X-ray generation from and spectroscopy of a thin liquid sheet

Jonathan Charles Tobit Barnard

Imperial College London

A thesis presented for the degree of

Doctor of Philosophy

Department of Physics

Imperial College London

Statement of originality

Unless otherwise noted, the work presented in this thesis is my own.

Jonathan Charles Tobit Barnard

Copyright declaration

Acknowledgements

There are many people without whom this thesis would not have occurred. First of all I would like to thank my supervisor, Professor Jon Marangos, for taking me on as a PhD student, providing me with advice and support throughout my time at Imperial and giving me the time I needed to write this all up. I would also like to thank our excellent technical staff; Susan Parker, Andy Gregory and Conor O'Donovan. I do not recall anything that I have asked them that they weren't able to do, and nothing in this thesis would have been possible without them.

I would also like to thank all of the students and postdocs in H012, of whom there have been too many to name. In particular though I would like to thank Mary Matthews and Esben Larsen for their unwavering support and encouragement and Oliver Alexander and Jacob Lee for keeping me sane during the last few months of attempted data collection mid-pandemic. I would also like to thank Judith Baylis and Marcia Salviato for their excellent advice, both on administrative matters and on life in general. Having a child halfway through my PhD would have been ten times harder if not for them.

Lastly I would like to thank my family. My parents, for being willing and able proofreaders, a constant source of moral support and for putting up with me for the past 27 years, and my wife Lizzy for everything she has done for me over my PhD. She has provided me with moral support, a sounding board for difficult decisions and kept me on the straight and narrow, all whilst being a good mother and finishing her degree. I don't know where I'd be without her, but it wouldn't be anywhere near as good.

I'd also like to mention my daughter Rosie. She didn't help — in fact she made it all much harder — but she has been the greatest blessing, and I look forwards to trying to explain what I've done here to her in the future.

List of papers that have come from this thesis

- Gediminas Galinis, Jergus Strucka, Jonathan C.T. Barnard, Avi Braun, Roland A. Smith, and Jon P.Marangos. Micrometer-thickness liquid sheet jets flowing in vacuum. *Rev. Sci. Instrum.*, 88(8), 2017
- Oliver Alexander, Jonathan C T Barnard, Esben W Larsen, Timur Avni, Sebastian Jarosch, Clement-Ferchaud, Andrew Gregory, Susan Parker, Gediminas Galinis, Alexandra Tofful, Douglas Garratt, Mary-Matthews, and Jon P Marangos. The mechanism of high harmonic generation in liquid alcohol. (Submitted to Nat. Physics).

Contents

List of Figures

List of Tables

Abstract

In this thesis a fan spray nozzle design is presented that is capable of producing single micrometer thick sheets of liquid, which are stable and flat. These sheets are characterised by various interferometric techniques, and some elementary fluid dynamics calculations are employed to understand the formation of the sheet. A new experimental vacuum chamber was designed and built, and the liquid sheet jet apparatus was installed into it and run successfully in the vacuum. Soft X-ray absorption measurements were attempted, but due to problems with a lack of X-ray flux and contamination of the experimental apparatus no results have yet been obtained. The sheet was also used as a source for high harmonic generation, generating harmonics up to 50 eV. The mechanism of high harmonic generation in the liquid is investigated, with the conclusion being reached that it is a similar mechanism to the semi-classical three step model. The effects of scattering are considered, and found to be less severe than anticipated due to correlation effects in the condensed phase.

Chapter 1

Introduction

1.1 The Holy Grail?

Recently, the official magazine of the Royal Society of Chemistry published a series about the 'Holy Grails' of Chemistry 25 years on [1]. These 'Holy Grails' are the subject of a series of papers published in 1995 by the Accounts of Chemical Research, each one on an area of Chemistry which was felt to be an important, exciting and active area of research. One of these papers was entitled 'Direct Observation of the Transition State', by John Polanyi and Ahmed Zewail [2]. The aim of this area of research is the direct observation and characterisation of the transition state - the in between configurations that are occupied during the course of a chemical reaction [3]. These transition states, which are comprised of unstable arrangements of atoms and electrons that quickly settle to more stable forms, often last for less than a picosecond. In order to observe the evolution of these states in time it is necessary to use a probe with a time resolution less than that of the lifetimes of the states themselves.

There are several reasons why it is desirable to probe the dynamics of these transition states, such as designing new ways to speed reactions up, finding ways to alter the outcome of reactions and even just simple curiosity at how chemistry works. Transition state theory was proposed in 1935, by Eyring [4] and seperately by Evans and Polanyi [5], and since then it has been foundational to our understanding of Chemistry. The theory (as shown in Figure 1.1) is that in order to progress from state A to state B, the molecules involved must progress along the reaction coordinate, which involves them passing through a local maximum of energy, known as the activation barrier, before they can relax down to the new energy minimum in state B.

As the molecules pass from state A to state B, they will deform into less energetically favourable configurations in order to accomodate each other, hence the increase in energy as the molecules move along the reaction coordinate. The transition state (which is denoted with the symbol ‡) is the highest energy configuration along this path, and according to transition state theory resembles the state closer to it in energy more than the state

further away. However, in order to obtain as much information about the reaction as possible it is desirable to observe all the dynamics of the system, rather than just the highest energy point.

Significant progress has been made on this front since the publication of the 'Holy Grail' paper. Using state of the art synchrotron sources complex electronic dynamics on a sub-picosecond time scale have been observed [6, 7], and using the technique of high harmonic generation (HHG) the whole journey through the transition state of photoactivated reactions have been observed [8, 9]. Attosecond X-ray pulses are now routinely available on a lab scale through HHG sources [10], and are able to be produced with significantly higher flux at large facilities via X-ray free electron lasers (XFELs) using novel operating modes [11]. Measurements with sufficient time resolution are clearly technically feasible, so we are now moving from proof-of-principle experiments to experiments with the goal of understanding the dynamics of important chemical systems.

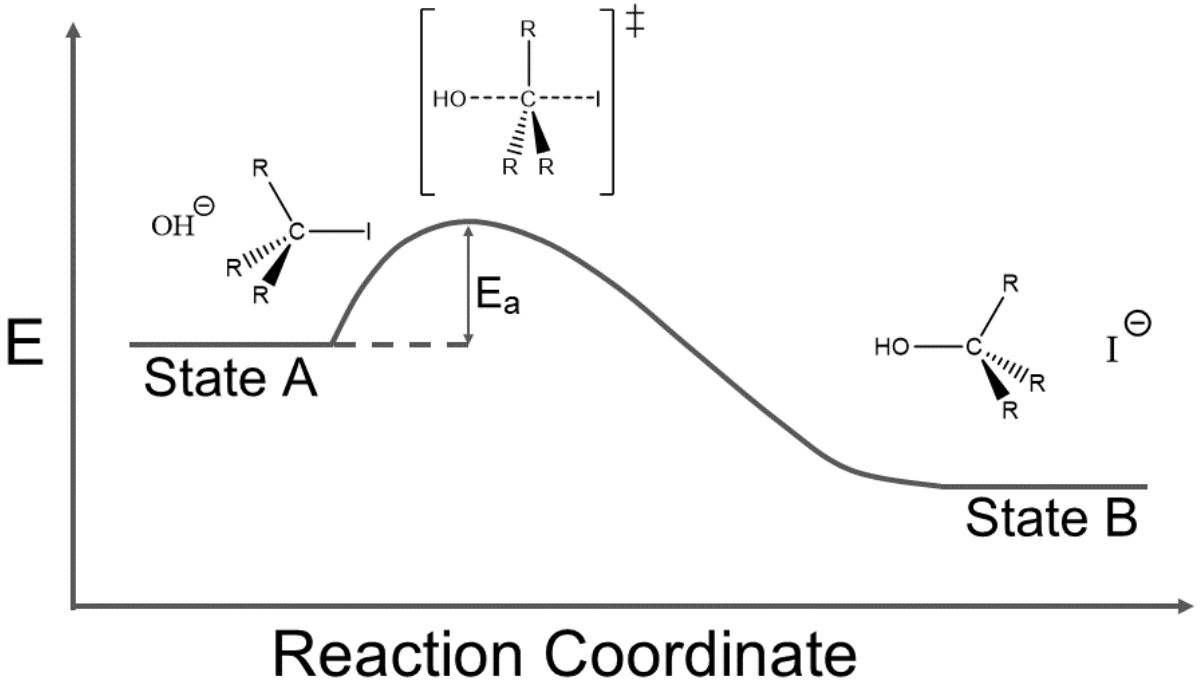

Figure 1.1: A simple diagram of the basic premise of transition state theory. As the system proceeds along the reaction coordinate (which in this case is the lengthening of the C-I bond and the shortening of the C-O bond) the overall energy of the system (E) increases to a maximum value, and then decreases smoothly to the new state. The transition state as defined by Eyring is the configuration of the system at the peak. As it is closer to state A than to state B the actual transition state in this case probably has a significantly shorter C-I bond than it does a C-O bond, but the whole journey over the potential energy surface (represented here by the grey line) is a smooth change in the length (and therefore strength) of the two bonds.

It is important when undertaking these measurements that we produce results that are useful for others down the line. For example, if we take the reaction shown in Figure 1.1, if the R groups are hydrogen atoms, then the activation energy (the energy required to cross the activation barrier, E_a) is small, as it is not too hard for the nucleophile (OH^-) to attack the carbon. However, the more hydrogen atoms are replaced by methyl

groups (CH_3) the harder it gets for the nucleophile to attack, and the higher the activation energy is for this pathway. In fact, in the case when all 3 R groups are methyl groups this reaction's barrier is so high in energy as to make it effectively impossible, and the reaction proceeds via an alternative reaction pathway. If we can observe the two reaction pathways clearly and accurately with our ultrafast measurement, then we can inform synthetic chemists or solar cell engineers about this, and they can proceed armed with new knowledge. Fairly complicated and important chemical processes have been elucidated by ultrafast measurement techniques [12], and the knowledge of transition states gained by these measurements along with the progress in theory and computational work have greatly improved the process of drug discovery [13].

Unfortunately however, the ultrafast measurements rarely take place in conditions that are similar to those of natural usage. Normally they take place in the gas phase, in a vacuum chamber, and are started with an intense burst of coherent radiation. Some of these constraints are necessary due to the techniques involved, but this thesis intends to show progress towards improving one of these — namely doing ultrafast X-ray absorption experiments in solution phase.

1.2 Overview of this thesis

This Chapter presents the background to the research contained in this thesis, and our motivation for performing these experiments. Chapter 2 discusses the fluid mechanics of generating a thin liquid sheet jet in more detail, as well as comparing the different methods of creating them. Chapter 3 lays out the relevant optical theory needed for the experiments in this thesis, and also discuss HHG and X-ray spectroscopy. Chapter 4 contains details of the experimental setups used for the experiments laid out in Chapters 5 and 6, as well as how they changed over time. Chapter 5 contains details of progress towards a static X-ray absorption measurement performed on solution phase targets and Chapter 6 contains details of the HHG from a liquid experiment that we performed over the course of my PhD. Finally Chapter 7 concludes, and offers suggestions for further work to pursue in this area.

1.3 Femtochemistry

The field of research we are aiming to contribute to is called Femtochemistry, as it is the study of chemistry on timescales on the order of $1–1000\times10^{-15}$ s (femtoseconds). This temporal resolution has become possible by the advent of ultrashort laser pulses which can now be routinely produced with a duration on the order of 10–15 fs or below [14]. Over the course of the past 25 years much progress has been made towards tracking and even controlling reactions in real time. It is now possible to capture transient states with short pulse X-ray scattering [15], and to observe the relative population and energies of species as reactions progress using various transition state spectroscopy techniques [16, 17]. It could even prove possible to resolve the pathways

of individual molecules over the course of a reaction, rather than viewing them as an ensemble [18]. Novel sample delivery techniques have also been invented, such as the 'molecular archery' technique which allows the experimenters to control the collision angle of the reaction, enabling more experimental control over the parts of the potential energy surface that are explored [19].

To perform our time-resolved experiments we use a nonlinear technique called high-order harmonic generation, which is capable of generating very broad bandwith single pulses of light in the XUV/soft X-ray spectral range, with a duration on the order of hundreds of attoseconds (10^{-18} s) [20]. These pulses can be created in the lab using a femtosecond laser system and a gas needle in a relatively simple arrangement, which allows us to perform femtosecond resolution experiments without needing the use of a big facility. We employ these pulses as a probe, to observe via X-ray absorption the state of the system at various points in time as it is evolving in response to a stumulus, in an experimental scheme called a *pump-probe* experiment (see Figure 1.2).

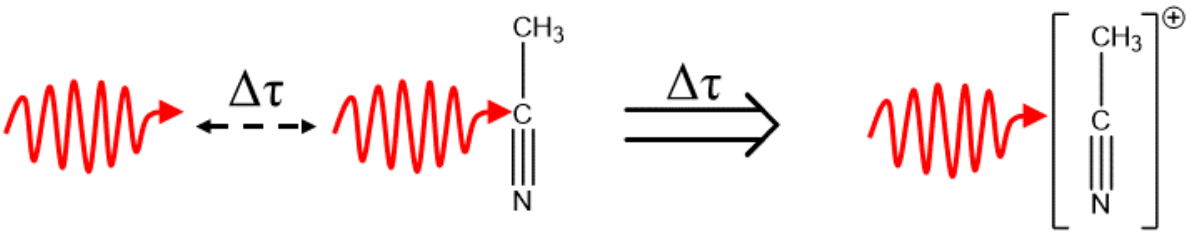

Figure 1.2: Illustration of an pump-probe experiment. The two pulses are separated by a time delay $\Delta\tau$. The first pulse reaches the target and starts the reaction of interest, in this case by ionising the acetonitrile. After $\Delta\tau$ seconds the second pulse reaches the target and performs a measurement, in this case by ionising again and causing a Coulomb explosion creating ion fragments which can be measured. By varying $\Delta\tau$ we can vary the amount of time the system has to evolve, and so we can build up a picture of the electronic and nuclear motion during this time.

A pump-probe experiment consists of two short pulses separated by a variable delay. It works on a somewhat similar principle to quenching experiments in reaction kinetics, where a reaction is stopped after a certain amount of time and the concentrations of reactants and products are measured. The first (*pump*) pulse induces the dynamics in the system that is to be studied, in effect acting as the start button on a stopwatch. The second (*probe*) pulse then arrives after a delay and produces an effect that can be measured, such as a change in absorption, current or voltage. As the system evolves under the dynamics induced by the pump during the delay, this probe pulse can take a measurement at any stage of the dynamics. By varying the delay, the progress of the dynamics of the system can be seen and conclusions drawn about the pathways of the reaction.

Figure 1.3 shows the results of a recent pump-probe measurement in our lab. The pump pulse delay is varied from −30 fs (where the probe arrives before the pump — this is done as a sanity check to see if there are any systematic variations between no pump and pump arriving after probe) to 100 fs before the probe. The resulting

spectra are divided by the unpumped spectrum (shown at the top of the figure), and the lineout is then plotted in the corresponding row in the figure. As can be seen, the measurement is sufficiently time sensitive to be able to capture a transient that exists for around 20 fs.

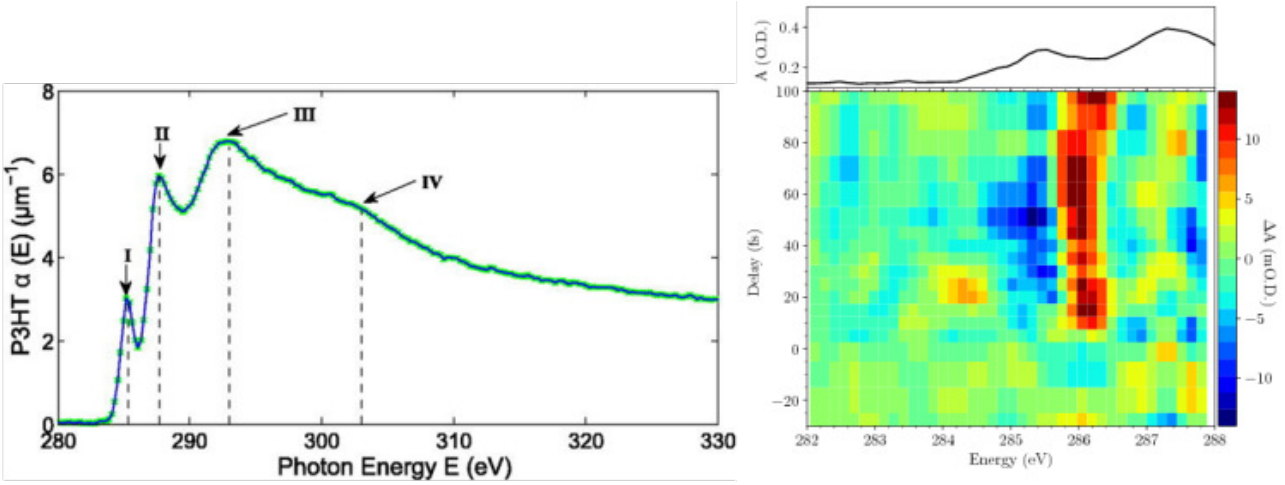

Figure 1.3: Two examples of X-ray absorption measuements in poly-3hexylthiophene (P3HT) performed in our lab. The first one is a static absorption spectrum from Johnson *et al.* [14]. The numbers indicate peaks that were identified, which will be explained further in Chapter 3. The second is a pump-probe measurement taken from the thesis of Douglas Garratt [21]. The figure is created by plotting the difference between the spectrum with the pump pulse present (at a particular delay) and not present. Red means an increase in signal, and blue means a decrease. The main conclusions from this figure is that there is a large increase in signal at positive delays around 286 eV, with a corresponding loss of signal next to it, and a small transient peak that appears between 15 and 30 fs at around 284 eV. Analysis is still ongoing as to the probable origin of these features.

1.4 X-ray absorption spectroscopy

A material's X-ray absorption spectrum (such as those shown in Figure 1.3) typically consists of a general decline in absorption with respect to energy, punctuated by steep rises at certain energies. These steep rises are known as edges, and they occur when the energy is at a resonance to promote an electron to the continuum (ionisation). Edges occur at characteristic energies for any given atom, and so by observing at which energies these edges occur one can tell which atoms exist within the sample. An advantage of using X-rays, as opposed to conventional optical absorption spectroscopy, is that the resonances are from core level states — orbitals with a high binding energy, which have their electron density concentrated near the nucleus and are therefore not involved in bonding — as opposed to higher lying valence states that are more diffuse orbitals, where the electron density is spread out and therefore is involved in bonding. X-ray absorption spectroscopy thus allows an atom specific probe of the dynamics in the reaction of interest [22] (providing there is only one environment in the molecule — if there are two different environments for a given element then the structure will be more complicated).

On top of these edges there will be additional features due to transitions from the core states to unoccupied valence states. These features (known as X-ray Absorption Fine Structure — XAFS) split into two regions of interest. The far edge region (Extended X-ray Absorption Fine Structure — EXAFS) from about 30 eV above the absorption edge is structured due to the scattering of high kinetic energy photoelectron waves off of neighbouring atoms, and therefore contains information about the short range environment of the atom [23, 24]. The near edge region (X-ray Absorption Near Edge Structure — XANES) is more complicated, but also contains structural information, as well as information about unoccupied valence states of the molecule [23, 25].

By observing the shifts in energy and amplitude of individual structure peaks as a function of delay (in effect tracking the appearance and disappearance of resonances) it is possible to observe the presence or absence of certain molecular orbitals, which indicate the presence or absence of chemical bonds [8, 22]. It is possible in this way to observe bonds breaking or forming in real time.

In addition to the location specificity of XANES, there are many other reasons why soft X-ray spectroscopy is good for this kind of experiment. Firstly it is more possible to get the bandwidth in the X-ray region to support pulses of less than 1 fs than it is in longer wavelength regions of the spectrum because the relationship between energy, wavelength and frequency is:

$$E = \hbar\omega = \frac{hc}{\lambda}, \tag{1.1}$$

Where E is energy, ω is the frequency, h is the planck constant, $\hbar$ is the planck constant divided by 2π, c is the speed of light in a vacuum and λ is the wavelength. The relationship between the spread of energies (or frequencies) and the temporal width is:

$$\Delta E \Delta t \geq \frac{\hbar}{2}, \tag{1.2}$$

i.e. the greater the bandwidth (spread of energies) you have then the shorter the pulse can be in time (t). The differential of Equation 1.1 with respect to λ is:

$$\frac{\mathrm{d}E}{\mathrm{d}\lambda} = -\frac{hc}{\lambda^2}. \tag{1.3}$$

Therefore as wavelength decreases the same change in energy is associated with a much smaller change in wavelength. So the same absolute range of wavelengths (*i.e.* 20 nm) would lead to a much greater energy bandwidth if the range is 1 to 21 nm, rather than 780 to 800 nm, which means that a shorter temporal pulse is allowed by the uncertainty principle (Equation 1.2).

Secondly, shorter wavelength pulses have a lower intrinsic pulse length limit due to the higher carrier frequency (the frequency of the central wavelength of the pulse). At 800 nm a single optical cycle pulse (the shortest pulse that can be made at that carrier frequency) has a duration of 2.7 fs. As the wavelength decreases

the carrier frequency increases, leading to shorter optical cycles and the potential of shorter pulses [26].

Thirdly, isolated attosecond pulses are available in the lab via the process of high harmonic generation (HHG) [20]. This has been developed into a reliable way of producing attosecond-duration XUV/Soft X-ray pulses on a tabletop [20], allowing time-resolved experiments to be conducted without needing expensive facilities [8, 9, 27].

1.5 High harmonic generation

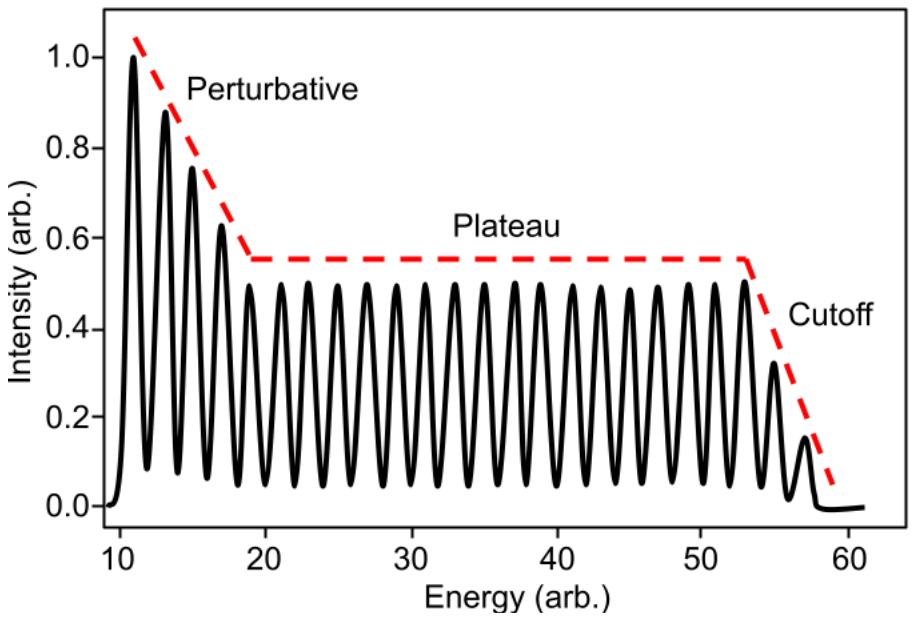

Figure 1.4: An example harmonic spectrum, based on [21]. The perturbative region is where the harmonics are still generated due to standard nonlinear processes (see Chapter 3). In the plateau region the field becomes too strong for this to hold, and the harmonics are generated by ionisation, acceleration and recollision as described in Chapter 3. The cutoff region occurs as the maximum possible harmonic energy is reached. No harmonics can be produced with an energy greater than this by the laser field, as it cannot impart any more energy to the recolliding electron, so the spectrum abruptly cuts off.

HHG was discovered in the early 1990s [28] by researchers attempting to use strong laser fields to create progressively higher and higher harmonics. Shortly after this discovery a theoretical semiclassical description of the mechanism was published independently by Corkum [29] and by Kulander, Schafer and Krause [30], and then developed by Lewenstien *et al.* [31]. The applicability of the method to producing short pulses was seen immediately by Corkum [32], who noted the short wavelength and the wide bandwidth could combine into a pulse as short as 10 attoseconds if the correct phase relations were achieved. The current record pulse length that has been measured is 43 as [33], and although most HHG pulses will not be as short as this they will still be on the order of hundreds of attoseconds — short enough to observe most chemical phenomena.

During the process of HHG, electrons are ripped from their parent ions by an intense laser field, accelerated in this laser field and then recombined with their parent ion, releasing their accumulated energy as a photon.

The spectrum produced by this process has a distinctive shape — the first few harmonics are much more intense, and decrease in intensity as the harmonic order increases. Subsequently a plateau region is reached, where the harmonics are all the same intensity, which stretches over a large energy range and then abruptly cuts off after a maximum energy (see Figure 1.4).

This technique has been developed extensively since its discovery. As stated above, it is a common probe source for lab based time-resolved experiments, even being expanded into facility level systems such as the new Extreme Light Infrastructure project [34], where their ELI-ALPS facility in Szeged is set up to generate attosecond pulses, as well as the ARTEMIS beamline in the Central Light Facility in the UK. In addition to gaseous systems, HHG has also been extensively studied in solid systems [35], which have their own interesting properties. The process of HHG itself has also been used as a probe due to its theoretically fairly simple nature, in many different techniques such as HHG spectroscopy [36] and RABITT (reconstruction of attosecond beating by interference of two-photon transitions) [37, 38].

HHG is a powerful technique, and it can produce pulses that are shorter than any other source in the world [10, 39]. However there are two major drawbacks — the efficiency of the process and the maximum photon energy that can be reached. HHG is a strong field effect, which means that in order to produce photons via this technique a very intense electric field is needed (on the order of $10^{13}\,\mathrm{Wcm}^{-1}$). However the flux produced by these high fields is about 10^8 times lower than that of the driving field. This can be countered by increasing the intensity of the driving field. However, there is a limit to this as well, as increasing the intensity of the driving field also increases the ionisation fraction of the target, and if this fraction gets too high then this will interfere with the process, decreasing the efficiency dramatically.

1.6 Soft X-ray absorption in condensed matter

The other problem with using HHG as a source of X-rays for experiments is absorption. HHG creates pulses in the XUV/Soft X-ray region — the highest energy photons we can make by HHG are around 600 eV [40]. Unfortunately, especially in the lower end of this energy range, these photons are absorbed by the majority of materials (see Figure 1.5). The reason for this is that these are the energies at which electronic transitions within atoms start to be accessible, mainly ionisation of the valence electrons [41]. This means that experiments have to be carried out in a vacuum chamber, as the air itself will tend to absorb the photons needed for measurements. This complicates experimental design, as everything we use needs to be able to survive operating in a vacuum.

This also applies to the design of targets for these experiments. Not only do the targets have to be able to be housed in a vacuum chamber, which imposes constraints by itself, but there is also a narrow range of thickness that the target can have (see Figure 1.5). Inspecting the transmission curves for two common solvents in the energy range that we would hope to be accessing with our HHG source, we can see that the target thickness

needs to be closer to 1 µm than 10 µm if we want to be confident of observing a change in the level of absorption, as the thicker the target is the fewer photons are transmitted, effectively saturating our absorption signal with background.

However there is a balancing act that needs to be performed here, as the thinner the target, the fewer opportunities there are for the relevant atoms to absorb a photon. Thus the target must be thick enough to ensure enough signal, but not too thick that the background overwhelms it. Previous solid phase experiments in our lab have shown that the signal to noise ratio is maximised for targets that are somewhere in the middle of the range of thicknesses we could manufacture [21], although the exact location of the peak on the signal to noise ratio curve depends on the specific system that is being studied.

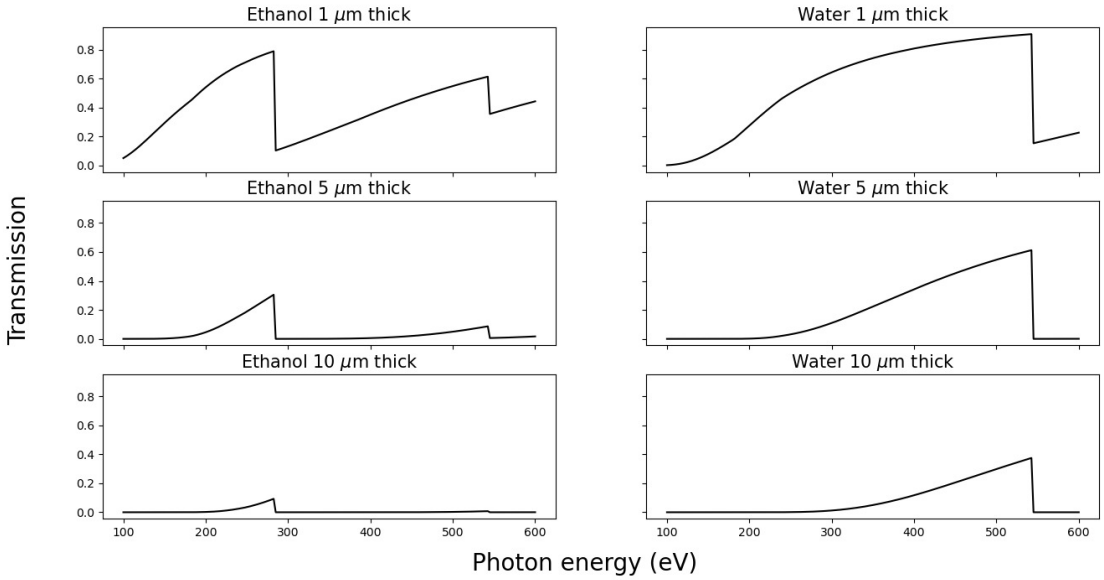

Figure 1.5: The transmission of soft X-rays through solvents of different thicknesses. Data from [42]. A 1 µm thick sheet of either solvent has a reasonable transmission rate, but as the thickness of both of them reach 5 or 10 µm the transmission drops to unuseable levels. This is the reason why we needed to create such a thin, flat sheet to use as a target.

1.7 Liquids: The third phase of matter

In the study of the physics of molecular systems on an ultrafast time scale, most efforts have been focused on the gas phase [8, 9] or the solid phase [21, 38]. This is due in part to the relative simplicity of these systems, and in part to experimental considerations. As we transition from gas phase (isolated molecules with no interactions between them) through crystalline and amorphous solids to liquids the size of the theoretical calculations necessary to capture the dynamics fully increases rapidly. The solid phase can have recurring units, which lead to the emergence of both intrinsic symmetry and band structure. These both simplify calculations by allowing one to perform a smaller calculation once, and then apply the results more widely. Even amorphous solids

can still have partial band structures, which can help somewhat to reduce the extent of required calculations. In the liquid phase none of these things are present, and often effects from molecules beyond the first set of nearest neighbours are important to the progress of the reaction [43] leading to a much greater cost of accurate calculation.

However chemical reactions and biological processes occur almost exclusively in the liquid phase. The complexity that makes liquid and solution phase reactions hard to model computationally is what leads to the rich variety of chemical and biological processes that occur in the world. Thus it is vital to fill this gap in ultrafast research so we can get an accurate picture of key processes that are relevant to other disciplines.

On the experimental side, time-resolved studies in the liquid phase using femtosecond laser pulses are well established [3, 44], as these can be performed at atmospheric pressure. However, there are other time-resolved techniques that have been limited in their application due to the difficulty of maintaining a liquid target in vacuum. Recently some time-resolved photoelectron spectroscopy results have been reported [45, 46], a method that gives direct access to the electronic structure of the reactants during a reaction by ionising them and then detecting the resulting electrons in a time of flight spectrometer [47]. This can then be used to track the shifts in the energies of the electrons directly over the course of a reaction using a pump-probe scheme.

Working in the X-ray range is advantageous, as at these energies we obtain direct access to the 'core' orbitals of the atoms — those not involved in bonding, which enables the making of element-specific measurements. Our measurement of choice is absorption spectroscopy rather than photoelectron spectroscopy because as the signal is the absorption of photons rather than the emission of electrons in principle more information can be obtained — not only the energy of the bound electron, but also the energies of higher lying unfilled bound and quasi-bound states. The trade off is that the XAS spectra can as a result become very difficult to interpret, sometimes requiring extensive modelling to confirm results (see Chapter 3 for a fuller discussion). In addition absorption spectroscopy is tolerant of higher pressure in the vacuum chamber than photoelectron spectroscopy and is not affected by space charge effects in the jet [47], which are both advantageous for the study of liquids.

Around the turn of the millennium, X-ray sources that were bright enough to perform scattering and absorption measurements in dense media with an ultrashort time resolution (here defined as sub nanosecond) were becoming available — due to advances in synchrotron technology [48] and some table top sources [49]. These sources offered a time resolution of the order of hundreds of picoseconds, and also did not require the use of targets in vacuum (as they produced hard X-rays — 1 keV or more). Absorption experiments have been carried out at synchrotrons in the last decade [6, 7], but the time resolution of synchrotrons is still limited to hundreds of femtoseconds.

In order to acheive the very high time resolutions needed to track both the electronic and the nuclear dynamics, we ideally want X-ray pulses which are sub-femtosecond in length (for context the atomic unit of time is equivalent to 24.2 as, and the shortest dynamics we might want to measure would occur on a scale of

tens of atomic units of time). Whilst facilities are starting to approach this, for instance the XLEAP program at the LCLS [11], the simplest way to achieve probe pulses of this length is through HHG [20]. The only drawback to this source is its inherent inefficiency, leading to flux that is many orders of magnitude less than that achievable at synchrotrons or XFELs. So the question becomes can we develop a target system that will allow us to perform these experiments in a liquid?

1.7.1 Experimental challenges

Liquids are difficult to handle experimentally in high vacuum. The extreme conditions affect liquids more than solids or gases, as at high vacuum we are usually below the triple point on the phase diagram — *i.e.* liquid does not exist in equilibrium under those conditions, so the target is inherently unstable (see Chapter 5, especially Figure 5.6). Evaporation also has the effect of unleashing large amounts of gas into the vacuum chamber, which raises the ultimate pressure achievable, whilst freezing may inflict disastrous results on delicate experimental apparatus.

Beyond issues with vacuum, we also have the problem that often the targets of interest are only sparingly soluble in the solvents available. This can then create an issue with the signal to noise ratio, as any solvent signal present will manifest as a huge background noise signal that can mask anything of interest in the solute signal.

This being said, there are also a number of advantages to using a liquid phase target in ultrafast measurements. The liquid is self-healing, especially if a jet is used (although a static liquid target is still self healing to an extent — the liquid can move around inside it, so as long as no permanent chemical change is effected there is still no need to raster the sample). This means that we can use more power in our laser pulses, without the worry of destroying the sample. In addition there is no possibility of irregularities in the structure or areas that are impure, as the liquid is randomly (and therefore presumably uniformly) arranged.

In order to overcome the challenges above, specialist equipment needs to be constructed to house a rapidly evaporating liquid sheet in vacuum. It is necessary to keep the pressure sufficiently low whilst at the same time protecting the integrity of the liquid target.

1.8 A target suitable for liquid phase ultrafast spectroscopy

The target needs to allow the soft X-ray spectroscopy measurements to be taken without too much absorption, whilst also remaining stable in the vacuum for the duration required for the measurement. There are two types of liquid target that are commonly used in this situation — a nanofluidic cell, where the target liquid is trapped between layers of transparent material and so protected from the vacuum, and a jet, where the liquid flows into and back out of the vacuum quickly enough that the evaporation doesn't compromise the vacuum too much.

Nanofluidic cells have many advantages to them, chiefly the fact that they can be made very thin very easily — all it takes is to fabricate the cell correctly (see Figure 1.6). For example, this approach was used for a time-resolved soft X-ray experiment by Kim *et al.* [50] in 2016 — noting the excessive solvent absorption in their free flowing jet in this energy range, they used a nanofluidic cell that forced the liquid target to be between 10 µm and 200 nm thick [51] successfully.

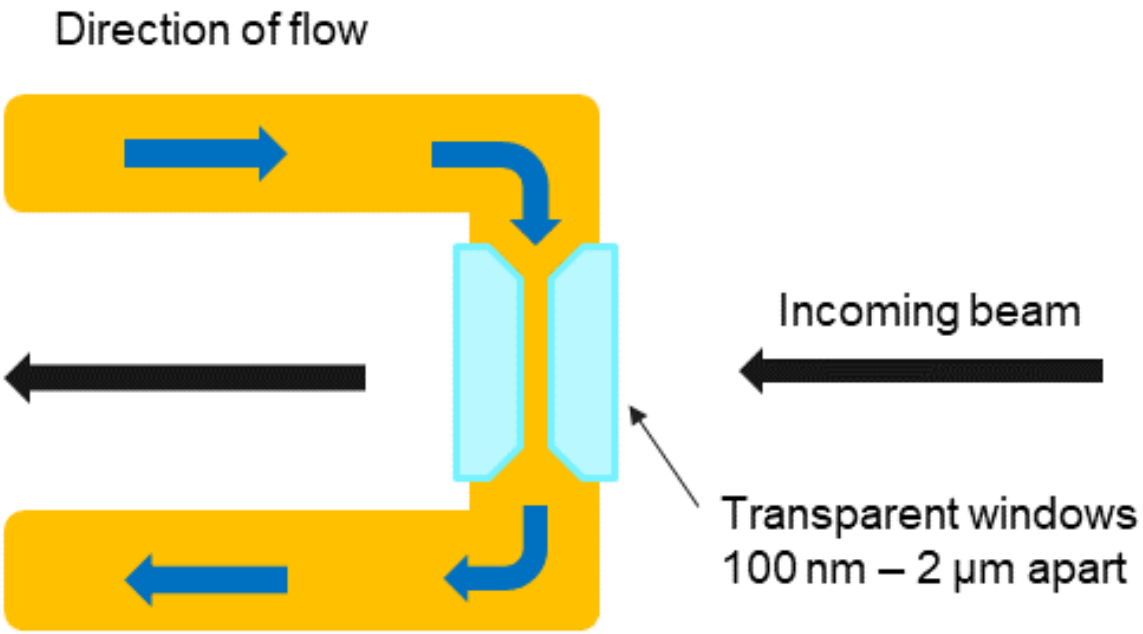

Figure 1.6: A schematic of a nanofluidic cell, based on drawings in [51]. The cell itself is created by sealing two very thin windows (typically of silicon nitride) together with a gap of 100 nm to 2 µm between them. Liquid is then flowed into the cell at one end, is forced through the gap between the windows and out the other side, creating a smooth, very thin flowing sheet of liquid.

However, there are drawbacks to this system. Firstly, the construction of such nanofluidic cells is far from trivial. Secondly, no matter how transparent the windows are, they are still windows and so will have some effect on the pulses used in the measurement. In particular, using an HHG source we are working with an X-ray flux many orders of magnitude less bright than the flux of a synchrotron, so the absorption of the window is best avoided.

The other target option is a free flowing jet, which is the most common way of creating a liquid vacuum target. This solution is windowless, but it has the obvious downside of therefore being vulnerable to evaporation into the chamber. In addition most liquid jets tend to be thicker than the 200 nm–10 µm quoted for the nanofluidic flow cell.

The most common liquid jet type to be placed in vacuum is a cylindrical jet. This is mostly based on a design reported in 1988 by Manfred Faubel *et al.* [57], where the liquid propagates through a hole, and is then caught in a liquid nitrogen trap. This design has been improved by a sophisticated nozzle and catcher design

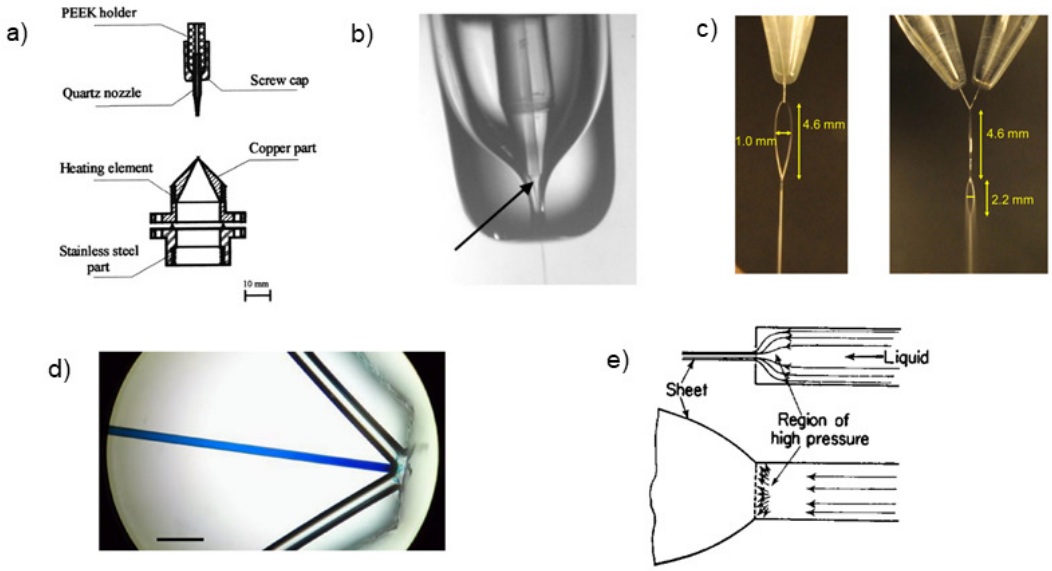

Figure 1.7: Different types of liquid jet delivery geometries. a) is a sketch of a cylindrical nozzle from Charvat *et al.* [52], showing a convergent nozzle made of quartz, which ensures long term stability, and also the catcher into which the jet is flowed. This catcher is heated to avoid the liquid freezing onto it. b) is a picture of a gas dynamic virtual nozzle from DePonte *et al.* [53]. Like a) the liquid comes through a convergent nozzle (indicated by the arrow on the photograph), but surrounding the nozzle is a gas channel, which creates a gas sheath, narrowing the jet and preventing it from freezing. c) are colliding jets producing a sheet from Ekimova *et al.* [54], d) is the tip of the gas dynamic virtual nozzle from Koralek *et al.* [55]. The central, blue dyed channel is the liquid channel, and the two side channels are the gas channels. The gas then provides the momentum to push the liquid into a sheet, but it is also possible to flow liquid through these outer channels to create a more conventional colliding jets type sheet. e) is a diagram of the fan spray nozzle geometry from Dombrowski *et al.* [56].

[52], the ability to recirculate the sample [47] and using gas to focus the stream to a smaller diameter and protect it from the vacuum [53, 58], which avoids the need for small diameter nozzles that can block easily (see Figure 1.7).

However, even the thinnest of these jets comes in at around 5–10 µm, which is at the very highest edge of the thickness that can be tolerated in our experiments (see Figure 1.5). Therefore, we turned to another common liquid jet type, which is the flat jet. The flat liquid jet has been widely used in dye lasers, which need a flat gain medium in their optical cavity to ensure optimum performance [59]. It has also been widely studied in the field of fluid dynamics, as the sheets produced by this technique have interesting breakup properties [60, 61].

The operating principle of a flat liquid jet is one of momentum transfer. In the most simple example, two jets of liquid collide head-on (as in *e.g.* Clanet and Villermaux 2002 [62]). This creates a circular sheet of liquid which thins as it spreads out as $\frac{1}{r}$, as the vertical momentum of the two jets is cancelled out and transformed into horizontal momentum (see Figure 1.8).

The two jets can be set at an angle to each other that is not 180°. This gives both jets a momentum component that adds up instead of being cancelled out. This causes the resulting sheet to be a segment of the circle that would oterwise have been produced if they were head-on. If the velocity is high enough, the resulting sheet is unstable, and spreads out in a fan-shape [63]. However if the velocity is lower the surface tension can compete with the momentum forcing the liquid to spread out, and the sheet will be drawn together into a leaf shape [64, 65]. Depending on the characteristics of the colliding flows this sheet can be very thin and flat — our sheet is around 1–2 µm thick with an optical flatness of $\frac{\lambda}{20}$ at 640 nm, which should be more than enough for our experiments.

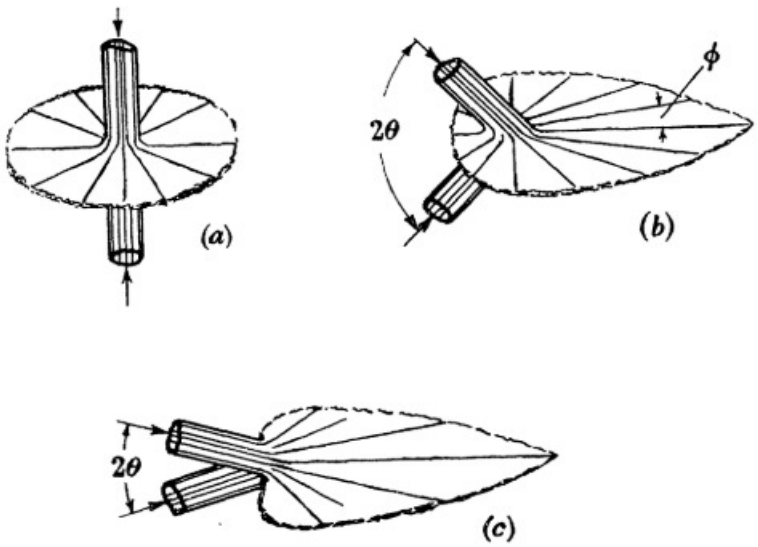

Figure 1.8: Drawings of the effect of different angles (2θ) on the resulting sheet of liquid produced. From Taylor [60]. In (a) the two jets are at 180° to each other, so the sheet spreads out as a circle around the axis of collision. As 2θ decreases through (b) and (c) the circle distorts as the transverse momentum components of the jets add together in one direction, ending up with a leaf shape. Although this is shown for colliding jets, the same process of momentum transfer holds for all different nozzle geometries.

1.8.1 Colliding jets

By far the most studied method of producing a thin flat sheet jet is with two colliding cylindrical jets. Two cylindrical jets collide at an angle to each other, supplying both the liquid for the sheet and the momentum to spread the sheet out. This setup is the simplest one to make, so has been the subject of numerous experimental studies [64, 65], theoretical studies [66–68] and was the subject of the classic series of papers on the subject by Taylor [69–71]. It was also the first geometry to be reported as having made sub 1 µm thick jets [54]. In this geometry the sheet is free to expand as it wants, leading to a much more rounded shape. It is simple to construct but tricky to align, which is why we explored other avenues of producing a thin liquid sheet.

1.8.2 Gas dynamic virtual nozzle

A gas dynamic virtual nozzle (GDVN) is a type of nozzle orginally designed for keeping cylindrical jets straight and avoiding them breaking up into spray [58]. It works by blowing gas through channels in the nozzle either side of the cylindrical jet, which then produces a vortex around the stream, keeping it from breaking up [58]. In 2018, Koralek *et al.* produced a type of nozzle that applies this idea to creating a thin liquid sheet [55]. The gas comes out of channels set at 45° to a liquid stream, and impart their momentum onto the stream, blowing it into the classic leaf shape. This enables the flow to be semi-independent of the properties of the sheet, as the momentum in provided by the gas. In addition, due to the presence of gas around the sheet the local pressure for the liquid should be higher, which should help the integrity of the sheet in the vacuum. It is also possible to flow liquid through the gas channels to make a geometry similar to colliding jets, above. This makes this type of nozzle more flexible.

1.8.3 Fan spray and similar nozzles

The third type of nozzle geometry considered here is a fan spray nozzle. This type of nozzle generates momentum transfer in a plane by having a slope at the desired angle along the direction of the flow of liquid, to cause the change in momentum [56]. The leaf that comes out of the nozzle is much thinner and longer than the colliding jets leaf, because the liquid is not free to spread out.

This is the type of nozzle used by dye lasers, due to ease of fabrication and use. Being a nozzle it is easy to use, as there is nothing to adjust and no alignment needed. Dye laser nozzles were typically made by crushing pipes into the desired shape [72], or using razor blades to create the nozzle [59]. We chose to use this type of nozzle due to these properties, however in order to create a sufficiently thin jet we needed to have greater control over the dimensions of the nozzle, so we fabricated them using a high resolution 3D printer, which gave us great control over the exact geometry of the nozzle, as well as enabling them to be easily reproduced [73].

1.9 HHG from a liquid

The other possibility that arises from a liquid sheet target in vacuum, especially in a lab dedicated to using HHG to produce X-ray pulses, is to use the liquid itself as a source of HHG. Novel HHG sources have risen in popularity over the past decade, with various solids offering different properties in the HHG that they can produce [35]. A liquid target offers at least the possibility of a self-healing condensed phase HHG source, as well as possibly having unique properties in its harmonic spectrum.

There has been a fair amount of scepticism over the possibility of creating harmonics from liquids in the past. The well established mechanism of HHG in gases involves the ionised electron travelling long distances in the vacuum, propelled by the strong oscillating electric field. In the condensed phase it therefore seems reasonable

to doubt that HHG can occur, as the density of atoms is so much higher than in the gas phase so it was thought unlikely that the electron could travel far without scattering, and therefore only perturbative harmonics could be generated from liquids [74, 75]. Following this logic, Flettner *et al.* and Kurz *et al.* performed pump-probe HHG experiments on water vapour, by sending a pulse to explode a water droplet, before generating harmonics from the dense gas that resulted [76, 77]. In addition, incoherent plasma radiation has been created from liquid targets by Kim *et al.* and Heissler *et al.* [78, 79], but this is of limited interest to us as it arises from a non-HHG mechanism.

The discovery of HHG from solid targets [80] suggests that HHG from a condensed phase target is not completely out of the question. Solid HHG relies on the existence of conduction and valence bands, which are not present to nearly the same degree in a liquid target due to the lack of long range order. The first experiment of this kind was performed by DiChiara *et al.* [75] in 2009, in which they only saw perturbative harmonics. We have performed a similar experiment, roughly concurrently with Luu *et al.* [81, 82], and both of us have seen high harmonics from a liquid, which leads to some interesting questions about the exact mechanism by which this can arise.

1.10 Conclusion

Having established the desire for and potential applications of a thin liquid sheet jet target, we now turn to Chapter 2 where the development of this target will be explained. As the majority of the work in this thesis was performed in conjunction with other members of our research team in the Imperial College extreme light consortium I have used the word 'we' to indicate work done by myself and others, and have used 'I' only when work was unambiguously done myself.

The following people have made significant contributions to the work presented in this thesis: Initial nozzle design was by Gediminas Galinis. Characterisation of the liquid sheet was performed by Jergus Strucka, Gediminas Galinis and myself. Manufature of the nozzles was performed by Gediminas Galinis,Oliver Alexander and myself. The liquid jet experimental apparatus was constructed by Sebastian Jarosch, Gediminas Galinis, Oliver Alexander, Jacob Lee, Clément Ferchaud and myself, with excellent technical assistance by Andy Gregory, Susan Parker and Conor O'Donovan. The liquid HHG measurements were taken by Sebastian Jarosch, Gediminas Galinis, Alexandra Tofful, Timur Avni, Mary Matthews, Esben Larsen, Oliver Alexander and myself. Data analysis was performed by Sebastian Jarosch, Timur Avni, Mary Matthews, Esben Larsen, Oliver Alexander, Clément Ferchaud and myself. The X-ray absorption measurements were taken by Oliver Alexander, Jacob Lee, Douglas Garratt and myself, with data analysis by myself.

Chapter 2

Creating thin sheets of liquid for use in XUV and soft X-ray experiments

2.1 Introduction

In order to perform liquid phase experiments with XUV and soft X-ray radiation it is necessary that the radiation can interact with the target with a manageable degree of absorption. Due to the high absorption cross sections most materials have for light at these energies, this means the target needs to be very thin. Figure 2.1 shows the normalised transmission of light of various X-ray energies in a slab of water that is just 1 µm thick. Here the unit of choice is eV, as is customary when dealing with X-rays as it makes the numbers more compact than either using Joules or converting to frequency or wavelength in this range of the electromagnetic spectrum. One eV is equivalent to 1 V multiplied by e, the charge on the electron (1.602×10^{-19} C). As can be seen, even within the 'water window' (a region where none of the constituent atoms of water have an absorption resonance) 1 µm of water is not very transparent to soft X-rays. Adding the effect of windows (for instance as in a microfluidic flow cell [83]) to this absorption will lower the transmission further, not to mention adding extra resonances to our X-ray absorption spectrum that might overlap with resonances we want to study, further complicating our measurements. Given that we do want to use XUV and soft X-ray wavelengths in our experiments the best solution to this problem seems to be a thin free flowing jet of liquid — this is because as jets are directed streams of liquid they do not need a container to hold them, so windows are not necessary. Jets can also be made very thin — cylindrical jets of width of 10s of µm are commonly available [47, 52, 79]. Before the work presented in this thesis Ekimova *et al.* developed a flat sheet jet based on two colliding cylindrical jets capable of producing a sheet with thickness of between 1.4 and 3 µm [54], and subsequently a modification on this geometry by Koralek *et al.* has produced a sheet which claims a thickness of about 0.5 µm [55]. At the time of writing this design is available to purchase from Micronit GmbH, and has been used for some of the measurements presented later

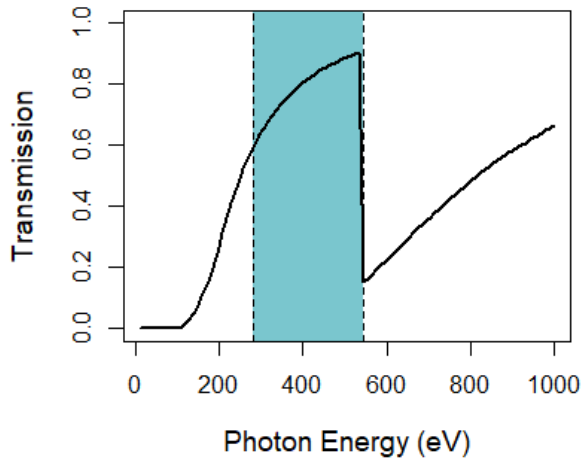

Figure 2.1: Incident photon energy plotted against transmission through 1 µm water, with the blue region indicating the water window (the region in which water is relatively transparent to soft X-ray photons). The large drop off at 543 eV is due to the oxygen K-edge. In the water window the transmission is about 0.87 — enough to avoid much solvent interference, but still not enough to have an arbitrarily thick water target. Data from B. L. Henke *et al.* [42]

on in this thesis.

Given the low transmission through only 1 µm of liquid, a flat sheet jet seems like a more appropriate choice than a cylindrical jet. Such a flat sheet jet is fundamentally created by momentum transfer into the plane of the sheet [55] (illustrated in Figure 2.2). The y-component of the momentum of the liquid coming from both sides is transferred during the central collision into the xz-plane, leading to a spreading out of a thin sheet of liquid (see section 2.2.1).

2.2 Physics of thin liquid sheets

Thin liquid sheets have been studied for many reasons — both as an investigation into the flow properties of liquids in general [84] and also with more practical applications in mind such as fuel injection [85], pesticide spraying [86] and dye lasers [59]. As such, a reasonably comprehensive body of work on this subject has accumulated, even if not all reports seem to be aware of the other work done (e.g. Watanabe *et al.* [59] reproduce the findings and methods of Dombrowski *et al.* [87] fairly exactly, without any seeming knowledge of the existence of this former paper).

A useful heuristic for understanding the behaviour of liquids (in any setting, including that of thin liquid sheets) is to consider a series of dimensionless numbers. Many of these numbers exist, dealing with relations between all kinds of quantities relevant to fluid dynamics. However, the two that are the most relevant to the behaviour of liquid sheets are the Reynolds and Weber numbers. It should be borne in mind in what follows

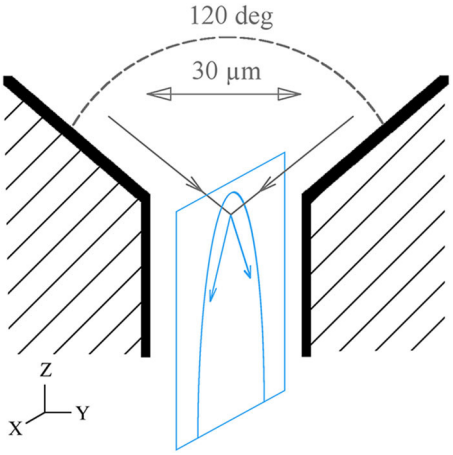

Figure 2.2: Schematic illustration of the process of momentum transfer in our sheet jet nozzle. Liquid flows down the slopes angled at 60° to the z-axis, and collides in the centre of the nozzle. This collision causes a region of high pressure, which redistributes the momentum component in the y-axis into the xz-plane, forming the sheet shown here. The exact details of our fan spray nozzle are a little more complex than this, and are shown later, however this is to first order what happens to produce a sheet. The colliding jets geometry that is often referred to in this chapter works on the exact same principle, except instead of using a nozzle to guide the flow it is created by two cylindrical jets oriented at an angle to each other. In the case where the angle of the jets is 180° to each other and gravity is negligible the sheet spreads out in a disc around the contact region. Figure from [73].

that as liquid flow is in general chaotic, these numbers and any conclusions we base on them are a general indication of regions of parameters that lead to different behaviours of the liquid, but that the exact boundaries for any given combination of liquid parameters may be hard to pin down.

Consider first the Reynolds number, which relates the inertia of the liquid (the tendency to resist changes in velocity measured in units of mass multiplied by area) to its viscosity (the tendency to resist deformation, measured in units of pressure multiplied by time). This number is usually used to estimate when a particular liquid flow will be laminar (the trajectories of the particles in the fluid are smooth lines that do not overlap) and when it will be turbulent (the trajectories of the particles in the fluid are chaotic and definitely overlap). Higher values indicate turbulence, as the viscosity of the fluid is insufficient to damp out the tendency of the particles to move chaotically, whilst lower values indicate laminar flow, as the viscosity of the fluid becomes more important, and the inertial forces driving the molecules in directions contrary to the flow are suppressed. The Reynolds number is defined as

$$\mathrm{Re} = \frac{\rho u L}{\mu}, \tag{2.1}$$

where ρ is the density of the liquid, u is the velocity of the liquid, L is a characteristic length which is defined by the parameters of the system under study and μ is the *dynamic* viscosity of the liquid. Often it is more convenient to use the *kinematic* viscosity $\nu = \mu/\rho$ instead, as this combines the two intrinsic properties of the liquid in question into one value, better enabling different liquids to be compared. This then renders the

equation

$$\mathrm{Re} = \frac{uL}{\nu}. \tag{2.2}$$

In relation to thin sheets of liquid, the Reynolds number is primarily a measure of the ability of capillary waves to form on the surface of the sheet (as more viscous liquids tend to damp such waves out — see section 2.2.4), which can contribute to the breakup or instability of the sheet.

The second dimensionless number to consider is the Weber number, which relates the inertia of the liquid with its surface tension (the tendency of liquids to minimise their surface, measured in units of force per unit length). This number is of great importance in the formation of liquid sheets, as the overall form of the liquid sheet is characterised by the interplay between the inertial forces of the particles spreading out into the sheet, and the force of surface tension pulling them back again. The Weber number is defined as

$$\mathrm{We} = \frac{\rho u^2 L}{\sigma}, \tag{2.3}$$

where σ is the surface tension of the liquid. Higher numbers indicate the inertial forces of the liquid dominate over the force of the surface tension, which thereby fail to pull the sheet into a fluid chain, leading to disintegration of the sheet (see section 2.2.4).

Two additional numbers exist that describe the effect of gravity on the sheet — the Froude number and the Bond number. The Froude number relates the intertia of the liquid to an external field. In the case of these experiments the only external field is gravity, so the Froude number is defined as

$$\mathrm{Fr} = \frac{u^2}{gL}, \tag{2.4}$$

where g is the acceleration due to gravity. In practice usually inertial forces dominate over the effects of the gravitational field, leading to high Froude numbers that indicate the effects of gravity can be neglected (for a flow rate of 3.4 mL/min the Froude number is around 83,000).

The Bond number

$$\mathrm{Bo} = \frac{L}{\lambda_\mathrm{c}} \tag{2.5}$$

where

$$\lambda_\mathrm{c} = \sqrt{\frac{\sigma}{\rho g}}, \tag{2.6}$$

denotes the relative effects of surface tension and the acceleration due to gravity. This is also small for the experiments (for isopropanol, the liquid with the lowest surface tension that we used, the Bond number is 0.0357), and so the effects of gravity can be neglected in this experiment.

2.2.1 Forming a thin liquid sheet

There are many ways to form thin liquid sheets, and many forms these sheets can take. In the context of a liquid sheet being a target for optical experiments however, the most useful form is a single liquid sheet extending in one direction. This form is easily created by colliding two cylindrical liquid jets [54], or by using a fan spray nozzle [73] (a nozzle consisting of a slit, with two sloping surfaces funneling liquid into the slit to produce the momentum transfer. This type of nozzle is used to create fan shaped sheets for spraying pesticides and other things, hence the name). As indicated in the introduction, the underlying principle of both of these methods is the same, even if the exact details are not. A schematic of the forces involved in creating the sheet is presented in Figure 2.2. In both methods, streams of liquid with momentum components in the z- and y-axes collide at a point, creating a region of high pressure. This high pressure region then converts the existing momentum in the y-axis into momentum in the xz-plane, spreading out the liquid and forming a sheet.

This sheet then spreads out radially until surface tension overcomes the radial velocity of the liquid, and it is pulled back in to a point, creating a sheet with a rather nice leaf shape (see Figure 2.3). As the fluid flows out radially in the sheet (before the contraction of the surface tension pulls it back to a point), the volume of the liquid must remain the same but the area covered increases proportional to the radius from the collision point (as it can be thought of as the circumference of a circle that is increasing in diameter out from this point). The thickness of the sheet must therefore be inversely proportional to the radius. Taylor [60] theorised that the thickness of the sheet formed by two colliding jets could be described in an equation of the form

$$h = \frac{\mathrm{K}(\theta, \phi)}{r}, \tag{2.7}$$

where h is the thickness, θ is the angle at which the jets collide, ϕ is the azimuthal angle around the collision point and $\mathrm{K}(\theta, \phi)$ is an arbitrary function of these two angles which describes the thickness.

In the discussion in this section the geometry being discussed is always colliding jets, unless otherwise mentioned. There has been much less work on fan spray nozzles, presumably because these are more complicated to make and do not lend themselves to changing the impact angle easily. Colliding jets are also easier to model theoretically, as inside a fan spray nozzle there is likely to be a complex interplay of forces which is harder to simulate. As mentioned above, the underlying principle of the creation of the sheet is the same, and once the sheets leave the nozzle they should thin out, form rims and disintegrate in the same way.

The exact nature of $\mathrm{K}(\theta, \phi)$ has been much discussed. Following the discussion in Choo and Kang [88], the first attempt to formulate an expression for the thickness of the sheet formed by two colliding jets (by Ranz [89]) was to predict the thickness distribution from the conservation of momentum, obtaining the equation

$$h = \frac{d^2}{4r}(1 + 2\cos\theta\cos\phi), \tag{2.8}$$

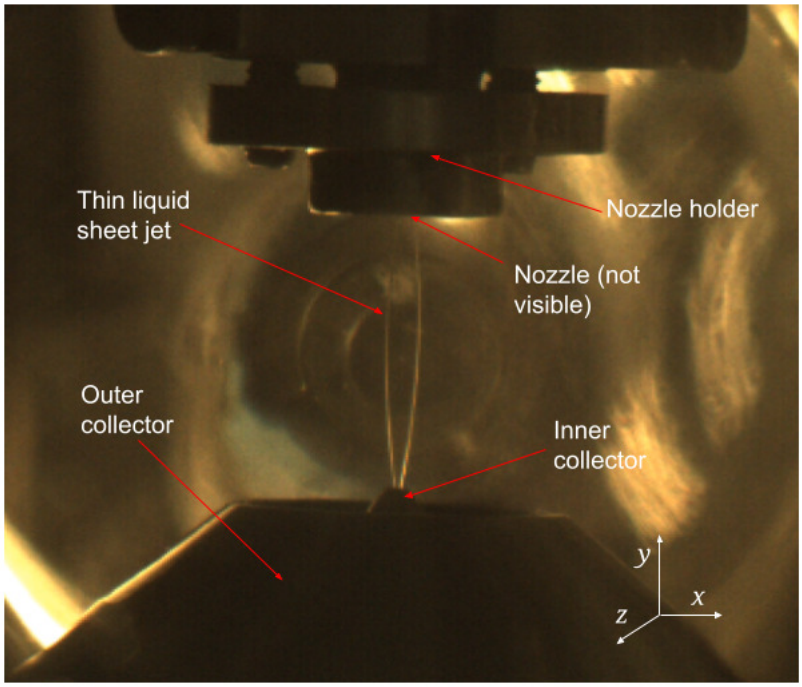

Figure 2.3: A face on picture of our thin liquid sheet. Note the initial expansion of the liquid, before surface tension starts to pull it in. The tip of the leaf is inside the hole of the collector (see section 4.5.1). Our jet is much thinner in the x-direction than the jets produced by the colliding jets geometry — this is probably because the nozzle only allows a small section of the sheet to actually form (see section 2.5). Picture taken by Gediminas Galinis

.

where d is the diameter of the jets that are colliding, the characteristic length for this system. Unfortunately, this equation predicts negative thickness at $\phi = 180°$ for θ less than 60°, so clearly needed to be refined.

Miller [90] suggested the equation

$$h = \frac{d^2}{4r}\left(\frac{1-\cos^2\theta}{1+\cos^2\theta}\right)\left(\frac{1}{1-\frac{2\cos\theta\cos\phi}{1+\cos^2\theta}}\right) \tag{2.9}$$

as an alternative solution to Ranz's balance of momentum and mass conservation, which fits the actual shape of the sheet produced much better.

Hasson and Peck [67] take a completely different approach, assuming that the liquid flux distribution in the colliding jets in a plane parallel to the sheet was an ellipse, and then assuming that the mass and momentum in a section of the ellipse is conserved as the jets collide to form the sheet, they obtain the equation (expressed in the same terms as those above)

$$h = \frac{d^2}{4r}\frac{\sin^3\theta}{1-\cos\phi\cos\theta}. \tag{2.10}$$

Comparing their results to Miller's, they note that while the equations are not identical, for values of θ greater than 60° they predict identical results. Below 60° the two sets of predictions diverge.

Ibrahim and Przekwas [66] note that Hasson and Peck's analysis is more accurate compared to experimental results than Miller's below $\theta = 60°$, but they take issue with the idea that the ellipse of the jet is preserved in the collision region. They adopt a semi-empirical approach, using an experimentally determined expression for the thickness of the sheet formed by one jet impinging on a wall advanced by Naber and Reitz [85] to obtain an expression for the initial thickness of the sheet at the point of impact of the jets of

$$h_{\mathrm{i}} = \left(\frac{\frac{d}{2}\beta \sin\theta}{\mathrm{e}^{\beta} - 1} \right) \mathrm{e}^{\beta(1-\phi/\pi)}, \tag{2.11}$$

(where β is a parameter determined from mass and momentum consideration) which can then be used to find the thickness as (again expressed in the same terms as above)

$$h = \frac{h_{\mathrm{i}} d}{2r \sin\theta}. \tag{2.12}$$

All of these approaches, different as they may be, have a couple of things in common. They assume a uniform velocity of liquid in the sheet, and they assume that the only factors affecting the thickness of the sheet are the impingement angle θ and the azimuthal angle ϕ. However, there is good reason for thinking this is not so. Choo and Kang [68, 91], and Bremond and Villermaux [65] have both shown that the velocity profile in both the jets that collide and the sheet is not uniform, and that the thickness profile of the sheet will change as a result of this non-uniform velocity. It also seems unlikely that the properties of the liquids themselves will not affect the thickness or the shape of the sheets produced.

The number of factors involved in forming the sheets can quickly overwhelm any model that is aiming for simplicity, and the situation only gets more complicated when we switch to a fan spray nozzle, and consider the interactions of the liquid inside the nozzle. Ha, DePonte and Santiago [92] used a full Navier-Stokes simulation of the velocities inside their nozzle, alongside a thorough parametric study of Reynolds and Weber numbers to obtain scaling laws for the thickness and dimensions of the thin sheet obtained. However it is unclear how valid their results are for nozzles with different internal geometries to their own. Dombrowski, Hasson and Ward [56], and following them Choo and Kang [88] have used monochromatic interferometry to determine $\mathrm{K}(\theta, \phi)$ experimentally (see section 2.4.1) — some combination of these two approaches may be the only way to properly characterise and describe the jet, although the necessary predictive modelling is both time and computing resource intensive.

I will now discuss three aspects of the behaviour of the liquid sheets that are important to understanding how our jet behaves — the relation between the sheet and the rim, the action of waves on the surface of the sheet and the mechanisms of the breakup of the sheet into droplets. After this I will describe the nozzle that we have designed and made, and the sheets that are produced from it, and then I will return to the models discussed here and try to apply these ideas to provide a quantitative description of the flow of the sheet.

2.2.2 Distinction between the sheet and the rim

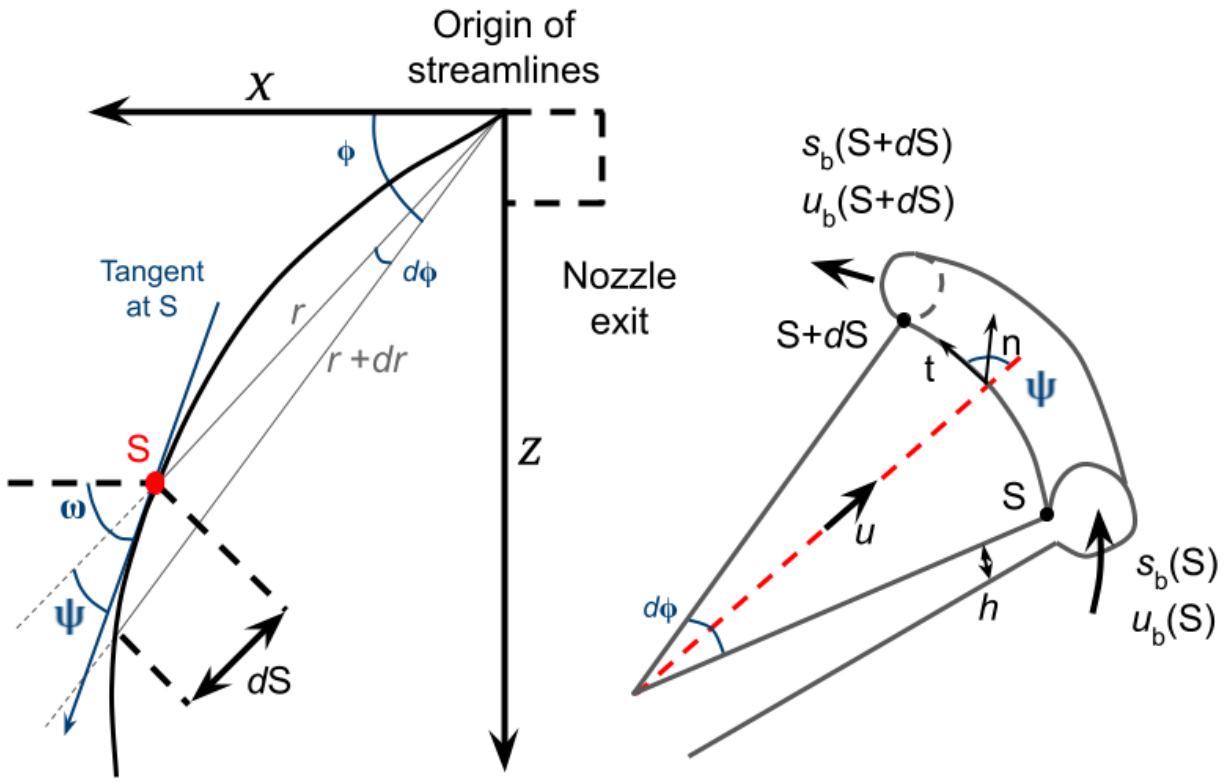

Figure 2.4: Illustrations of the coordinate system used in this section. The leftmost drawing is tracing out the arc of the rim from its origin inside the nozzle (based on Figure 2 in Clark and Dombrowski [93]). The rightmost drawing is the section from the origin between the points S and dS on the rim (based on Figure 4 in Bermond and Villermaux [65]). At the point S the area of the rim is s_{b}(S) and the velocity of the flow in the rim is u_{b}(S). The velocity of the sheet flowing into the rim is u, and the thickness of the sheet is given by h. h varies with r, and so the thickness at the edge of the sheet is denoted by h_{e} (not shown on the figure — equivalent to the thickness at the point S). The angle between the x-axis and the radius r is given by ϕ, the angle between the x-axis and the tangent to the rim at the point S is ω, and the angle between r and the tangent at point S is ψ.

A stable liquid sheet is characterised by a thin flat leaf shaped sheet portion, surrounded by a thicker rim. This rim transports the liquid that has flowed out of the central sheet towards the tip of the jet, whereupon it collides, either forming a second sheet or breaking up into droplets depending on how fast the liquid is going in the rim. Even in an unstable sheet the rim can often play a role in the breakup of the sheet into droplets [94], whilst in a stable sheet the size and shape of the sheet are determined by where the rims form.

In Figure 2.4 the system of coordinates used in this analysis is outlined with respect to drawings of the sheet and the rim. In order to calculate the trajectory that the rim takes as liquid flows into it from the sheet, as well as the size of the rim and the velocity of the flow at any given point along this trajectory, we need to consider conservation of both momentum and flux. In what follows I have used the discussions in Bush and Hasha [64], Bremond and Villermaux [65] and Clark and Dombrowski [93] — all of them follow the same basic scheme, but with different emphases.

To start with, as the rims are stationary, the flux flowing out of a section of the arc dS must be equal to the flux flowing into that section from the sheet

$$\frac{\rho d(s_{\mathrm{b}} u_{\mathrm{b}})}{d\mathrm{S}} = \rho u h_{\mathrm{e}} \sin\psi, \tag{2.13}$$

where s_b is the cross-sectional area of the rim, u_b is the velocity of the fluid along the centerline of the rim, u is the velocity along the streamline intersecting dS, h_e is the thickness of the sheet just before the rim and ψ is the angle between the streamline intersecting dS and the tangent to the rim at that point. The geometrical definition of ψ is such that

$$\sin\psi = r\frac{d\phi}{d\mathrm{S}}, \tag{2.14}$$

i.e. the length along the tangent at S that is swept out along dS, which for a sufficiently small area is equal to the the length of the sheet bordering the rim in this section.

We can split the momentum conservation up into two axes — tangential to the rim and normal to it. Tangential to the rim, the momentum change over dS must be equal to the additional momentum added by the flux flowing into the rim

$$\frac{\rho d(s_{\mathrm{b}} u_{\mathrm{b}}^2)}{d\mathrm{S}} = \frac{\rho d(s_{\mathrm{b}} u_{\mathrm{b}})}{d\mathrm{S}} u\cos\psi. \tag{2.15}$$

Using Equation 2.13 we get

$$\frac{\rho d(s_{\mathrm{b}} u_{\mathrm{b}}^2)}{d\mathrm{S}} = \rho u^2 h_{\mathrm{e}} \sin\psi\cos\psi. \tag{2.16}$$

Normally to the rim the forces must balance out, as the rims do not move laterally. Therefore the forces pulling into the sheet (surface tension) must equal the forces pulling out (centripetal force and the normal component of the momentum of the flux flowing into the rim). Using Equation 2.13 again we get

$$2\sigma = \rho s_{\mathrm{b}} u_{\mathrm{b}}^2 \frac{d\omega}{d\mathrm{S}} + \rho u^2 h_{\mathrm{e}} \sin^2\psi, \tag{2.17}$$

where $\frac{d\omega}{d\mathrm{S}}$ is the change in the angle between the x-axis and the tangent of the rim, which is the curvature of the section dS.

Bush and Hasha add an additional term to each of the momentum conservation equations, corresponding to an additional inwards force due to the curvature of the centerline of the rim in the normal momentum conservation, and due to tangental gradients in rim curvature in the tangental momentum conservation. This gives equations of

$$2\sigma + 2\sigma(\pi - 1)R\frac{d\omega}{d\mathrm{S}} = \rho s_{\mathrm{b}} u_{\mathrm{b}}^2 \frac{d\omega}{d\mathrm{S}} + \rho u^2 h_{\mathrm{e}} \sin^2\psi \tag{2.18}$$

and

$$\frac{\rho d(s_{\mathrm{b}} u_{\mathrm{b}}^2)}{d\mathrm{S}} = \rho u^2 h_{\mathrm{e}} \sin\psi cos\psi - \pi R^2 \sigma \frac{d}{d\mathrm{S}}\left(\frac{1}{R}\right), \tag{2.19}$$

where R is the radius of the rim and σ is the surface tension of the liquid. These terms can be regarded as correction terms for the main equations for situations where the rim thickness and curvature are sufficiently large, and where the surface tension exerts a significant enough force to affect the tangental flow.

To transform this into a system of equations that can be solved numerically we note Equation 2.14 can be rearranged to give

$$\frac{d}{d\mathrm{S}} = \frac{\sin\psi}{r}\frac{d}{d\phi}, \tag{2.20}$$

that the relation between ϕ and r can be expressed in terms of ψ as

$$\frac{dr}{d\phi} = \frac{r}{\tan\psi} \tag{2.21}$$

and the insight by Taylor at the start of section 2.2.1

$$h = \frac{\mathrm{K}(\theta,\phi)}{r}. \tag{2.22}$$

These all can be substituted into Equations 2.13, 2.18 and 2.19, to yield the following set of differential equations

$$\begin{aligned}
\frac{d(s_\mathrm{b}u_\mathrm{b})}{d\phi} &= u\mathrm{K}(\theta,\phi)\\
\frac{d(s_\mathrm{b}u_\mathrm{b}^2)}{d\phi} &= u^2\cos\psi\mathrm{K}(\theta,\phi)\\
s_\mathrm{b}u_\mathrm{b}^2\left(\frac{d\psi}{d\phi}+1\right)\sin\psi &= \frac{2\sigma r}{\rho} - u^2\sin^2\psi\mathrm{K}(\theta,\phi)\\
\frac{dr}{d\phi} &= \frac{r}{\tan\psi}.
\end{aligned} \tag{2.23}$$

In order to be as general as possible we can define a characteristic length scale (d_j - the diameter of the colliding jets, or equivalently the diameter of the point behind the nozzle orifice where the streamlines emanate from) and velocity scale (u_j, the velocity of the fluid in the sheet). Dividing through by these, we are left with the following dimensionless equations

$$\begin{aligned}
\frac{d(\tilde{s}_\mathrm{b}\tilde{u}_\mathrm{b})}{d\phi} &= \tilde{\mathrm{K}}(\theta,\phi)\\
\frac{d(\tilde{s}_\mathrm{b}\tilde{u}_\mathrm{b}^2)}{d\phi} &= \cos\psi\tilde{\mathrm{K}}(\theta,\phi)\\
\tilde{s}_\mathrm{b}\tilde{u}_\mathrm{b}^2\left(\frac{d\psi}{d\phi}+1\right)\sin\psi &= \frac{2r}{\mathrm{We}} - \sin^2\psi\tilde{\mathrm{K}}(\theta,\phi)\\
\frac{dr}{d\phi} &= \frac{r}{\tan\psi},
\end{aligned} \tag{2.24}$$

where the tildes indicate dimensionlessness and $\mathrm{We} = \frac{\rho d_\mathrm{j} u_\mathrm{j}^2}{\sigma}$ is the Weber number. These equations can be numerically integrated to yield the shape and profile of the rims of the sheet. This will be done in section 2.5.

2.2.3 Capillary waves

Capillary waves are waves that form on the boundary between two different phases, such as that of the surface of our thin sheet of liquid and the surrounding air. Taylor [60, 70] produced the definitive analysis of waves on a thin sheet in 1959, as part of his series on the dynamics of thin liquid sheets. As the liquid sheets are bounded on two sides by air, the dynamics are more complicated than that of a body of liquid with only one air-liquid boundary. In this case, the waves can be described by the interference of two simple harmonic waves propagating on the surface in opposite directions. When there are two parallel boundaries that are close enough to affect each other there are two pairs of waves that have to be accounted for, and the relations between them lead to two distinct types of capillary waves — what Taylor called symmetric and antisymmetric waves (see Figure 2.5).

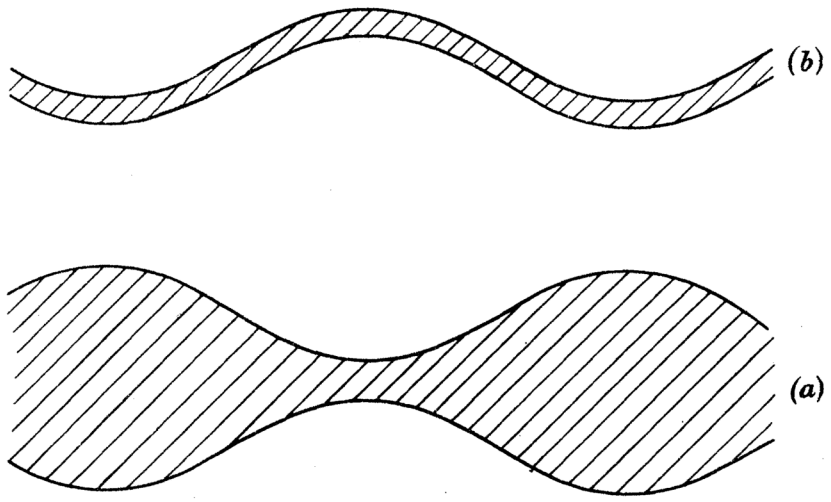

Figure 2.5: A sketch of symmetric (b) and antisymmetric (a) waves by Taylor [70]. The displacements of the top and bottom surfaces are in phase in (b), but antiphase in (a). Drawing from [70]

The disturbance of the sheet can be regarded as being made up of these two types of waves. Symmetric waves are those in which the front and back surfaces of the sheet oscillate in phase, and antisymmetric waves are those in which the front and back surfaces of the sheet oscillate in anti-phase.

Symmetric waves play a part in the Kelvin-Helmholtz instability [95] (see section 2.2.4 below), but antisymmetric waves are important in governing the shape and extent of the thin liquid sheet and also to play some part in the breakup of these sheets, so we will confine ourselves to them here.

The expression given for the velocity of the antisymmetric wave is

$$w_\alpha^2 = \frac{2\sigma}{\rho h}\frac{1}{2}kh\coth(\frac{1}{2}kh), \tag{2.25}$$

where w_α is the speed of the wave and $k = \frac{2\pi}{\lambda}$ is the wavenumber of the wave. When kh is small (as in a thin

liquid sheet) $\frac{1}{2}kh\coth(\frac{1}{2}kh) \approx 1$ and therefore the speed of the wave is given by

$$w_\alpha^2 = \frac{2\sigma}{\rho h}. \tag{2.26}$$

Taylor then suggested that as a result the lines of constant phase on an expanding sheet which is thinning as $\frac{1}{r}$ will be at rest at an angle Φ to the radii if

$$\sin^2\Phi = \frac{2\sigma}{\rho u^2 h} = \frac{4\pi\sigma r}{\rho u Q} = \frac{r}{R_{\text{rim}}}, \tag{2.27}$$

where

$$R_{\text{rim}} = \frac{\rho u Q}{4\pi\sigma} = \frac{d_{\text{j}}\text{We}}{16} \tag{2.28}$$

and Q is the volume flowing into the sheet per second. R_{rim} is therefore a radius beyond which no waves can remain at rest and, as we shall see in the next section, therefore defines the maximum extent of the stable sheet formed.

2.2.4 Breakup and instability

One of the primary reasons that thin sheet jets have been studied over the years is the ease with which they break up into droplets. As the sheet disintegrates it is able to cover a wide area with relatively evenly distributed droplets, which is useful for many applications, including pesticide spraying and jet fuel injection. Our application, which requires a thin sheet target to sit stably in a vacuum chamber, is in some ways diametrically opposed to this as ideally we do not want any droplets to form whatsoever. However it is still important to have some idea of how liquid jets break down into droplets, if only so that we can design our jets so that they don't.

There are two main regions of instability and breakup in a liquid sheet, characterised primarily by the Weber number of the sheet [61, 66, 96]. At high Weber numbers the sheet breaks up due to waves caused by the Kelvin-Helmholtz instability, whilst at lower Weber numbers the breakup is caused by cardioid antisymmetric waves.

The Kelvin-Helmoltz instability occurs when there is a velocity difference between two fluids. This velocity difference causes waves with a short wavelengths $\lambda = \frac{2\pi}{k}$ to occur on the surface and when the effects of surface tension and viscosity are neglected they amplify exponentially proportional to e^{kt} [97], which causes a catastrophic instability in the surface. Surface tension causes these waves to decay, as it acts as a force opposing the wave motion, whilst the viscosity of the liquid acts as a damper, causing the growth of the waves to be limited [97]. This then explains the Weber number dependence of this instability — it is only past a critical value of the Weber number that the relative velocities of the surrounding air and the liquid become significant

enough to overcome the surface tension of the liquid and the waves can start to amplify, breaking up the jet.

Villermaux and Clanet have undertaken a thorough study of this instability regime [98]. They observed that the transition from a smooth liquid sheet to what they described as a 'flag-like' flapping occurred at Weber numbers roughly greater than $\frac{40}{\sqrt{\alpha}}$, where $\alpha = \frac{\rho_a}{\rho}$, the ratio of the densities of the air and the liquid. The radius of the sheet was also observed to contract with increasing Weber number in this regime, decreasing as described in [61]:

$$\frac{R}{d_\mathrm{j}} \approx \alpha^{-2/3}\mathrm{We}^{-1/3}. \tag{2.29}$$

For our purposes this regime represents an upper limit on the stability of the jet, in particular on the speed that the liquid can flow through the jet. This is confirmed by the parametric study of Bush and Hasha [64] — above a certain value of Weber number, no stable sheet can form. Their boundary between these two regions is curved, indicating a slight dependence on Reynolds number and therefore viscosity, but the overwhelming physical factor that affects this instability is surface tension. In a vacuum, one would expect this instability not to occur so readily, so feasibly if this was the limiting factor on our jet then running in vacuum may fix it.

The second breakup regime is more interesting, as in this Weber number range a stable sheet can form. Its appearance however can vary significantly, from an oscillating stream of liquid [64, 99] to a single sheet with a rim that instantly dissociates into droplets [60, 65]. Partly this is due to the variation of the Weber number — Huang [96] identified several regions of different Weber numbers below the critical one at which Kelvin-Helmholtz instability becomes important, each characterised by slightly different behaviour with regards to waves forming on the surface. In general though, the breakup of the sheet in this regime is not total, and consists of droplets being ejected from the rims of the sheet, rather than the sheet itself disintegrating.

Clanet and Villermaux [62] have performed an extensive study on this regime as well (the two papers form a pair). Above a minimum Weber number needed to form the sheet (below which the pressure is not enough to overcome the force of the surface tension and spread out into a sheet) and below the critical Weber number after which the Kelvin-Helmholtz instability sets in, the sheet is smooth and relatively unperturbed by waves. It increases in radius according to

$$\frac{R}{d_\mathrm{b}} = \frac{\mathrm{We}}{16}, \tag{2.30}$$

which is the radius predicted by the stationary cardioid waves. The sheet in both papers was created by two liquid jets at 180° to each other, so there was no prospect of stable rims forming a fluid chain as discussed in section 2.2.2 above. What they observed however is important to this discussion as it concerns what happens when a stable rim cannot form, and the liquid is ejected in droplets. They observed indentations on the rim of the sheet, like Huang [96] before them. These indentations are not stable, and a time averaged photo shows a circular sheet. In both cases these are theorised to result from waves on the sheet. As the maximum radius corresponds to that of the stationary antisymmetric waves (see section 2.2.3) it would make sense that these

waves are cardioid in shape and they approximately are, although Clanet and Villermaux found some deviations from this. To first order then the profile of the sheet formed can be thought of as interfering cardioid waves, and Ibrahim and Przekwas [66] fit the shape of the sheet formed by two jets colliding at an angle to measured sheet profiles with a reasonable degree of accuracy.

Droplets form on the rim of the circular sheets used by most of the experimenters quoted above through the Rayleigh-Plateau instability [62, 96, 100]. In this, small fluctuations in the rims of the sheet get amplified with a speed that depends on how their wavenumber $k = \frac{2\pi}{\lambda}$ relates to the radius of the rim. The fluctuation will eventually become sufficiently large that it overcomes the force of the surface tension, detaching a droplet from the rim of the sheet. The size of the droplet depends on the wavenumber that dominates, and thus on the size of the rim.

When we turn to leaf shaped sheets, we find a similar behaviour in the region of single sheets with unstable rims [60, 64, 65, 94]. The results of Bremond and Villermaux [65] are especially interesting, as they form a jet in both ethanol and water. The former has stable rims, and forms a nice fluid chain at all the Weber numbers they tested, whilst the latter, water, never forms a sheet with stable rims. This cannot be due to a difference in Weber number, as water has a higher surface tension than ethanol which would act against the Rayleigh-Plateau instability. Instead the authors attribute it to an increased Reynolds number (which doubled in the case of water). This makes sense, as the lower viscosity of water allows the antisymmetric capillary waves to form much more easily, and consequently with larger amplitudes. This increased disturbance affects the rims of the sheet, favouring droplet formation over a stable rim.

Another instability that has been observed to affect leaf shaped sheets is the fishbone instability, discovered by Bush and Hasha [64] and extensively studied by Jung *et al.* [101]. This instability arises due to a slight asymmetry in the speeds of the colliding jets in the positive and negative y-directions, and consists of a shrinking of the sheet, with waves of droplets being thrown out of the bottom in lines. The whole effect is slightly reminiscent of a fish skeleton, hence the name. This instability is transitory — it occurs for a certain range of Reynolds and Weber numbers, and only if there is a velocity mismatch in the jets that collide, and then disappears again once those parameters have been passed through. This indicates to me that some kind of resonance or beating is occurring to form these structures, which seems less likely to occur in a fan spray nozzle due to the increased variety of momenta in the liquid entering the nozzle.

In any case, both the Weber and Reynolds number need to fall under a certain critical value in order for a stable smooth sheet with stable rims to be formed, which is the type of sheet we require. Unfortunately, the numbers are linked (they both depend on a length, and they both depend to some extent on velocity, Reynolds linearly and Weber quadratically), and the parameters that are independent are intrinsic properties of the liquids, and are not easily changed.

2.3 Nozzle design

The design of the nozzles we use is a fan spray nozzle, similar to the one presented in Watanabe *et al.* [59] and Dombrowski, Hasson and Ward [87], among others. Although the material compositions of their nozzles differ (a commercial porcelain nozzle for the latter and a crushed steel pipe and a pair of razorblades for the former), the principle behind the formation of the thin liquid sheet is the same. Assuming the jet spreads out in the x-direction, the cross-section of the nozzle in the xz-plane is effectively a triangle, allowing the streamlines of the sheet to spread out. In the yz-plane however, there is a slope, directing the streamlines in this plane towards each other. This aids the momentum transfer from the y-direction to the x-direction, as from an initial compression point upstream the liquid with momentum oriented along the preferred sheet direction is free to spread out, whilst liquid with momentum oriented perpendicularly to the preferred sheet direction is redirected into a collision point, causing the sheet to spread out further.

This philosophy has been followed in the design of our jet (illustrated in Figure 2.6). From a channel of width 1.59 mm, the liquid is narrowed down into a channel of diameter 70 µm, before spreading out again. At this point the geometry of the xz- and yz-planes differ. The xz-plane continues spreading, and then the nozzle wall stops, allowing the sheet to spread out without any interference from the nozzle. Meanwhile, in the yz-plane the walls of the nozzle cut back in at an angle of 60° to the z-axis, and then remain there for a further 50 µm to ensure the momentum along the y-axis is totally dispersed. In this way the most efficient transfer of y-momentum to momentum in the xz-plane is achieved (this is shown schematically in Figure 2.6).

2.3.1 Manufacture and wear of a fan spray nozzle

Our nozzles are created by 3D-printing onto a specially designed substrate, using a 3D printer that is designed for fine detail work (Nanoscribe, by Nanoscribe GmbH). The nozzles are printed using a polymer called IP-S (also made by Nanoscribe GmbH), which is designed for printing structures of the size of our nozzles. This 3D printing has allowed for rapid prototyping in order to iterate through many different designs to come up with an optimal one. The substrate the nozzle is printed on is shaped like an ellipse (to ensure that the jet is facing in a reproducible direction — the slit is printed perpendicular to the major axis of the ellipse). On top of the ellipse is an inverse cone, to allow the printed part of the nozzle to attach to the substrate. In the cone is drilled a hole, through which the liquid flows during the experiment to get to the more finely detailed structure of the nozzle where the sheet production takes place.

When making these nozzles it is very time consuming to 3D print every part of the nozzle — printing a nozzle in this way takes around 8 hours, which is prohibitively long if multiple different nozzles need to be manufactured and tested. Instead the skeleton of the nozzle shape is traced out in the IP-S monomer using a laser, trapping the monomer in a shell made of polymerised IP-S. This procedure takes only 2 hours, a quarter

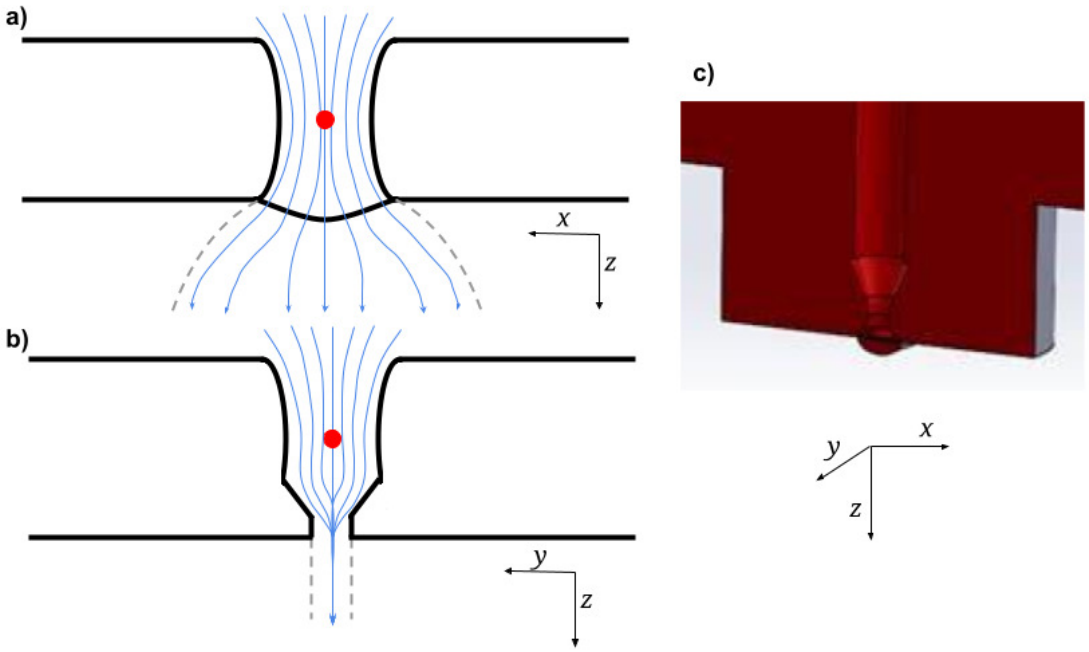

Figure 2.6: The xz- (**a**) and yz- (**b**) plane cross sections of our nozzle. **c** shows a cut through of a 3D model of the nozzle. The red dot indicates the approximate position of the compression point from which all streamlines in the sheet behave as if they originate from. The streamlines (shown in blue) are the directions of travel of particles in the jet — if a particle starts at the tail of a streamline it will follow the curve of it down to the tip. In the xz-plane the streamlines are allowed to spread out freely, whilst in the yz-plane they are collided by a slope at 60° in order to transfer their y-direction momentum component to the xz-plane. The grey dashed lines show the extent of the sheet. 3D drawing by Clement Ferchaud.

of the time. The nozzles are then washed with propylene glycol methyl ether acetate (PGMEA) and isopropanol to remove excess monomer from the outside, and baked in an oven for an hour to polymerise the monomer that is trapped inside the shell. IP-S polymerises in a two photon process at the wavelength of the laser (780 nm), so there is only sufficient intensity for polymerisation at the focus of the printer's laser [102]. This gives a very fine resolution (less than 1 µm in the settings we use), which is good for the detailed structure of the nozzle (for a smooth sheet it is important that the edge of the nozzle be very flat), whilst allowing nozzles to be printed in a reasonable timeframe.

As these nozzles are intended to be used over an extended period, we were concerned about wear. IP-S is a proprietary mixture manufactured by Nanoscribe GmbH, and data on the solubility or reactivity of polymerised IP-S with many common solvents is hard to come by. Therefore we decided to test the susceptability of IP-S to several common solvents, selected on the basis of which solvents we were likely to want to use in future experiments. Six solvents were tested — isopropanol, acetone, DMSO, ethanol, water and a saturated solution of water and NaCl. Seven cuboids of 1850 µm × 1000 µm × 1000 µm were fabricated on ITO glass — glass with a conducting layer on one side to facilitate detection of the surface for printing. Six of them were then immersed

in 40–50 mL of the solvents for a week, whilst the seventh one was left in air as a control. They were then removed, washed briefly with isopropanol to remove any residue solvent and air dried. Photographs were then taken of them under a 10× microscope to determine if any damage had occurred.

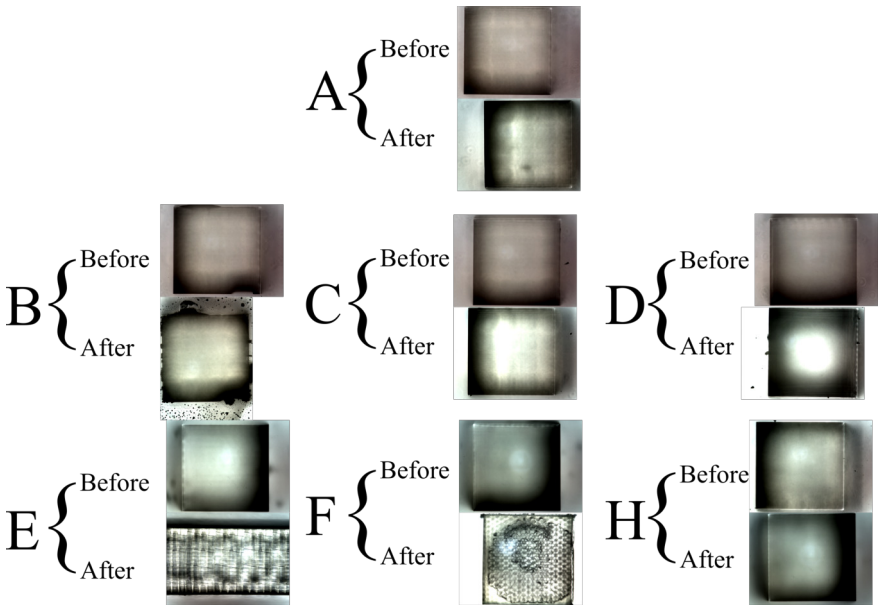

Figure 2.7: The cubes before and after being immersed in solvent for a week. A was left in air and suffered no damage. B, C and D were immersed in the alcohols and H was immersed in water — none of these showed any signs of wear. E and F were immersed in DMSO and acetone respectively, and do show signs of wear — E was only photographed on its side, as it fell off the ITO glass and F shows significant signs of wear.

Most of the solvents appear to have had negligible effect on the cubes. The only two solvents to have exhibited any sort of damage were DMSO and acetone, which makes some sense as they have a similar molecular structure so the damage mechanism could be similar. The damage exhibited on the cubes was only minor surface damage, which in the case of acetone led to some eating away of the edges and in the case of DMSO caused no visible surface damage but did enable the cube to detach from the ITO substrate.

Subsequently five of the solvents have been jetted through the nozzle — isopropanol, ethanol, water (with and without NaCl) and DMSO. All of these solvents have undergone significant running, and the results from the long periods of use confirm the results of the wear test above. When running either alcohol or water the jet lasts for a very long time — upwards of three months of constant use without showing any signs of wear, and around 6 months before the degradation gets severe enough that it warrants the replacement of the nozzle. When running DMSO an individual nozzle lasted for an average of two weeks before becoming completely unusable. A comparison between a new and a worn nozzle is shown in Figure 2.8. A worn nozzle causes the sheet to become unstable, as noted by Letouzey *et al.* [103]. The uneven surface of the nozzle causes more surface waves, which are accordingly less damped by the viscosity than with a smooth nozzle, and so therefore destabilise more and more of the sheet, leaving only a small part at the top which is stable enough to be used

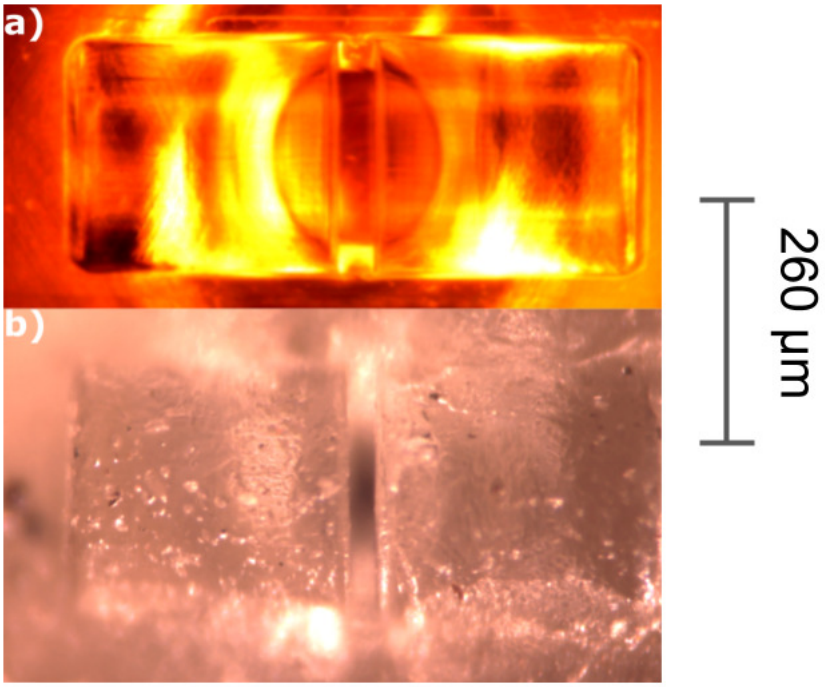

Figure 2.8: The edges of the new nozzle (a) are very clean and sharp, whilst that of the worn nozzle (b) are rougher and more rounded, having been worn away by erosion from the jet.

for experiments.

2.4 The liquid sheets produced by our fan spray nozzle

Once the sheet has been made, we need to know how thick it is. The measurement of the thickness of a liquid sheet is not a trivial exercise, firstly because it is so thin (the difference between 1 and 5 µm is not a distinction it is possible to make by eye, but matters a great deal when it comes to absorption of XUV and soft X-ray photons) and secondly because any mechanical measurement of the sheet will disrupt it. People have tried to measure the flux of fluid as a function of angle from the point of impact (the nozzle head in our case) [60, 64] by collecting fluid with a razor blade, but while this is useful for other considerations it is a very impractical and roundabout way of measuring the thickness of the sheet, and cannot give the full thickness profile.

The method (or family of methods) that have proved the most effective for measuring the thickness profile of thin liquid sheets is interferometry — measuring with light [55, 59, 87, 88]. The insight behind this method is that the liquid sheet consists of two surfaces separated by a layer of liquid. When light is incident upon a surface it can either reflect off it or travel through it. The incident light can therefore reflect off of both the front and back surfaces of the sheet, and then be detected. The two reflected light waves interfere with each other as the second has travelled the extra distance from the front of the sheet and back again, and therefore has accumulated some extra phase. From the interference pattern it is therefore possible to work out the extra

distance travelled, and so the thickness of the sheet. This is the same physics that causes the well known rainbow pattern on an oil film on top of water.

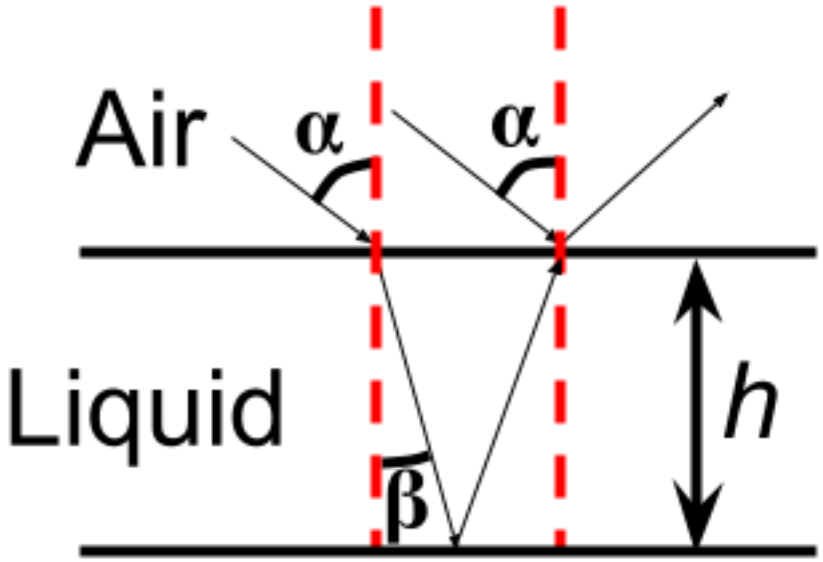

Figure 2.9: A schematic drawing of the principle of interferometry. α is the incident angle of the light on the surface, β is the internal angle caused by refraction of the light and h is the thickness of the sheet. The best results for interferometry are obtained when α is very nearly 0.

2.4.1 Monochromatic interferometry

Monochromatic interferometry is interferometry performed with monochromatic light. As this light is all of one wavelength the interference pattern produced is easy to see, with bands of colour interspersed with dark bands across the sheet (see Figure 2.10). The optical path difference between the reflections from the two different

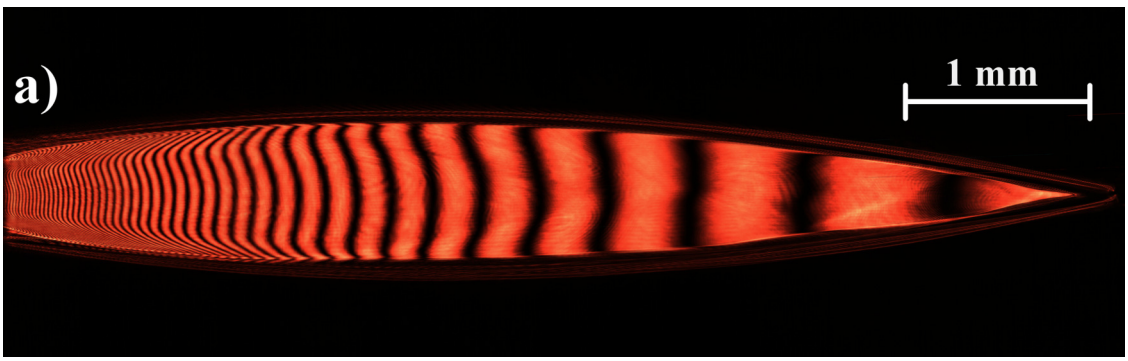

Figure 2.10: A monochromatic interferogram of our jet. Note that the pattern does not include the rims, which show up as red lines tracing the outside of the sheet due to the sudden discontinuity in thickness.

surfaces is

$$D_{\text{opt}} = 2hn\cos\beta \tag{2.31}$$

where h is the thickness of the sheet, n is the refractive index of the liquid and β is the angle of refraction of the light. Using Snell's law of refraction

$$\frac{\sin\beta}{\sin\alpha} = \frac{n_\alpha}{n_\beta}, \tag{2.32}$$

where α is the angle of incidence of light onto the sheet and n_α and n_β are the refractive indices in air and the sheet respectively, and taking n_α to be 1, we can write this in terms of α as

$$D_{\text{opt}} = 2h(n^2 - \sin^2\alpha)^{1/2}. \tag{2.33}$$

The bright fringes are created when D_{opt} is such that they are in phase. This condition should recur once in every optical cycle — implying that the difference in optical path difference between one bright fringe and the next one is λ — the wavelength of the light. This leads to an equation for the change in thickness of the sheet of

$$\Delta h = \frac{\lambda}{2(n^2 - \sin^2\alpha)^{1/2}}. \tag{2.34}$$

In order to move from the monochromatic interferogram to a measurement of the thickness of the sheet multiple methods can be used. The more general method, which was the one employed in our paper [73], is to retrieve the phase profile of the light and unwrap it. (As phase repeats every optical cycle it is not possible to say *a priori* what absolute phase relationship one part of the phase profile has with another. Phase unwrapping algorithms seek to establish this). This then gives a relative thickness profile, which can be anchored to actual thickness by other measurements.

The phase profile of our interferograms was retrieved and unwrapped using software written by Jason Cole, details of which can be found in his thesis [104]. The retrieval algorithm is called a continuous wavelet transform [105], which is a generalisation of a Fourier transform, in that it projects a function f(x) onto a set of known functions. However, these functions are not necessarily mutually orthogonal. They are instead a group of functions $\psi_{ab}(x)$, which are related to a defined 'mother' function $\psi(x)$ by

$$\psi_{ab}(x) = \psi\left(\frac{x-b}{a}\right). \tag{2.35}$$

With an appropriate choice of $\psi(x)$ it is possible to use this to find the frequency at every position b along the image in a similar way to a Fourier transform, and from this the phase at each point can be retrieved.

The phase unwrapping is then performed using an algorithm by Goldstein *et al.* [106]. The basis for this is that in a smooth phase map, the sum of the phase differences around a given set of pixels should be zero —

i.e. there should be no discontinuities in the phase. This algorithm calculates this sum for each pixel, and then connects together pixels with positive and negative sums. It then spirals outwards from a pixel that it knows to be good on a path that avoids all known bad pixels, and then having established a phase map for most of the image it can slot the discontinuities in sensibly.

Dombrowski, Hasson and Ward [56], and following them Choo and Kang [88] employ a simpler technique, using knowledge about the expected thickness profile of the sheet. As thickness varies with $\frac{1}{r}$ according to Equation 2.7, the variation in thickness between two points along a radial streamline r_1 and r_2 will be

$$h_1 - h_2 = \mathrm{K}(\theta, \phi)\left(\frac{1}{r_1} - \frac{1}{r_2}\right). \tag{2.36}$$

If the thickness at a point on the sheet is known, then the thickness at any other point can in principle be established by counting fringes from that point. In this case the difference in thickness between the two points can be written as

$$h_1 - h_2 = (N_2 - N_1)\frac{\lambda}{2(n^2 - \sin^2\alpha)^{1/2}} \tag{2.37}$$

We can combine the two above equations into

$$\frac{1}{r_2} = -\frac{\lambda(N_2 - N_1)}{2\mathrm{K}(\theta, \phi)(n^2 - \sin^2\alpha)^{1/2}} + \frac{1}{r_1}. \tag{2.38}$$

$\mathrm{K}(\theta, \phi)$ can therefore be worked out experimentally from the slope of the line of $(N_2 - N_1)$ vs $\frac{1}{r_2}$, with only the thickness at one point needed to anchor it to the experiment, as with the more general method above. Both groups that use this method treat K as a constant for a given value of θ, which is not a reasonable assumption — the thickness of the sheet will have some variation with the azimuthal angle ϕ. However this is not a problem for the method in general, as all that has to change is the recognition that the slope will be dependent on ϕ.

Both methods ultimately provide the same insight though — they both provide a means of finding a relative thickness profile for the whole jet, but they need a fixed point for which the thickness is absolutely known to anchor them in reality and give the actual thickness of the jet. That method is white light interferometry.

2.4.2 White light interferometry

This method of interferometry is more complicated than monochromatic interferometry, as instead of using only one wavelength, we use a very broad spectrum, encompassing most of the visible range. The idea is that at any given point, some wavelengths will constructively interfere, and others will destructively interfere, under the same principle as monochromatic interferometry discussed above. By observing which wavelengths these are, the absolute thickness of the sheet at that point can be determined. This was used as a complement to the monochromatic interferometry described above, as it can provide us with absolute reference points to anchor

the thickness profile given by monochromatic interferometry.

Peaks in the white light spectrum indicate wavelengths of light for which their optical path difference is equal to a half integer number of wavelengths

$$D_{\text{opt}} = \frac{2m+1}{2}\lambda. \tag{2.39}$$

Two adjacent peaks will therefore differ in the number of wavelengths of light their optical path difference corresponds to by one. This implies that

$$\frac{D_{\text{opt},1}}{\lambda_1} - \frac{D_{\text{opt},2}}{\lambda_2} = 1. \tag{2.40}$$

Using Equation 2.33 and solving for h (which has to be the same in both cases) we get an equation for determining the thickness of the sheet [59]:

$$h = [\frac{2(n_1^2 - \sin^2\alpha)^{\frac{1}{2}}}{\lambda_1} - \frac{2(n_2^2 - \sin^2\alpha)^{\frac{1}{2}}}{\lambda_2}]^{-1}. \tag{2.41}$$

Therefore all we have to do to to determine the thickness of the sheet at a point is to determine the peak wavelengths in the white light interference spectrum and enter them into the equation. A white light interference spectrum is shown in Figure 2.11. As can be seen, not all the pairs of peak wavelengths corresponded to the same thickness of sheet. This could be for many reasons — the spectrum is noisy, so determining the exact peak wavelength is hard, some peaks were better resolved than others and there may have been some other optical distortion effects not taken into account in this analysis. Therefore in order to obtain the thickness value that fit the data the best the measured positions of all the peaks were input into an optimisation routine, which then found the thickness value that predicted peaks as close to these measured peaks as possible. The error on this measurement was taken to be the difference between the highest possible and lowest possible thickness from any pair of peaks in the spectrum. The measured white light thicknesses with errors are overlaid on the thickness curve found by monochromatic interferometry in Figure 2.11.

2.4.3 Sheet observations

As can be seen in Figure 2.11, the thickness profile of our sheet obeys the $\frac{1}{r}$ dependency, and at its thinnest point is around 1 µm thick. However, it is not just the thickness of the sheet that matters. The profile of the surface can also matter — if the surface is not flat relative to the wavelength of light incident upon it then this can lead to spatial distortions in the signals generated. Happily, relative thickness profiles across the jet can be obtained from the monochromatic interferograms in much the same way as the the profile down the center of the jet was obtained (see Figure 2.12). What we can see from this is that the jet, particularly the lower regions

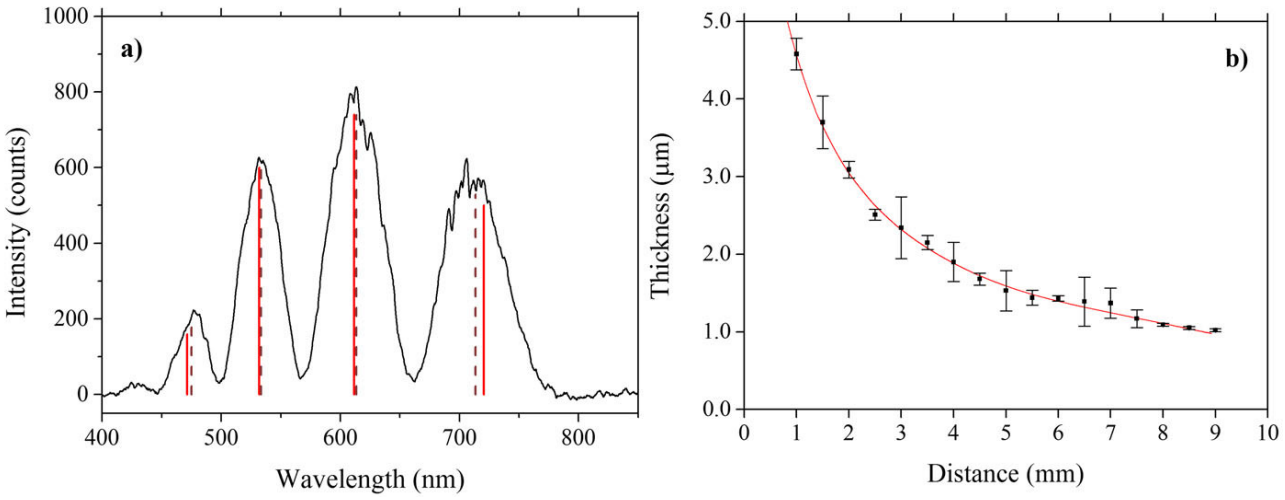

Figure 2.11: Figure 7 from Galinis *et al.* [73]. Panel a) shows the white light interference spectrum at a point halfway down the jet. The dotted gray lines are the positions of the peaks found by our peakfinding routine, and the solid red lines are the peaks predicted by the thickness value found by our optimisation. Although the match is not perfect it is fairly good, and gives a reasonable confidence in the thickness of this point. Panel b) shows the relative thickness slope overlaid on top of the points found by white light interferometry — the two slopes agree very well, and confirm a $\frac{1}{r}$ dependence for the sheet thickness.

of the jet, is optically flat (defined as variations in surface thickness of $< \frac{\lambda}{20}$) at the wavelengths that we are likely to be using over a region of 100 μm diameter. The thickness of the rims of the sheet is not included on this, as the large variation in thickness due to the rims leads to the interference pattern not being visible, and as such we don't have a good measurement of the thickness of the rims.

The two interferograms shown in Figure 2.12 were taken at two different gas pressures (for more on how the liquid is pressurized and sent through the nozzle, see Section 4.5.1). The pressures were 2 and 3 bar, which corresponded to flow rates of 7.2 and 9.1 mL/min. As the pressure increases, the velocity of the fluid through the nozzle increases, which correspondingly increases the Weber number of the flow in the sheet. This then counteracts the surface tension, leading to an increase in the size of the sheet produced (consistent with the predictions in section 2.2.3).

Another feature to note is the ridges on the surface of the sheet, especially on the larger one. These are anti-symmetrical waves of the type explored in section 2.2.3. As the jet velocity increases, the Reynolds number also increases along with the Weber number. This means that the viscosity of the fluid is relatively less important, and so surface waves like this are not damped out as readily, and instead form this standing wave structure on the surface of the jet. As can be seen by the thickness profile measurements, this standing wave does not affect the surface profile enough to disrupt the optical flatness.

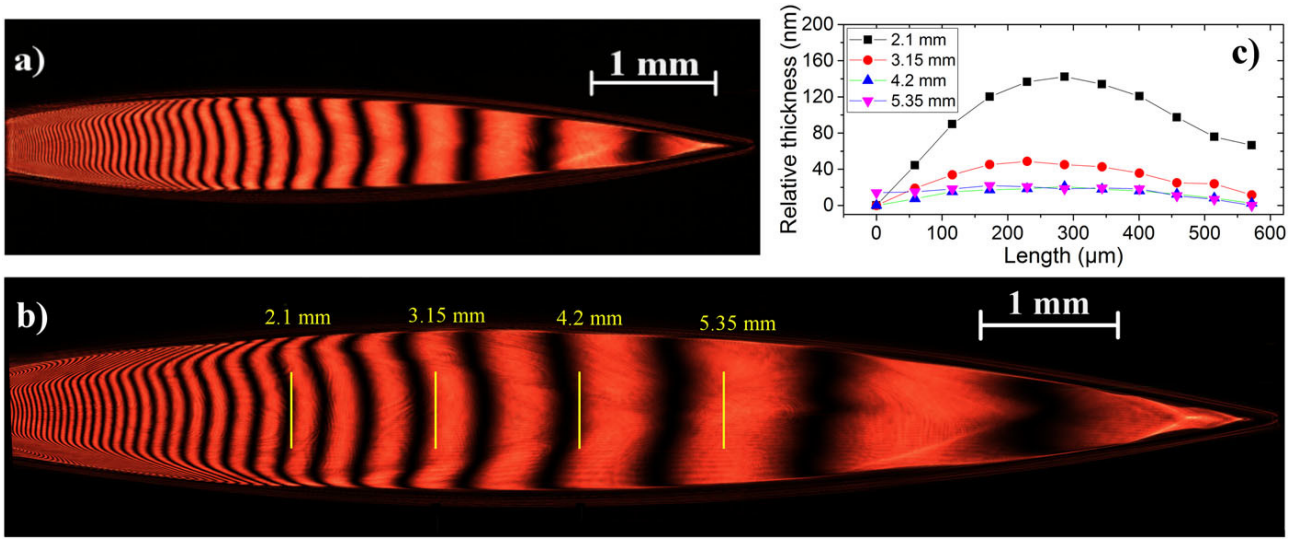

Figure 2.12: Figure 6 from Galinis *et al.* [73]. Two interferograms taken at pressures of 2 bar (7.2 mL/min, panel a)) and 3 bar (9.1 mL/min, panel b)), shown to scale, showing the increase of the size of the jet with increasing velocity (and so Weber number). Panel c) shows the horizontal thickness profiles taken at the points indicated by yellow lines — note that even when these cross one of the surface waves the change in thickness is not significant.

2.4.4 Variation of jet dimensions with liquid and nozzle properties

The fan spray nozzle was developed initially using isopropanol as the liquid, due to early experiments that showed it was much easier to create a large thin sheet in an alcohol than in water (due to the difference in surface tension of the two liquids, which will be covered below). Having produced this sheet however, there was a question about whether other liquids (especially water), could be used with this nozzle to make a stable sheet.

Runge and Rosenberg [72] noticed in 1972 that only certain liquids that they ran through their nozzle (which was roughly a fan spray nozzle - the construction is a little odd, but the general principle should be the same) would produce a stable sheet. They linked this to viscosity (dynamic rather than kinematic), observing that below $\mu = 1 \times 10^{-2}$ Pa·s, instead of a stable flat sheet, they obtained one that was destabilised by surface waves. Unlike the surface waves observed in section 2.4.3 above, these waves were non-stationary and caused a noticeable deterioration in the surface quality of the sheet.

We have noticed a similar liquid dependence in our jet, most especially in the case of water. It was very difficult to get a stable sheet of water out of our jet, and sheets when they did form tended to be fairly small. The viscosity of water is fairly low, around 0.89×10^{-3} Pa·s at 20°C, but isopropanol and ethanol form stable sheets with our jet, and their viscosities are 2.3×10^{-3} Pa·s and 1.2×10^{-3} Pa·s at 20°C respectively - certainly much closer to that of water than that of ethylene glycol, which Runge and Rosenberg required. The more likely explanation is that Runge and Rosenberg's nozzle (which was made out of a piece of pipe with holes cut in it) had some surface roughness that they were unable to remove, which was what was causing the surface waves. Our nozzle is smooth (and not formed by machining), which means we do not have as much of a problem with

this effect.

However, what we do have a problem with is surface tension. Water's surface tension at 20°C is $72.8 \times 10^{-3}\,\mathrm{Nm}^{-1}$, whilst that of the two alcohols is $22.1 \times 10^{-3}\,\mathrm{Nm}^{-1}$ for ethanol and $23.0 \times 10^{-3}\,\mathrm{Nm}^{-1}$ for isopropanol — around a quarter of that of water. Thus the Weber number for the sheet will be roughly four times smaller for water than for the alcohols, requiring more velocity to be imparted to the liquid to beat out the surface tension to form a sheet. However, at these high velocities the Reynolds number crosses a threshold, and the viscosity of the liquid is unable to damp out capillary waves in the sheet, causing the rims to become unstable as in section 2.2.4. This instability is not noticeable by eye, but showed up as a blurring of the monochromatic interferometry lines in the water jet. The main issue that we have had with it is that the droplets from the rims made the jet harder to run in vacuum (see Chapter 5).

The other question that occurred to us was 'how does the nozzle affect the sheet dimensions'? To this end we obtained some gas dynamic virtual nozzle (GDVN) chips from Micronit (descibed in [55]). These nozzles consist of three channels, one normal to the bottom of the nozzle, and two at 45° (see Figure 2.13). The normal mode of operation is to run gas through the two outside channels and liquid through the central one, thus transferring the momentum from the gas to the liquid, decoupling flow rate from liquid velocity in the sheet and enabling very thin sheets with low flow rates. However it is also possible to run it by flowing liquid through the two outside channels, creating a more conventional colliding jets setup, which is what we were doing.

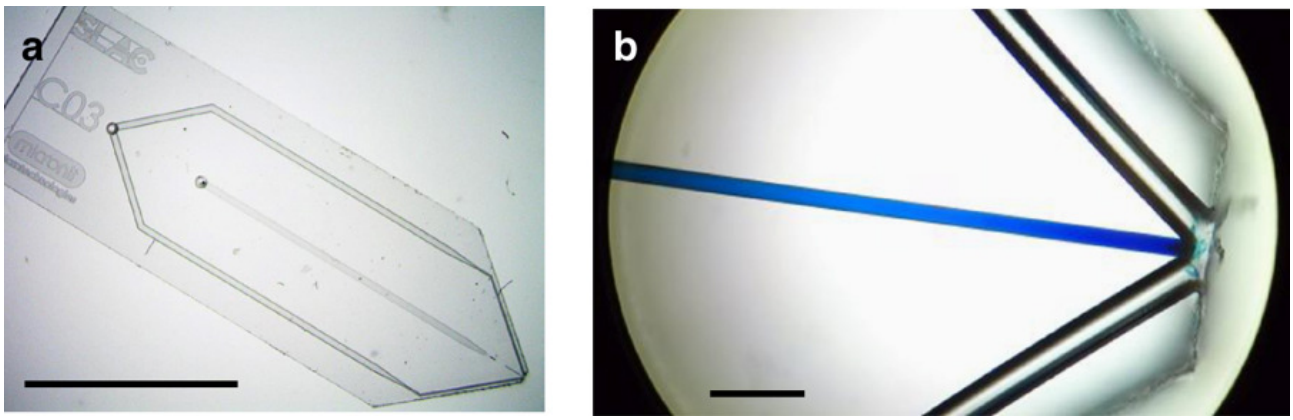

Figure 2.13: Figure 1 from Koralek *et al.* [55]. The channels are engraved onto a borosillicate glass chip by photolithography. In panel b) the liquid channel is coloured in blue, whilst the gas channels are left clear.

Water was pumped through the chip at varying flow rates, from two to four mL/min, the results of which are shown in Figure 2.14. As the flow rate increases the cohesion of the jet decreases, but the thickness also decreases, whilst the size of the sheet increases. At the lowest flow rates several links in the fluid chain below the first sheet are clearly visible. Then, as the main sheet elongates and thins out the subsequent links do not form as well, as the fluid has sufficient energy to cause it to break up into droplets at the end of the first sheet. Then as the flow rate is increased even further the main sheet elongates and thins, until it crosses a threshold

and breaks apart into several streams.

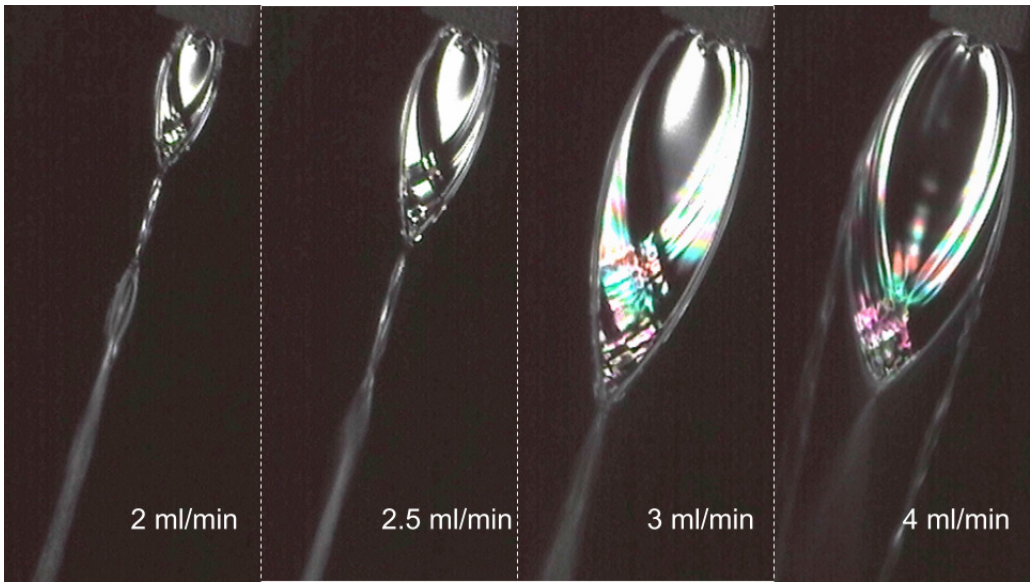

Figure 2.14: A coarse scan of flow rates through the GDVN in colliding jets mode. As the flow rate increases the jet grows in size before splitting apart, and thins out (indicated by the bands of colour at 3 and 4 mL/min — this is jet wide white light interferometry in which only one wavelength interferes constructively, indicating that the jet at this point is of comparable thickness to one or two wavelengths of visible light. The images have been brightened a bit as the originals were too dark to see the spray properly.

The sheet produced is much shorter and wider than that produced by the fan spray nozzle, probably because there is not as much guidance for the liquid, so it can spread out around the collision point to a greater extent. The other major difference is that the sheet disintegrates at much lower flow rates — with water in the fan spray nozzle we were able to increase the flow rate up to 7 mL/min before the sheet started showing signs of instability, whereas with this type of nozzle instability set in at half of that value. This is probably due to the geometry of the fan spray nozzle ensuring a much more laminar flow, although other instabilities may be present due to the geometry.

In order to check this the crossover point was examined more closely, shown in Figure 2.15. As the flow rate increases, more waves are visible on the surface of the sheet and the sheet starts to twitch at the bottom, as it is destabilised by the kinetic energy of the waves running across its surface. As the flow rate increases further a small hole opens up at the tip of the sheet, which widens and widens, until the breakup of the sheet is complete. Interestingly, at 4 mL/min it looks like the rims are continuing on in a straight line without the rest of the sheet to hold them together. This supports the idea that this is due to increase in the Reynolds number, rather than the Weber number, because if the velocity was dominating over surface tension then we would expect to see a uniform transition into spray caused by the Kelvin-Helmholtz waves, whereas if the rims were breaking up due to surface waves driving the Rayleigh-Plateau instability then we would expect the rims to effectively carry on straight as streams of droplets.

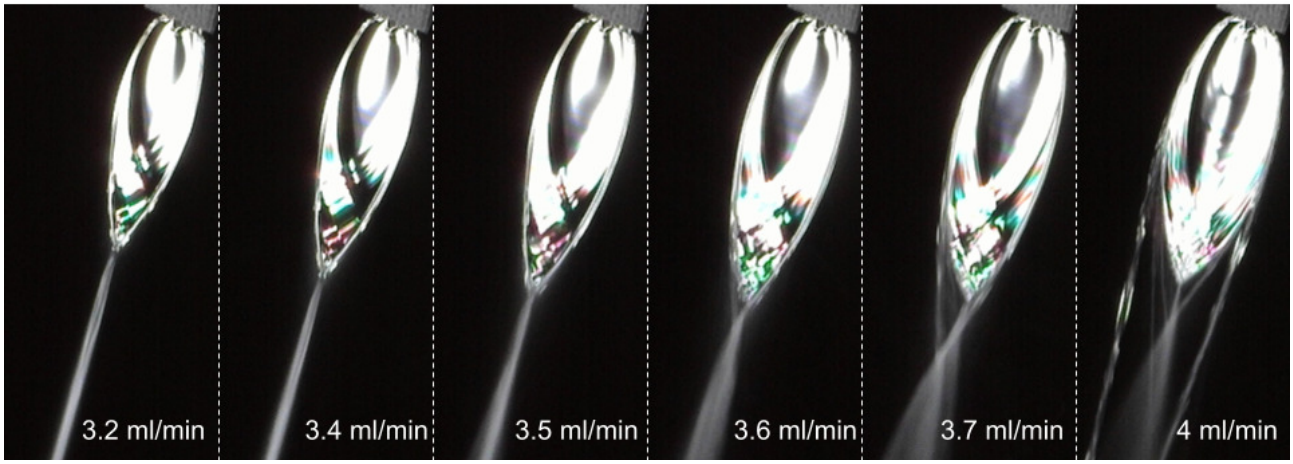

Figure 2.15: A finer scan of flow rates. The jet was starting to twitch at 3.5 mL/min, and we can see the tip start to split apart as the flow rate is increased to 3.7 mL/min.

This breakup condition of the jet looks a little bit like the fishbone structure that Bush and Hasha observed [64] (see section 2.2.4). As they note there, it is possible to create this kind of instability by driving the sheet with an acoustic wave (as per the beautiful photo in Van Dyke's Album of fluid motion, p. 87 [107]). It is possible that this was what was occurring as these jets were created using a high pressure liquid chromatography (HPLC) pump with no vibration damping (see section 4.5.1), although then the question arises as to why it was only above a certain flow rate that it seemed to be affected. In addition similar behaviour was observed using a gas driven system and a fan spray nozzle, which makes vibrations in the system an unlikely cause.

Jung *et al.* [101] observed that the fishbone instability occurs when the flow rates in a colliding jet system are slightly asymmetric, leading to a deformation of the sheet that causes these large oscillations, destabilising the jet. It is possible that some small asymmetry in the way that we have mounted the chip has caused a differential in the flow rate, leading to the breakup shown. The surface waves on the sheet are much more prominent than on the sheets produced by the fan spray nozzle, which is another indication that there may be another factor involved in the breakup of these sheets compared to what would happen for an 'ideal' colliding jets system.

2.5 Comparison of theory with the observed liquid jets

The reason for considering the fluid mechanics of sheet formation in the detail I have done in this chapter was because I wanted to understand the effect that the properties of the liquid (particularly the viscosity and surface tension) would have on the sheet created by our nozzle. As mentioned above, we believed that it was not possible to create a stable sheet of water that was large and thin enough to be useful as a target. It turned out that the instabilities caused by the easy breakup of the rims were only significant to the experiment in a secondary sense (droplets freezing, see Chapter 5 for more details) and so no changes to the nozzle were necessary, but the question remains — how well does the analysis presented above correspond to the properties

of the sheets produced?

The data presented in this section was obtained from images taken with either a Sumix MX-7 CCD camera (the fan spray nozzle images) or a Dinolite microscope camera (the GDVN chip images), and the relevant distances were measured using ImageJ image processing software. The measured lengths in pixels were converted to mm by comparison with the measured length of an object of known length in the image (a screw head for the fan spray nozzle and the GDVN chip itself respectively).

2.5.1 Dimensionless number regions

According to Huang [96] there are three regions of sheet formation determined by the Reynolds number of the jet — from the first formation of a sheet (around 100 or so) to around 500 is a stable sheet region, with no visible waves on the surface. From 500 to 2000 is the region they identified as a transition region, between a stable sheet and an unstable one. In this region cardioid waves increase more and more in visibility, and the extent of the sheet swaps over between the unstable and stable radii identified by Clanet and Villermaux [62, 98]. Above 2000 the Kelvin-Helmholtz instability takes over and the sheet starts to flap violently.

In Bush and Hasha's [64] parametric study the regions are not so easy to define. They have a different definition of the Reynolds and Weber numbers, depending on the total flux in the sheet, and using the radius of the colliding jets rather than the diameter. From their numbers the transition between stable and unstable sheets occurs at a Weber number of around 5000–6000, and the fishbones occur in a region of around 500-4000 Weber numbers and 300–1200 Reynolds numbers.

2.5.2 Determining the dimensionless numbers of the jet

The normal definition of the characteristic length for a flow through a circular pipe is the diameter of the pipe. In order to generate comparable measures for non-circular pipes a concept called the hydraulic diameter is used. This is equal to four times the area of the pipe, divided by the total wetted perimiter, that is the total amount of the pipes surface that the liquid is in contact with. For a rectangular pipe this is

$$d_{\mathrm{h}} = \frac{2ab}{a+b}, \tag{2.42}$$

where a and b are the long and short sides of the rectangle respectively. In the case (such as with the final section of the nozzle) that the rectangle is open at both ends, so that it is just two surfaces then the wetted perimeter changes to become just twice the long section, *i.e.*

$$d_{\mathrm{h}} = \frac{4ab}{2b} = 2a \tag{2.43}$$

.

Occasionally, when the pipe is sufficiently short the correct characteristic length to use is the length of the pipe, rather than its diameter. In this case the length of the section is actually very close to both hydraulic diameters (50 µm long vs 54 µm and 60 µm for the first and second hydraulic diameters respectively). For the purposes of comparison with the existing literature (all of which use the diameter or radius of the colliding jets or fan spray nozzle as the characteristic length) I will use the hydraulic diameter without sides, as that is the best representation of the actual nozzle.

In Figure 2.15 we can see a stable water sheet form, and then split apart as the flow rate is increased. As explained above the geometry is colliding jets, so the characteristic length is the diameter of the jets, which is 50 µm. The Weber number range for this was from around 500 to around 800, so we are in the start of the transition regime identified by Huang, where cardioid waves start to appear. The Reynolds number range is between 1300 and 1700, with the split of the leaf coming at around 1500. This is a little lower than the 2000 Huang predicted for an unstable sheet, but it is not too far off — especially if there is another source of oscillations as speculated above. It certainly is not due to the velocity outweighing the surface tension, as the Weber numbers are fairly small. As a comparison, the Weber and Reynolds numbers for water flowing through the fan spray nozzle at 7 ml/min are 170 and 864 respectively. The comparatively higher velocities required to obtain similar flow rates with the GDVN chip are what cause the higher numbers, and this seems like a reasonable explanation for the earlier breakup of the jet.

2.5.3 Simple theory prediction of jet dimensions

Equation 2.7 suggests that the thickness of the jet as it spreads out radially is proportional to $\frac{1}{r}$, with the constant of proportionality K depending on the initial conditions of the colliding jets and the angle in which the liquid is spreading. However, this clearly doesn't tell the whole story, because the jets we observe are finite in extent, and bounded by thick rims. There must therefore be another factor that determines the spatial extent of the jet.

Taylor has provided a first-order solution to this problem, by noting that when two jets collide, the movement of the flux out from the collision point follows Equation 2.27. Solutions to this equation cannot exist outside the radius

$$R = \frac{\rho\, u^2\, \mathrm{K}(\theta, \phi)}{2\sigma}, \tag{2.44}$$

(where $\mathrm{K}(\theta, \phi)$ has been used in place of Q to account for an angle dependence of the flux), because $\sin^2(\Phi)$ cannot exceed one, so therefore the maximum radius will be equal to one divided by the constants of proportionality. I have presented three approximations to $\mathrm{K}(\theta, \phi)$, and Figure 2.16 shows the results of using Equation 2.44 to predict the length of the jet and Equation 2.7 to predict the jet thickness along it's major axis, compared to measured data from the images taken of our jets running.

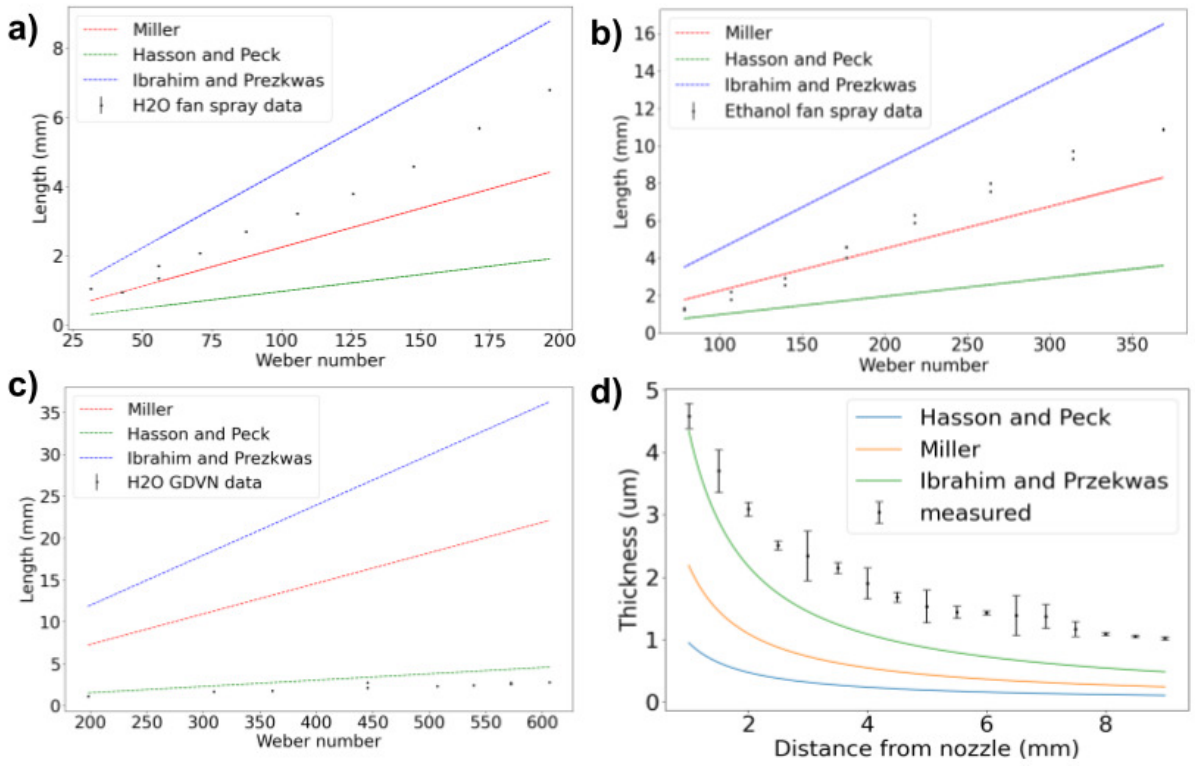

Figure 2.16: Comparison of cardioid waves analysis to measured data. Panels a–c show predicted length against the Weber number of the jet at generation in a) a water jet made by out fan spray nozzle, b) an ethanol jet made by our fan spray nozzle and c) a water jet made by the GDVN chip in colliding jets mode. Panel d) shows the comparison of the thickness predicted to the measured thickness data published in [73].

Unfortunately, not only do the predictions not agree with the measured data, they don't even agree with themselves well. This is not especially surprising, as the geometry of both of these nozzles heavily obstructs flux in the opposite direction to the desired orientation of the leaf which will have an effect on its shape not accounted for here. For instance, the fan spray nozzle obstructs flow everywhere outside its desired area, which leads to a much longer and thinner leaf to the GDVN in colliding jets mode which only obstructs flow in the area above the collision point, leading to a shorter and wider leaf.

A similar case exists for the thickness data, with all of the models underestimating the thickness of the jet, presumably because they are accounting for flux moving in directions that is is obstructed from moving in in the actual measured data. These K values are approximations in any case, but in the case of a fan spray nozzle they appear to be really inaccurate when used on their own. A more sophisticated analysis is needed to account for the observed dimensions of the jet.

2.5.4 Differential equations prediction of jet dimensions

Fortunately, such an analysis is accesible, in the form of the coupled differential equations presented in Equation 2.24. I solved these equations using the fourth order Runge-Kutta method both with and without the correction

terms proposed by Bush and Hasha for all three K models, and the results are presented in Figure 2.17. The

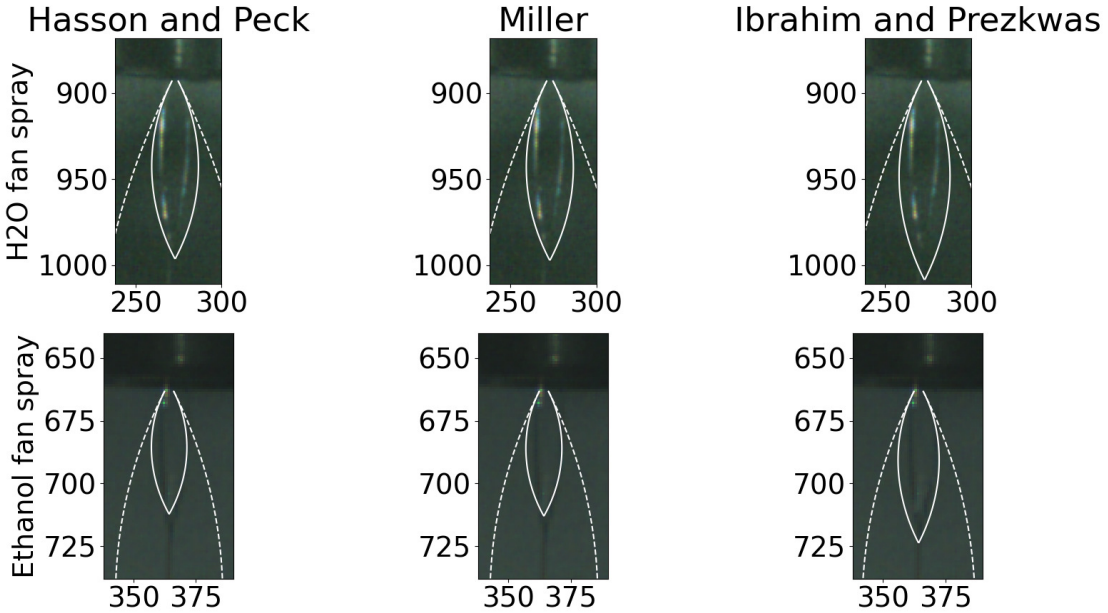

Figure 2.17: Predicted shapes of the jet with (solid line) and without (dashed line) the correction terms proposed by Bush and Hasha. The initial rim thickness was determined by measuring using ImageJ and converting pixels to mm in the same way as presented in Figure 2.16, although the resolution of the images was such that the rims were 2 and 3 pixels for ethanol and water respectively, which leads to a greater uncertainty in the results. The x- and y-axes are in pixels — due to experimental considerations we were unable to take higher resolution pictures of these jets.

differential equations with the correction terms added give a fair reconstruction of the shape of the jet. The predicted length is mostly accurate, but the predicted width is consistently too wide. Without the correction terms the prediction is hopelessly inaccurate, but one might suggest that the situation where surface tension is important is exactly the situation where the rims pull back together to form a stable sheet, and so this should not be a surprise.

Unfortunately it appears that the liquid flat jet system is a complex system, and that one can only get so far with simple models before one needs either experimental data (in the form of a measured K or initial conditions for the differential equations in Equation 2.24) or a full Navier-Stokes simulation of the system. One halfway approach could be that taken in Ha *et al.* [92], where the initial conditions of the jet leaving the nozzle are computed using full Navier-Stokes, and then the propagation from then on is done using a much less computationally expensive method.

2.6 Conclusion

We have developed a nozzle that is capable of producing micrometer thick sheets of liquid in several common solvents, and I have established a simple theoretical framework to attempt to predict how the parameters of the

nozzle will effect the sheet. This gives us control of the liquid sheet, an important element of our experiments.

We now turn to the other important experimental piece — the laser. In the next chapter I will explain the theoretical concepts used to understand the interaction of light with matter, which includes high harmonic generation and X-ray absorption spectroscopy, and then in the following chapter I will describe the laser system used and how we integrated this liquid sheet into it.

Chapter 3

Ultrafast optics and strong field optical response

3.1 Introduction

In this chapter we will cover the relevant optical concepts used in the rest of this thesis. In order to generate coherent X-ray radiation from a liquid sheet, or to track photodissociation in real time, it is necessary to use very specialised laser pulses, and therefore an understanding of optical theory is necessary to give a complete picture of what is happening over the course of these experiments.

We will first cover the linear optics required to describe a laser pulse with a phase and a polarisation before turning to nonlinear optics in general, including low order harmonic generation, optical parametric amplification and the intensity dependent refractive index. We will then cover high harmonic generation in more detail, and look at different descriptions of the mechanism of generation dependent on the medium in which the generation is happening. Finally we will cover X-ray absorption spectroscopy, and in particular how it can be applied to the field of femtochemistry to provide information rich, high time resolution measurements of systems of great chemical interest in their natural environment.

3.2 Description of a laser pulse

For the following it will be useful to have a mathematical description of a laser pulse, which will allow easy reference to the various pulse features in order to explain relevant optical phenomena. An electromagentic wave, such as light, consists of an electric and a magnetic field oscillating at mutually perpendicular angles, and both perpendicular to the direction of propagation. In the cases we are considering, the effect of the magnetic field is very much weaker than that of the electric field, so it is common to describe light as just an oscillating electric

field propagating along z, the propagation axis:

$$\mathcal{E}(t,z) = A(t,z)\cos(\omega t), \tag{3.1}$$

where $\mathcal{E}(t,z)$ is the time dependent oscillating electric field, $A(t,z)$ is the time and space dependent amplitude of the field and ω is the characteristic frequency of the oscillation. For simplicity we can split the amplitude variation $A(t,z)$ into a slowly varying envelope function $\tilde{A}(t,z)$ and a plane wave function e^{-ikz} [108], giving:

$$\mathcal{E}(t,z) = \tilde{A}(t,z)\mathrm{e}^{-ikz}\cos(\omega t), \tag{3.2}$$

where k is known as the angular wavenumber, and is equal to $\frac{\omega n(\omega)}{c}$, where n is the refractive index at frequency ω and c is the speed of light in a vacuum. For convenience the cosine term is often split into exponentials, giving rise to what is known as the analytic form [108, 109]

$$E(t,z) = \tilde{A}(t,z)\mathrm{e}^{i(\omega t - kz)} \tag{3.3}$$

where

$$\mathcal{E} = \frac{1}{2}(E(t,z) + E^*(t,z)). \tag{3.4}$$

We will now use this description of a pulse to cover two linear optical topics that are important to this thesis — phase and polarisation

3.2.1 Phase

The phase of a wave is the $(\omega t - kz)$ part of the above expression indicating the position along the optical cycle the wave is at the present moment, and is often denoted ϕ. In the complex representation above we can use Euler's formula and similar relations,

$$\mathrm{e}^{i\pi} = \mathrm{e}^{-i\pi} = -1,\ \mathrm{e}^{i2\pi} = 1,\ \mathrm{e}^{\pm\frac{\pi}{2}} = \pm i, \tag{3.5}$$

and see that this can be represented graphically on a diagram with the x-axis as the real numbers, the y-axis as the imaginary numbers and the phase as an angle (see Figure 3.1). As the phase increases from 0 to 2π the real part of the wave goes from $+A$ down to $-A$ and back again, like the cos function. The phase can be split into the spatial component kz and a temporal component ωt. The spatial phase controls the distribution in space of the pulse and the temporal phase controls the phase of the oscillations in time. The temporal phase becomes important for the generation of high order harmonics, as only certain parts of the optical cycle can produce them, but otherwise is fixed for a given frequency and so most of what follows is focused on the spatial phase.

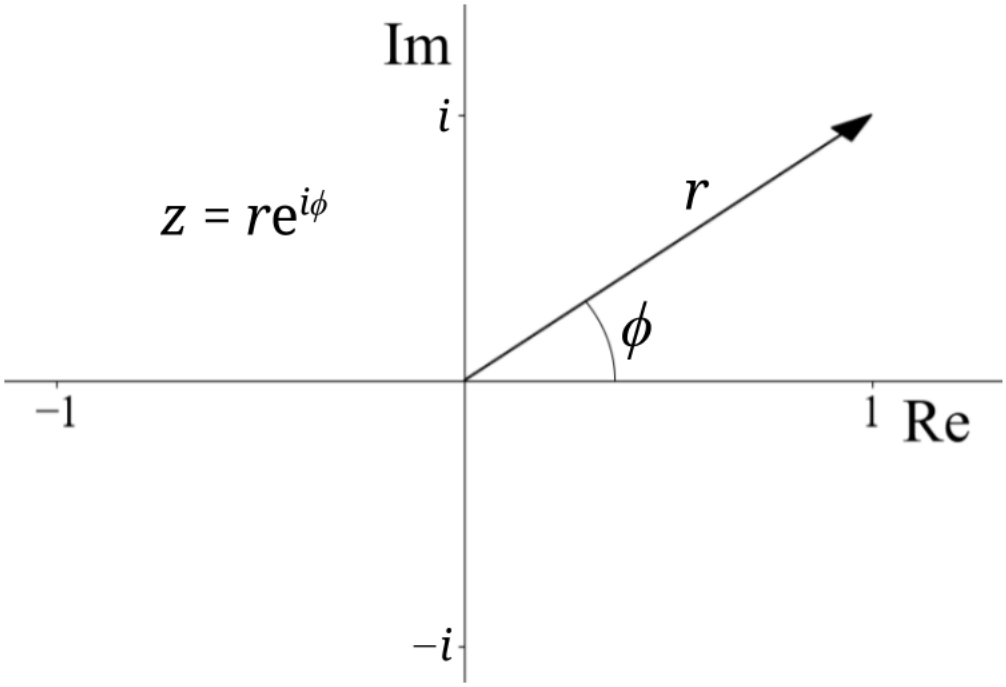

Figure 3.1: An Argand diagram. The y-axis scale is the imaginary numbers, whilst the x-axis scale is the real numbers. A complex number $z = x + iy$ can be represented on this diagram in terms of a magnitude r and an angle (or phase) ϕ as $z = re^{i\phi}$. To get back to the real part of the number we can take a projection onto the real axis of this complex vector, by adding the vector to it's complex conjugate and dividing by two (as in Equation 3.4)

This is important for many nonlinear phenomena as they rely on keeping the spatial phases of various beams as closely matched as possible to enable energy transfer from one to the other, a process know as *phasematching* (see Section 3.3.2 below).

3.2.1.1 Carrier envelope phase

When the envelope of the pulse is at a peak we can define the phase of the pulse at that point as the carrier envelope phase offset, known as the carrier envelope phase or CEP [109] (see Figure 3.2). This phase is normally not particularly important, as in a pulse longer than a few cycles of the field the peak is sufficiently flat that it is immaterial what the phase at the exact peak is. However, in pulses consisting of only one or two optical cycles it can become critically important as processes that only happen over a threshold intensity will only be able to occur for certain regions of the optical cycle, leading to measurably different results depending on which parts of the pulse pass the threshold, which is controlled by the CEP [110].

We can include the CEP in our analytical expression of the electric field fairly easily, by just including a constant phase term. Defining $\phi_0 = 0$ as the phase where the electrical field peaks at the same time as the

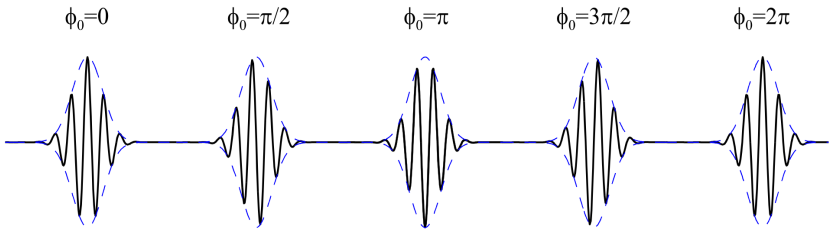

Figure 3.2: A demonstration of the carrier envelope phase, by HartmutG reproduced under the creative commons licence v.3.1. As the CEP (ϕ_0) varies between 0 and π the peak of the cycles of the electric field moves further away from the peak of the overall pulse envelope, which is the blue dashed line.

envelope, we can write the expression as

$$E(t,z) = \tilde{A}(t,z)\mathrm{e}^{i(\omega t - kz + \phi_0)}, \tag{3.6}$$

which will reproduce the behaviour shown in Figure 3.2 for the corresponding values of ϕ_0.

3.2.1.2 Phase velocity and group velocity

As a pulse propagates, the fast oscillations of the electric field can effectively move at a different velocity to the pulse envelope, causing the CEP of the pulse to shift as it propagates. The velocity of the field oscillations is called the phase velocity

$$\nu_{\text{phase}} = \frac{\omega}{k} = \frac{c}{n(\omega)}, \tag{3.7}$$

and the velocity of the envelope is called the group velocity [108]

$$\nu_{\text{group}} = \frac{d\omega}{dk}. \tag{3.8}$$

As the phase velocity is dependent on the refractive index, and the refractive index is frequency dependent, as the pulse propagates along z the different frequencies propagate at different velocities, leading to a smearing out of the pulse in time. This process is referred to as chromatic dispersion, and the smeared out pulse is referred to as chirped, as when this process happens to a sound wave it can sound like a bird chirping. The analagous process for the group velocity also occurs, but as the group velocity is the derivative of the phase velocity it is somewhat harder to visualise.

The chromatic dispersion experienced by a pulse can be expressed by taking a Taylor expansion of $k(\omega)$ around the central wavelength of the pulse ω_0 [111]:

$$k(\omega) = k_0 + \frac{dk}{d\omega}(\omega - \omega_0) + \frac{1}{2}\frac{d^2k}{d\omega^2}(\omega - \omega_0)^2 + \frac{1}{6}\frac{d^3k}{d\omega^3}(\omega - \omega_0)^3 + \dots. \tag{3.9}$$

The zeroth order term describes a common phase shift for all the frequencies. The first order term describes an overall time delay that does not lead to any change in the pulse shape. The second order term contains the group velocity dispersion, and is the term that is usually the most important one in controlling the chromatic dispersion of a pulse. By multiplying this term by the length of material that the pulse propagates through one can obtain the group delay dispersion (GDD), which can be used to estimate with good accuracy how the pulse will stretch in time upon traveling through the material. The third order term contains the third order dispersion. This term only really becomes important for very short pulses (30 fs or below [111]), but if we want to use pulses this short then we will need to provide compensation for this source of chromatic dispersion as well.

3.2.2 Polarisation

So far we have only considered an electric field oscillating along one axis. This case is called linear polarisation. However, this does not have to be the case. It is very possible to have an electric field that rotates on the xy-plane as it propagates along the z-axis, or some other more complicated scheme. To understand this within the mathematical picture we have, consider two light waves propagating along the z-axis with their electric fields along the x and y axes respectively, and a phase difference of ε [112]:

$$\begin{aligned} \vec{E_x}(t,z) &= \hat{\mathbf{i}}\tilde{A}_x \mathrm{e}^{i(\omega t - kz)}, \\ \vec{E_y}(t,z) &= \hat{\mathbf{j}}\tilde{A}_y \mathrm{e}^{i(\omega t - kz + \varepsilon)}. \end{aligned} \tag{3.10}$$

Where $\hat{\mathbf{i}}$ and $\hat{\mathbf{j}}$ are the unit vectors along the x- and y-axes respectively. According to the principle of superposition the resultant wave from the overlap of these two waves is the vector sum of the two [112]:

$$\vec{E}(t,z) = \vec{E_x}(t,z) + \vec{E_y}(t,z) \tag{3.11}$$

i.e.

$$\vec{E}(t,z) = (\hat{\mathbf{i}}\tilde{A}_x + \hat{\mathbf{j}}\tilde{A}_y \mathrm{e}^{i\varepsilon})\mathrm{e}^{i(\omega t - kz)}. \tag{3.12}$$

The polarisation of the resultant electric field depends on the value of ε. If ε is 0 or a multiple of π then the phase multiplying the y-component is entirely real, and the electric field is linearly polarised, at an angle θ to the x-axis given by trigonometry as

$$\tan(\theta) = \frac{\tilde{A}_y}{\tilde{A}_x} = \frac{\tilde{A}\sin(\theta)}{\tilde{A}\cos(\theta)}. \tag{3.13}$$

If ε is a multiple of $\frac{\pi}{2}$ then the phase multiplying the y-component is entirely imaginary ($\pm i$). To understand what this means we can use Equation 3.4 to turn the analytic form into the real electric field. Gathering terms

along the x- and y-axes and converting from exponential into sine and cosine form gives

$$\mathcal{E} = \hat{\mathbf{i}}\tilde{A_x}\cos(\omega t - kz) - \hat{\mathbf{j}}\tilde{A_y}\sin(\omega t - kz). \tag{3.14}$$

In other words the polarisation is time dependent, and rotates in the xy-plane as the field propagates in time and space. In the case where $\tilde{A_x} = \tilde{A_y}$ this is circular polarisation as the peak of the field traces a circle, otherwise it is some form of elliptical polarisation. When ε is not equal to one of the two cases above then the added phase is complex, and a linear combination of the two results, which is an elliptical polarisation (the electric field traces out an ellipse in the xy-plane).

3.2.2.1 Waveplates and polarisers

Waveplates and polarisers are optics that manipulate the polarization of the light waves that pass through them. A waveplate is a plate made up of one or more birefringent crystals. It has two axes, perpendicular to each other and also to the intended laser propagation direction. Light travels along one of these axes (the *fast axis*) faster than along the other one (the *slow axis*), which means that the components of the light pulse oriented along these axes emerge from the plate with a different phase with respect to each other than they had at the beginning [113]. Assuming them to have been in phase at the beginning, we can treat the phase added as ε above. Two types of waveplate are used in this thesis, a half waveplate and a quarter waveplate.

The half waveplate retards the slow axis by an integer multiple of $\frac{\lambda}{2}$, thus introducing a phase shift of $\mathrm{e}^{i\pi}$. This leads to a final wave of the form

$$\vec{E}(t,z) = (\hat{\mathbf{f}}\tilde{A}\cos(\theta) - \hat{\mathbf{s}}\tilde{A}\sin(\theta))\mathrm{e}^{i(\omega t - kz)}, \tag{3.15}$$

where $\hat{\mathbf{f}}$ and $\hat{\mathbf{s}}$ stand for the components along the fast and slow axes, and the beam is at an angle θ to the fast axis. This corresponds to linearly polarised light with the amplitude along the slow axis negated. Using Equation 3.13 and the fact that $-\tan(\theta) = \tan(-\theta)$ we can see that the effect of the half waveplate is to rotate the linear polarisation of light from an angle of θ to the fast axis to an angle of $-\theta$.

The quarter waveplate retards the slow axis by an integer multiple of $\frac{\lambda}{4}$, which corresponds to a phase shift of $\mathrm{e}^{i\frac{\pi}{2}}$. This leads to a final wave of the form

$$\vec{E}(t,z) = (\hat{\mathbf{f}}\tilde{A}\cos(\theta) + i\hat{\mathbf{s}}\tilde{A}\sin(\theta))\mathrm{e}^{i(\omega t - kz)}. \tag{3.16}$$

If $\theta = 45°$, $\tilde{A}\cos(\theta) = \tilde{A}\sin(\theta)$ and this corresponds to circularly polarised light. If not then the electric field instead traces an ellipse, who's degree of ellipticity depends on θ and thus the relative sizes of the amplitudes along the axes. Thus if linearly polarised light is transmitted through the quarter waveplate it is made more

and more elliptical as the angle to the fast axis increases towards 45° at which point it is circular.

A polariser is a device that only lets through light that is polarised along one axis. There are several methods of achieving this, which don't concern us here. The important thing is to consider what happens to linearly polarised light when it passes through a polariser that only admits light that is polarised along the x-axis. Our initial wave again has the form

$$\vec{E}(t,z) = (\hat{\mathbf{i}}\tilde{A}\cos(\theta) + \hat{\mathbf{j}}\tilde{A}\sin(\theta))\mathrm{e}^{i(\omega t - kz)}, \tag{3.17}$$

where the overall amplitude of the wave is $\tilde{A}$, and the polarisation of the light is at an angle of θ to the x-axis. We then lose the component of the electric field that is along the y-axis (in reality no polariser is perfectly efficient, but this doesn't affect the point here). We are then left with a wave of the form

$$\vec{E}(t,z) = \hat{\mathbf{i}}\tilde{A}\cos(\theta)\mathrm{e}^{i(\omega t - kz)}, \tag{3.18}$$

i.e. the electric field that passes through has an amplitude of $\tilde{A}\cos(\theta)$. We can square this to get the intensity, which gives Malus' law

$$I(\theta) \propto I_0 \cos^2(\theta). \tag{3.19}$$

Therefore we can use a polariser to control the intensity of the light that passes through itself by rotating the linear polarisation of this light.

3.3 Nonlinear Optics

Naturally occurring interactions of light and matter are in general linear (the response produced is proportional to the strength of the input). It is only when strong enough perturbations are introduced that the response becomes nonlinear (being proportional to the input raised to some power other than one, or a combination of such power terms) [108]. As these pertubations require strong electric fields or very intense beams of light optics proceeded for a long time without awareness of these phenomena. The Kerr and Pockels effects (discovered in 1875 [114] and 1894 [115] respectively) notwithstanding, nonlinear optics became generally important with the invention of the laser in 1960 [116], with the first nonlinear optics experiment using a laser published one year later, in 1961 [117].

In linear optics the reaction of a given medium to the electric field of a light wave is assumed to be [108]:

$$P = \epsilon_0 \chi^{(1)} \mathcal{E}, \tag{3.20}$$

where P is the polarisation of the medium (dipole moment per meter squared), ϵ_0 is the permittivity of free

space (a constant), $\mathcal{E}$ is the electric field of the light as discussed above and $\chi^{(1)}$ is the linear susceptibility of the medium, a measure of how much the system responds to a given light field. However, when the amplitude of the electric field is sufficiently large or there is an external field present then a deviation from linearity can be observed. In most cases this can be understood in terms of small perturbations to the linear polarisation of the medium, represented by [108]:

$$P = \epsilon_0(\chi^{(1)}\mathcal{E} + \chi^{(2)}\mathcal{E}^2 + \chi^{(3)}\mathcal{E}^3 + \ldots), \tag{3.21}$$

where $\chi^{(i)}$ is the nonlinear susceptibility associated with the $\mathcal{E}^i$ term, and the series can continue on to infinity, although usually only the first three terms are needed as further nonlinear susceptibility terms become vanishingly small.

There is a point to note here about the symmetry of this equation. In Figure 3.3 it can be seen that the second order terms distort the linear polarisation in one direction on both ends of the scale, but the third order terms distort in opposite directions. This means that the second order term produces a different effect depending on the sign of the electric field (positive is an increase, negative is a decrease in amplitude). This is only possible in a material which has some difference which would make it possible to tell whether the field was positive or negative — in other words the material cannot possess inversion symmetry. Most of the materials for which this is true are highly ordered crystals, which will be covered in section 3.3.3 below. Another way of thinking about this, from [108], is to imagine the perturbation of the material due to an nth order process, $P = \epsilon_0\chi^{(n)}\mathcal{E}^n$. Under inversion symmetry when $\mathcal{E}$ changes sign P must as well. But when n is even it does not, so under inversion symmetry these terms must be 0. The point to keep in mind is that this means that second order (and fourth, sixth, etc.) processes can only take place if there is some means of breaking the inversion symmetry of the system present.

According to Equation 3.4 we can write the electric field as the sum of plane waves like so

$$\mathcal{E} = \frac{1}{2}\sum_n \tilde{A}_n \mathrm{e}^{i(\omega_n t - k_n z)} + \text{c.c.}. \tag{3.22}$$

By analogy we can write the polarisation due to this wave as a sum of plane waves as well

$$P = \frac{1}{2}\sum_n \tilde{P}_n \mathrm{e}^{i(\omega_n t - k_n z)} + \text{c.c.}. \tag{3.23}$$

According to Equation 3.21 above we can split the polarisation into linear contributions and nonlinear contributions

$$P = P^L + P^{NL}, \tag{3.24}$$

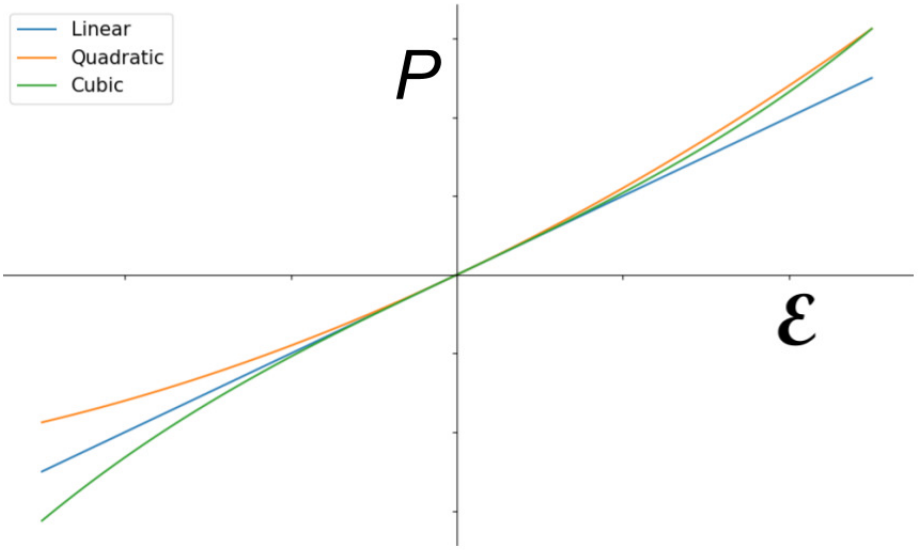

Figure 3.3: A diagram plotting the polarization of the medium against the electric field strength. The three lines are just the linear response, the linear response plus the quadratic response and the linear response plus the cubic response. Note that the cubic response has inversion symmetry through the origin, whereas the quadratic response does not.

with which we can write the wave equation for the electric wave travelling in the z direction as [108]

$$-\frac{\partial^2 \mathcal{E}}{\partial z^2} + \frac{1}{c^2}\frac{\partial^2 E}{\partial t^2} = -\mu_0 \frac{\partial^2 (P^L + P^{NL})}{\partial t^2}. \tag{3.25}$$

We can then use the fact that the linear polarisation envelope is related to the electric field by $\tilde{P}_n^L = \epsilon_0 \chi^{(1)} \tilde{A}_n$, considering only the analytic form of the field and using the slowly varying envelope approximation to remove the $\frac{\partial^2 \tilde{A}_n}{\partial z^2}$ term from the expansion of $\frac{\partial^2 \mathcal{E}}{\partial z^2}$ term to get an expression for the relation of the envelope of the electric field of frequency ω_n to that of the nonlinear polarisation at frequency ω_n

$$2ik_n \frac{\partial \tilde{A}_n}{\partial z} + (k_n^2 - \frac{\omega_n^2}{c^2}(1 - \chi^{(1)}))\tilde{A}_n = \frac{\omega_n^2}{c^2 \epsilon_0}\tilde{P}_n^{NL}, \tag{3.26}$$

where $c^2 = \frac{1}{\mu_0 \epsilon_0}$ has been used.This can be simplified by setting $k_n = \frac{\omega_n}{c}\sqrt{1 + \chi^{(1)}} = \frac{\omega_n n_n}{c}$ (as $\sqrt{1 + \chi^{(1)}} = n_n$, the refractive index at frequency ω_n), so that the middle terms cancel, leaving us with

$$\frac{\partial \tilde{A}_n}{\partial z} = -i\frac{\omega_n^2}{2c^2 n_2 \epsilon_0}\tilde{P}_n^{NL}. \tag{3.27}$$

Using this equation, we can thus describe the change of the field at ω_n by finding the components of the polarisation at this frquency.

3.3.1 Second order nonlinear optics

To show how this can work the second order polarisation term due to a monochromatic field may be found by substituting the equation

$$\mathcal{E} = \frac{1}{2}(A(t,z)\mathrm{e}^{i\omega t} + \text{c.c.}) \tag{3.28}$$

into the term in $\chi^{(2)}$ in Equation 3.21:

$$P^{(2)} = \epsilon_0\chi^{(2)}\mathcal{E}^2 = \frac{1}{2}\epsilon_0\chi^{(2)}A(t,z)A^*(t,z) + \frac{1}{4}\epsilon_0\chi^{(2)}(A^2(t,z)\mathrm{e}^{2i\omega t} + \text{c.c}). \tag{3.29}$$

This equation describes the generation of a DC (non-oscillating) electric field — a process known as optical rectification — and the generation of a second wavelength of light of frequency 2ω (the second harmonic). In fact nonlinear processes are required to convert one frequency of light to another, and it is in this capacity that they are mostly used in our lab. This treatment of second order nonlinear effects can be generalised to reflect this and include the effects of multiple different wavelengths mixing in a nonlinear medium. For this we will use not two but three electric fields with frequencies related by $\omega_1 + \omega_2 = \omega_3$, for reasons that will become clear soon. We start with the field

$$\mathcal{E} = \frac{1}{2}(A_1(t,z)\mathrm{e}^{i\omega_1 t} + A_2(t,z)\mathrm{e}^{i\omega_2 t} + A_3(t,z)\mathrm{e}^{i\omega_3 t} + \text{c.c.}), \tag{3.30}$$

and place this into the second order polarisation term as in Equation 3.29 above. Among the many terms that arise from doing this we can identify those related by the frequency relation above, and disregard the rest as not being anywhere close to phasematching [108] (see Section 3.3.2 below). This gives us three nonlinear polarisations, one corresponding to each frequency

$$\begin{aligned}
P_1^{(2)} &= \frac{1}{2}\epsilon_0\chi^{(2)}(A_3(t,z)A_2^*(t,z)\mathrm{e}^{i\omega_1 t} + \text{c.c}) \\
P_2^{(2)} &= \frac{1}{2}\epsilon_0\chi^{(2)}(A_3(t,z)A_1^*(t,z)\mathrm{e}^{i\omega_2 t)} + \text{c.c}) \\
P_3^{(2)} &= \frac{1}{2}\epsilon_0\chi^{(2)}(A_1(t,z)A_2(t,z)\mathrm{e}^{i\omega_3 t} + \text{c.c}).
\end{aligned} \tag{3.31}$$

We can then separate out the spatial component of these equations into the envelope and the plane wave component and substitute them into Equation 3.27

$$\begin{aligned}
\frac{\partial\tilde{A}_1}{\partial z} &= -i\frac{\omega_n^1}{4c^2n_1}\chi^{(2)}(\tilde{A}_3(t,z)\tilde{A}_2^*(t,z)\mathrm{e}^{-i\Delta kz}) \\
\frac{\partial\tilde{A}_2}{\partial z} &= -i\frac{\omega_n^2}{4c^2n_2}\chi^{(2)}(\tilde{A}_3(t,z)\tilde{A}_1^*(t,z)\mathrm{e}^{-i\Delta kz}) \\
\frac{\partial\tilde{A}_3}{\partial z} &= -i\frac{\omega_n^3}{4c^2n_3}\chi^{(2)}(\tilde{A}_1(t,z)\tilde{A}_2(t,z)\mathrm{e}^{i\Delta kz}),
\end{aligned} \tag{3.32}$$

where $\Delta kz = k_3 - k_2 - k_1$, the phase mismatch parameter. This leads us into a discussion of phasematching. Why does second harmonic generation not occur whenever the intensity is high enough? Even in non-centrosymmetric crystals this does not automatically happen. The answer depends on this phase mismatch parameter.

3.3.2 Phasematching

As it sounds, the phase mismatch parameter is effectively a measure of how well matched the spatial phases of the different harmonics are. In order to illustrate this, let us consider the case of second harmonic generation. In terms of the three frequencies mentioned above, in this case

$$\begin{aligned} \omega_1 = \omega_2 = \omega_\mathrm{f}, \\ \omega_3 = 2\omega_\mathrm{f} = \omega_\mathrm{h}. \end{aligned} \tag{3.33}$$

The coupled wave equations above thus read

$$\begin{aligned} \frac{\partial \tilde{A}_\mathrm{f}}{\partial z} &= -i\frac{\omega_\mathrm{f}}{2cn_\mathrm{f}}\chi^{(2)}(\tilde{A}_\mathrm{h}(t,z)\tilde{A}_\mathrm{f}^*(t,z)\mathrm{e}^{-i\Delta kz} \\ \frac{\partial \tilde{A}_\mathrm{h}}{\partial z} &= -i\frac{\omega_\mathrm{h}}{4cn_\mathrm{h}}\chi^{(2)}\tilde{A}_\mathrm{f}{}^2(t,z)\mathrm{e}^{i\Delta kz}. \end{aligned} \tag{3.34}$$

The asymmetry in the prefactors comes in this case from the fact that both ω_1 and ω_2 are included. We can assume for now that the fundamental beam will not be depleted by the second harmonic beam, which allows us to ignore the first of these equations (as we are assuming the amplitude of the fundamental does not change with z) and treat $\tilde{A}_\mathrm{f}{}^2(t,z)$ as a constant term for the same reason. The solution of the remaining equation then becomes

$$\tilde{A}_\mathrm{h} = -i\frac{\omega_n^3}{4c^2n_\mathrm{h}}\chi^{(2)}\tilde{A}_\mathrm{f}{}^2\frac{\mathrm{e}^{i\Delta kz}-1}{i\Delta k} = -i\frac{\omega_n^3}{4c^2n_\mathrm{h}}\chi^{(2)}\tilde{A}_\mathrm{f}{}^2 2\mathrm{e}^{i\Delta kz/2}\frac{\sin(\Delta kz/2)}{\Delta k}. \tag{3.35}$$

We can use the relation

$$I_n(z) = \frac{1}{2}c\varepsilon_0 n_n|\tilde{A}_n(z)|^2 \tag{3.36}$$

to write the generated harmonic intensity as

$$I_\mathrm{h}(z) = \frac{(\omega_\mathrm{h}\chi^{(2)})^2}{2c^3\varepsilon_0 n_\mathrm{h}n_\mathrm{f}^2}I_\mathrm{f}^2\frac{\sin(\Delta kz/2)}{\Delta k}. \tag{3.37}$$

From this solution we can see that the intensity of the generated second harmonic rises and falls periodically with z. We can define a characteristic length called the coherence length,

$$L_\mathrm{coh} = \frac{\pi}{\Delta k}, \tag{3.38}$$

which is the distance that the polarisation can propagate before the harmonic intensity peaks, and energy starts to be transferred back to the fundamental beam. The key to efficient harmonic generation then is reducing Δk as much as possible, so that the harmonic intensity is always increasing with z until it leaves the generation medium. Later we will see that calculating Δk can get fairly involved, as multiple frequencies and the tensor response of the medium complicate it further and that the phasematching condition can also be affected by changes in the medium, which change n, and geometric effects due to the shape of the light beam.

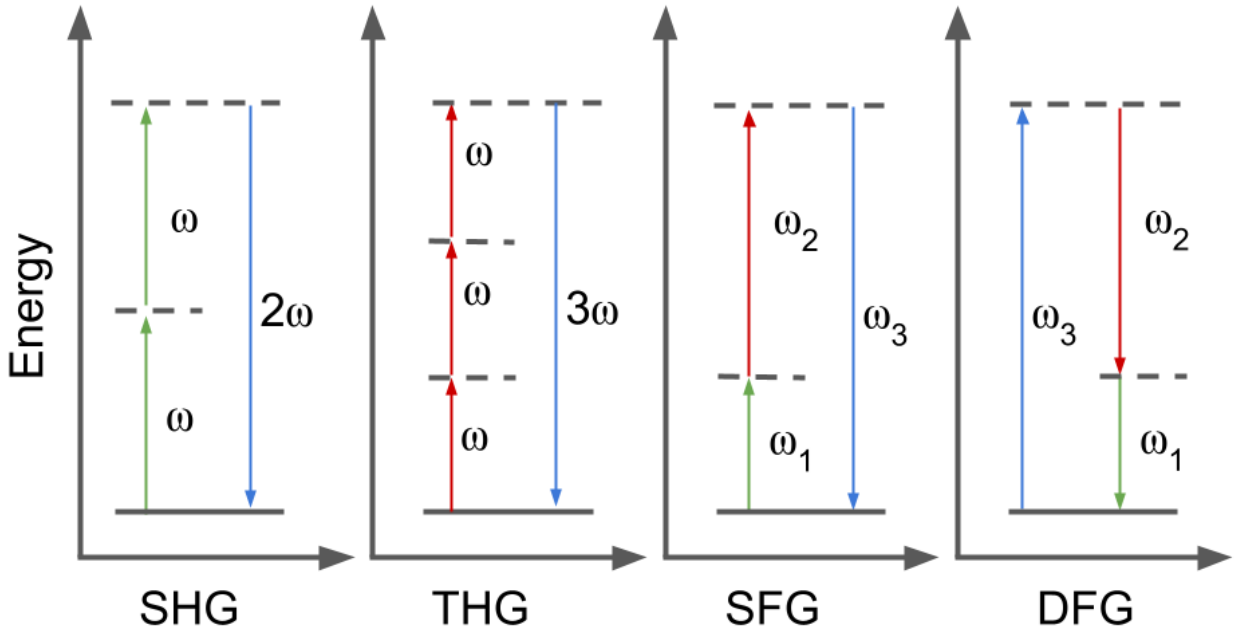

Figure 3.4: A diagram illustrating some of the harmonic generation processes covered in this section. The processes are second harmonic generation (SHG), third harmonic generation (THG), sum frequency generation (SFG) and difference frequency generation (DFG). All of these processes can be viewed as the simultaneous excitation of the medium to a virtual state (represented by the dashed lines) by various photons, followed by de-excitation and the release of a photon of frequency equal to the sum of the absorbed photons, with sum frequency generation as a general case and the nonlinear order of the process equivalent to the number of photons absorbed. The only process that deviates from this is DFG, in which one photon is absorbed and two are generated, with frequencies that sum to the absorbed photon's frequency.

3.3.3 Low order harmonic generation

The first type of nonlinear process I want to consider here is low order harmonic generation (LHG), named as such in contrast to high order harmonic generation (HHG), which is very important and is covered in Section 3.4. For the purposes of the discussion here, I define LHG as harmonic generation which is accomplished by the process laid out above, a small nonlinear perturbation to the linear response of the medium. Although qth harmonic generation is a qth order process and therefore depends not on $\mathcal{E}^2$ but on $\mathcal{E}^q$, the process always

consists of q photons of frequency ω adding together to form a photon of frequency $q\omega$ and so a generalised phase mismatch parameter for qth harmonic generation can be given as

$$\Delta k = qk_\omega - k_{q\omega} = \frac{q\omega}{c}(n_\omega - n_{q\omega}). \tag{3.39}$$

As with second harmonic generation above, the phase mismatch depends on the difference in the refractive index of the medium between the fundamental and the qth harmonic. This makes sense intuitively, as the refractive index is an indicator of how fast light of a given wavelength travels through a medium, and if the disturbance generating the qth harmonic travels at a significantly different speed to the harmonic itself then it will quickly find itself out of phase. The phase mismatch is usually sufficiently large that the coherence length is of the order of tens of nanometers.

The solution to this problem is to change the refractive indices such that they are more closely aligned. The most common way to do this is to use a birefringent crystal. A light wave travelling through a transparent material propagates by driving the bound electrons in the substance. These electrons are excited, but then failing to reach any new energy level (in this case they stay in the excited level, and we speak of the photon being absorbed) they fall back to the ground state and reradiate the light. This scattering effectively slows down the light by some amount, which is the physical meaning of the refractive index ($n = \frac{c}{\nu}$, where ν is the phase speed of light within the medium). In the case where the electron is more strongly bound in one direction than the others then light polarised in this direction will propagate at a different speed to light polarised along any other direction, thus this axis will exhibit a different refractive index to the others. A material with this property in known as a birefringent material, and the two axes are known as the *extraordinary* axis, for the changed refractive index, and the *ordinary* axis for everything else.

In order to use this for phasematching we need the refractive index of one of the waves along the e-axis to equal the refractive index along the o-axis of the other wave. This can be achieved by orienting the crystal relative to the incoming wave such that the combination of components oriented along both axes make the refractive indices match. For instance, if there is only one extraordinary axis, at an angle of θ to it the refractive index can be calculated as

$$n = \frac{1}{\sqrt{\frac{\cos^2(\theta)}{n_\mathrm{o}^2} + \frac{\sin^2(\theta)}{n_\mathrm{e}^2}}}. \tag{3.40}$$

If (as is often the case) all three axes have different refractive indices then this can get even more complicated. However, the point is that one can match the refractive index of the generated wave to the generating wave by rotating the crystal correctly. A birefringent crystal also usually has the benefit of being non-centrosymmetric, allowing the generation of even as well as odd harmonics.

3.3.4 Optical parametric amplification

Returning to second order processes only for the moment, the coupled wave equations presented in Equation 3.32 do not of course have to be solved only for second harmonic generation. They describe more generally the process

$$\hbar\omega_3 \leftrightharpoons \hbar\omega_1 + \hbar\omega_2. \tag{3.41}$$

The reversible symbol has been used to indicate that this process can run either way — a photon of frequency ω_3 can be split into two photons of frequencies ω_1 and ω_2 (difference frequency generation, DFG), or equally a photon of frequncy ω_1 and one of frequency ω_2 can combine into one of frequency ω_3 (sum frequency generation, SFG). This principle lies behind the technique of optical parametric amplification and optical parametric amplifiers (both abbreviated to OPA). In this the pump beam is split into two photons of lower frequency. In order for this process to take place one of the other photons needs to be present, if only in a small quantity, and then the remaining photons are generated by DFG. In order to have a photon of the correct wavelength present a little of the pump is used to generate what is called a *seed* beam by generating a white light continuum. In this continuum many different wavelengths of light are present, which gives the OPA great tuneability, as it can use any of them for DFG, but also makes it very sensitive to alignment, as it is necessary to phase match very precisely to phase match the generation of the wavelengths desired.

The two frequencies produced by the OPA are known as the signal and idler. The signal beam is the one that is seeded by the white light continuum, and the idler is the beam that is not initially present. As the white light continuum extends outwards from the pump pulse, the signal beam is almost always produced with a higher frequency than the idler beam. The idler beam also has the property that it is (or at least can be, if the generation is performed correctly) passively CEP locked — i.e. it is not affected by the jitter of the CEP in the pump beam, but always has a fixed CEP.

The fact that the idler could be made passively CEP stable was proposed by Baltuska *et al.* in 2002 [118]. The argument runs as follows. The electric field of pulse i in the pulse train of the laser can be defined as

$$\mathcal{E}_i(t) = \mathcal{E}_0(t)\mathrm{e}^{i\psi_i} \tag{3.42}$$

where $\mathcal{E}_0(t)$ is the electric field of a pulse with a CEP of 0, and ψ_i is the phase offset of the CEP. The phase offset of the frequencies generated in the white light continuum is [118]

$$\psi_{\omega 4} = \frac{\pi}{2} + \psi_{\omega 1} + \psi_{\omega 2} - \psi_{\omega 3}. \tag{3.43}$$

(This extra frequency is a product of self phase modulation, a third order process, hence it is viewed as four wave mixing. See section 3.3.5 below). The phase offset of the signal beam then is $\frac{\pi}{2}$ ahead of the pump beam,

but otherwise mirrors its CEP offset. The phase offset of the signal and idler beams can then be written

$$\begin{aligned} \psi_{\mathrm{S}} &= \psi_{\mathrm{P}} + \frac{\pi}{2} \\ \psi_{\mathrm{I}} &= -\frac{\pi}{2} + \psi_{\mathrm{P}} - \psi_{\mathrm{S}} \\ \psi_{\mathrm{I}} &= -\pi, \end{aligned} \tag{3.44}$$

where P denotes the pump beam, S the signal beam and I the idler beam. This shows that the idler always ends up with the same CEP despite shot to shot fluctuations in the pump pulse, as the fluctuations in the pump are reproduced exactly in the signal beam, and end up cancelling each other in the production of the idler.

3.3.5 Intensity dependent refractive index

Another process that is important for the generation and manipulation of laser pulses is the phenomenon of intensity dependent refractive index (IDRI, often denoted by n_2). A third order effect, this phenomenon lies behind a large number of nonlinear processes, from self-focusing and Kerr-lens modelocking, to white light continuum generation as we saw above in the OPA. As a third order process the polarisation of the medium is dependent on $\mathcal{E}^3$, and as such the mathematics is very involved, and only the results will be presented here. In general, a third order process can be viewed as the combination of three frequencies into a fourth, just as second order processes combined two frequencies into a third

$$\hbar\omega_4 \leftrightharpoons \hbar\omega_3 + \hbar\omega_2 + \hbar\omega_1. \tag{3.45}$$

In the case of IDRI, one of these frequencies is the negative of the other two (a case known as degenerate four wave mixing), giving

$$\hbar\omega = \hbar\omega - \hbar\omega + \hbar\omega. \tag{3.46}$$

As this process only involves one frequency it is automatically phase matched, and will always happen when the intensity is high enough. The polarisation amplitude of this process can be written in the case of linearly polarised light [108]

$$P_x(\omega) = \frac{3}{4}\epsilon_0 \chi^{\mathrm{IDRI}} \left|A_x(\omega)\right|^2 A_x(\omega) \tag{3.47}$$

where $|A_x(\omega)|^2$ indicates the amplitude multiplied by its complex conjugate and χ^{IDRI} is the nonlinear susceptibility associated with this process. As the refractive index is defined as

$$n = \sqrt{1 + \chi^{(\omega)}} \tag{3.48}$$

where $\chi^{(\omega)}$ is the nonlinear susceptibility of the terms in the polarisation with a linear dependence on the electric field, this then manifests itself as a refractive index term dependent on the square of the electric field, that is the intensity. Treating $\chi^{(\omega)}$ as much less than one, we can then write

$$n = n_0 + \left(\frac{3\chi^{\text{IDRI}}}{4n_0^2 c\epsilon_0}\right) I = n_0 + n_2 I \tag{3.49}$$

as a definition of the new intensity dependent refractive index.

3.3.5.1 Self focusing

In a beam the intensity will vary in space — typically it will be higher in the center than at the edges. Thus the effect of IDRI will be higher in the center than at the edges, and as n_2 is positive for most materials [108] this leads to a slowing down of the light in the center, causing a spatially dependent phase in the wavefront of the beam which leads to a narrowing of the profile. At lower intensities this process is small, and is outweighed by the diffraction due to the medium which causes the wavefront to spread out. However at higher intensities the term becomes bigger and outweighs diffraction, leading to the beam focusing in the medium, which can cause catastrophic damage if it is allowed to go too far.

However, as long as it is accounted for, self focusing can be of great benefit. In the process of Kerr-lens modelocking of an oscillator the gain medium is pumped such that the center of the beam is able to extract more energy from it, thus increasing the intensity in the center of the beam, increasing n_2 and leading to self focusing. The beam will also narrow in time due to a similar process (self phase modulation, see below), creating a very small, short pulse travelling in the oscillator cavity (see Figure 3.5). This narrowing is only limited in principle by the maximum bandwidth of the gain medium and the dispersion of the cavity, enabling the production of a femtosecond pulse train which is used as the basis of all the pulses produced by our laser [119].

3.3.5.2 Self phase modulation

Self phase modulation is the analagous process to self focussing in the spectral domain. Using Equations 3.47 and 3.27 we can write the propagation of the electric field like so

$$\frac{\partial \tilde{A}}{\partial z} = -i(\frac{\omega n_2 I}{c})\tilde{A}. \tag{3.50}$$

This implies that the amplitude of the electric field varies with z according to

$$\tilde{A}(z) = \tilde{A}(0)\mathrm{e}^{i\Phi(z)} \tag{3.51}$$

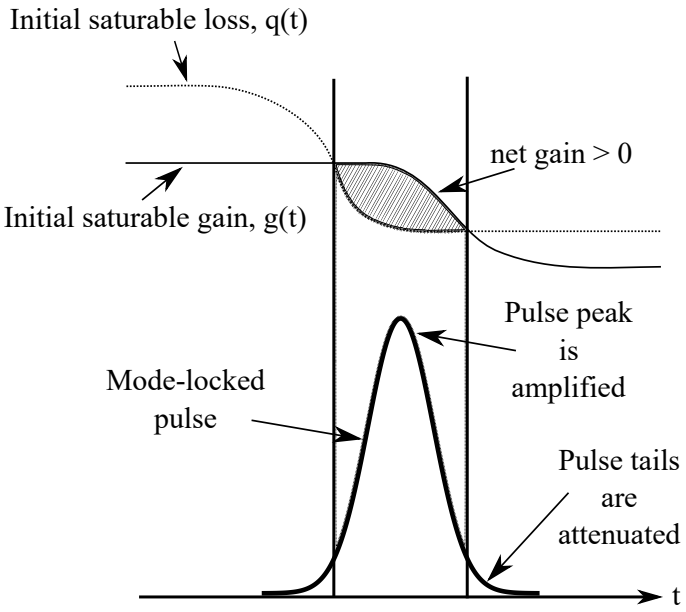

Figure 3.5: An illustration of the principle of saturable absorbtion in modelocking. The gain medium acts as a saturable absorber (dotted line) and a saturable amplifier (solid line). The x-axis is the time that the pulse passes through the gain medium, and so it is best to think of the pulse as travelling from right to left in this case. Initially it is both easy for the pulse to be absorbed and for the pulse to be amplified. At the leading tail of the pulse absorption outweighs gain. Once the first vertical line is reached enough stimulated emission is produced to start to outweigh the losses due to absorption. By the peak of the pulse the absorption has been saturated, leading to the maximum amount of gain. The gain then also starts to saturate, dropping down towards the absorption, until at the second vertical line they cross over again and the absorption outweighs the gain. Thus the pulse is narrowed in time. Diagram adapted from Ducasse *et al.*[119]

where

$$\Phi(z,t) = -\frac{\omega n_2 I(t) z}{c} \tag{3.52}$$

is the phase change of the peak of the field. As the field evolves in time the phase shift of each part of the temporal envelope of the field is given by

$$\frac{\partial \Phi}{\partial t} = -(\frac{\omega n_2 z}{c})\frac{\partial I}{\partial t}. \tag{3.53}$$

This phase shift results in a frequency shift of the pulse, where the leading edge redshifts as the intensity is increasing with time leading to a negative frequency shift, and the trailing edge blueshifts as the intensity is decreasing with time which leads to a positive frequency shift (see Figure 3.6). This causes a broadening of the spectrum of the pulse, which enables the pulse to support a shorter time duration due to the time bandwidth product, which states that the product of the time and the bandwidth of any given pulse $\Delta t \Delta \nu \geq K$, where K is a number of order unity that depends of the shape of the envelope of the pulse. For Δt to become as small as possible therefore requires a greater $\Delta \nu$. Note that a large bandwidth does not automatically mean a short pulse — the pulse still has to be compressed (see chapter 4 for more details).

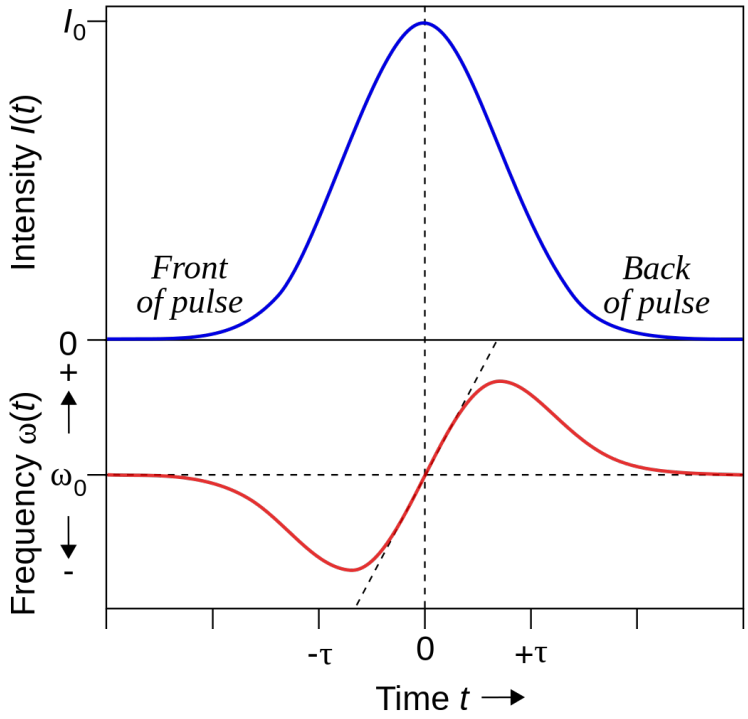

Figure 3.6: A diagram of a pulse undergoing self phase modulation, by Emannuel Boutet, licenced under the Creative Commons Attribution-Share Alike 3.0 Unported licence. The blue curve is the pulse envelope and the red curve is the phase of the pulse. The red curve is the differential of the blue one. As the intensity increases this causes a larger and larger phase shift, initially towards red as intensity increases, but then towards blue as it starts to decrease again. The earlier part of the pulse therefore has a longer wavelength than the later part, which is positive chirp.

If the effects of the IDRI are sufficiently strong, the refractive index that the pulse envelope itself experiences can become intensity dependent [108]. Thus the peak of the pulse will travel slower than the wings of the pulse, leading to a steepening of the trailing edge of the pulse. This process is called self-steepening.

3.3.6 Use of nonlinear optics in the creation of ultrashort laser pulses

The generation of the ultrashort 1.8 µm laser pulses described in Chapter 4 depends very heavily upon the concepts outlined in this Section. First of all the fact that the laser is pulsed at all depends on the self focusing caused by the intensity dependent refractive index. The selection of only certain pulses in our pulse train depends on the Pockels effect (a second order effect which occurs in the presence of a DC electric field). More details on this and on the exact method of the generation and amplification of our laser pulses can be found in Chapter 4. The conversion of our 800 nm pulses to the desired central wavelength takes place via DFG and the compression of the resulting infrared pulses uses self phase modulation in a gas filled fiber and self steepening in fused sillica to compress the pulse and correct for third order chromatic dispersion.

3.4 High order harmonic generation

The discussion of nonlinear optical processes so far has proceeded under the assumption that the electric field amplitude is sufficiently small that the nonlinear terms constitute a small perturbation to the linear polarisation wave. However this is not necessarily the case. When the amplitude of the electric field of the laser pulse is comparable in size to the potential binding the electrons to the atoms then many processes that otherwise would not be possible suddenly become feasible [120], as the perturbation due to the electric field is no longer small. The most relevant of these for this thesis is high order harmonic generation (HHG). This process can generate harmonic orders up to several hundreds or even thousands [121], and is regularly used as a source of soft X-rays for laboratory-scale X-ray absorption measurements (see Section 3.5) [9, 14]. It is also possible to use the production of harmonics themselves as a form of spectroscopy, as signatures of different molecular states can be found inside the spectra produced from molecules [36].

3.4.1 Single atom response

The most simple form of HHG to consider is HHG from an atomic gas (for instance one of the noble gases, typically Ar or Ne). Although the complete mathematical description of this process requires a full quantum mechanical treatment under what is known as the strong field approximation [31, 120], for this simple system a semi-classical approach known as the three step (or recollision) model produces very good results [29], and is laid out in graphical form in Figure 3.7. The three steps in this model are:

1. Tunnel ionisation
2. Propagation in the laser field
3. Recollision and emission of harmonics

3.4.1.1 Tunnel ionisation

The first step in the three step model is tunnel ionisation, ionisation of an electron via tunneling through the distorted potential energy well in the electric field. As the electric field oscillates this causes the potential well that the electron feels inside the atom to oscillate as well. This results in a potential energy barrier which looks more like a small bump than an infinite well, which gives the electron a nonzero probability to tunnel out of it (see Figure 3.7). This is a purely quantum mechanical step, as under classical mechanics the electron does not have sufficient energy to climb out over the barrier, no matter how small the bump is.

Tunnel ionisation was first described by Keldysh as part of a scheme to understand the phenomenon of ionisation of atoms by strong laser fields with a frequency too low to ionise in the normal way. He described two mechanisms, multiphoton ionisation and tunnel ionisation that explain this phenomenon. The difference

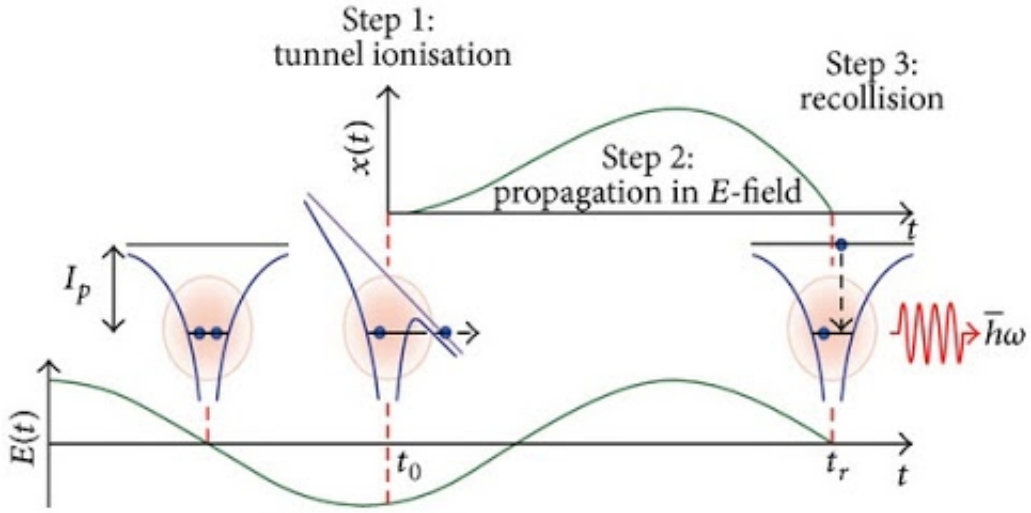

Figure 3.7: A schematic illustration of the recollision model of HHG. The oscillating electric field is shown along the bottom of the figure. When the magnitude is 0 the atomic potential remains undistorted. As the magnitude of the electric field increases the atomic potential distorts more and more, until tunnel ionisation becomes possible. note that although tunnel ionisation could happen on the other side of the peak in the electric field magnitude as well, in this case the electron would not recollide with the parent ion. From time t_0 to time t_r the electron propagates in the continuum, with a trajectory shown in the insert at the top. Finally at time t_r the electron recollides with the parent ion, emitting a high order harmonic photon. Reproduced under creative commons attribution licence 3.0, Image by Rashid Ganeev

between the two is given by the Keldysh parameter [122, 123], which in atomic units ($m_e = e = \hbar = 4\pi\epsilon_0 = 1$) is

$$\gamma = \sqrt{\frac{I'_p}{2U'_p}} \qquad U'_p = \frac{\mathcal{E}'^2}{4\omega^2}, \tag{3.54}$$

where I_p is the ionisation potential of the atom and U_p is the ponderomotive potential of the electric field — the cycle-averaged kinetic energy imparted to an electron oscillating freely in the electric field — and the primes indicate that they are in atomic units. A small Keldysh parameter indicates that the distortion of the electric field will be capable of making a barrier small enough for the electron to tunnel through. As it increases in size relative to U_p the distortion decreases in significance, and the dominant mechanism becomes multiphoton ionisation (many photons combining together to provide enough energy to overcome I_p). In our experiments we assume U_p is large compared to I_p as we have a long laser wavelength, and so we treat the ionisation process as tunnel ionisation

For us the main concern is with the rate of ionisation due to the strong laser field. This can be expressed according to the formalism created by Ammosov, Delone and Krainov (ADK) [124]. This model is based upon a hydrogen-like potential, and has been further expanded by including the work of Perelomov, Popov and Terent'ev

(PPT) on the long range coulomb potential of the atom [125] and by Yudin and Ivanov to include the effect of sub cycle dynamics [123]. The equation is very complicated, but it can be expressed as a prefactor (which depends on the electric field, ionisation potential and the quantum state of the electron) and an exponential term

$$\Gamma(t) = \mathrm{N}(t)\exp\left(-\frac{2(2I'_{\mathrm{p}})^{3/2}}{3\mathcal{E}'(t)}\right), \tag{3.55}$$

where again primes indicate that atomic units are used. Using this equation we can calculate the ionisation rate of the electron at different times in the laser cycle, and thus weight the probability of the electron following a particular trajectory.

3.4.1.2 Propagation in the laser field

The trajectory the electron takes after tunnel ionisation can be treated classically as a first approximation. The electric field can be assumed to have the form

$$\mathcal{E}(t) = \hat{\mathbf{i}}A(t)\cos(\omega t) + \alpha\hat{\mathbf{j}}A(t)\sin(\omega t), \tag{3.56}$$

where propagation is neglected as we are only interested in one or two optical cycles and α can vary between 0 and ± 1 as a way of changing the ellipticity of the field. When α is 0, the field is linearly polarised in the x-axis, and when α is ± 1 the field is either right- or left-circularly polarised. Treating both the electron and the nucleus as point particles, with the nucleus as fixed at $(0, 0)$ and the electron free to move in this electric field, we can write the classical equations of motion of the electron in the xy-plane as [29, 108]

$$\begin{aligned} \ddot{x}(t) &= \frac{-eA(t)}{m_{\mathrm{e}}}\cos(\omega t) \\ \ddot{y}(t) &= \frac{-\alpha eA(t)}{m_{\mathrm{e}}}\sin(\omega t) \end{aligned} \tag{3.57}$$

In this we have neglected the contribution of the ion to the electron motion, assuming that the electric field is strong enough to render this contribution negligible - this is known as the strong field approximation, and it is very important in the quantum mechanical treatment of this system[120].

Taking first the case of linearly polarised light (so $\alpha = 0$), we can integrate Equation 3.57 twice to get the velocity and path in space of the electron

$$\dot{x}(t) = \frac{-eA}{m_{\mathrm{e}}\omega}(\sin(\omega t) - \sin(\omega t_0)) = \frac{-eA}{m_{\mathrm{e}}\omega}\mathrm{V}(t) \tag{3.58}$$

$$x(t) = \frac{eA}{m_{\mathrm{e}}\omega^2}(\cos(\omega t) + \omega t\sin(\omega t_0) - \cos(\omega t_0) - \omega t_0\sin(\omega t_0)), \tag{3.59}$$

where the time variation of the envelope has been neglected and the constants of integration have been evaluated

in order to set $x(t_0)$ and $\dot{x}(t_0)$ to 0, where t_0 is the birth time of the electron.

The case of linearly polarised light means that if the electron crosses $x = 0$ again then it will (in the classical model) collide with the nucleus, but — as can be seen in Figure 3.8 — the electron takes different paths in the continuum depending on its birth time, only some of which lead to re-collision with the parent ion and subsequent emission of high harmonics.

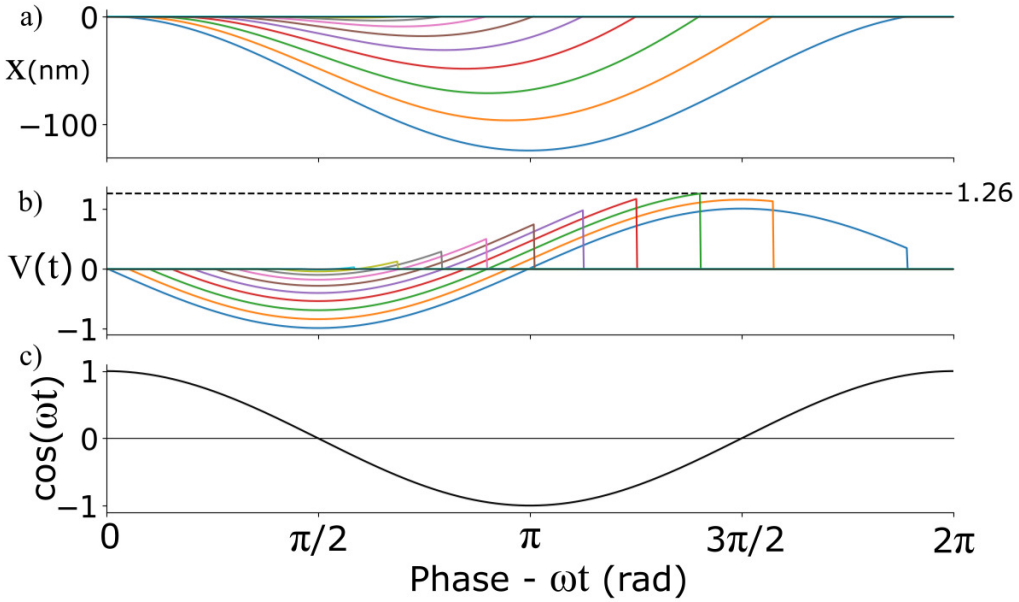

Figure 3.8: a) trajectories of the electron in the continuum for different values of ωt_0, the birth phase with an electric field of 1.8 µm wavelength and strength 1 VÅ^{-1}. b) the cosine sum part of the velocity for each trajectory in panel a) (see Equation 3.58). This value is dimensionless, and is shown here to show that there is a common limit to the cutoff energy due to the physics of the process. The vertical lines indicate the point at which the electron recollides, with a kinetic energy of $2U_\mathrm{p}\mathrm{V}(t)^2$. The maximum value that $\mathrm{V}(t)$ can take is 1.26, indicated by the dotted line. c) shows the value of the cosine part of the electric field during this process.

When α is varied from 0 the picture gets worse from a recollision perspective, as it very rapidly gets to the point where there is no chance that the electron will recollide with the parent ion, as can be seen in Figure 3.9. It is for this reason that one of the signatures of HHG under this recollision model is its very strong dependence on the ellipticity of the driving field of the laser.

3.4.1.3 Recollision and emission of harmonics

The times of recollision of the electron can then be calculated by setting $x(t) = 0$, and then solving for t for different birth times. The varying recollision times can be fed into the equation for $\dot{x}(t)$ to find the velocities of the electrons at the time of recollision. The kinetic energy of these electrons can then be calculated as $\frac{1}{2}m_\mathrm{e}\dot{x}(t)^2$, and so — assuming that the electron recombines into the bound state it occupied before and all the energy

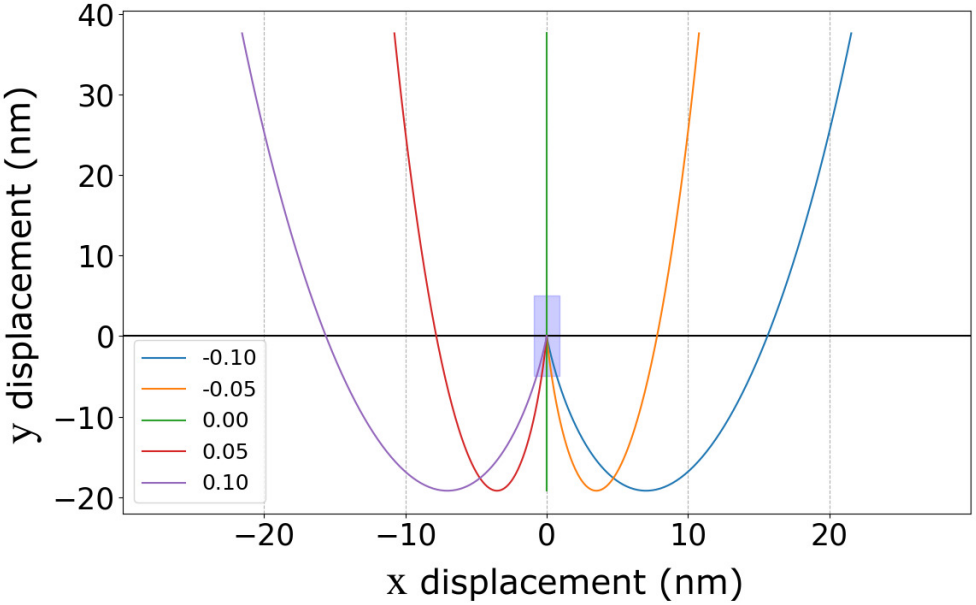

Figure 3.9: trajectories in the xy-plane of electrons in a slightly elliptical electric field. α is given in each case by the legend in the bottom, and the calculation is for a 1.8 µm electric field of strength 1 VÅ^{-1}. The blue shaded area around 0 indicates the atomic diameter of Kr, according to [126]. Even a very small deviation from linearity is enough to make the electron miss its parent ion.

released is released as one photon — the energy of a photon released at time t_r will be

$$\hbar\omega = \frac{1}{2}m_\mathrm{e}\dot{x}(t_\mathrm{r}) + I_\mathrm{p}, \tag{3.60}$$

where I_p is the ionisation potential of the atom. The ponderomotive potential of the laser field (U_p) can be defined as the cycle averaged kinetic energy of an electron moving freely under the influence of the field

$$U_\mathrm{p} = \frac{1}{2}m_\mathrm{e}\overline{\dot{x}(t)^2} = \frac{e^2A^2}{2m_\mathrm{e}\omega^2}\sin^2(\omega t) = \frac{e^2A^2}{4m_\mathrm{e}\omega^2}. \tag{3.61}$$

The kinetic energy of a recolliding electron can then be written in terms of this value as

$$K.E. = 2U_\mathrm{p}\mathrm{V}(t)^2, \tag{3.62}$$

where $\mathrm{V}(t)$ is the cosine sum of $\dot{x}(t)$ without the prefactor, and hence is dimensionless. From Figure 3.8 we can see that the maximum value of $\mathrm{V}(t)$ at recollision is 1.26, which leads to the famous equation for the cutoff energy of the harmonic spectrum produced

$$E_\mathrm{cutoff} = I_\mathrm{p} + 3.17U_\mathrm{p}. \tag{3.63}$$

The important points to take away from this equation are that there is a very definite limit to the energy of the harmonics that can be produced by HHG, and that in the recollision model the cutoff energy scales linearly with the amplitude of the electric field squared (in other words, linearly with the intensity of the laser pulse) and inversely with the square of the frequency. Therefore in order to increase the energy of the harmonics produced by this model we need to either increase the laser intensity or increase the laser wavelength.

The actual situation is not quite as simple as this picture however. Increasing intensity leads to increased ionisation of the medium. As we will see in section 3.4.2.1 this leads to bad phase matching of the harmonic generation, re-absorption of the harmonics and saturation of the medium as no atoms are left to ionise. In addition not all points on the laser cycle have equal probablility of birthing an electron into the continuum, which requires the equation mentioned above in the tunnel ionisation section to weight the probability of the electron being born. Another factor is the wavepacket spreading in the continuum. The electron is not actually a point particle, instead it should be understood as a quantum wavepacket, and as such when it propagates in space it spreads out. This wavepacket spreading has been measured to scale the harmonic yield by a factor of around λ^{-5} [127], although this has been disputed and the picture may be more complicated than a simple scaling law [128, 129].

3.4.1.4 Strong field approximation

The three step model is useful for understanding the process of HHG, and for predicting simple concepts like the cutoff. However, HHG is a fundamentally quantum mechanical process, and so in order to predict the full harmonic spectrum produced with some degree of accuracy a more quantum mechanical process is required.

Obviously, the most accurate simulation would involve solving the time dependent Schrödinger equation for the interaction of the laser field with the atom or molecule in question. This however is very computationally expensive, so in order to make this more tractable one must make some approximations about the nature of the system under consideration [120, 130].

The approximation used most commonly is known as the strong field approximation (SFA), and was introduced in a paper by Lewenstein *et al.* [31]. The approximation assumes a single active electron interacting with the laser field (which is linearly polarised in one axis), so the effect of the other electrons can be considered as just a potential felt by the active electron. It also assumes that U_p is comparable to or greater than I_p (hence the name — the field is strong compared to the ionisation potential of the atom), a regime where tunnel ionisation dominates.

The theoretical concepts in the derivation of the master equation of the SFA are quite involved and the details are beyond the scope of this thesis. However as this equation will be used to predict a harmonic spectrum later on in this thesis I will present the equation here and explain the general concepts involved.

In the SFA, the time dependent dipole moment experienced by the electron (x(t)) is given by[31]

$$x(t) = i\int_0^\infty d\tau \left(\frac{\pi}{\varepsilon + \frac{i\tau}{2}}\right)^{3/2} d_x^*(p_{\rm st}(t,\tau) - A_x(t)) \times d_x(p_{\rm st}(t,\tau) - A_x(t-\tau)) E\cos(t-\tau) \times {\rm e}^{-iS_{\rm st}(t,\tau)} + c.c. \quad (3.64)$$

where $d_x(p-A)$ is the transition dipole matrix element from the ground state to the continuum with a total momentum p and the vector potential of the laser field A, E is the amplitude of the laser field, τ is the time the electron spends in the continuum and S is the quasiclassical action, which describes the classical trajectory the electron takes in the continuum.

Equation 3.64 is obtained using the saddle point method to integrate over the momentum of the electron. This is a method of approximating integrals of the form f(x)e$^{ag(x)}$, where one would expect the integral to be dominated by the largest stationary point in g(x). One then re-expresses the integral in terms of x_0, the location of the largest stationary point, which then allows the integral to be expressed in a soluble form. Thus the momentum and quasiclassical action terms in Equation 3.64 are only the stationary points, and are given by[31]

$$p_{\rm st}(t,\tau) = E(\cos(t) - \cos(t-\tau))/\tau \quad (3.65)$$

and

$$S_{\rm st}(t,\tau) = (I_{\rm p} + U_{\rm p})\tau - 2U_{\rm p}(1-\cos(\tau))/\tau - U_{\rm p}(\sin(\tau) - 4\sin^2(\tau/2)/\tau)\cos(2t-\tau). \quad (3.66)$$

The HHG spectrum is then just the Fourier transform of the time dependent dipole moment (as time and frequency are a Fourier pair). This method can be extended to accommodate molecules, other dynamical effects and even a second active electron [120], but what I have presented here is the basis for all of this.

3.4.2 Bulk effects

HHG does not of course only consist of harmonics generated from a single atom. Like LHG mentioned above it requires that the harmonic wave be generated in phase from a vast number of emitters, which require a phasematching condition. In addition as the medium that is generating the harmonics becomes more dense it becomes no longer safe to assume that the electron can propagate in the continuum without meeting another body and thus scattering off of it. The XUV and soft X-ray pulses produced by HHG are at a wavelength where there are a lot of core-valence state transitions, so absorption will also be a big factor in the harmonic yield produced as the density of the medium is increased.

3.4.2.1 Phasematching HHG

As HHG is a form of harmonic generation we start with Equation 3.39, the phase matching equation for q-th harmonic generation:

$$\Delta k_q = qk(\omega_1) - k(\omega_q) \tag{3.67}$$

where $\omega_q = q\omega_1$. However, as the mechanism of HHG is different to that of LHG an extra term has to be added to account for the phase mismatch between the extra phase accumulated by the electron (and thus the harmonics) due to it's excursion in the continuum and the electric field, which has not accumulated this extra phase [131–133]

$$\Delta k_q = qk(\omega_1) - k(\omega_q) - \Delta K_{q,\text{dipole}}. \tag{3.68}$$

This dipole phase can be approximated in atomic units as [131, 133]

$$\phi_{\text{dipole}} \approx -U_\text{p}\tau \approx -\alpha_q I, \tag{3.69}$$

where τ is the electron travel time and α_q is a slope that depends on both the travel time and the harmonic order. The dipole phase then can be said to scale as

$$\Delta K_{q,\text{dipole}} \approx -\alpha_q \frac{\partial I(z)}{\partial z}. \tag{3.70}$$

This value can either be positive or negative, depending on whether the intensity is increasing with z (the focus is after the interaction region) or decreasing (the focus is before the interaction region). This implies that we can adjust the phase matching of HHG by adjusting the position of the focus relative to the interaction medium.

The remaining non dipole part of the phase mismatch parameter can be rewritten (as it is in Equation 3.39) as

$$qk(\omega_1) - k(\omega_q) = \frac{q\omega_1}{c}(n(\omega_1) - n(\omega_q)), \tag{3.71}$$

the phase mismatch due to the different dispersions of the medium at the different frequencies. In general the refractive index of a medium at soft X-ray wavelengths is much lower than that at optical or IR wavelengths, and therefore this expression is likely to be fairly large and positive, and thus the process will be badly phasematched. However, there are two corrections to this that have to be taken into account before conclusions can be drawn [131, 134]. These are the Gouy phase and the phase shift due to the plasma.

The Gouy phase is a phase shift in the beams due to their focusing geometry. When a Gaussian beam is focused it obtains an extra spatial phase along the z-axis of[134]

$$\phi_{\text{Gouy}} = \arctan\left(\frac{\lambda z}{\pi w_0^2}\right) = \arctan\left(\frac{2z}{b}\right) \tag{3.72}$$

where w_0 is the beam waist in the focal plane and $b = \frac{2\pi w_0^2}{\lambda}$ the confocal parameter, equal to twice the Rayleigh length. When $z \ll b$ (a small region around the focus, which is where our interest lies) we can write the wave vector associated with this as

$$k_{\text{Gouy}} = \frac{d\phi_{\text{Gouy}}}{dz} \approx \frac{2}{b}. \tag{3.73}$$

The confocal parameters of both beams are the same, this wave vector does not change depending on the frequency. We can therefore write its mismatch as

$$\Delta k_{\text{Gouy}} = k_{\text{Gouy}}(q-1) = \frac{2(q-1)}{b}. \tag{3.74}$$

This contribution is always positive in this region around the focus.

The phase shift due to the plasma is caused by the ionisation of the medium. Only a few electrons that ionise successfully recombine with their parent ion and release a harmonic photon. The rest scatter or otherwise miss, and thus the continuum will have an appreciable number of free electrons in it for a time that is long compared to the laser pulse. These electrons will oscillate with the electric field, with a resonance frequency of [134]

$$\omega_{\text{r}} = \sqrt{\frac{e^2\rho_{\text{e}}}{\epsilon_0 m_{\text{e}}}} \tag{3.75}$$

where ρ_{e} is the electron density in the medium. This polarisation of the plasma leads to a refractive index of the form

$$n_{\text{plasma}}(\omega) = \sqrt{1 - \left(\frac{\omega_{\text{r}}}{\omega}\right)^2}. \tag{3.76}$$

As ρ_{e} increases ω_r approaches ω, and light of frequency ω becomes increasingly absorbed by the plasma, which will kill the HHG. Therefore we want to keep ionisation as low as possible, in which case $\frac{\omega_r}{\omega}$ will be small and we can approximate Equation 3.76 using the binomial approximation as

$$n_{\text{plasma}} \approx 1 - \frac{1}{2}\left(\frac{\omega_{\text{r}}}{\omega}\right)^2. \tag{3.77}$$

using Equation 3.39 we can then write

$$\Delta k_{\text{plasma}} = \frac{q\omega_1}{c}(n_{\text{plasma}}(\omega_1) - n_{\text{plasma}}(\omega_q) = \frac{\omega_{\text{r}}^2(1-q^2)}{2qc\omega_1}. \tag{3.78}$$

This is always negative, so the effect of the plasma on the phase mismatch parameter always acts counter to dispersion.

Taking this all into account, the whole equation can then be written

$$\Delta k_q = \Delta k_{q,\text{dispersion}} + \Delta k_{q,\text{plasma}} + \Delta k_{q,\text{Gouy}} - \Delta K_{q,\text{dipole}}. \tag{3.79}$$

The dispersion term is usually positive and the plasma term is usually negative. The Gouy phase term is positive, but can be reduced by moving the focus relative to the interaction medium. The dipole phase can be positive or negative depending on where the focus is relative to the interaction medium. Dispersion can also be affected by changing the pressure (and thus the density of emmiters) and temperature of the interaction medium [131, 135]. Thus it is possible to phase match the HHG process, despite the unfavourable dispersion relation.

3.4.2.2 Scattering probabilities in the continuum

As the electron propagates through the continuum there is a finite probability that it will collide with a particle other than its parent ion, and thereby be scattered and (effectively) lost from the point of view of generating high order harmonics. The probability of an electron not colliding with a particle and scattering after the j-th small section of its trajectory Δx can be considered as $P(NS_j|NS_{j-1})$, the probability that the electron has not scattered after section j, given that it has not scattered after any of the previous sections either. The recursion goes back to the first section of the trajectory, $j = 1$, and so this can be thought of in terms of a tree diagram, where at each step the electron is either scattered and lost or not scattered and continues on. According to the normal rules of tree diagrams this can be expressed as

$$P_{\text{n.s.}} = \prod_{j=1}^{N} 1 - p_j \Delta x, \tag{3.80}$$

where N is the number of steps taken and

$$p_j = p_{\text{scat}}(t_0 + j\Delta x) \tag{3.81}$$

is the probability of scattering in section j of the trajectory, in units of m^{-1}. For sufficiently large N Equation 3.80 can be expressed in terms of sums as

$$P_{\text{n.s.}} = 1 + \sum_{j=1}^{N} \left(-\frac{1}{j!} \sum_{k=1}^{N} p_k \Delta x \right)^j, \tag{3.82}$$

as the additional terms which arise due to this approximation become vanishingly small. This has the form of the Taylor series of $\text{e}^{-\langle n \rangle}$, where

$$\langle n \rangle = \sum_{k=1}^{N} p_k \Delta x. \tag{3.83}$$

This can be understood as the average number of collisions that an electron would experience as it carries on on it's trajectory. We can express the distance travelled Δx in terms of a time step Δt as

$$\Delta x = |v(t)|\Delta t \tag{3.84}$$

where |v(t)| is the speed of the particle at time t (as we care about total distance travelled, rather than displacement from the origin). For an infinitesimally small time step this can then be expressed as

$$\langle n \rangle = \int_{t_0}^{t_{\mathrm{traj}}} p_{\mathrm{scat}} |v(t)| \, dt, \tag{3.85}$$

and so we can think about this in terms of integrating the probability of the electron scattering per unit time.

The probability of scattering can be modelled as a function of the number density of molecules ρ_{N} and the scattering cross-section of those molecules σ_{scat}

$$p_{\mathrm{scat}} = \rho_{\mathrm{N}} \sigma_{\mathrm{scat}}, \tag{3.86}$$

which gives a total probability of not scattering over the course of the electron trajectory of

$$P_{\mathrm{n.s.}} = \mathrm{e}^{-\rho_{\mathrm{N}} \int_{t_0}^{t_{\max}} \sigma |v(t)| \, dt}. \tag{3.87}$$

3.4.2.3 Reabsorption of emitted harmonics

One more factor needs to be taken into consideration when talking about the amount of harmonic emission that occurs, and that is re-absorption of the harmonic radiation by the medium. In the soft X-ray range most materials have strong resonances, which leads to large values for the cross-section for absorption at these wavelengths (σ). The absorption length for a given medium can be defined as [136]

$$L_{\mathrm{abs}} = \frac{1}{\sigma \rho} \tag{3.88}$$

where ρ is the density of the medium. Constant *et al.* found a relationship between this length, the coherence length of the harmonic generation due to the phasematching conditions (L_{coh}), the length of the medium (L_{med}) and the number of harmonic photons emitted at the end of the medium (N_{out}) for the q-th harmonic as [136]

$$N_{\mathrm{out}} = \rho^2 A_q^2 \frac{4L_{\mathrm{abs}}^2}{1 + 4\pi^2 \left(\frac{L_{\mathrm{abs}}^2}{L_{\mathrm{coh}}^2}\right)} \left[1 + \exp\left(-\frac{L_{\mathrm{med}}}{L_{\mathrm{abs}}}\right) - 2\cos\left(\frac{\pi L_{\mathrm{med}}}{L_{\mathrm{coh}}}\right) \exp\left(-\frac{L_{\mathrm{med}}}{2L_{\mathrm{abs}}}\right)\right]. \tag{3.89}$$

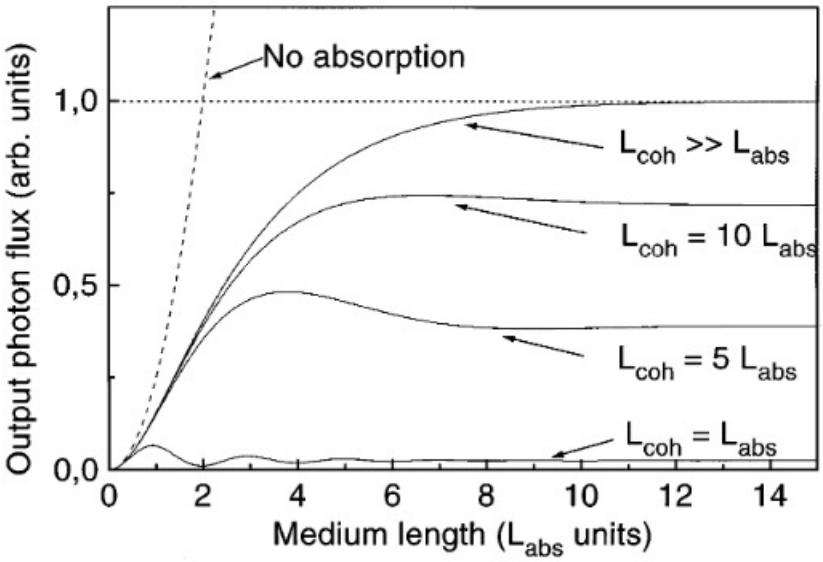

Figure 3.10: Figure of the output photon flux for different ratios of coherence length to absorption length, plotted as a function of the medium length (in units of absorption length). The dotted line indicates what the flux would be in the case that there was no re-absorption of the harmonic photons. Figure taken from Constant *et al.* [136].

This puts an additional constraint on the medium - even with perfect phasematching absorption will limit the amount of flux possible at a given harmonic. Constant *et al.* suggest from this the limiting conditions

$$\begin{aligned} L_{\text{med}} &> 3L_{\text{abs}} \\ L_{\text{coh}} &> 5L_{\text{abs}} \end{aligned} \tag{3.90}$$

required to ensure that the harmonics produced are above 50% of their theoretical maximum value (see Figure 3.10).

Absorption will also be important outside of the harmonic generation medium. In this case we can use the Beer-Lambert law to calculate the absorption of the background gas:

$$\tau = \sigma \int_0^L n(z)dz \tag{3.91}$$

where $n(z)$ is the number density of the absorbing species and τ is the optical depth of the medium. This is related to the transmitted flux by

$$T = 10^{-\tau}. \tag{3.92}$$

3.4.3 Solid phase high order harmonic generation

Moving from a gas to a solid the picture becomes somewhat different, although as we will see the recollision picture is still a good description of one of the possible mechanisms. The ionised electron is no longer able

to move freely with the electric field, as the electronic structure of the densely packed solid ensures that the electron is always affected by the Coulomb potentials of the nuclei around it [137]. The density of the solid ensures that the atomic orbitals of the constituent atoms overlap significantly, which allows for a significant amount of delocalisation of both the valence electrons and electrons in the conduction band — the band structure equivalent to the lowest unoccupied molecular orbital in a molecule [35]. This is a big difference to gas phase HHG, as the excited electron can in principle combine at any of the sites it was originally delocalised over instead of having to recombine with only one parent ion [138], although the electron is also limited in the extent of where it can travel by the band structure which constrains the final recombination energy [139, 140].

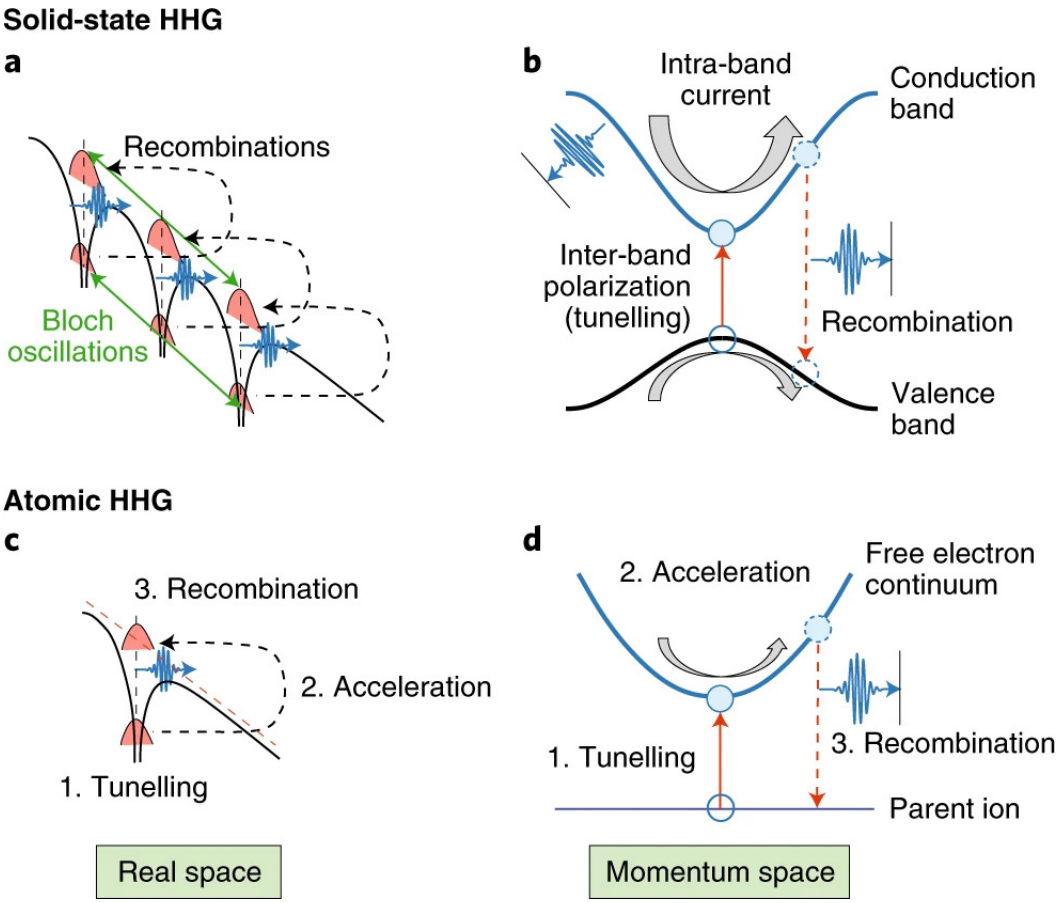

Figure 3.11: Illustration of the mechanisms of solid state and gaseous HHG in both real and momentum space. The most instructive comparison is the momentum space one (panels b and d). The recombination pictures in both mediums can be seen to be essentially the same — the electron tunnels up to a higher energy state, where it is driven by the electric field, acquires some extra momentum and then recombines back to the ground state, releasing a photon with energy equal to the band gap. The difference in this case is the shapes of the upper and lower energy states. The intraband mechanism is very different, as the radiation occurs due to the electrons being driven along the conduction band causing Bloch oscillations. This has no counterpart in the atomic case, as the electrons are not localised in the continuum and therefore cannot oscillate in the required manner. Figure taken from a review of solid state HHG by Ghimire and Reis [35].

There are two proposed mechanisms for HHG from a solid, known as the intra- and inter-band mechanisms (see Figure 3.11). The first mechanism, intra-band oscillation, is only available to a solid as it relies on having

a large delocalised band over which the electons can oscillate. The electrons are excited to the conduction band by tunnel ionisation, and then once in the conduction band they are driven by the strong field, leading to Bloch oscillations [80, 137]. This leads to a cutoff energy of

$$E_{\text{cutoff}} = n\hbar\omega_{\text{B}}, \tag{3.93}$$

where n is the number of different sites involved in the oscillation and ω_{B} is the Bloch frequency corresponding to the peak electric field:

$$\omega_{\text{B}} = \frac{e\mathcal{E}_{\text{peak}}d}{\hbar}, \tag{3.94}$$

where d is the lattice spacing. The implication of this is that harmonics generated via this intraband mechanism should have a cutoff frequency that scales linearly with the electric field, rather than with the square of the electric field as in the recollision model above. Another feature of this oscillatory mechanism is that due to the fact that all the harmonics are radiated by the electrons oscillating in the conduction band, there is no time delay between the production of different harmonics. As previously described, this time delay is present in harmonics generated by the recollision model above, as different harmonics are generated by different length trajectories that are born and recollide at different times, leading to a chirp in the pulse produced called the attochirp. This will not be present in intraband harmonics.

The second proposed mechanism (of interband recollision) is much more similar to recollision HHG. The electron tunnels into the conduction band as before. The electron and the hole that are thereby created are then accelerated within their respective bands by the strong electric field. They then recombine at some later time, radiating a photon with an energy equal to the band gap for the momentum that they have acquired during their journey [139, 140]. The cutoff harmonic energy therefore scales in the same way as the band gap scales with the amount of momentum which the electric field is able to impart to the electron-hole pair. The band gap can have a somewhat complicated relationship with momentum, however there is an approximately linear region for most band structures, which only deviates near the bottom and top [140]. The expected cutoff scaling is therefore linear as well, which is what is observed in experiment [80] and predicted by diverse theories [138, 140]. There is also a maximum cutoff in this mechanism — no matter what the field is, the highest harmonic generated cannot exceed the maximum band gap energy.

This interband mechanism therefore predicts roughly the same energy cutoff scaling as the intraband mechanism. Harmonics produced by interband recombination should however exhibit attochirp, since just as different trajectories spend different amounts of time in the continuum different momenta require different amounts of time to be gained. In a model based upon ZnO, Vampa *et al.* suggest that the interband mechanism dominates for driving pulses between 1 and 4 μm in wavelength, beyond which the intraband mechanism starts to dominate [140]. However in SiO_2 the observed harmonic generation at a wavelength on the shorter edge of that range

by Garg *et al.* showed almost no chirp, consistent with an intraband mechanism [141]. In principle then both mechanisms can be responsible for harmonics from solids.

The main difference in observable macroscopic response between solid high harmonics and gas ones thus is the scaling of the cutoff, with it being quadratic with electric field in the case of gas, and linear with electric field in a solid, due to the effects of the band structure. In principle we can also differentiate between a recollision-type model and a Bloch oscillations-type model in any medium by observing the chirp of the produced pulses. A further potential means of probing the electronic structure of the high order harmonic medium is in its dependence on ellipticity [138]. Ghimire *et al.* found a relatively weak ellipticity dependence in harmonics from ZnO [80]. This was backed up by a theoretical study by Vampa *et al.* [139], which found a slightly stronger ellipticity dependence, although still weaker than that found in gas. Contrary to this, an experiment by Ndabashimiye *et al.* in solid phase Ar found a similar level of ellipticity dependence to gas phase Ar [142]. The ellipticity dependence can be understood as being inversely related to the area over which an electron is capable of recombining with its hole, this can be used as a probe of the delocalisation of the ground state — i.e. how 'band-like' or 'atomic-like' it is.

3.5 X-ray absorption spectroscopy

Having used HHG to generate a soft X-ray pulse, we want to put this pulse to use. There are many things that HHG pulses are useful for, but their very large bandwidth and reach into the X-ray region of the spectrum make them ideal for X-ray absorption spectroscopy (XAS) measurements. Not only are the pulses short in duration, enabling high time resolution experiments, the broad bandwidth enables multiple X-ray absorption edges to be probed at the same time [9, 27]. This potentially allows for a very large amount of information to be encoded in the X-ray absorption spectrum.

The advantage that XAS has over absorption spectroscopy at conventional optical wavelengths is that at X-ray energies most atoms have transitions from core states to unoccupied valence states or the continuum (that is, ionisation). Core states here are defined as states which, in a molecule, do not extend far enough from the nucleus of the atom to have significant overlap with any other atom. These states are thus localised to their atom, and so give information about the electronic state of their atom only. This element specificity is very useful in interpreting the spectra — as long as there is only one environment in the molecule in which the element in question can be found then any transitions can be confidently identified with transitions to valence states involving this element, and so the electronic structure in the vicinity of this element can be probed. These transitions also occur at very different energies for different elements and different core states within the same element, preventing transitions due to different elements overlapping in energy and confusing the spectrum. Unfortunately if there is more than one environment the XAS shape will be a linear combination of

the transitions in these environments, which can very quickly render the spectroscopy confusing and hard to interpret without extensive theoretical support [23]. However, if the theoretical support is in place this can be used to gain information about complex materials that would otherwise be inaccessible.

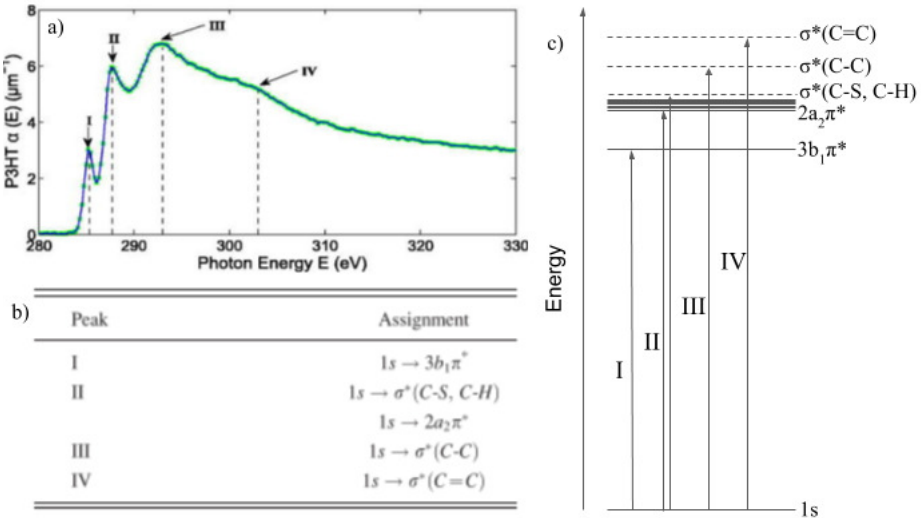

Figure 3.12: Illustration of an XAS spectrum. a) is a spectrum of the carbon K-edge of an organic polymer called P3HT, taken with our HHG source [14]. The electron transitions have been marked on it, with assignments taken from [143]. These assignments are written out in b). Note that the carbon K-edge is at around 289 eV, which is roughly the same energy as peak II, hence the uncertainty as to whether the transition is to a π^* state or a σ^* state. The π^* states show up as pre-edge features, whilst the σ^* states are actually higher in energy than the ionisation potential of the polymer, so they show up as shape resonances. c) is an energy level diagram to show the transitions that are occurring. The transitions that cause peaks I–IV are labelled on the arrows. As the quasi-bound states are not stable bound states of the molecule they have been shown as dotted lines.

An X-ray absorption spectrum of a given material consists of a general decline in absorption with energy punctuated by steep rises at the atomic edges, as the core level excitations are known. The actual edge *per se* is the core-continuum transition, but in the vicinity of these edges there will be additional features due to transitions to other states. This is split into two regions — X-ray absorption near edge structure (XANES) and extended X-ray absorption fine structure (EXAFS) [24, 25]. The EXAFS region is far beyond the edge, and manifests itself as oscillations in the descent down from the edge. This emerges due to electrons being ionised into the continuum with a non-zero kinetic energy. These electrons then move off into the continuum, where they can encounter other nuclei and scatter backwards towards their parent ion. The forwards and backwards propagating electrons then interfere, producing the oscillatory structure in the spectrum. This oscillation therefore encodes information about the local structure of the molecule in the region of the element.

The XANES region is effectively everything else. We can split it up further — before the edge there is the pre-edge structure, and after the edge there is post-edge structure (see Figure 3.12). The post-edge structure is due to quasi-bound states, states where the electron has an energy greater than the ionisation potential of the atom, but is hindered from entering the continuum due to the shape of the ionisation potential. These are also

known as shape resonances. The pre-edge structure consists of transitions from the core state to bound states, and are where our main interest lies in this thesis. These enable the spectroscopist to probe changes in energy of bound states and to observe the appearance and disappearance of these states corresponding to electronic and structural changes in the molecule during the experiment. By observing the shifts in energy and amplitude of individual structure peaks as a function of delay (in effect tracking the appearance and disappearance of resonances) it is possible to observe the presence or absence of certain molecular orbitals, which indicate the presence or absence of chemical bonds [8, 22]. It is possible in this way to observe bonds breaking or forming directly.

The edges have different nomenclature depending on the principal quantum number of the orbital that the electron is being ionised from. $n = 1$ is known as a K-edge, $n = 2$ is an L-edge, $n = 3$ is an M-edge, and so on alphabetically. Due to the fact that not all orbitals with the same principal quantum number are degenerate in energy there will be multiple different versions of a particular edge, which are denoted with subscript numbers. For example, the L-edge is split into three, with the L_1 edge being the ionisation from the 2s atomic orbital, whilst the L_2 and L_3 edges are the ionisation from the $2p_{1/2}$ and $2p_{3/2}$ atomic orbitals respectively, with the degeneracy lifted due to spin-orbit coupling.

Edge	Energy (eV)
Carbon K-edge	282
Nitrogen K-edge	397
Oxygen K-edge	533
Silicon $L_{2,3}$-edge	101,100
Sulfur $L_{2,3}$-edge	164,163
Chlorine L_1-edge	237
Chlorine $L_{2,3}$-edge	204,202
Argon $L_{2,3}$-edge	247,245
Zirconium $M_{4,5}$-edge	187,184

Table 3.1: A table of X-ray edges used in this thesis and their energies, taken from [41]. Edges split due to spin-orbit coupling have been presented on the same line.

Apart from the edges, the XANES and EXAFS oscillations contain the information that we are looking for about the electronic environment of the atom and the local structure the atom finds itself in. The difference between these two regions is the final destination of the photoelectron — in XANES the electron enters a bound or quasi-bound state, whereas in EXAFS the electron is unbound, and therefore is free to scatter off neighbouring atoms and provide structural information. Both cases can be in principle described by Fermi's golden rule [144]

$$\sigma(\omega) \propto \sum_f |\langle f|\,\hat{\mathrm{d}}\,|i\rangle|^2 \delta(E_f - E_i - \hbar\omega), \tag{3.95}$$

according to which the absorption cross section at frequency ω is proportional only to the square of the transition matrix element. f is the final state, i is the initial state and the corresponding $E_{f,i}$ terms are the energies of

these states, whilst $\hat{\mathrm{d}}$ is the transition dipole operator.

The EXAFS part of the spectrum is the easiest to interpret. As the photoelectron is released from the absorbing material it spreads out in a wavelike manner. It can then scatter back off neighbouring atoms, interfering with itself in the process. For the transition of a highly localised core orbital to an unbound photoelectron, the transition matrix element will be the overlap integral between the core and the wavefunction of the photoelectron. This will vary depending on whether the photoelectron experiences constructive or destructive interference at the site of the absorbing atom. As the phase gained by the photoelectron reflecting straight back off an atom is $2kR$, where R is the distance to the atom in question and $k = \frac{1}{\lambda}$, where λ is the De Broglie wavelength of the photoelectron (which depends on the energy of the photoelectron), this transition matrix element will vary with the energy of the absorbed photon, leading to the oscillatory structure in the spectrum.

This oscillatory structure (χ) is isolated from the edge absorption by subtracting a theoretical absorption cross section with all backscattering effects 'switched off' from the actual measured absorption cross section. This is then divided by a normalisation constant based on the size of the absorption edge in order to isolate the effects of the backscattering on the absorption. This oscillatory structure can be modelled by the EXAFS equation [24], which contains a sinusoidal term dependent on $2kR$ (as well as other phase shifts experienced by the photoelectron), a term containing the backscattering amplitude and two exponential terms containing decay terms which limit the structural information that it is possible to obtain to only the immediate vicinity of the atom.

The XANES part of the spectrum is altogether harder to interpret, because it is formed by transitions from the core to empty high lying bound and quasi-bound states. These vary not only from atom to atom, but also vary greatly with the different environments that the atoms find themselves bound in. This however is the source of the incredible richness of these spectra, as this variation contributes to a great sensitivity to the surrounding environment. Typically in order to reproduce a XANES spectrum a full calculation of the molecule in question is required, including the influence of the newly created core hole on the energies of the states available to be excited to, as the spectra observed will depend strongly on the exact electronic environment of the molecule. This makes it a very sensitive probe [22, 144, 145].

3.6 Conclusion

Having covered the science behind the apparatus used in the lab, we now tun to the details of the experiments performed. I will describe in detail our laser system, the design and modification of a new target chamber built in order to house our liquid jet target and then finally the experimental schema employed in order to perform HHG from our liquid target, and an attempt to take XAS measurements on solvated molecules within it.

Chapter 4

Experiments on liquid phase HHG and X-ray spectroscopy with a thin sheet jet in vacuum

4.1 Introduction

Having produced the thin liquid sheet jet, we now turn to using it as a target for experiments. In this chapter we will cover the experimental details and methods, whilst in the next two chapters we will look at the results from each experiment and discuss our interpretation of them. This chapter will detail the production and measurement of the ultrashort 1.8 µm pulses that are used to perform all the experiments in this thesis, the equipment used to detect the HHG pulses that we generate, the design and installation of a new target chamber and our flat liquid jet system within this chamber and finally the experimental setup used in the two experimental runs presented in the next two chapters.

4.2 Laser system

4.2.1 Titanium sapphire chirped pulse amplification laser

The driving source of the laser pulses used in all experiments in this thesis is a modified KM Labs Red Dragon CPA Ti:Sapphire laser. This laser is capable of producing 8–9 mJ 35 fs duration pulses with around 790 nm central wavelength. It consists of a passively modelocked oscillator, a grating pulse stretcher, two amplification stages and a final pulse compressor stage. We will discuss all of these in detail in turn.

4.2.1.1 Oscillator

The oscillator stage of our laser is a diode pumped Ti:Sapphire laser cavity which uses passive modelocking to form a femtosecond pulse train (KM Labs Griffin model oscillator). Passive modelocking (also called Kerr-lens modelocking) uses self-focusing and self phase modulation to induce a fixed phase relationship between cavity modes, 'locking' them together to form a pulse train. The cavity modes are standing optical waves that can exist inside a laser cavity, and when they are locked into a fixed phase relationship they form a pulse train due to the interference between the modes of different frequencies. The pulses in this pulse train are more intense at their peak than the continuous wave operation of a single cavity mode, and so they can extract sufficient gain from the laser crystal to overcome the losses in the cavity due to dispersion, absorption etc. (Note that the gain at this stage is limited by the photon flux incident on the crystal). This net gain enables these pulses to grow in each round trip through the cavity. The self-focusing and self phase modulation then ensure that the pulses in the cavity are short by favouring the central part of the pulse (in space and time) over the wings. This cavity operation can grow from noise to stable operation very quickly, and provides a reliable source of short (30 fs or so) pulses to be amplified in the rest of the system.

4.2.1.2 Chirped pulse amplification

This pulse train is amplified by a process known as chirped pulse amplification (CPA), whose inventors (Donna Strickland and Gérard Mourou [146]) won the Nobel Prize for physics in 2018. In any large amplification system there is a trade off between producing a higher energy pulse and the damage threshold of the amplification system. CPA is a way of mitigating the damage to the amplification system, whilst allowing high pulse energies to be produced. It does this by stretching the pulses in time to several hundreds of picoseconds, using a grating pair to add a frequency dependent group delay to the pulse, retarding some frequencies and speeding some up in a uniform manner and separating the frequencies contained in the pulse in time — a process known as *chirping*. This greatly increases the temporal width of the pulse, which can now be amplified without obtaining the high peak intensities that damage the optics in the system due to a corresponding reduction in peak power [146, 147]. Once the pulse has been amplified the chirp that is stretching it is compensated by another grating pair (known as the compressor), to produce a pulse of approximately 35 fs duration [14].

There are two main ways in which optics can be damaged — by a pulse with a peak power over the damage threshold, directly damaging the optic itself, and by the cumulative heating effect of multiple pulses with peak power lower than the damage threshold, each of which impart some energy which doesn't have enough time to dissipate before the next pulse, leading to heating and damage. The first damage mechanism is dealt with by chirping the pulse, but the second is still a risk, especially as the pulses leave the oscillator at a rate of around 80 MHz (one every 12.5 ns). If the pulses are to be stretched to around 0.1 ns then there will be very little time to dissipate the heat before the next one arrives.

The solution is to use a Pockels cell to pick out one pulse every millisecond, turning the 80 MHz pulse train into a 1 kHz one. The Pockels effect is a second order nonlinear effect, which in the presence of a DC electric field induces a change in the refractive index of a material. The electric field

$$\mathcal{E} = A_0 + A_1 \cos(\omega t - kz) \tag{4.1}$$

leads to a polarisation amplitude of

$$P = \frac{1}{2}\varepsilon_0 \chi^{(2)} A_0 A_1. \tag{4.2}$$

Like the intensity dependent refractive index this is a polarisation amplitude that oscillates with the frequency of the fundamental electric field, so induces a change of refractive index which scales linearly with A_0, the DC field. This effect only involves waves of the fundamental frequency, and so it is automatically phase matched (assuming the existence of a DC electric field), and it is a second order process so it is only possible to induce in non-centrosymmetric materials.

This process is used to make what is effectively an electrically switched half waveplate. When the pulse to be picked out travels through the Pockels cell a DC field of the correct value is applied to the cell, and the new refractive index then causes the polarisation of the pulse to rotate. The pulse train then passes through a polariser, which blocks every pulse except the one that has been rotated. In this way the frequency of the pulses in the pulse train is reduced. A similar process is used after the first amplification stage, in order to remove pulses generated by amplified spontaneous emission (ASE).

4.2.1.3 Amplifiers

Our laser system has two amplification stages, both Ti:Sapphire crystals pumped by diode lasers at 532 nm. The first stage amplifies the pulses to around 1.3 mJ, whilst the second one amplifies them up to around 16 mJ. We have two as a compromise between adjustability and cost. Amplifying higher energy pulses requires different parameters to amplifying lower energy pulses, not least that more energy needs to be imparted to the gain medium. One stage would put too much load on the crystal, whilst three is unnecessary for the energies we are using in our experiments.

The first amplifier is a ring type amplifier, in which the pulses pass through the gain medium several times in a ring pattern, slightly displaced horizontally from each other each time until finally being picked off by another mirror and sent onwards (see Figure 4.1). The number of passes through the crystal depends on how the mirrors in the ring are aligned, and this has been varied between 9 and 13. This type of amplifier is good for enabling a high number of passes through the gain medium, which is necessary for amplifying lower energy beams, where the initial power per pulse is sufficiently low that doubling it does not come close to saturating the gain of the amplifier and so many successive doublings can be performed before reaching this limit. However

the output mode of the beam is very sensitive to alignment, so needs careful monitoring to ensure the beam that is sent onwards can be amplified well and won't burn anything. In addition the amplifier produces a significant amount of ASE, as a spontaneously emitted photon can travel round the ring and be amplified fairly easily, and so it requires a second Pockels cell to cope with this.

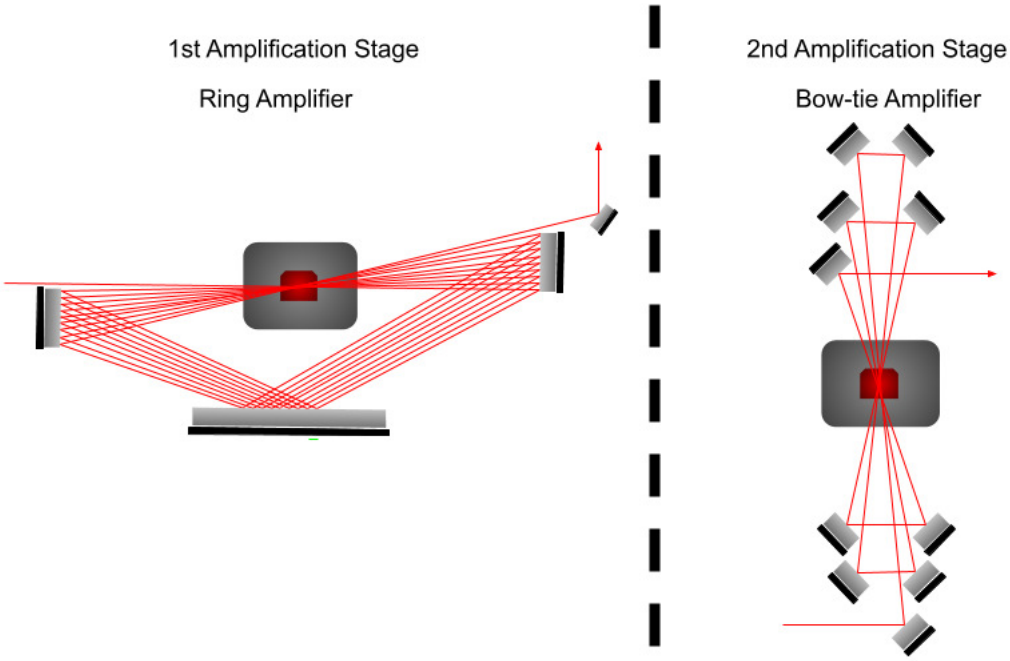

Figure 4.1: Diagram of the two amplification stages in the Red Dragon. The first stage is a ring amplifier, where the pulse passes around in a ring, as shown. The optics on either end of the crystal are more complicated than just a single mirror, but only one mirror is shown for simplicity. the long mirror at the bottom is a large rectangular mirror, which is what allows the ring to operate as a ring. Misalignment of this mirror can throw the whole ring out of alignment. Nine passes are shown, but the number can vary depending on need. The second stage is a bow-tie amplifier, in which the beam passes five times through the medium before being picked off. It is a mark of the relative simplicity of this amplifier that the optical layout in this has not been simplified at all. This simplicity is due to the higher gain in this stage — the pulse already has lots of photons in it, each of which can cause a stimulated emission event, so each pass represents a much larger increase in the absolute energy of the pulse than in the first stage amplifier.

The second amplification stage is a bow-tie configuration with five passes (see Figure 4.1). This configuration is a bow-tie shape that widens each pass through, until an outside mirror picks the pulse off. This second amplifier is pumped by two diode lasers as opposed to the first stage which is pumped by one, and is also where most of the energy is added, going from 1.3 mJ of energy to 16 mJ per pulse. Since the pulses have many more photons per pass, fewer passes are possible before the amplifier saturates, so in contrast to the ring amplifier in the first stage the second amplifier is arranged with only one path through the optics and out the other side. This makes the optics much easier to align, and consequently makes the mode much nicer out of this amplifier.

4.2.1.4 Pulse compressor

The pulses are then split with a beamsplitter, with about 1 mJ sent to other laser systems in the lab and the remaining energy sent into a grating compressor. This is the opposite to the stretcher, in that it applies a frequency dependent group delay in the same way, but it does so to counteract the stretching, by compressing the pulse back down to somewhere near the transform limit (*i.e.* the theoretical shortest pulse, according to the bandwidth theorem (see section 3.3.5)). This optical system is therefore exposed to higher peak intensities than the rest of the system, and correspondingly burns more often. To compensate for this the beam area is increased before it is compressed to lower the peak intensity. After compression and the losses accumulated after the second amplification stage the final output pulse has around 8–9 mJ of energy.

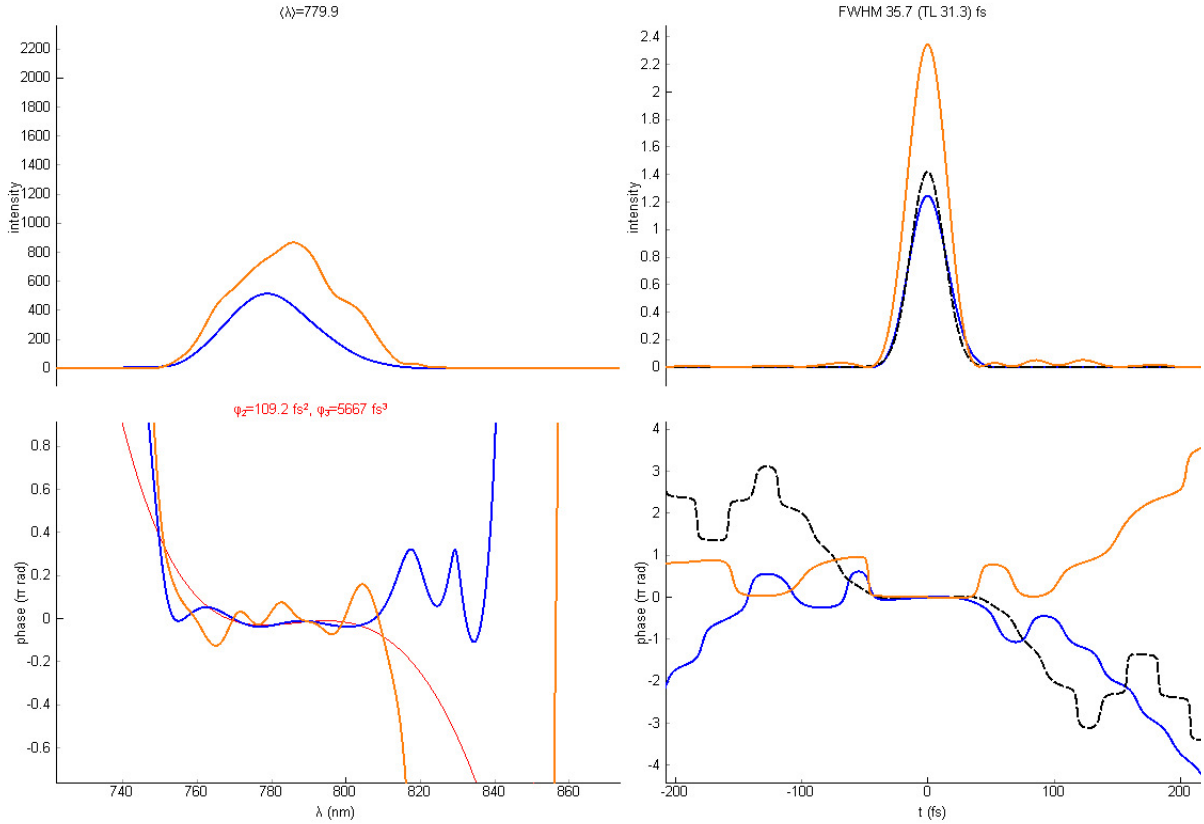

Figure 4.2: The SPIDER traces of the compressed pulse out of the Red Dragon, measured on 16/7/2021. The top row are measurements of intensity against wavelength and time, and the bottom row are measurements of phase against those two. In all measurements the orange line is a reference and the blue line is the measured signal. In the temporal measurements the profile of the calculated transform limited pulse is shown as a dashed black line, whilst in the phase/wavelength measurement the fitted phase profile is shown as a red line. The phase measurements outside the intensity profiles are essentially noise. Our pulse is fairly close to transform limited, with a relatively flat phase over the wavelengths actually present in the pulse.

A small part of the energy of the Red Dragon output pulses is split off from the rest by a 99:1 beamsplitter and is sent into a SPIDER (spectral phase interferometry for direct electric-field reconstruction) pulse diagnostic system in order to measure the output pulses in real time. This allows us to be sure that the pulses we are

using to seed the rest of the system are what we think they are. A typical SPIDER measurement is shown in Figure 4.2, and more details on the operation of SPIDER are given in section 4.3.1.

4.2.2 OPA system

The pulses produced by the Red Dragon have a central wavelength of around 790 nm. However, as was discussed in Section 3.4, the longer the wavelength the higher the harmonics that can be produced in theory, as U_p is much greater. In order to convert our laser pulses to a longer wavelength (and therefore lower frequency) we use an optical parametric amplifier (an HE-TOPAS by Light Conversion). As detailed in Section 3.3.4, optical parametric amplification splits a higher frequency wave of light into two lower frequency waves, a signal beam and an idler beam.

The HE-TOPAS has four stages, a seed generation stage and three amplification stages (see Figure 4.3). The seed is generated by focusing a part of the pump beam into a sapphire crystal, producing a beam of white light by self phase modulation [148]. As we saw in section 3.3.5, the intensity dependent refractive index can induce a frequency shift in different parts of the pulse, proportional to the rate of change of intensity with time in that part of the pulse. As the rate of change of intensity can be large, this frequency shift can also be quite large, generating a white light supercontinuum containing a broad range of frequencies. This seed generation stage is necessary because in order to transfer energy from the pump beam to the signal beam there needs to be light of the correct frequency already present [149]. This white light is then chirped to make the selection of the signal frequency easier.

This seed is then fed through the three amplification stages, which each split off a portion of the pump to amplify the beam — there are three for a similar reason to the fact that there are two amplification stages in the Red Dragon, *i.e.* each pass through the gain medium may need different parameters in order to optimise the efficiency. The first two crystals are optimised for gain, to increase the amount of the signal beam present. The idler beams that are generated in this way are dumped into beam blocks. The final stage is optimised to convert as much of the pump as possible into the signal and idler beams, which then forms the output of the system.

The wavelengths of products of this process are tunable, as the signal frequency generated will be the one with the most efficient transfer of energy, *i.e.* the one for which momentum is the closest to being conserved, in which the wave vectors of the three beams obey the phase matching condition:

$$\vec{k}_\mathrm{p} = \vec{k}_\mathrm{s} + \vec{k}_\mathrm{I}. \tag{4.3}$$

Note that this is a vector sum because the beams are not necessarily colinear. The frequencies for which this is true depend on the angle the optical axes of the crystal make with the pump beam, the temperature of the

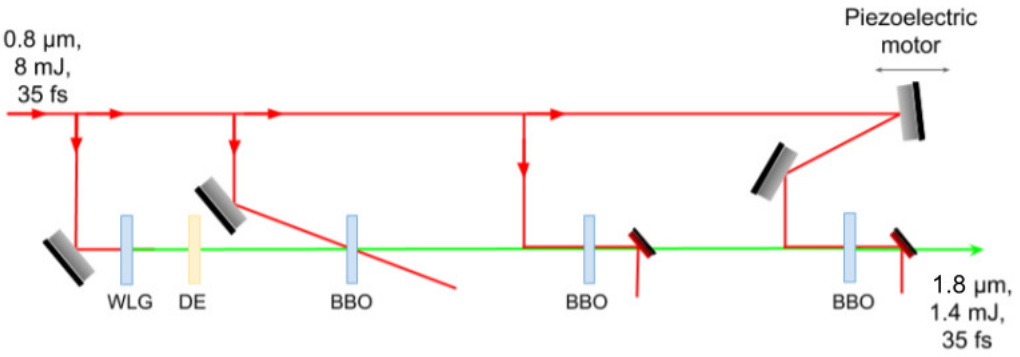

Figure 4.3: Diagram of the optical layout of our HE-TOPAS. A small proportion of the pump beam is sent into a white light generation plate, where it undergoes massive spectral broadening. This supercontinuum is then sent into a dispersive element — a plate that imparts a linear chirp to the supercontinuum, to facilitate tuning the wavelength of the signal. The signal wavelength is then amplified in the first two barium borate (BBO) nonlinear crystals, one non-colinearly and one colinearly, with a dielectric mirror to remove the excess pump beam. The third BBO crystal then generates the signal and idler pulse, and again a dielectric mirror is used to remove the excess pump beam. The output characteristics are quoted for the idler beam, as that is the one that is used in the rest of the system. The signal beam is around 1.4 µm, 2 mJ, 35 fs.

crystal and other effects such as voltage being applied to the crystal [150]. In our case we phasematch the signal beam to be at around 1400 nm, giving an idler beam of around 1800 nm. We are able to get around 3.4 mJ out of both beams combined. They are then split, and after the splitter we have around 1.4 mJ in the idler, and 2 mJ in the signal.

The idler is then sent onwards to the hollow core fibre to be broadened, whilst the signal is blocked and not used further. The reason for using the idler over the signal is twofold — ease of generation of HHG at longer wavelengths and passive CEP stability. It is easier to make the idler the longer wavelength, as then the supercontinuum that the OPA is seeded with does not have to extend all the way down to the desired wavelength. The two wavelengths used are comfortably within the operating range of our OPA, and so we can use this long wavelength light whilst still extracting maximum gain from our OPA.

As explained in section 3.3.4, the idler of a white light seeded OPA is passively CEP stable — its CEP does not fluctuate from shot to shot like that of the pump beam. This means that using the idler in an experiment makes it possible to control the CEP of the driving field. This is desirable when generating high order harmonics, as it is well established that HHG using a short pulse has a cutoff energy that varies with the CEP of the driving field [20, 40, 110]. In order to ensure stability of the CEP over longer periods of time (when even the nominally CEP stable idler can drift due to thermal fluctuations causing subtle differences in path length) we use a 2f–3f interferometer, which is connected to the motorised mirror inside the TOPAS, which then finely tunes the pump delay to keep the CEP of the idler beam constant.

In a 2f–3f interferometer the spectrum of the pulse to be measured first needs to be broadened so that it

spans at least $\frac{2}{3}$ of an octave ($3\omega_{\text{low}} = 2\omega_{\text{high}}$, where low and high denote the highest and lowest frequencies present in the pulse). In our case this is done via self phase modulation in a hollow core fibre (see section 4.2.3). With this accomplished the third harmonic of the low frequency part of the spectrum is overlapped with the second harmonic of the high frequency part, and the resulting interferogram is observed on a spectrometer. This is capable of finding, and therefore stabilising, the relative CEP offset in the pulses due to the method with which the pulses were originally generated. Kerr-lens modelocking locks together a wide variety of cavity modes with frequencies equally spaced by the pulse repetition rate [151]

$$\omega_{\text{r}} = \frac{\nu_{\text{g}}}{2L}, \tag{4.4}$$

where ν_{g} is the group velocity of the pulse envelope and L is the cavity length. However the frequency of each mode is in general not an integer multiple of this round trip frequency, as the electric wave travels not at the group velocity, but at it's phase velocity (see section 3.2.1). There will therefore be an offset in CEP of $\Delta\phi$ every round trip. (This is the origin of the CEP jitter that is eliminated in the generation of the idler). The frequency of each mode in the pulse will then contain a CEP offset term given by [152]

$$\omega_{\text{CEP}} = \frac{\Delta\phi}{2\pi}\omega_{\text{r}} \tag{4.5}$$

giving an equation for the nth frequency in the pulse of

$$\omega_n = \omega_{\text{CEP}} + n\omega_{\text{r}}. \tag{4.6}$$

Returning to the 2f–3f interferometer, the interference of the two frequencies gives a beat frequency of[153]

$$\omega_{\text{beat}} = 3\omega_{\text{low}} - 2\omega_{\text{high}} = 3\omega_{\text{CEP}} + 3\omega_{\text{low}} - 2\omega_{\text{CEP}} - 2\omega_{\text{high}} = \omega_{\text{CEP}}, \tag{4.7}$$

as $2\omega_{\text{high}} = 3\omega_{\text{low}}$, which can then be used to calculate out $\Delta\phi$, and in the event that it drifts from the set value the interferometer software will send a signal to the piezomotor in the HE-TOPAS, to adjust the pump timing a little and bring $\Delta\phi$ back to the desired value.

In our interferometer the second harmonic is generated in a BBO, whilst the third harmonic is present already, as it is generated in the hollow core fibre. The two beams are projected onto the same axis using a polariser, and the time delay between the two frequency components of the pulse is corrected using a ZnSe plate, which has very high dispersion in this wavelength range.

4.2.3 Hollow core fibre

In addition to wanting to use the longest possible wavelength to generate high order harmonics, we also want each driving pulse to be as short as possible. This is both because a shorter pulse has a higher peak intensity, which increases U_p and thus cutoff frequency, and also because, with a sufficiently short pulse, we can generate only one burst of harmonics as only one half cycle of the electric field has sufficient intensity to generate harmonics, a process known as amplitude gating [20, 154]. With a longer pulse each half cycle generates its own burst of harmonics, leading to a pulse train. This is undesirable for a pump-probe type experimental scheme, as the measurement needs to take place at a definite point in time which requires there to be only one pulse.

In order to ensure that the pulse is short enough to generate isolated attosecond pulses we need to compress it. According to the bandwidth theorem, the product $\Delta t \Delta \nu \geq K$, and so in order to compress our idler pulse in time the spectrum needs to be broadened. This is achieved by propagating the beam through a 1 m long Ar filled hollow core fibre (HCF). The existence of the noble gas in the fibre means that third order nonlinear effects can occur (second order being ruled out due to the centrosymmetricality of the gas). The refractive index experienced by the pulse thus becomes intensity dependent, and so the spectrum is broadened by both self phase modulation and self steepening [155]. This results in a very broad spectrum (due to the SPM contribution), asymmetrically shifted towards shorter wavelengths (due to the self steepening contribution). This incidentally is why we use a 2f–3f interferometer rather than the more conventional f–2f one — the asymmetry in our spectrum coupled with the large wavelength range for an octave at these wavelengths means that the condition for the 2f–3f interferometer is easier to achieve consistently.

In general a pulse broadened in this manner would have to be compressed using chirped mirrors (mirrors that have been specially made to impart a fixed group delay to light in a certain wavelength range). These are expensive and add path length and complexity to the laser system, and don't always compress the pulse to its transform limit, as they cannot compensate for the third order dispersion in the pulse [156]. However happily (and this was another reason for using the 1.8 µm idler) above 1.3 µm fused sillica glass has negative group delay dispersion, so the group delay caused by the broadening can be compensated by a simple pair of fused silica wedges. This also has the beneficial property of compensating the third order dispersion of the pulse in addition, due to the asymmetric phase profile caused by self steepening [156].

This compressed pulse is then sent to the experiments. The pulse duration and phase was measured by a second harmonic frequency resolved optical gating (FROG) experiment (see section 4.3.2). The trace of this pulse is shown in Figure 6.4 in Chapter 6, indicating that we can compress the pulse to around 13 fs (although the pulse is not as clean as those out of the Red Dragon).

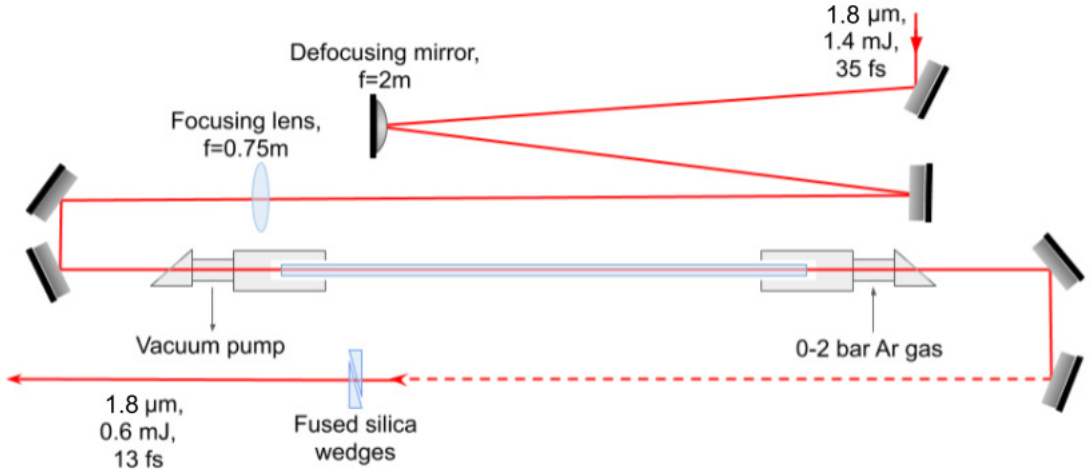

Figure 4.4: Diagram of the hollow core fibre optical layout. The idler is expanded by a defocusing mirror in order to give as tight a focus as possible, and then is focused into the entrance of the hollow core fibre. The fibre is housed inside a differential pumping stage, where a variable pressure of Ar is input into the fibre at the exit, but at the entrance a vacuum pump is used to evacuate the fibre. This creates a pressure gradient across the fibre, which is necessary to stop self-focusing distorting the pulse before it can propagate through the whole fibre. This way at the start of the fibre the pulse can couple into it with little disruption, and then as it propagates the third order effects slowly ramp up. The broadened spectrum is then sent onwards to the experiment, where it is compressed with a pair of fused silica wedges.

4.3 Pulse Diagnostics

In order to be sure that the experiments we are performing are the experiments we want to perform, we not only have to generate the requisite optical pulses, but we have to verify that we have generated them correctly by measuring them as well. This is not a trivial exercise, as the usual way of measuring the length in time of an optical pulse is to use a shorter event to measure it. However, when measuring pulses with durations of tens of femtoseconds, there is not really a shorter event available to measure the pulses with. In consequence, both of the techniques presented here are some variation on measuring the pulse with itself, and the process of measuring the pulse can appear to be as challenging as performing the experiment itself.

In order to fully characterise an electric field (such as the field of the pulses we are measuring) it is necessary to specify not only the amplitude of the electric field but also the phase of the electric field [157]. There are two main techniques used to characterise pulses during the experiments presented in this thesis, SPIDER [158] and FROG [159].

4.3.1 SPIDER

SPIDER (spectral phase interferometry for direct electric-field reconstruction) is a pulse measurement technique that uses a spectral phase shear (separating out the spectral phases in time by a known amount) to back out the original phase of the pulse under investigation.

SPIDER works by creating three replicas of the pulse to be measured, using D-mirrors (half of a circular mirror, designed to provide a straight edge allowing a laser beam to be divided cleanly in half) to divide the wavefront of the input pulse. One of these replicas is stretched by a grating in a similar manner to the Red Dragon pulses in section 4.2.1.2, to the extent that in the period of time that one of the other two pulse replicas would overlap with the stretched pulse it is essentially monochromatic. These two other replicas are then delayed relative to each other by a time delay τ, and as such overlap different frequencies of the stretched pulse (ω and $\omega - \Omega$, where Ω is the spectral shear, the difference between the frequencies overlapped by the two test pulses. For a stretched pulse of spectral width $\delta\omega$ and duration T, Ω is equal to $\frac{\tau\delta\omega}{T}$). This overlap takes place in a nonlinear crystal which is oriented so that the overlapping pulses undergo sum frequency generation. The resulting upconverted pulses are overlapped in time and produce an interference pattern, the intensity of which is given by [158]

$$S(\omega) = |\mathcal{E}(\omega) + \mathcal{E}(\omega - \Omega)\mathrm{e}^{i\omega\tau}|^2, \tag{4.8}$$

where the extra $\mathrm{e}^{i\omega\tau}$ is the extra phase obtained due to the time delay. The cross term of this expression can be obtained by spectral filtering[160], and has a phase modulation of

$$\theta(\omega) = \phi(\omega) - \phi(\omega - \Omega) - \omega\tau. \tag{4.9}$$

The $\omega\tau$ due to the time delay can be found by calibration and accounted for, giving a frequency dependent phase of

$$\theta'(\omega) = \phi(\omega) - \phi(\omega - \Omega), \tag{4.10}$$

from which the spectral phase $\phi(\omega)$ can be backed out by a number of methods. This technique is very sensitive to correct alignment and calibration, as the interference signal that gives the phase is superimposed on top of other interference signals, especially that due to the time delay.

4.3.2 FROG

Frequency resolved optical gating (FROG) is a pulse characterisation technique based on intensity autocorrelation [157, 159], a crude method of measuring a pulse using itself. In both autocorrelation and FROG, the pulse to be measured is split into two replicas, delayed relative to each other by τ, which is then varied. The replicas are then overlapped in a nonlinear medium. Virtually any nonlinear process with an instantaneous response will do [159], although it does matter for the retrieval which one is chosen. The signal field is a spectrally resolved autocorrelation of the two pulses [157]

$$I_{\mathrm{sig}}(\omega, \tau) = \left| \int_{-\infty}^{\infty} \mathcal{E}(t)\mathrm{g}(t - \tau)\mathrm{e}^{i\omega t} dt \right|^2, \tag{4.11}$$

where $g(t-\tau)$ is a gating function. This function is different for different nonlinear processes. For SHG FROG it is just the electric field of the delayed pulse $\mathcal{E}(t-\tau)$, whereas for cross-correlation FROG (XFROG) it is the electric field of the reference pulse $\mathcal{E}_{\text{ref}}(t-\tau)$, which is a different pulse to the one that is being measured. Whatever form the gate function takes, as τ is scanned across the range the signal intensity changes, revealing the overlap of the gate function with $\mathcal{E}(t)$, producing the nonlinear signal only in the delay region where the two pulses overlap. In an autocorrelation we stop here, and estimate the pulse duration from the signal produced. In order to retrieve $\mathcal{E}(t)$ complete with spectral phase from this autocorrelation an iterative algorithm is used to perform a two-dimensional phase retrieval on the FROG spectrogram given by Equation 4.11, which is a solved problem that always has a unique, determinable solution provided the mathematical form of the signal field is known [157, 159].

4.4 Detection of harmonics

In both of the experiments that will be described at the end of this Chapter it is necessary to detect photons in the XUV and soft X-ray ranges. Detection of photons in these wavelength ranges can be a challenge, as they are absorbed easily by just about every material. Instruments capable of registering these photons in a useful manner also tend to be fairly expensive as they require careful manufacture. The absorption problem means that all optics used in the spectrometer need to be grazing incidence (as that is the geometry that leads to the least amount of absorption of photons), which leads to most X-ray optical setups being very large, and difficult to fit in a laboratory.

The best solution we have to this problem is a spectrometer based on a flat-field grating. This type of grating is both a diffractive and a focusing optic, designed to make X-Ray spectrometers more compact, and thus cheaper and easier to make and use [161–163]. The grating is slightly curved (a typical radius of curvature is on the order of several meters) in order to serve as a grazing incidence focusing optic. In addition, the spacing between the grooves of the grating is varied along the length in order to take into account the radius of curvature, to produce a focus in which a certain range of wavelengths are linearly spaced along the detector [161–163]. The two gratings that were used are both 1200 grooves/mm gratings made by Hitachi. The only major difference between them is their radii of curvature, which lead to different focus positions, suiting each grating to the geometry of the spectrometer that houses it. Apart from this, the spectral response of the two gratings is very similar.

The spectrum observed is the first order diffraction from the grating. The second order is in most cases too weak to be observed, and the zeroth order, which anyway contains no spectral information and also contains the most energy, is blocked to prevent it from scattering and blinding the detector. As can be seen in Figure 4.5, at an incidence angle of 87° to the normal (the conventional way of quoting angles onto flat field gratings)

the field is flat on both gratings between at least 50 eV to 150 eV. Above 150 eV the deviation is only very slight, and although the deviation from flat field is larger in the region below 50 eV (which would preclude us from resolving sub-eV features, for instance) it is flat enough for our purposes.

Our X-ray detectors are mounted on translation stages which allow them to be moved easily to view different positions of this spread of X-rays, and facilitates the measurement of the harmonic spectrum across the full photon energy range of the generation. However, this means that for different detector positions, the same pixel value can correspond to different X-ray energies. In order to have a reliable conversion of pixel value to wavelength, we needed to find the values of features of known wavelength on the detector whilst the detector is at the same position as it will be for the measurements we wish to take. These fixed points, along with the knowledge that the spread of wavelengths is linear, allow us to perform a linear fit and therefore map every pixel value to a wavelength.

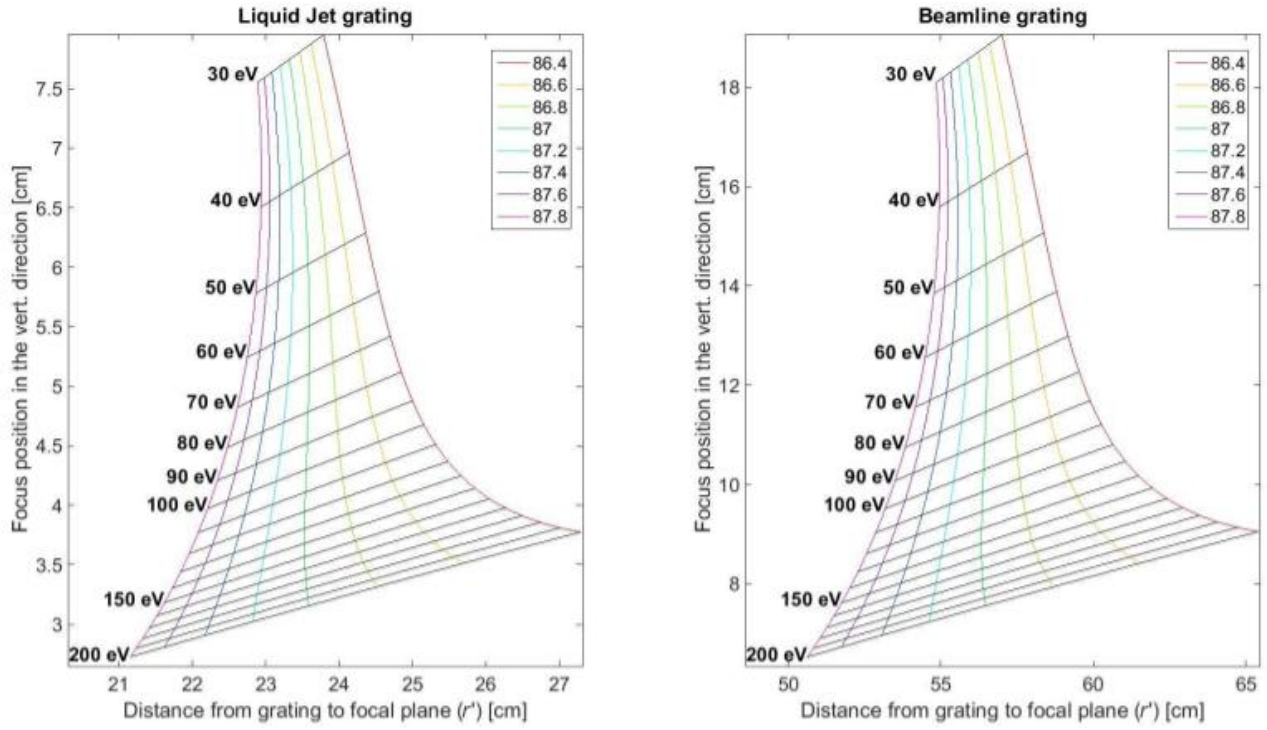

Figure 4.5: Simulated curves of focal plane vs energy for the two gratings used for the experiments, laid out in a similar manner to [161, 162]. The x-axis is the distance from the grating to the focal plane of the light at a given energy, and the y-axis is the vertical position that the light of that energy will focus at. The black lines across the curves are lines of constant energy. The different curves are the focusing results of different incidence angles, indicated in the legend. The regions where the coloured lines are vertical are the flat field regions, which give a well resolved energy separation which is linear with wavelength.

Two types of detector were used to detect the harmonics that are scattered off the gratings, a micro-channel plate (MCP) and an Andor XUV camera.

4.4.0.1 Microchannel plate

A microchannel plate is made of some resistive material with a regular pattern of thin channels running diagonally through the plate. A large voltage bias is applied between the surfaces of the plate, and when a photon is incident upon the surface of one of these channels it triggers an electron cascade through the channel. Our detector has two MCPs in a stack, with the channels at 180° to each other, so the electron cascade from one plate is amplified more in the second plate. The electrons then strike a phosphor screen, which glows at the place that it is struck. A PCO pixelfly CCD camera captures images of the screen.

The MCP-phosphor-camera detector system is mounted on a vertical translation stage to allow it to detect different photon energy regions diffracted by the flat field grating. It is 4 cm wide, which allows for a fairly wide range of energies to be detected, making it a good detector for capturing the entire harmonic spectrum. However due to the nature of the electron cascade it is impossible to get a quantative relationship between the intensity of the signal on the phosphor screen and the intensity of the beam incident on the MCP stack, so we cannot measure an absolute photon flux with this device. The MCP also produces a significantly larger number of dark counts than the camera, and this increases if the voltage is increased to make the detector more sensitive.

4.4.0.2 X-ray CCD camera

The other detector used is an Andor Newton SO XUV CCD camera. This detector has a very low noise floor, which can be reduced further using a built in cooling circuit that can cool the detector chip down to −60°C, although this requires a cold trap to be used in the spectrometer chamber to prevent vapour condensing on the chip and breaking it (see Chapter 5). As a CCD camera, it is also produces a signal that is proportional to the number of photons incident upon the chip, so can be used for photon counting. This detector is therefore more suited to measurements where small changes in signal need to be observed, and can be used to quantify the absolute photon flux produced by our HHG apparatus.

4.5 Liquid sheet target

In addition to the detector, the other common feature of the two experiments is the liquid sheet jet that was used as a target. The details of how our liquid sheet is created can be found in Chapter 2, here I will concentrate on the infrastructure necessary to use this sheet as a target in experiments. This infrastructure has a developmental history based on the particular needs of the experiments at the time. First, during the development of the nozzle, a variant on the high pressure liquid chromatography (HPLC) pump driven system was used, which is described in Galinis *et al.* [73]. However, implementing this system in the lab proved difficult, for reasons I will go into in the discussion below. We therefore moved to the hydrostatic gas pressure driven system for the liquid jet

harmonics experiments, which worked well. However, concurrently with finishing the liquid jet harmonics work, I was designing and building a new target chamber to house the jet for the transient absorption work, and certain aspects of the hydrostatic gas pressure system would render that experiment infeasible. We therefore rebuilt a new version of the HPLC driven system for use with the new chamber.

4.5.1 Running the jet

In order to create a liquid jet from our nozzle, the liquid has to be forced through the nozzle with sufficient momentum to create a liquid sheet of the desired size and thickness. In practice this requires the liquid to be pressurised, as gravity or even the negative pressure of a vacuum chamber is not capable of pulling the liquid through at sufficient speed. In the lab two different methods of pressurising the liquid have been used — static gas pressure forcing the liquid through the nozzle, and an HPLC pump pumping the liquid through.

The other element to consider when running the jet is what to do with the liquid after it has passed through the nozzle. The two different pressurizing methods have different methods of dealing with the liquid, which stem ultimately from whether the system as a whole has the ability to recirculate the liquid whilst the jet is running. However there is a common element that can be dealt with before we turn to the specifics of the two systems, which is the collector tip. The liquid after it has passed through the nozzle is then flowed into an aperture in the tip of a cone, based on the design in Charvat *et al.* [52]. This aperture is larger than the one in that paper, being around 500 µm in diameter, which reflects the larger size of the liquid jet that needs to enter it. However, this larger aperture allows more gas to backstream into the chamber, which can have the effect of destabilising the liquid sheet, causing it to flap and spray all over the place. Due to this a second cone with a much bigger aperture was added around the first cone (see Figure 4.6), which reduced the amount of gas flowing from the reservoir back through the primary aperture.

4.5.1.1 Hydrostatic pressure driven system

The Hydrostatic pressure driven system is the simpler arrangement of the two, and is easy to set up and use. As shown in Figure 4.7, a gas bottle is used to pressurise a 2 L reservoir of the liquid to be used. This forces the liquid up through the nozzle, resulting in a liquid sheet jet. Anywhere between 2 and 5 bar of pressure can be used to form the sheet, since as the nozzle quality degraded more pressure was needed to form a sheet of similar dimensions. A new nozzle typically formed a nice sheet at around 2 bar pressure, but a severely worn one might form a similar sheet at 5 bar, at which point the Reynolds number is almost certainly high enough to lead to instabilities in the sheet [64]. Nozzles typically have a lifetime of a few months before reaching this point. The pressure required for a sheet also depends on the pressure in the vacuum chamber housing the jet — lower chamber pressure leads to a larger sheet for the same gas pressure.

Loading a liquid sample into this setup is somewhat more complicated than forming the liquid jet, as a large

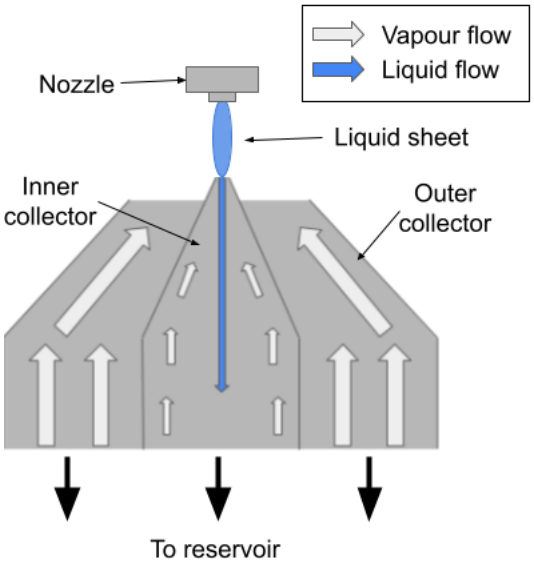

Figure 4.6: An illustration of the idea behind the double walled collector. The white arrows represent the backflow from the reservoir into the liquid jet target chamber, and the blue arrow represents the flow of the liquid into the reservoir. It is much easier for the vapour to stream out of the outside collector than it is for the vapour to get out of the inner one, which lowers the pressure differential between the reservoir and the vacuum chamber, and effectively prevents significant backstreaming out of the inner collector, which can disrupt the jet. Unfortunately, depressing the pressure in the reservoir leads to significant evaporation, so this system was only used when evaporation of the liquid was not an issue.

volume of sample needs to be loaded and as the reservoir is not designed to allow gas to escape (in fact it is designed to not let gas escape) the excess gas in the reservoir needs to be removed by a vacuum pump. The high pressure gas inlet is blocked off, and excess pressure over atmosphere is released via an outlet valve. The reservoir is then evacuated to a very rough vacuum (around 1 mbar or so, depending on the amount of liquid left in the reservoir — the vapour pressure of isopropanol is around 50 mbar at 20°C, so this affects the ultimate pressure possible in the reservoir a lot) and then the liquid is loaded in batches via a funnel and a tap (see Figure 4.7), taking care to only allow in liquid, and not the lab air as well.

Due to this disruptive loading procedure it is not possible to recirculate the liquid whilst using this system to create a liquid sheet. Reclamation of the used liquid is possible, as it is collected in a reservoir under the chamber. However, a high vapour pressure liquid like isopropanol has a tendency to evaporate when stored in a chamber that is exposed to vacuum, especially when the tap to the secondary aperture is open, allowing more vapour back into the vacuum chamber. It often was not worth the effort of attempting to reuse the isopropanol as by the time we came to try and reclaim it there was nothing left. One of the primary motivations for moving back to the HPLC pump driven system was an attempt to avoid inefficiencies like this and the funnel loading system, as we wanted to use a solvated molecule which was much more expensive than low grade isopropanol, and reclaiming it in crystalline form from the reservoir is not ideal.

In the end, the big advantage of this system is its simplicity. It is easy to use (turn the gas on and the jet

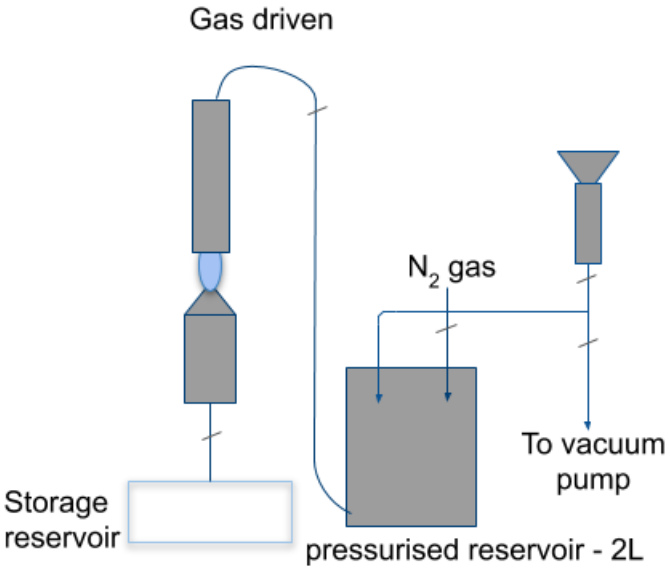

Figure 4.7: Illustration of the gas driven liquid jet system. Between 2 and 5 bar static pressure of inert gas (typically nitrogen) is applied to a 2 litre steel reservoir, which then forces the liquid out and through the nozzle. The liquid is then collected in a storage reservoir under the chamber, from which it can be collected and re-used. To refill the reservoir, the reservoir must be blocked off from the gas and pumped out using a vacuum pump. The liquid can then be added via a funnel. This system can run for 2–3 hours without needing a refill, but the need to stop to refill makes longer scans infeasible. The fact that the storage reservoir is exposed to the vacuum also is a problem, as this leads to a lot of evaporation which raises the base chamber pressure and wastes sample.

comes out) and it is simple to maintain as the only points of failure in the three years it has been used have been the nozzle itself, or the vacuum pumps if a large amount of liquid has made its way into them. With a 2 L reservoir the jet can be run for around two hours before running out of liquid, which is enough for an experiment that has a large signal to noise ratio.

However, the big disadvantage of this system is that it cannot recirculate. This means that either the experiment can only last for two hours, or must be paused two hours in and then restarted, a tricky operation that often requires re-optimisation of the jet position. This arrangement also requires a large volume of liquid to be used, which is not only wasteful but could be a hazard and expensive as well. Ultimately it was a compromise between the ideal system and a system that could be made to work quickly, and for that it worked admirably. However, the ultimate aim of the jet was to do transient absorption measurements on solvated molecules, and for that a more sophisticated system was needed. This hydrostatic system was used for the liquid jet harmonics experiments, whilst the HPLC driven system was used for the transient absorption experiments.

4.5.1.2 HPLC driven system

This more complex system has been implemented twice, as mentioned above. It was the original system shown in our paper [73], but when we tried to migrate it across to the laser lab the HPLC pump failed to build up

sufficient pressure to recirculate the liquid. At the time we were unable to find the cause, and time pressure encouraged us to press ahead with a simpler system. On reflection, the cause was probably that the lower reservoir, where we were collecting the liquid for recirculation, was both sealed and too long and flat.

Inside an HPLC pump is a piston, connected to a chamber with two one way valves (see Figure 4.8). This piston is driven in and out of the chamber by a cam. As it is drawn out it sucks liquid through one of the valves into the chamber, and then as it is pushed back in it squeezes liquid out of the other valve, and onwards into the system. It is a pump that is capable of driving relatively small amounts of liquid at very high pressure, which is perfect for the original application of high pressure liquid chromatography. However, if the reservoir the liquid is drawn from is sealed, the liquid will be prevented from entering the chamber as that would create a vacuum, and if the reservoir is flat (or worse, actually underneath the pump as it was in this case) then there is no gravity assistance for the liquid to be pushed into the chamber. Therefore no liquid is drawn up, and no pressure is generated to push the liquid around the system.

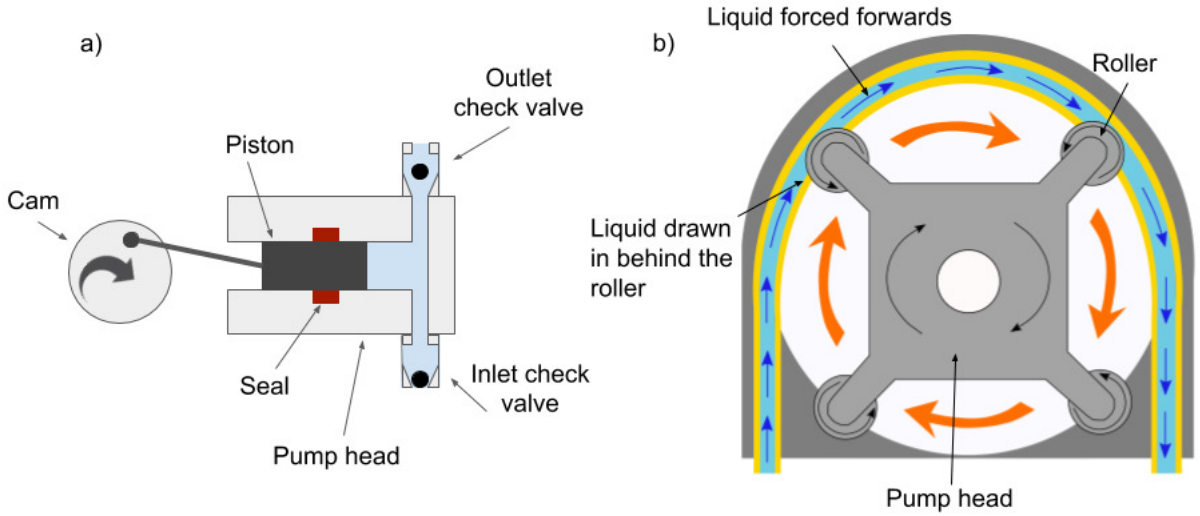

Figure 4.8: Diagrams of the mechanisms of an HPLC pump (a) and a peristaltic pump (b). In the HPLC pump, a cam rotates and pushes a piston in and out of a chamber. Liquid can only enter this chamber through the lower check valve, and can only exit through the upper check valve, so the net effect of this motion is to pump liquid from below the pump to above it. Liquid is prevented from entering the rest of the pump by a seal around the piston. In the peristaltic pump, the pump head rotates turning a series of rollers which squash a flexible piece of tubing. As the roller moves around, liquid in front of it is trapped by the pinching of the tubing, and is forced forwards. Meanwhile the tubing that was previously pinched by the roller opens up, and liquid behind the roller is drawn into it in order to prevent a vacuum. Thus liquid is drawn into and pumped out of the pump head.

When the system was rebuilt, to avoid this problem we followed the solution of Riley *et al.* (who themselves followed Charvat *et al.*) [47, 52]. The reservoir which the HPLC pump draws from is kept at atmosphere. This ensures that the pump can always generate the required pressure to keep the system operating at full capacity, whilst also having the beneficial side effect of being able to load the system whilst the jet is running. This is fed by a peristaltic pump, connected to the primary aperture of the collector. A peristaltic pump works by

squeezing a tube with a roller (see Figure 4.8). As the roller moves along the tube any liquid trapped inside is forced along, whilst in the wake of the roller liquid is induced to flow to fill the gap. Whilst this pump also suffers from the same problems that the HPLC pump does in respect to evacuated reservoirs, it has the advantage of being smaller and less complicated. This has allowed us to position the peristaltic pump directly below the collector, thus using gravity to directly feed the pump whilst simultaneously removing the liquid from the evacuated region as soon as possible, to guard against evaporation.

The ease of recirculation also depends heavily on the vapour pressure of the liquid, For a low vapour pressure liquid like water the rate of evaporation is slow, so a recirculation system will work inside a fairly broad range of parameters as most of the liquid will reach the pump. However for a high vapour pressure liquid like the alcohols we used the rate of evaporation is high, so if the liquid has to travel far to reach the peristaltic pump due to the geometry of the vacuum chamber, then most of it will have gone by the time it gets to the pump, rendering recirculation difficult.

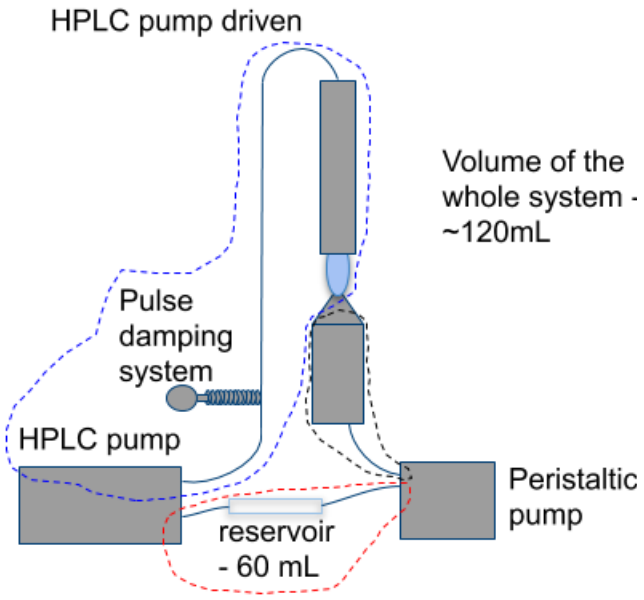

Figure 4.9: Illustration of the HPLC driven recirculation system. Liquid is forced by the HPLC pump up and through the nozzle, which is then collected and fed into the peristaltic pump. This pump then pushes the liquid into the atmospheric pressure reservoir, from which it is pulled by the HPLC pump. A pulsation damping system consisting of a pressure guage and a length of corrugated tubing is attached to the high pressure side in order to prevent pulsing of the jet [164]. This system should be much less vulnerable to evaporation, as much less of the liquid is exposed to the vacuum at any one time. In addition, the reservoir can be exposed to the atmosphere, allowing sample to be topped up whilst the experiment is running. It was also designed to be as small volume as possible, to minimise the amount of whatever target molecule we are using that needs to be used. The three pressure regions are shown on this figure by dashed lines — the areas with red dashed lines are at atmospheric pressure, the areas with blue dashed lines are at high pressure, and the black dashed lines indicate the low pressure region of the pumping system.

The system then has three pressure regions: 1) a high pressure one between the HPLC pump and the nozzle where the liquid is driven to form the sheet, 2) a low pressure one between the nozzle and the peristaltic pump where evaporation is an issue (especially if high value solute is used, as with evaporation there is always a risk of the solute crashing out and crystallising on the walls of the chamber) and therefore this length is minimised as much as possible, and 3) an atmospheric pressure region between the peristaltic and HPLC pumps, which is where the liquid is stored and pumped from. All this is laid out in Figure 4.9.

An issue that arose in the original system was pulsation of the sheet due to the action of the HPLC pump. The periodic squeezing of liquid in the chamber of the pump causes a periodic pressure variation, which was reflected very visibly in the jet, which could reduce in size by up to a third. This was alleviated in both iterations of the system using a technique proposed by Ventura *et al.* [164], which uses a length of corrugated steel pipe with a pressure gauge at the end situated at right angles to the main feed pipe (see Figure 4.9). The idea is that pressure fluctuations now have two directions in which they can travel, and the pressure gauge acts as a damper to wash out the fluctuations that hit it. This is very effective at damping the oscillations (they are no longer detectable with the naked eye, a big reduction). However, it adds complication to changing the sample, requiring the pulsation damping system to be flushed out separately.

The main aim of rebuilding this setup, apart from and in fact dependent on allowing proper recirculation of the liquid, was to reduce the volume of liquid inside the system as much as possible. This was done to reduce both risk and cost, as the risk of exposure to dangerous chemicals and the cost of buying said chemicals both rise with the amount of these chemicals necessary for the experiment. The total volume of the system is estimated to be around 120 mL at most, which compares favourably with over 2L for the previous system. In addition the loading funnel has been replaced with a Luer lock system for loading the sample with a syringe, which should reduce spillage, lowering risk and cost further.

4.5.2 Integrating the liquid sheet jet with the vacuum system

Working at vacuum is necessary for the experiments reported in this thesis because light at the wavelengths generated by HHG is very easily absorbed by all materials, including air. Therefore in order to have any signal at all to detect it is necessary to house the experimental apparatus downstream of the HHG source in vacuum chambers, which significantly increases the complexity of performing experiments. In particular, it presents a massive headache when it comes to using liquid targets.

Unlike solids, the vapour pressure of liquids is usually much larger than the background pressure we would like the vacuum chambers to be held at, leading to rapid evaporation of the liquid. On the other hand, liquids are both denser and harder to pump with conventional vacuum pumps than gasses, leading to a rapid rise in pressure due to this evaporation that is difficult to counteract by merely sticking a larger pump on the chamber (although with a large enough pump this would work). A large rise in pressure, in addition to increasing the

absorption of the soft X-rays, can also play havoc with sensitive vacuum optics by condensing on their surface.

The vacuum chamber therefore has to be able to keep the pressure in itself as depressed as possible, whilst protecting fragile optics that are housed in adjacent vacuum chambers. These objectives are achieved using a cold trap and differential pumping apertures respectively. Two different vacuum chambers were used for these experiments (one for the liquid jet HHG and one for the liquid jet transient absorption) and whilst both chambers were slightly differently set up, they are fairly similar in operation. I will therefore only discuss the target chamber that was used for the transient absorption experiments, as I spent significant time designing and building it for this purpose.

4.5.2.1 New target chamber

This chamber is a large six-way cross, with 13 ¼" diameter arms. One horizontal arm is taken up with the main pump for the chamber, an Adaixen ATH1603M turbo pump by Pfeiffer. It has a pumping speed of 1400 L/s for N_2, and is backed by a Leybold Leyvac 80 screw pump, with a pumping speed of 22.2 L/s. Together they are capable of achieving a base pressure in the chamber of mid 10^{-8} mbar (with no liquid present). In each of the horizontal arms a breadboard is placed, for ease of use for placing diagnostics, stages and other useful things inside the chamber. The flange on the arm opposite the turbo pump has a hatch with a large viewport attached to it. Provided by Kurt J. Lesker, this hatch arrived to us slightly scuffed, probably due to rough handling during shipping. Despite this it seals well, as the aforementioned base pressure is with the hatch attached.

The liquid jet assembly is mounted to the top flange of the chamber by an xyz-translation stage (Vacgen Miniax), and the collector assembly is mounted to the bottom flange by another xyz-translation stage (Thermionics Northwest EC series). This allows us to translate the position of the jet with respect to the laser, as well as realigning the jet position relative to the collector if necessary. A breadboard is mounted outside the viewport with two cameras and a powerful LED mounted on it. One camera images the jet from an angle of 90° to the laser path, whilst the other one images it via a rectangular mirror from an angle that is almost parallel to the beam path. Having images in both the yz- and xz-planes allows us to align the jet to the collector tip with almost no difficulty.

4.5.2.2 Differential pumping

The formula for working out the conductance of a pipe in the limit of molecular flow ($d \cdot \bar{p} < 10^{-2}$ — the regime in which we should be operating, where d is the diameter of the pipe and $\bar{p}$ is the average of the pressures at either end of the pipe) is [165]:

$$C = 12.1 \cdot \frac{d^3}{l} \text{L/s},$$

where l is the length of the pipe. What this means is that in order to minimise the conductance between the target chamber and other parts of the vacuum system that are sensitive to high pressure it is therefore important

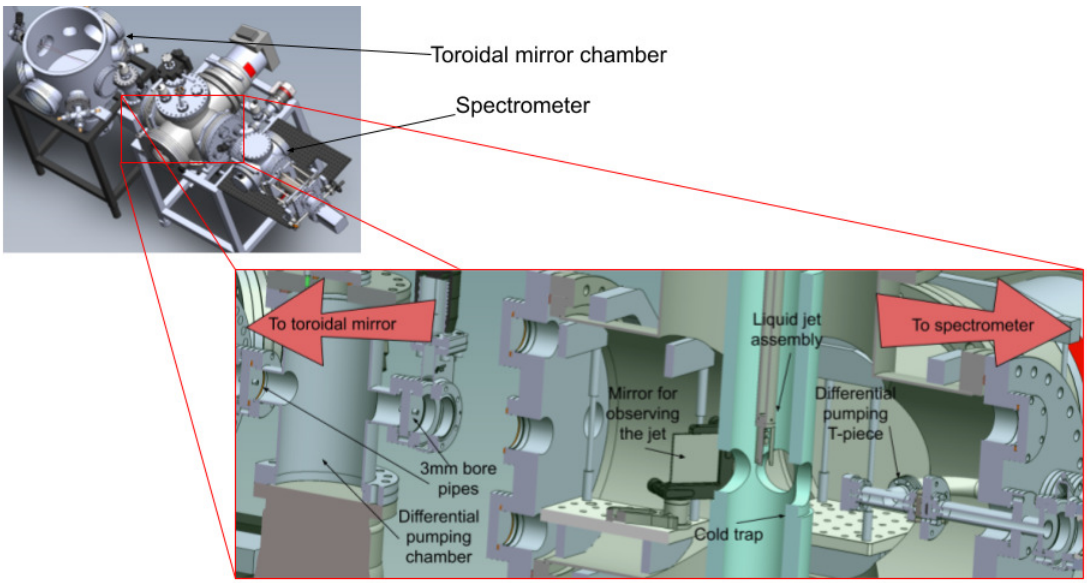

Figure 4.10: Cross section of the new chamber, showing the liquid jet nozzle housing, cold trap jacket and differential pumping apertures. The laser propagation goes from left to right across the image. The two pipes either side of the left hand most chamber are the differential pumping apertures in front of the jet, with a small chamber with a turbo pump attached between them to aid in removing any gas that does get through. On the right hand side of the liquid jet there was no room to put a similar chamber, so a small T-piece was mounted inside the chamber with an external turbo pump attached via a bellows, and the differential pumping apertures were attached to that. Image adapted from a drawing by Clement Ferchaud

to minimise the diameter of the aperture, and to increase the length — i.e. in order to protect the integrity of the vacuum in neighbouring vacuum chambers we needed to install thin pipes between them. We settled on a pair of 3 mm inner diameter 4 cm long pipes with a small chamber between them on either side of the target chamber as the best solution to this problem (see Figure 4.10). This set of specifications is sufficient to maintain a large pressure differential between the chambers (under testing pressures of below 10^{-6} mbar were observed in the chambers either side of the target chamber whilst the jet was running, with a pressure of around 10^{-3} mbar in the target chamber), whilst also remaining sufficiently open to allow for the laser to pass through with a minimum of fuss.

On the side closest to the toroidal mirror (the left hand side in Figure 4.10, see section 4.7) this design fits perfectly, however on the other side there was not sufficient room. This was due to the fact that the grating used in the detector requires a distance of 56.4 cm from the source of the light to the grating in order for the flat field to work optimally. There simply was no room for another chamber, it would have made the path from the source (the jet) to the grating too long, so we mounted our differential pumping apertures on either side of a T-piece, and attached a turbo pump to the T-piece, in a similar manner to the setup in Lorek *et al.* [166]. To aid with alignment through all four differential pumping apertures the front aperture on the chamber was made able to move normally to the laser propagation axis, to allow a certain degree of alignment of the differential pumping onto the pre-existing laser path.

4.5.2.3 Cold trap

Whilst we have the differential pumping sections to protect the vacuum optics in the adjacent vacuum chambers to our liquid jet target chamber, they will not work if the pressure in the target chamber rises too high due to the evaporation of liquid. In order to counteract this we have a cold trap, which is also mounted on the top flange of the chamber, along with the liquid jet assembly. (This incidentally is how we made the target chamber to be multipurpose — all the (admittedly bulky) liquid jet apparatus has been mounted on the top flange, so to swap to a gas or solid phase experiment all that is needed is to exchange the top flange). The conventional pumping on the chamber (turbo pumps backed by a screw pump) does not pump liquid vapour very well at all, and the pressure in the target chamber with the jet running but without the trap cooled very quickly rises to the millibar or tens of millibar range, even with all the pumps on.

The cold trap is a steel cylinder with hollow walls, 320 mm long, with an inner diameter of 73 mm and an outer diameter of 122 mm. It can be filled with liquid nitrogen in order to function as a cryopump, 'pumping' the isopropanol vapour by freezing it to the surface of the trap. The difference between the temperature of liquid nitrogen and the melting point of isopropanol is such that a very thick layer of ice can be formed before the cold trap loses effectiveness all together (see Figure 4.11), although it does need to be topped up frequently to maintain cryogenic temperatures and avoid the ice melting. After a couple of hours however the ice gets too thick for the rate of freezing to outweigh the rate of melting and evaporation, and the pressure in the chamber creeps steadily upwards.

The pumping speed of a cryopump where the proportion of molecules that hit the surface can be assumed to be trapped is 100% can be estimated by [165]

$$S = A\sqrt{\frac{RT_{\mathrm{G}}}{2\pi M}},$$

where A is the inside area of the cold trap, R is the gas constant $8.314\,\mathrm{J\,mol^{-1}K^{-1}}$, T_{G} is the temperature of the gas and M is the molar mass of the gas. Assuming that the liquid vapour has a temperature of 295 K (roughly room temperature) the pumping speed of our cold trap is $3.42{\times}10^3\,\mathrm{Ls^{-1}}$, which is just over twice that of the turbo pump.

4.6 HHG from a liquid

The first experiment that was carried out using the liquid jet was high order harmonic generation from a liquid sheet. Harmonic generation from gases is well known and understood, and harmonic generation from solids is becoming better known and understood (see section 3.4). However, there has been little investigation of high order harmonic generation in liquids. The first reported harmonics generated from a liquid was Zürl and

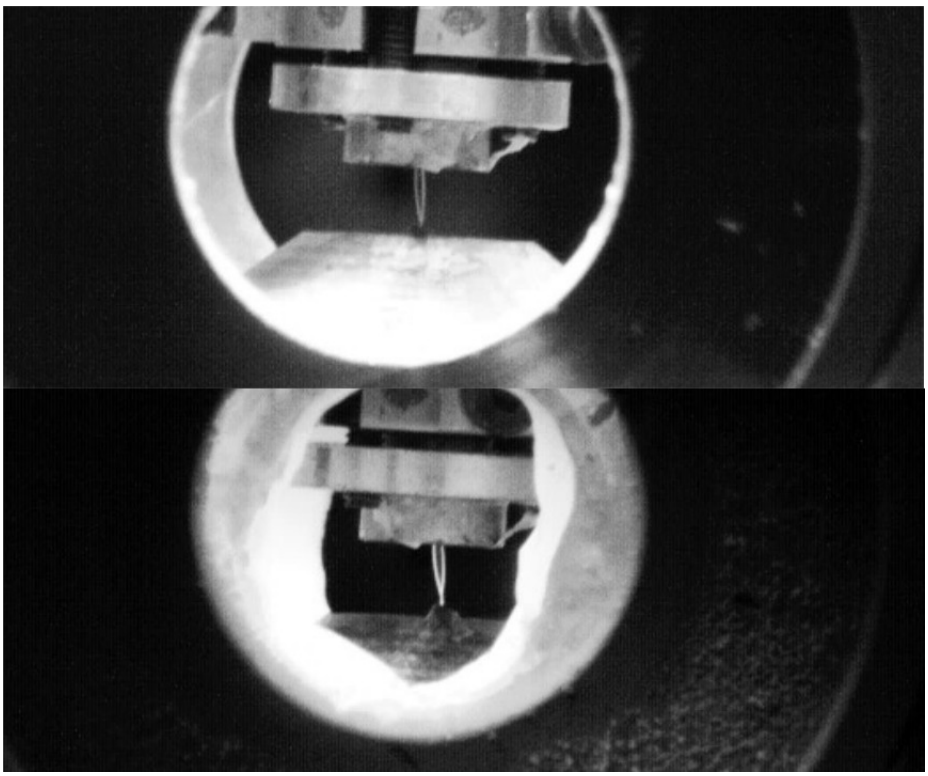

Figure 4.11: Before and after image of a fill cycle of the cold trap. The ice can build to a fairly thick layer before the trap starts to lose effectiveness. Images taken by Sebastian Jarosch.

Graener in 1998 [167], who generated up to the 7th harmonic of roughly 3 µm light in chloroform, apparently accidentally whilst attempting to probe its vibrational resonances. Subsequntly DiChiara *et al.* in 2009 [75] generated harmonics from H_2O and D_2O with a 3.66 µm laser in an attempt to follow this up and see if HHG proper was possible in a condensed phase target. They did not observe any convincing sign of HHG — up to the 9th harmonic they could explain the generation as LHG, and the 11th and 13th harmonics were observed only when the ionisation started to saturate, and were explained by indirect generation by wave mixing processes.

Heissler *et al.* generated high order harmonics by coherent wake emission from a cylindrical liquid jet in 2014 [79], but as this is a different mechanism to any form of HHG that we could generate it is not especially relevant for our purposes. Kurz *et al.* in 2013 [77] and Flettner *et al.* in 2003 [76] both generated high order harmonics in water vapour produced by vapourising water droplets with a second laser. Again, as this is gaseous water it is not especially relevant for investigating the properties of liquid HHG. However, it is interesting to note that one of the reasons given by Kurz *et al.* for using the droplet setup is that they believed that recollision HHG would be damped out by the density of a liquid phase sample [77], a sentiment echoed by DiChiara *et al.* [75].

Finally, during the course of the experiments I will present here Luu *et al.* published a paper detailing their attempts at generating high order harmonics in a thin liquid sheet [82]. Unlike the others above, and similar to the results to be presented in Chapter 6, Luu *et al.* do find the signature of the HHG mechanism in their results, although there are some important differences to our results. This will be discussed further in Chapter

6.

In the light of the above background (excepting Luu *et al.*), we decided that a good use for our liquid sheet jet would be to attempt to generate high order harmonics from it. After all, we have a lot of experience with HHG from gasses, both atomic and molecular [40, 168], so if it was possible to generate high order harmonics then we should be able to do it. The laser scheme we use for generating high order harmonics is relatively simple and was easily adapted to generate harmonics from the liquid target, with the overall layout as indicated in Figure 4.12. The broadened pulses from the hollow core fibre are compressed by a pair of fused silica wedges and then focused by a 50 cm focusing lens into the liquid jet, which is kept at 90° to the laser propagation direction. An iris is kept partially closed in front of the chamber to clean up the beam and reduce intensity, which is important due to the ease with which large amounts of plasma can be generated. This is easily visible, as the plasma is shot backwards into the vacuum chamber as a large white plume. There is typically a small 'sweet spot' in intensity just before the plume starts to form, where the maximum number of harmonics can be obtained but increased plasma doesn't disrupt the phasematching of the generation.

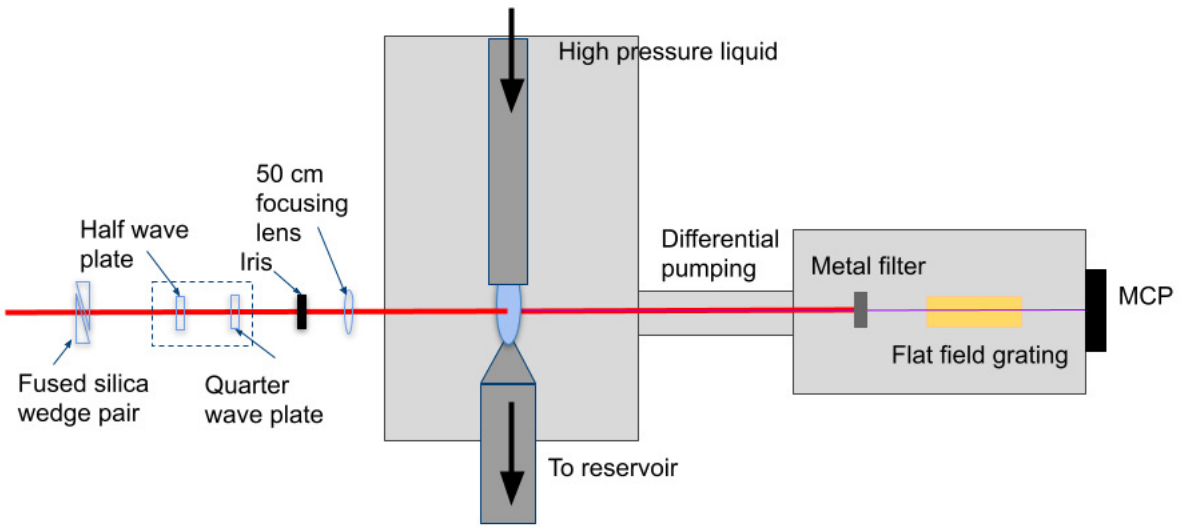

Figure 4.12: Diagram of the experimental set up used. The spectrally broadened pulses from our hollow core fibre are compressed to around 12 fs by a fused silica wedge pair. The waveplates in the dotted line are not always present, but are used in some experiments to alter the polarisation of the pulses. The beam is then focused into the jet to create harmonics. The iris just before the chamber is used to clean up the beam and reduce the intensity to the optimal point. There is then a filter wheel (not shown) containing several metal filters to block the fundamental beam. The harmonics generated are spatially separated by a flat field grating, which projects photon energy (in the region we are interested in) onto space, with a multichannel plate/phosphor screen to pick them up.

Having observed harmonics that bore the signatures of HHG, the next step was to test their properties. As discussed in section 3.4 there are two possibilities for condensed phase HHG, an intra-band mechanism based on Bloch oscillations, and an inter-band mechanism analogous to the recollision model in gaseous HHG. The

difference between them should show up in a scan of CEP — intra-band oscillations generate all frequencies up to the maximum at all points in the laser cycle, whilst the inter-band recombination should have a similar CEP dependence to gas phase HHG [20, 110]. Solid phase HHG scales linearly with the amplitude of the electric field, whilst gas phase HHG scales linearly with the intensity of the electric field, due to the fact that the electron is still bound in a conduction band in the solid, so it can only be given the momentum that the conduction band allows it to have. In the continuum the electron is not so restricted, and so can receive the full ponderomotive energy. The scaling of the cutoff energy with increasing intensity should then be able to tell us whether liquid phase harmonics are more solid or gas like. Finally the ellipticity dependence of HHG reveals the extent of delocalisation of the ground state of the electron. *A priori* we would not expect there to be significant delocalisation of the ground state, as that would seem to require a fixed relationship between isopropanol molecules when one of the contrasting features of a liquid vs a solid is the lack of any local ordering. However, on the time scales of our driving laser the nuclei hardly move at all, so it is not clear on this basis whether there should be much difference between a liquid and an amorphous solid. In addition Panman *et al.* have suggested that alcohols may be more locally ordered than may otherwise be suspected due to their ability to form chains of hydrogen bonds [169]. All of these were measured, and the results will be discussed in Chapter 6. Here the procedure followed in these experiments is outlined.

4.6.1 CEP scan

The CEP scan is the easiest to describe, as it requires no extra optical components. As established in section 3.3.4, the CEP of the idler is consistent pulse to pulse, but its actual value at any point along the optical path is arbitrary and depends on the path length, dispersion and other factors. It is therefore possible to scan through 2π radians of CEP at the focus of the pulse by progressively adding or removing small amounts of glass to the beam path, thereby changing the dispersion by a small amount each time. The easiest way to do this is with a pair of wedges, such as the fused sillica ones being used to compress the pulse. As covered in section 3.2.1, the group velocity (speed of the envelope) and the phase velocity (speed of the propagating electric field) are not usually the same. As the effective length of the path taken through a material of refractive index n by a pulse of light is the thickness L multiplied by n, the change in CEP due to the glass is

$$\Delta\text{CEP} = \frac{2\pi L}{\lambda}(n_\text{g} - n_\text{p}),$$

where λ is the wavelength of the light, and n_g and n_p are the effective refractive indices experienced by the envelope and the carrier respectively. In fused silica at 1.8 µm the refractive indices for the phase velocity and group velocity respectively are $n_\text{g} = 1.4648$ and $n_\text{p} = 1.4409$ [170], which leads to an additional thickness of 77 µm to add 2π radians to the CEP. The group delay dispersion due to this extra thickness is around $-5\,\text{fs}^2$,

which leads to a negligible amount of broadening of the pulse.

The wedges are therefore scanned backwards and forwards over a range large enough to give the required 2π radians of CEP, whilst the CEP is being actively stabilised with the 2f–3f interferometer. The detector is set to observe the cutoff harmonics, which can be observed to move with a π radian dependency.

4.6.2 Intensity scan

In order to observe the effects of changing the intensity of the electric field on the harmonic cutoff two methods were used, closing the iris in order to clip the beam and using a waveplate and a polariser to vary the intensity according to Malus' law. Two methods were used because it was not clear that either of them would give the effect of changing the intensity uncoupled to any other effects. Aperturing the beam to clip the outlying energy and so decrease the intensity adds no extra dispersion to the pulse, but it does add diffraction effects which change the pulse profile and so change the focus geometry, whilst using the half waveplate-polariser technique maintains the integrity of the pulse spatial profile, but adds a lot of extra dispersion to the pulse, affecting the pulse duration. Using both techniques should provide a good correction for both of these effects — if the same relationship is seen for both measurements then we can have confidence in it.

4.6.2.1 Intensity scanning via a variable aperture in the beam

The iris scan was performed by partially closing the iris in front of the chamber in order to clip the beam. The power of the beam after it was clipped by the iris was measured, and was used to determine the extent to which the iris was closed — we aimed for an even spread of powers. The beam was then focused into the jet, and the harmonic spectrum at each power was recorded. Because the iris reduces the power by progressively clipping more and more of the beam, the calculation of the intensity of the beam at the focus not only has to include possible dispersion effect of the optics and liquid vapour in the chamber, it also has to account for the changing focus shape due to this diffraction. In order to simulate this, we multiplied a Gaussian pulse distribution with a circular mask representing the iris aperture, and then Fourier transformed the resulting distribution to obtain the focus spatial profile. From this we can obtain the peak intensity of the apertured beam in the absence of dispersion. The effect of dispersion is then included by considering the group delay dispersions (GDDs) of the optics the beam propagates through, which stretch it a little bit to around 15 fs.

4.6.2.2 Intensity scanning via waveplate and polariser

The waveplate scan was performed using a half waveplate and a wire polariser. First of all the orientation at which the light was aligned along the polariser was found by rotating the waveplate such that the power coming through the polariser was at a maximum. Subsequently the waveplate was then rotated in small steps until the intensity was sufficiently low that no harmonics were observed.

The half waveplate is made of a 2.67 mm thick mixture of layers of quartz and MgF_2, which has a GDD of around $-140\,fs^2$ at 1800 nm and will stretch the pulse to around 33 fs. Along with the other optics in the system, this will reduce all the intensities at the focus by a factor of around $\frac{1}{3}$ due to the increased pulse length and correspondingly decreased peak power. We have therefore ended up with two intensity scans that probe different intensity regimes — the iris scan which probes the intensity in wide steps between massively overdriven harmonics to the threshold where no harmonics are produced at all, and the waveplate scan which probes the intensity in much finer steps in the lower intensity regime.

4.6.3 Ellipticity scan

The ellipticity scan was performed using a quarter waveplate. As discussed in Chapter 3 this can convert the polarisation of light from linear to circular or elliptical, depending on the angle of the fast axis to the linear polarisation of the beam. The initial angle of the waveplate was set by observing the angle which gave the maxiumim harmonic signal (which was assumed to occur when the polarisation was linear), and then it was rotated such that the ratio of the minor to the major axis was varied from 0 (linear) to around 0.2, after which point no more harmonics could be observed. The ellipticity dependence of the signal was found by observing the dependence of the intensity of each harmonic on the angle of the quarter waveplate. The quarter waveplate is thinner than the half waveplate, at 1.4 mm thick, and it has a GDD of around $-60\,fs^2$ at 1800 nm, and will stretch the pulse to only around 18 fs, which meant that the intensity at the focus was more similar to the unstretched intensity and allowed us to observe higher energy harmonics than we were able to in the intensity scans.

4.7 X-ray spectroscopy of solvated molecules

In order to perform X-ray spectroscopy of solvated molecules we need to generate X-rays (using our HHG apparatus), and we need to be able to deliver solvated molecules to the interaction region via our liquid jet apparatus. Over the course of this experimental run we encountered a number of issues, which I will discuss more fully in the next chapter, so here I will outline the plan for this experiment and leave the modifications to the next chapter.

In order to investigate electronic and nuclear dynamics on their characteristic timescales it is necessary to be able to resolve a time delay on a finer timescale. In practice the way this is done is by doing a pump-probe experiment, as discussed in Chapter 1. The time resolution of such an experiment is in theory only limited by the temporal width of the optical pulses used, although in practice many other factors can limit this. Our HHG beamline was set up for performing pump-probe experiments, as shown in Figure 4.13. For these initial absorption experiments we were not using the pump arm of this set up, only the probe arm to generate the

X-rays for our spectroscopy measurements.

The probe pulse is generated by HHG in a neon gas cell. The gas cell is a hypodermic needle that has been crushed at the end and glued to form a gas tight seal. The needle is attached to a pipe which is fed Ne gas via a regulator. In order to ensure the integrity of the vacuum in the chamber in which the needle is mounted, a two stage differential pumping jacket is used to house the needle [40, 171], which is shown in an insert in Figure 4.13. It consists of two stainless steel jackets, one nested inside the other, with perspex windows with a hole for the laser to pass through. The inside jacket is directly pumped by a screw pump of a similar specification to that used to back the liquid jet target turbopump. This pumping ensures that the pressure in the outside jacket is below 5×10^{-1} mbar, whilst the pressure in the needle is on the order of several bar. The outside jacket is pumped by a small turbo pump, which is enough to ensure the main chamber in which the needle target assembly is housed can be maintained at below 1×10^{-4} mbar with the gas needle in operation. In order to allow the gas to flow (and the laser to propagate through the gas), the laser is used to drill a hole through the needle. This avoids problems of aligning the laser to a pre-drilled hole, and ensures the hole is only ever as big as it needs to be.

The driving pulse is the same pulse from the hollow core fibre that was used for liquid HHG in section 4.6. It is compressed with fused silica wedges, and steered into the needle target with a motorised mirror inside the vacuum chamber. The compression and pointing of the driving beam is optimised by observing the harmonics produced on the spectrometer, with no sample present. The assumption on a day to day basis is that a spectrum with the highest cutoff possible, high flux and no visible harmonic peaks indicates that an isolated attosecond pulse has been generated. In particular the lack of harmonic peaks is a key indicator, as the more pulses there are in a pulse train, the more the frequencies in the pulses show up as distinct peaks. The lack of such peaks is an indicator that there are not many, or only one pulse generated. However, none of these heuristics are conclusive, and so additional checks will need to be made, such as a CEP scan (if the driving pulse is long enough to generate a pulse train then varying the CEP should not affect the spectrum).

After the pulse is generated it passes through a hole that has been drilled in a mirror, and is focused into the target by a toroidal mirror. This mirror is a section of a torus which can focus light that impacts it at a grazing incidence, in our case 5°, with a focal length of 1 m. Having passed through the target, the focus is then imaged by the flat field grating onto the XUV camera as described in section 4.4 above.

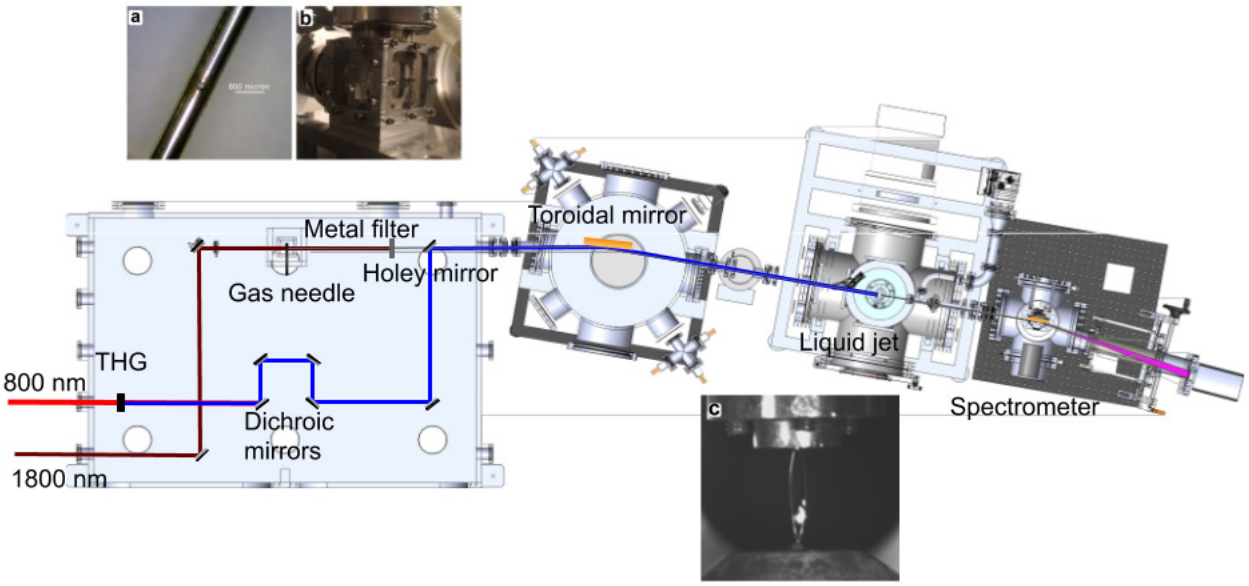

Figure 4.13: The experimental layout of our pump-probe system, including a proposed pump generation line based on the existing layout. The pump beam would be taken directly from our 800 nm laser system, bypassing the OPA system. It then would used to generate 267 nm via third harmonic generation in either a crystal or a gas cell (represented as a black box on the figure). The fundamental beam would then be removed by a set of dichroic mirrors from Eksma coated specifically to reflect 267 nm and transmit 800 nm whilst adding as little dispersion as possible. These dichroic mirrors could also be used as a delay stage, if no delay stage is able to be placed before the third harmonic generation. The probe beam is generated in a gas needle target by the compressed 1800 nm idler beam. Insert **a** shows a needle target after use with the hole drilled in it by the laser, whilst insert **b** shows the double differential pumping jacket used to keep the pressure in the chamber down whilst gas is running through the needle. Inserts a and b are taken from [40]. The idler is removed by a thin metal filter, and then the HHG and UV beams are combined by a UV-enhanced aluminium holey mirror, which is a mirror with a hole in the center. This ensures that the HHG and UV beams are colinear, which greatly eases the alignment of the rest of the system. The beams are then focused at grazing incidence (5°) by a toroidal mirror into our liquid jet target (shown in insert **c**). The X-ray absorption spectrum is then detected by our flat-field spectrometer.

The two molecules we used were an organic molecule called 2-Aldrithiol (dissolved in ethanol) and NaCl (dissolved in water) — for more details of why we used these molecules see Chapter 5. In order to be able to switch back and forth between pure solvent and solvent plus solute we installed an extra reservoir and attached it to the HPLC pump via a three way valve, so that we could select whether the pump was pumping from the recirculating system (which had the solute in it) or the extra reservoir with the pure solvent. We also used this system to switch between ethanol and water in some of the measurements reported in Chapter 2, and from the time it took the jet to change dimensions after making the switch we estimate that it took around 8 minutes for the new liquid to make its way to the jet, due to the lengths of the tubes involved.

Whilst we were taking measurements we were careful to record spectra with solute and no solute, as well as

spectra without the jet present at all. Ideally these would have been taken with the X-rays passing next to the jet whilst it was running, to account for absorption due to vapour, but this proved practically challenging so we settled for measurements before and after the jet had been switched on.

Another challenge was aligning the jet to the X-rays, especially initially. As there was no visible sign that the X-rays were striking the jet we had to align on a decrease in the overall flux of the X-rays. However, as the signal from the X-rays was fairly weak during this experimental run we had to integrate for 30 seconds to get a usable signal, which is not a response time that is easy to use to align anything. In the end we used the green alignment laser to place the nozzle holder into the beam path, then raised the jet apparatus so that the sheet was in the place the nozzle holder was, and then switched back to using the main laser to generate harmonics.

4.8 Conclusion

We constructed and installed a new target chamber, designed by myself in collaboration with others, into our HHG XAS beamline. We installed our flat liquid jet source inside it, and were able to run this jet with a rough vacuum ($\sim 1 \times 10^{-3}$ mbar) inside the chamber, whilst maintaining high vacuum in the surrounding chambers ($\sim 1 \times 10^{-6}$ mbar) due to differential pumping sections. We were able to use this jet as a source for high harmonic generation, and we were also able to use it as a target for X-ray absorption spectroscopy, albeit with some difficulty. We now turn to the results of these two experiments, starting with the X-ray absorption spectroscopy.

Chapter 5

Towards X-ray transient absorption in a liquid

5.1 X-ray transient absorption overview

Since the field of femtochemistry was made possible in the 1980s by advances in ultrashort laser technology, researchers have sought to observe chemical reactions on the shortest possible timescales [3]. However, the processes that can be directly observed by these techniques are limited by the energies of the photons that can be generated by these lasers. The recent advent of ultrashort pulse X-ray sources has allowed experimentalists to access the element specific electronic and structural information contained within X-ray absorption spectroscopy (XAS, see Chapter 3) on a timescale limited only by the length of the X-ray pulses themselves [172]. The most common technique used is pump-probe (see Chapter 4), where an optical pump initiates the dynamics, and the X-ray pulse arrives after a variable delay allowing us to map out the dynamics over the time range of the delay.

Initially the shortest X-ray sources available were synchrotrons (circular electron accelerators, which store electrons by accelerating them round in a circle at relativistic speeds, and then generate X-rays by causing a sudden change in direction in the electron bunch), which are able to make pulses as short as 30 ps. These pulses can be 'sliced' — where the energy of the generating electron bunch is modulated by a femtosecond laser pulse so that the X-ray radiation is released within as short a window as possible. This leads to the ability to make pulses as short as 100 fs, although some fluence is sacrificed [173]. Due to this pulse length limit studies carried out at synchrotrons were not able to observe very fast electron rearrangements, as even the femtosecond sliced pulses were unable to resolve the fastest dynamics [6]. Charge transfer and fast nuclear motion such as ligand rearrangement can take place on a timescale of tens or even single femtoseconds [174]. However synchrotrons have been around for much longer than the new alternatives (the first synchrotron was built in the 1940s), so for a long time they were the gold standard and they continue to be used for longer time scale experiments as

a lot of them have been built, so they are relatively accessible.

Progress in X-ray free electron lasers (XFELs) and lab based techniques have now led to studies with much shorter time resolution being feasible. The lab based technique that we use is HHG, which has been described in greater detail in Chapter 3, and which is now being used in time-resolved experiments with resolutions of single femtoseconds [27]. XFELs work in a similar way to synchrotrons (quickly changing direction of relativistic electrons to generate X-rays), but the electron bunch is accelerated linearly instead of in a ring. It then is wiggled side to side by magnets, generating the X-rays. This configuration allows the X-ray pulses to be shorter and brighter than those from synchrotrons, and using a similar technique to slicing the synchrotron beam (known as XLEAP), XFELs are now capable of delivering pulses with a time resolution of less than 1 fs [11].

Whilst synchrotron and XFEL sources are sufficiently bright that absorption rates are not a problem when working in a condensed medium, HHG sources are a different story. The soft X-rays generated by the HHG source are very readily absorbed by the target, and so the thickness of the target quickly becomes an issue to any XAS experiments performed with this source, and especially transient absorption experiments where the signal that is being sought can be only a small change in an already small feature. If the limiting factor of the experiment is the photon flux, then a large background absorption will decrease this already limited resource, lowering the number of photons that can be absorbed by the target of interest and so decreasing the signal obtained.

On the other hand, an advantage that our method of HHG XAS has over XFELS and synchrotrons is that we can detect a very broad bandwidth simultaneously. At these large-scale X-ray sources the normal mode of operation is to monochromatize the X-ray beam, measure the absorbance at each energy, then move on to the next energy and repeat. This can stretch the pulse duration, and in addition leads to problems of comparison where one has to be careful to ensure the delay points measured at each energy are the same so as to allow sure conclusions to be drawn. Our measurements have no such problems, as with the correct calibration of detector position and grating angle we have the ability to observe multiple absorption edges simultaneously, whilst preserving the duration of the X-ray pulse entering the interaction region.

As indicated in Chapter 1, the sample state of principal interest to ultrafast chemistry and biology is the liquid phase (which for our purposes includes solutions and suspensions), and so therefore being able to bring TR-XAS techniques to bear in a liquid sample would represent a big step forwards in being able to observe ultrafast reactions in as close to their natural state as possible. We therefore designed and built a thin liquid sheet jet, which is thin enough to allow sufficient transmission of XUV radiation for a signal to be observed (details in Chapter 2). We then designed and built apparatus to enable this jet to be run in the vacuum conditions required for XUV radiation transmission (details in Chapter 4).

Figure 5.1 shows the signal generated by our HHG source during the measurements in this Chapter. It was not the highest flux recorded during the time of performing these experiments, but it was recorded on the same

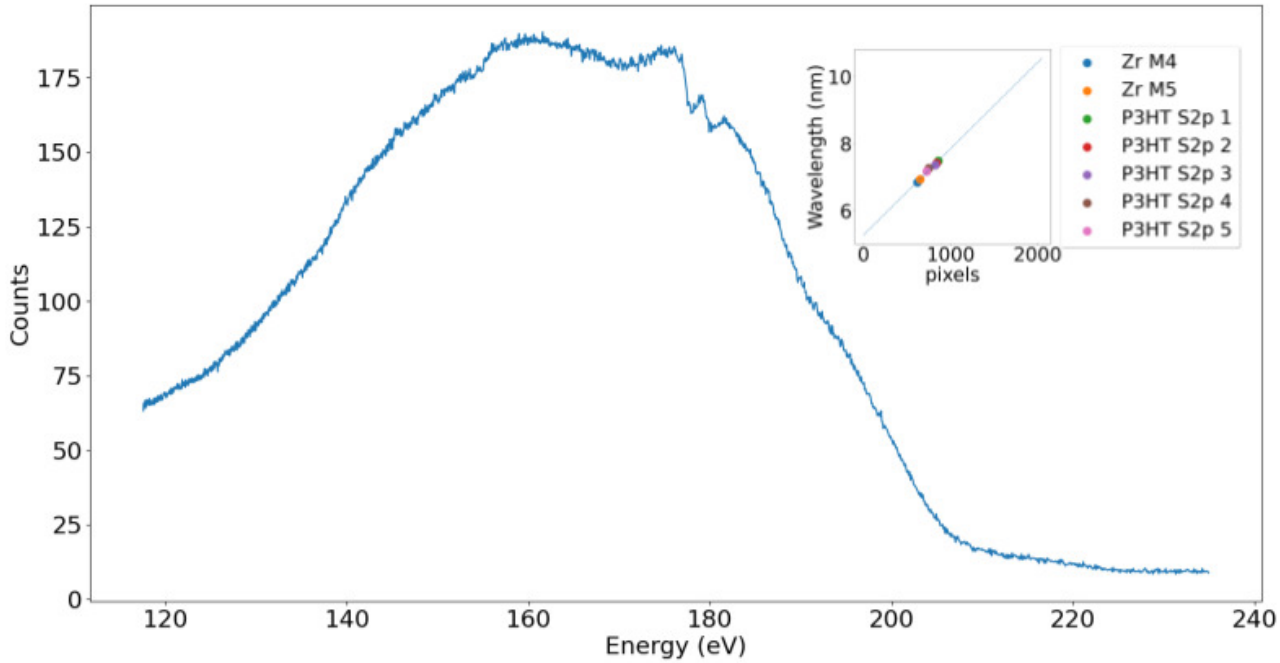

Figure 5.1: X-ray spectrum recorded during the X-ray absorption experiments detailed in this chapter. This spectrum was obtained by integrating for 30 seconds. The two absorption features around 180 eV are the Zr M_4 and M_5 edges, from the filter that we use to block the IR beam after the harmonics have been generated. The x-axis has been calibrated by taking a measurement of poly-3hexylthiophene (P3HT), and fitting a line to the known energies of the observed absorption features (see insert)

day as the data presented in this Chapter and therefore is as accurate a record of the flux in these measurements as can be presented. The maximum number of photon counts at a given energy recorded in 30 s is around 185. If the quantum efficiency of the camera is around 90% [21], this amounts to approximately 7 photons per second per pixel at the most. This is several orders of magnitude lower than the synchrotron and XFEL sources, but it is also several orders of magnitude below the previously reported flux in our beamline [21].

In this Chapter I will outline our preliminary attempts to obtain X-ray absorption spectra of solvated molecules, what problems we encountered along the way and what lessons we can take from these issues to enable us to perform these experiments more succesfully in the future.

5.2 Target molecules

We had identified two candidates as target molecules for these experiments, Aldrithiol and NaClO. Our criteria for deciding on the targets were based on several things. We considered the availability of the target molecule, as well as its toxicity to humans, as our liquid jet system is not yet as efficient as we would like and so the loss of target is quite likely. Therefore a rare or toxic target would be an unwise idea, as a rare target would be difficult or expensive to get hold of, so wasting it to inefficiencies could prove financially unsustainable, and a target that is sufficiently dangerous would pose an unacceptable risk to experimentalists in the event of a leak.

We also considered the dynamics that it was possible to induce in the target. In our lab we have developed a few-fs tuneable deep UV source based on resonant dispersive wave emission [175]. Our plan was to implement this system as a pump source that could induce dynamics in many different organic systems (as most of them have a resonance in the deep UV) without the loss of time resolution that would be caused by the stretching induced by a birefringent crystal if the pump was generated by third or fourth harmonic generation from our 800 nm beam. We would also have the flexibility to tune our pump beam to exactly the resonance we wanted to excite.

Both Aldrithiol and NaClO have fast electron rearrangements upon excitation, the outcomes of which will drive the slower dynamics of the molecule, and ultimately the final outcome of the reaction. Although we were initially aiming just to do an X-ray absorption measurement, our choice of targets was informed by the desire to make a time-resolved measurement in the future.

5.2.1 Aldrithiol

The first target we tried was an organic molecule called 2,2-dithiodipyridyl, otherwise known by its brand name Aldrithiol. Aldrithiol, shown in Figure 5.2, consists of two pyridyl rings joined by a disulfide bridge. The cleavage of disulfide bridges has attracted some interest in the ultrafast experimental community recently, due to their biological relevance [9, 176–178]. Our HHG beamline should be able to resolve the electron dynamics involved in this bond cleavage, as shown by Schnorr *et al.* [9], and we believe that the dynamics involved would be more similar to that experienced in a biological system.

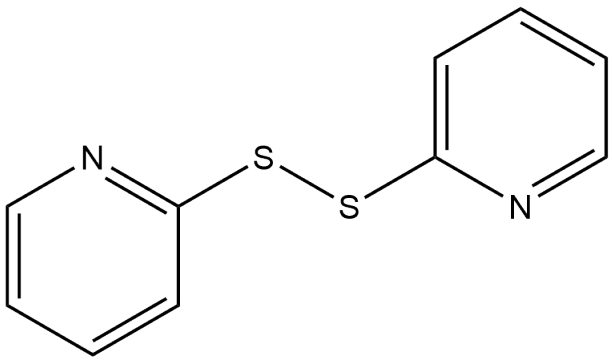

Figure 5.2: A drawing of Aldrithiol-2. The 2 refers to the fact that the nitrogens in the pyridyl ring are in the 2 position rather than the 4 position, which is a different molecule which is also available. The convention of numbering the positions starts by counting round from the carbon attached to the 'highest priority' group (in this case the sulfur). Thus the nitrogen is right next to this carbon, so position 2. Position 4 would be the para position, opposite to carbon 1.

As Aldrithiol has a delocalised π-system the radical monomer should be somewhat more stabilised, which should lead to longer recombination times and possibly some charge transfer around the ring which we could observe. As biological disulfide bridges are found in proteins (which tend to have large π-systems, capable of

stabilising radicals), this behaviour should reflect the cleavage of those disulfide bridges better than dimethyldisulfide (two methyl groups joined by a disulfide bridge — the target Schnorr *et al.* investigated) while still being small enough for the electron dynamics to be tracked. It would also allow us to observe in real time the effect of the aromatic ring when a radical is formed.

The fact that our sample is dissolved in ethanol (at 150 mM concentration) will almost certainly play a large part in the observed dynamics as well, as it would in a biological sample (although the presence of a protein is very capable of altering these things again, and the solvent in that case would be water). The solvent certainly will play a large part in the recombination of the radicals, as if the radicals end up in a solvent cage with each other then they will almost certainly recombine back to the dimer. However if the two radicals end up in different solvent cages then there is little chance of them recombining, especially if the solvent reacts with the radical readily rather than merely stabilising it. We were also planning to analyse the solution after an experimental run, which could have provided us with some more clues with which to make a complete picture of the dynamics in solution.

5.2.2 NaClO

The second target molecule was the ClO^- ion. In the choice of this molecule we were inspired by a paper by Lim, Gnanakaran and Hochstrasser [179], in which the authors perform a UV pump, UV probe experiment on this molecule, with a time resolution of around 120 fs. They observe a sharp peak in the (UV) transient absorption data with a width of around 60 fs, and then a fast and slow decay channel ($\tilde{2}$30 fs and several ps respectively) corresponding to the dissociation or not of the ion.

Our UV source should be able to offer much higher time resolution than this, and as we have access to X-ray spectroscopy as well we could obtain multiple sets of data in the UV and at the Cl L-edge which could correlate into a much finer time resolution than previously achieved. In particular it would be interesting to observe what occurs during the first 60 fs, as it is on this timescale that the negative charge of the molecule needs to shift to the O from the Cl, and the solvent cage is supposed to rapidly disappear as a result.

This molecule was our second choice, as it required an aqueous environment and is also a strong base. This would have led to practical difficulties, as we need to avoid corrosion and the health hazards associated with this molecule make handling much harder. The aqueous jet was also relatively new at this stage, so we preferred to use a target that had successfully run in vacuum before.

5.3 Challenges of performing XAS on a liquid sheet jet

Unfortunately, we never managed to bring these planned experiments to a satisfactory conclusion in the time we had available due to several problems encountered during the course of this experimental run. As already

noted in this thesis, obtaining a stable liquid sheet jet in vacuum is a significant challenge on its own, and we were unable to solve the additional challenges presented by using this target to perform precise measurements of X-ray absorption on solvated molecules in the time we had allotted for this experiment. However, we did learn a lot about the process of performing such an experiment with this type of target, which hopefully will enable us to build on this platform and perform an experiment similar to these in the future.

We faced three major challenges in the course of attempting these experiments — poor signal, an unstable jet and contamination of our vacuum system by the sample molecule. I will explain how we encountered these three challenges and what we did to try and overcome them. Then in the conclusions section I will outline what we could have done differently to avoid these challenges, and some potential improvements to the experimental system which could be tried in the future.

5.3.1 Signal to noise ratio

This was the first challenge that we encountered. The sample used during this experimental run was Aldrithiol dissolved in ethanol. We ran the jet with just ethanol, and then with ethanol containing 150 mM Aldrithiol (4% by mass) in order to obtain a differential absorption spectrum.

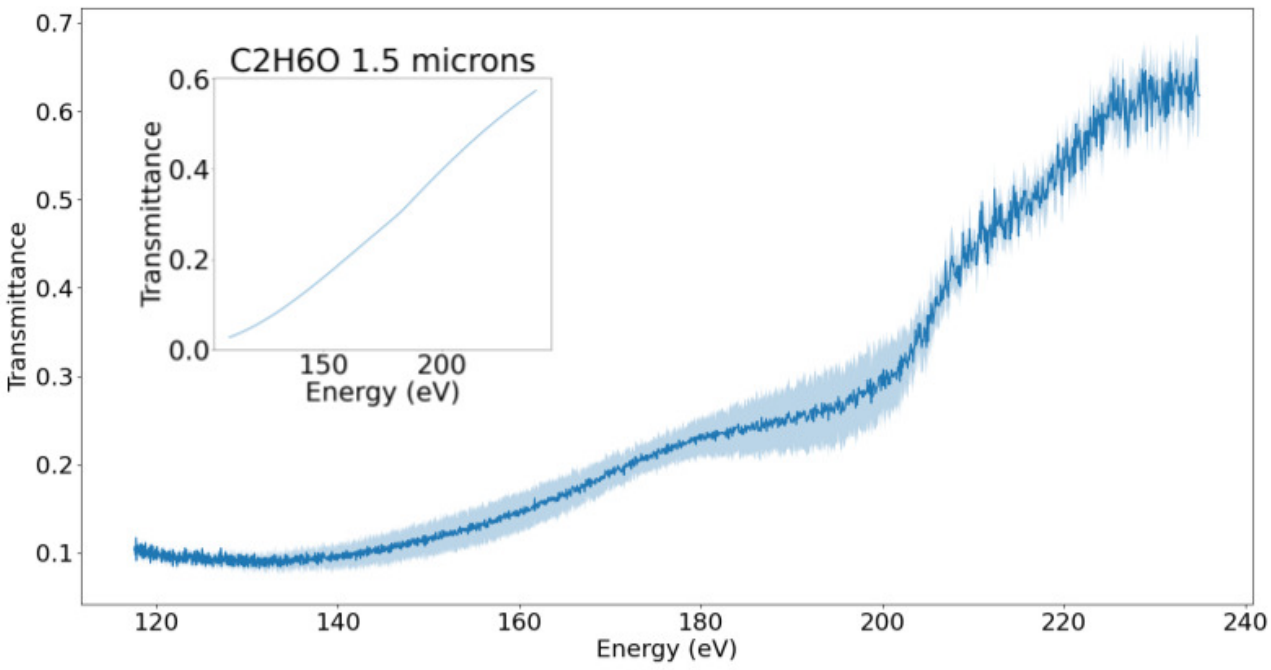

Figure 5.3: The average transmittance of our ethanol sheet. The spectra taken in transmission were averaged, and then divided by the spectrum taken in the absence of the target, shown in Figure 5.1. The shaded area is the standard deviation of this average transmittance. The inset shows data from Henke *et al.* [42] showing the expected transmittance of a 1.5 µm thick sheet of ethanol — it corresponds fairly well to our measured transmittance.

Figure 5.3 shows the transmittance spectrum of our thin flat ethanol jet. By comparison with the Henke data we estimate the jet thickness to be around 1.5 µm, which accords with the other measurements we have

made. Already with a relatively thin jet we see a very large amount of absorption, especially at lower energies. As mentioned above this large absorption background makes these experiments much harder, because we are already limited by the amount of flux our HHG source produces. If a large proportion of these photons are absorbed by the solvent then the flux that the solute 'sees' will be reduced by this proportion, further reducing the signal.

Figure 5.4 shows both the transmittance of the jet containing Aldrithiol, and then the ratio of this spectrum and the average ethanol spectrum. In both differential spectra there is a broad peak between 160 and 170 eV. This is in roughly the correct place for it to be the signal that we are looking for (see Table 5.1 and Figure 5.5), but it is so weak and broad that we cannot draw any firm conclusions from its presence. There is certainly no sign of any absorption peaks in the transmittance spectra on their own.

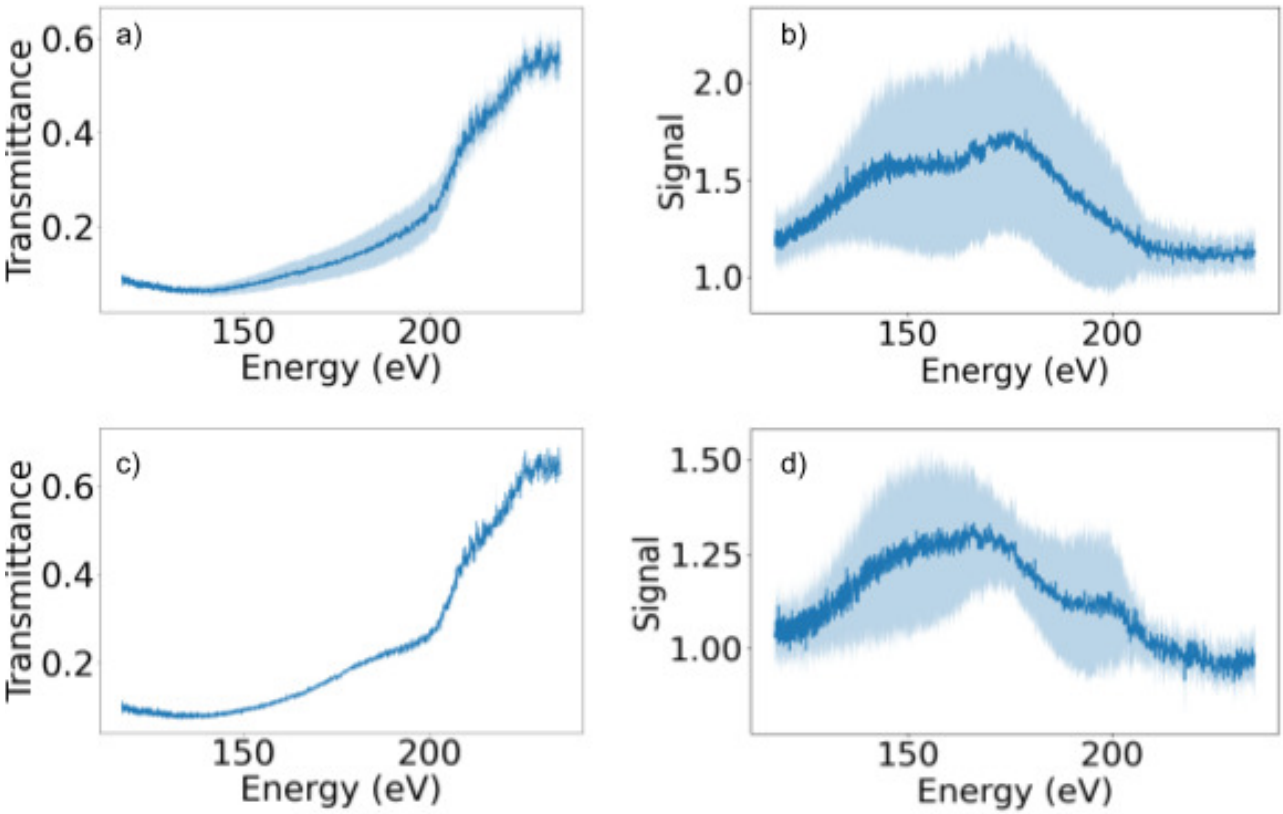

Figure 5.4: Transmittance spectra of the jet with 150 mM of Aldrithiol inside. a) and c) are two different experimental runs, and b) and d) are the difference of Figure 5.3 and a) and c) respectively (subtracted in this way so as to make the signal positive). The two sets of data were taken under slightly different conditions (we took a break in the middle to stabilise the jet), so I have presented them seperately. Regardless, both difference spectra show a peak between 160 and 170 eV. This is at the correct energy, but it is very weak and washed out.

In an attempt to understand why the signal is so weak we calculated the oscillator strengths of the transitions we were hoping to excite in the Aldrithiol. My colleague Jacob Lee calculated them using the ADC(2) method

with core valence separation.

Energy (eV)	Oscillator strength	Character
226.328	0.031	2s $\rightarrow \sigma^*_{\mathrm{CS}}$
224.149	0.051	2s $\rightarrow \sigma^*_{\mathrm{SS}}$
169.189	0.006	2p $\rightarrow \sigma^*_{\mathrm{CS/ryd.}}$
167.178	0.021	2p $\rightarrow \sigma^*_{\mathrm{CS}}$
166.982	0.013	2p $\rightarrow \sigma^*_{\mathrm{CS}}$
166.981	0.013	2p $\rightarrow \sigma^*_{\mathrm{CS}}$
164.952	0.015	2p $\rightarrow \sigma^*_{\mathrm{SS}}$
164.834	0.009	2p $\rightarrow \sigma^*_{\mathrm{SS}}$

Table 5.1: Oscillator strengths of transitions in aldrithiol. Oscillator strengths of calculated transitions of the same energy have been added together. The top two transitions are from the S 2s orbitals to the SS and CS σ^* orbitals, whilst the bottom six are from the S 2p orbitals to these σ^* orbitals along with a Rydberg state that also had some CS anti-bonding character. The energy splitting can be attributed to spin-orbit coupling and the p orbitals being less core-like than might otherwise be assumed.

The oscillator strength is a measure of how likely a particular transition between states is, defined as the ratio of the absorption rate of the system to that of a classical single electron oscillator [180]. In this case the transitions are between n=2 orbitals and the σ^* orbitals of the sulfur sulfur bond and sulfur carbon bonds in Aldrithiol. The oscillator strength is related to the absorption cross section of these transitions by

$$\int_{-\infty}^{\infty} \sigma_{\mathrm{abs}}(\omega)\, d\omega = \frac{\pi e^2}{2 m_{\mathrm{e}} c \epsilon_0} f_{\mathrm{abs}} \tag{5.1}$$

where $\sigma_{\mathrm{abs}}(\omega)$ is the absorption cross section at frequency ω, e is the charge on an electron, m_{e} is the mass of an electron, c is the speed of light in a vacuum, ϵ_0 is the permittivity of the vacuum and f_{abs} is the oscillator strength of the absorption.

The absorption cross section at frequency ω can be expressed as [181]

$$\sigma_{\mathrm{abs}}(\omega) = \frac{q^2}{4mc\epsilon_0} \frac{\gamma}{(\omega - \omega_0)^2 + (\gamma/2)^2}, \tag{5.2}$$

where ω_0 is the resonant frequency of the transition in question, q and m are the charge and mass of the oscillating particle respectively and γ is a damping constant reflecting the lifetime of the state. In what follows I treated this as just the core hole lifetime of the sulfur L-edge (around 6.6 fs [182]).

The oscillator strength as mentioned above can be thought of as the ratio of the oscillation rate of the system with that of a classical single electron oscillator, which can be expressed as

$$f_{\mathrm{abs}} = \left(\frac{m_{\mathrm{e}}}{m}\right) \left(\frac{q}{e}\right)^2 \tag{5.3}$$

We can therefore express the absorption cross section in terms of the oscillator strength as

$$\sigma_{\text{abs}}(\omega) = \frac{e^2}{4m_e c\epsilon_0} \frac{\gamma}{(\omega - \omega_0)^2 + (\gamma/2)^2} f_{\text{abs}}. \tag{5.4}$$

The intensity change due to absorption of species i is given by the Beer-Lambert law

$$I_i(\omega) = I_0(\omega) \mathrm{e}^{-n_i \sigma_{\text{abs},i}(\omega) L}, \tag{5.5}$$

where $\mathrm{I}_i(\omega)$ and $\mathrm{I}_0(\omega)$ are the final and initial intensities at frequency ω respectively, n_i is the number density of species i and L is the length of the medium through which the flux passes. The absorption of multiple species can be accounted for by adding the exponents, or equivalently multiplying the transmittances (the exponential part).

A concentration of 150 mM is equivalent to $150\,\mathrm{mol\,m^{-3}}$, which can then be converted to number density by multiplying by Avogadro's number. The jet is 1.5 µm thick, and the frequency dependent cross section can be calculated using the above method. This gives us the tools to calculate the expected transmittance spectrum, which is shown on panel a) of Figure 5.5.

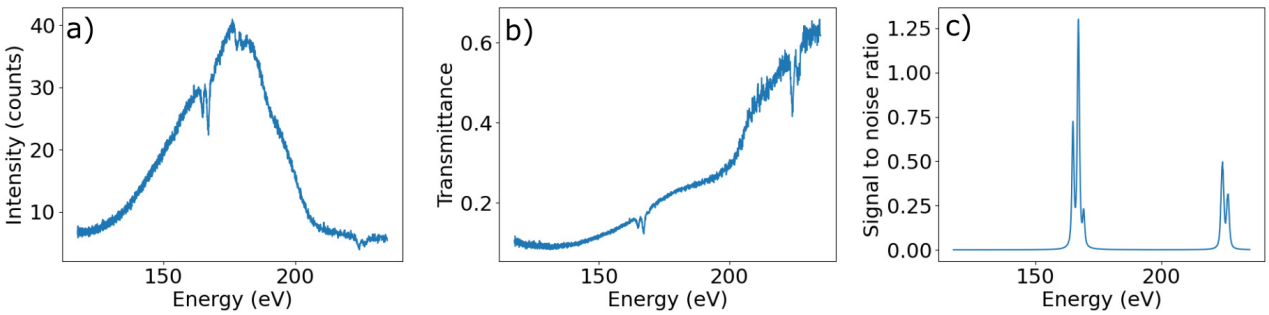

Figure 5.5: Panel a) shows the projected intensity of the ethanol + aldrithiol mixture, calculated by multiplying the measured ethanol intensity profile by the calculated transmittance of aldrithiol. Panel b) shows the result of the product of Figure 5.3 and the calculated transmittance of aldrithiol. The absorption cross sections in both panels were convolved with a Gaussian with FWHM of 0.5 eV to account for the resolution of our spectrometer. Panel c) shows the calculated signal to noise ratio of these peaks for the flux shown in Figure 5.1.

Panels a) and b) suggest that we had used a sufficient concentration of Aldrithiol in order to see a signal — there should be a significantly greater chance of a photon being absorbed at the peak energies than there is when only ethanol is present. To explain why there was no signal the suspicion therefore then turns to our flux — were there enough photons in those 30 seconds for the absorption of the solute to have produced a discernable signal?

The signal that we are looking for is an increase in the ratio $\frac{I_{\text{eth}}}{I_{\text{sample}}}$, which indicates the presence of extra

absorption due to the sample. This can be expressed as [183]

$$S = \frac{\partial {}^{I_{\text{eth}}}/_{I_{\text{sample}}}}{\partial \sigma(\omega)} \Delta\sigma(\omega) = \Delta\sigma(\omega) L \frac{I_{\text{eth}}}{I_{\text{sample}}}, \tag{5.6}$$

where $\Delta\sigma(\omega)$ is the change in absorption cross section due to the addition of the sample. The noise can be expressed using poissonian statistics as [183]

$$N = \frac{I_{\text{eth}}}{I_{\text{sample}}} \sqrt{\frac{1}{I_{\text{eth}}} + \frac{1}{I_{\text{sample}}}}. \tag{5.7}$$

The ratio of the signal to the noise is then easily obtained (using the Beer-Lambert law to express all intensities in terms of the incoming photon flux) as

$$\text{SNR} = {}^{\Delta\sigma(\omega)L}\Big/\sqrt{\frac{1}{I_0 \mathrm{e}^{-\sigma(\omega)L}} + \frac{1}{I_0 \mathrm{e}^{-(\sigma(\omega)+\Delta\sigma(\omega)L}}}. \tag{5.8}$$

The results of this calculation with I_0 taken from Figure 5.1 are shown in panel c), which roughly agrees with the simulated results in panels a) and b). According to these calculations we should have seen a much clearer signal than we did.

A possible reason for this is instability of the jet, which may have smeared out the signal. The jet was not very stable on the day that we took this data. We know from other groups that jet instabilities can give you a very bad resolution on the absorption features from the jet itself.

It is also possible that the simulated absorption features are overestimating the strength of the signal - possibly the absorption cross sections are smaller than calculated, or the lifetime of the excited states is longer than anticipated due to stabilisation by neighbouring atoms.

Finally the flux from our HHG beamline was quite poor during this experiment. Comparing the flux in Figure 5.1 to historical measurements for this same beamline (for instance in a recent thesis by Douglas Garratt [21]) shows that the flux that we were using during this experiment was around two orders of magnitude lower than is possible. However, at the time we had not performed this analysis and did not know whether the flux required for a clear signal would be achievable with our beamline at all, no matter how long we integrated for. We therefore elected to attempt these measurements with our second choice target instead.

5.3.2 Stability of the water jet

The second experiment intended to use NaClO as a target. However, as ClO^- is a strong base we did not want to use it for the initial experiments, to minimise any potential damage to the vacuum and pumping systems. We therefore performed these test experiments with NaCl instead.

Both of these targets are salts, and so are not soluble in alcohols. They are, however, very soluble in water,

so that was the solvent that we used for these experiments. As laid out in Chapter 2, we initially thought we could not make a stable sheet in water, but subsequent testing proved this to be false. We therefore made the switch from ethanol to water without much trepidation.

Unfortunately, running a water flatjet in vacuum did not prove to be as tractable as we had hoped. One of the other main differences between the physical properties of water and alcohols is that the critical point of water occurs at a much higher temperature and pressure. In actual fact, as shown by the phase diagram in Figure 5.6, liquid water cannot survive in equilibrium below around 6 mbar.

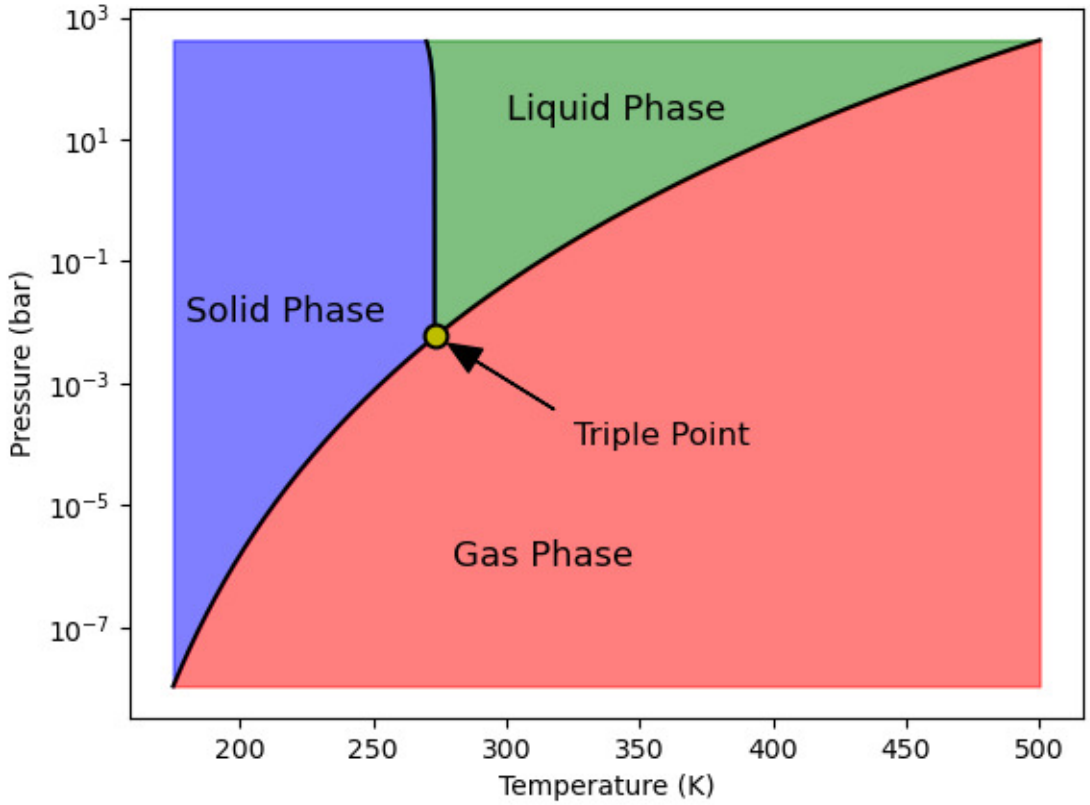

Figure 5.6: Phase diagram of water based on power laws found in [184] accessed 19/11/2020. The triple point at 0°C and 6.117 mbar is indicated by the yellow dot at the intersection of the three phase boundaries. At any given pair of pressure and temperature, the equilibrium phase of water is that shown on this figure. Whilst outside of equilibrium water can exist in other forms it will not be stable. At pressures below the triple point liquid water is not stable at all. In order to avoid freezing it is therefore important that the temperature of the liquid be maintained above the solid/gas boundary, so that the liquid will evaporate away.

Ordinarily this would not be an issue, as the jet is not required to stay in equilibrium in the vacuum chamber — all we need is for it to pass through without substantial disruption, and with a cylindrical jet this is very achievable [53, 57]. However, due to the high surface tension of water, flat sheets of water tend not to have stable rims [60] — the tendency to form droplets due to Rayleigh instabilities is very high. This leads to a large

amount of spray generated from the rims, which is prone to freezing and can lead to a cascade of ice build up. If the liquid in the jet goes through the central collector hole then this problem should be avoided, but due to the spray effects it does not.

The first solution we tried to make the water jet run properly was to heat the water. If the water is hot, the theory goes, it will be on the gas side of the phase diagram, rather than the solid side, and the problem is avoided. However, the scale of the evaporative cooling makes this difficult to maintain, and in addition there are other non-equilibrium thermodynamic effects that are important.

The enthalpy of vaporisation of water is fairly large, at $45\,\mathrm{kJ\,mol^{-1}}$. The flow rate of the water is controlled by the pump, and it typically run at around 3.8 ml/min, which is equivalent to $3.52 \times 10^{-3}\,\mathrm{mol\,s^{-1}}$. Assuming all of this water lands on the cone and evaporates instantly this leads to a power loss of 0.159 kW.

The catcher cone itself weighs around $1.44 \times 10^{-3}\,\mathrm{kg}$, and the specific heat capacity of stainless steel is $490\,\mathrm{kJ\,kg^{-1}K^{-1}}$. This leads to a temperature drop of $225\,\mathrm{K\,s^{-1}}$, which is very large. However, the assumption that all of the water lands on the cone and evaporates is very likely false, as when the jet is aligned the majority of the water should go through the hole, spray notwithstanding. Assuming only a tenth of the water lands on the cone and evaporates leads to a cooling of $22.5\,\mathrm{K\,s^{-1}}$, which while still large is actually quite similar to what we observe when a thermocouple is attached to the cone.

Therefore, in order to counteract this heat loss, and ensure the liquid in the vacuum chamber evaporates rather than freezes, we knew that we needed to provide around 16 W worth of heating to the tip of the catcher. We had been heating the catcher before this point, but from the outside of the chamber. The thermal conductivity (k) of stainless steel is $13\,\mathrm{Wm^{-1}K^{-1}}$, which relates to the power transmitted through the material (Q) by

$$Q = \frac{kA(T_2 - T_1)}{L}, \tag{5.9}$$

where $(T_2 - T_1)$ is the difference in temperature between the two ends of the material, A is the area of the material the heat is flowing through and L is the length of material the heat is flowing along. Unfortunately, for a meter long tube of stainless steel, of 4 cm diameter and 1 mm thick walls, transmitting 16 W of power requires a temperature differential of over 10,000 K, which is infeasible.

We redesigned the collector tip (shown in Figure 5.7) to counteract this. We added in slots to fit a heating cartridge, capable of delivering 25 W, and a thermocouple to monitor the temperature. We also made the tip out of beryllium copper, which has a much better thermal conductivity than stainless steel (the exact value depends on the composition of the particular beryllium copper alloy used, which we don't know as it was old stock, but quoted values tend to be at least ten times that of stainless steel), whilst being much more corrosion resistant than regular copper. We also heated the liquid in the reservoir and added a slot for another heating cartridge to the nozzle holder.

Despite all of this, we still encountered catastrophic freezing. We were able to vaporise the liquid around

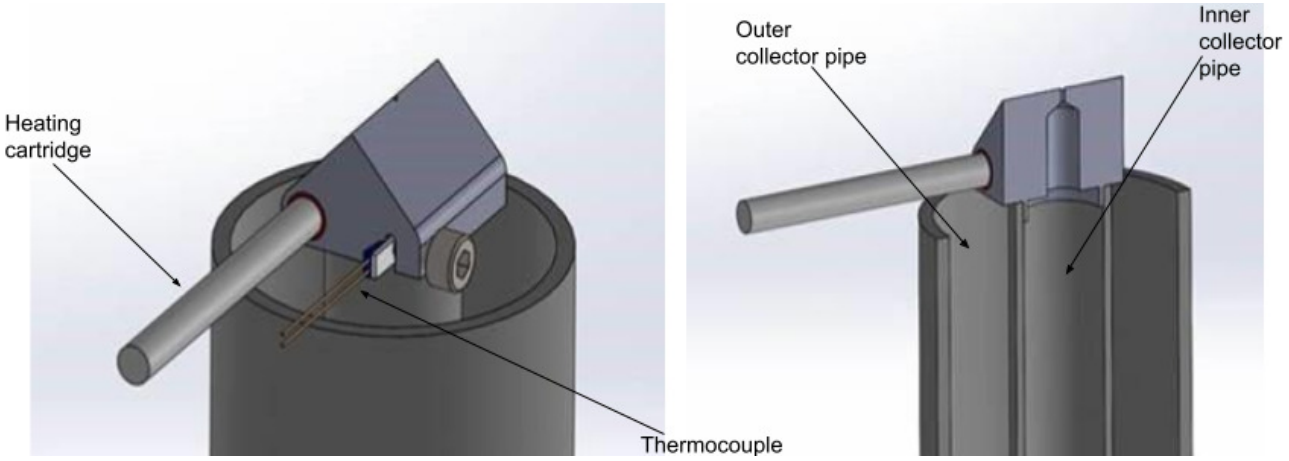

Figure 5.7: The plans for the redesigned collector tip. The part that is made of beryllium copper is the lighter grey prism shaped tip that is inserted into the inner collector pipe. With this new design we were able to have a consistent source of heat and temperature measurement right at the point where it was needed. Design drawn by Clement Ferchaud

the nozzle, but further out the liquid froze, forming an ice shell around the vapour. Acting on advice from a member of the Fielding group, we lowered the pressure very slowly with the jet running. By doing this we were able to avoid this catastrophic freezing issue.

This is probably due to the fact that lowering the pressure slowly allows the decrease in pressure to be an isothermal process, rather than an adiabatic one. When the pressure is lowered quickly, the temperature of the remaining gas drops as there is no time for the remaining particles to gain energy from their surroundings to maintain it. Thus no matter how much heating there is, quick drops in pressure will cause catastrophic freezing due to the large temperature drop — no heat can get to the molecules to compensate. On the other hand, slow lowering of pressure allows the system to maintain its temperature as the remaining gas particles can gain energy to make up for the loss of gas density. This then causes no temperature drops and therefore, along with the heating to avoid temperature loss due to evaporation, avoids catastrophic freezing.

5.3.3 Beamline contamination

Having overcome the issues caused by using water as the solvent of choice for this experiment, we encountered another problem. As we were setting up the X-rays to make a measurement at the Cl L-edge we observed a very strong absorption signal at exactly this wavelength. Comparing the signal to literature showed that it was probably Cl, as nothing except Ar has absorption edges in this range [41] and the harmonics were generated in Ne, although some key features that we would expect from the Cl absorption spectrum were missing (see Figure 5.8). Regardless, we proceeded on the assumption that this was contamination by NaCl, as there were no other candidates that could plausibly have been a source of contamination.

This contamination has proved very difficult to shift, due to uncertainty about which optics were contaminated. There are no optics in the target chamber or either differential pumping chamber that the X-rays make

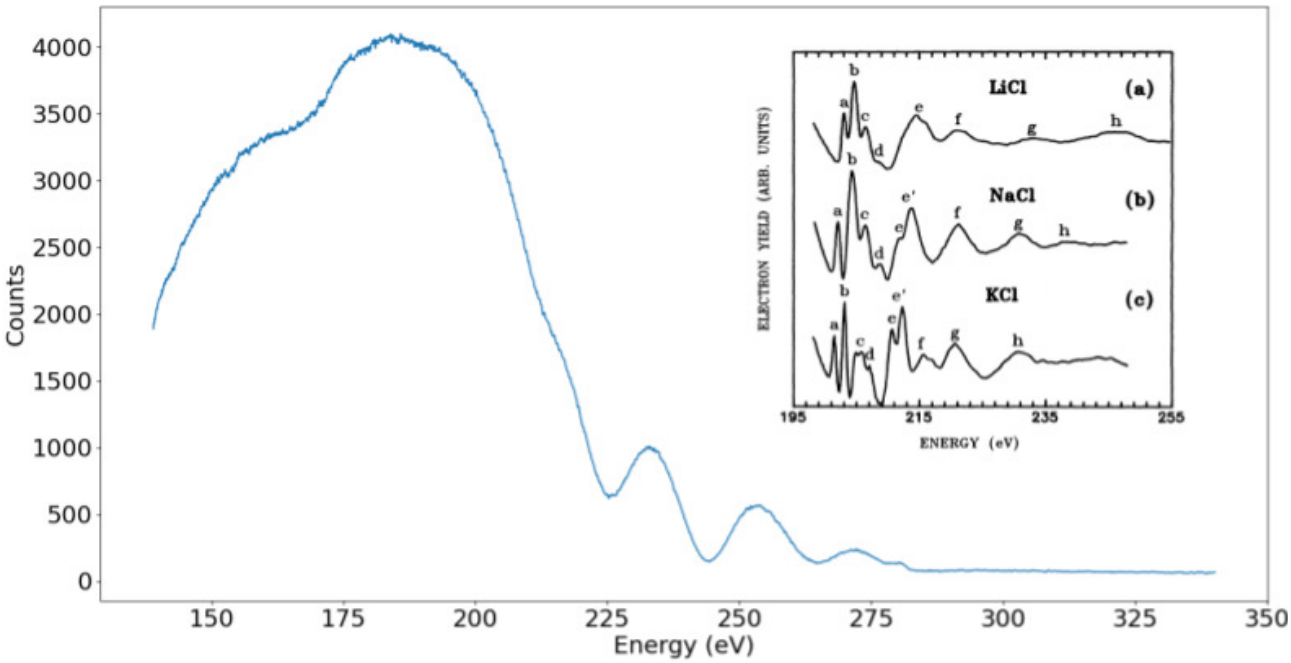

Figure 5.8: 5 minute long integration of the contamination signal. The grating parameters had been adjusted since the Aldrithiol measurements, but the x-axis was calibrated in the same way using the Zr $M_{4,5}$ edges and the P3HT L-edge resonances. There is a very large modulation starting at around 210 ev, which completely kills any signal beyond this point. The inset is a figure from Kasrai *et al.* [185] showing the measured X-ray absorption spectrum of NaCl. The trough at 225 eV (which would be a peak in the insert spectrum) approximately corresponds to peak f in the insert. The trough at around 240 eV is roughly at the correct energy for the Cl L_1 edge, which is not shown on the insert. Missing, however, is the structure around 200 eV which is characteristic of the Cl $L_{2,3}$ edge. There is already carbon contamination in our beamline, which can account for the lack of signal above 280 eV, and due to all this contamination it was difficult to optimise the signal for maximum cutoff so the lack of a rise in signal after the Carbon K-edge is unsurprising. The oscillations could also be an EXAFS signal after the Cl edge.

contact with, which means the differential pumping failed to keep the salt from reaching either the toroidal mirror or the grating/detector (and possibly both). As we were unsure why this had happened (which will be discussed below), we had to assume that all the optics that could be contaminated were indeed contaminated, and clean each one.

Fortunately the contamination was NaCl, which is very soluble in water, and we had a supply of pure water for the experiment we were attempting anyway. We were therefore able to clean the grating and toroidal mirror in water, avoiding any harsh chemicals that could have destroyed the coatings on these optics, or the optics themselves in the case of the grating which has very fine grooves on it which could be worn away. As there was also a question mark over the ANDOR X-ray CCD camera, and we were not sure whether running water over the CCD chip would end up destroying the camera, we sent the camera off for refurbishment and proceeded using an MCP.

Unfortunately after this round of cleaning the contamination was still present. Furthermore, subsequent tests of the MCP we used in a different HHG setup (one that should have had no NaCl contamination at all)

showed contamination on the MCP itself, which suggests that whatever was causing the NaCl contamination was still causing it. We therefore embarked on another much more thorough round of cleaning. This round of cleaning showed that there was a large build up of NaCl in the pipes backing the turbo pumps. It is not yet clear whether this source caused any secondary contamination or not, as the vapour pressure of NaCl is very low so one would assume that when it is crystallised on the side of a pipe it would tend to stay there. However, it would be exposed to a lot of water vapour in this position, so it is possible that microsolvated Na^+ or Cl^- ions could have been produced from this reservoir (see discussion below).

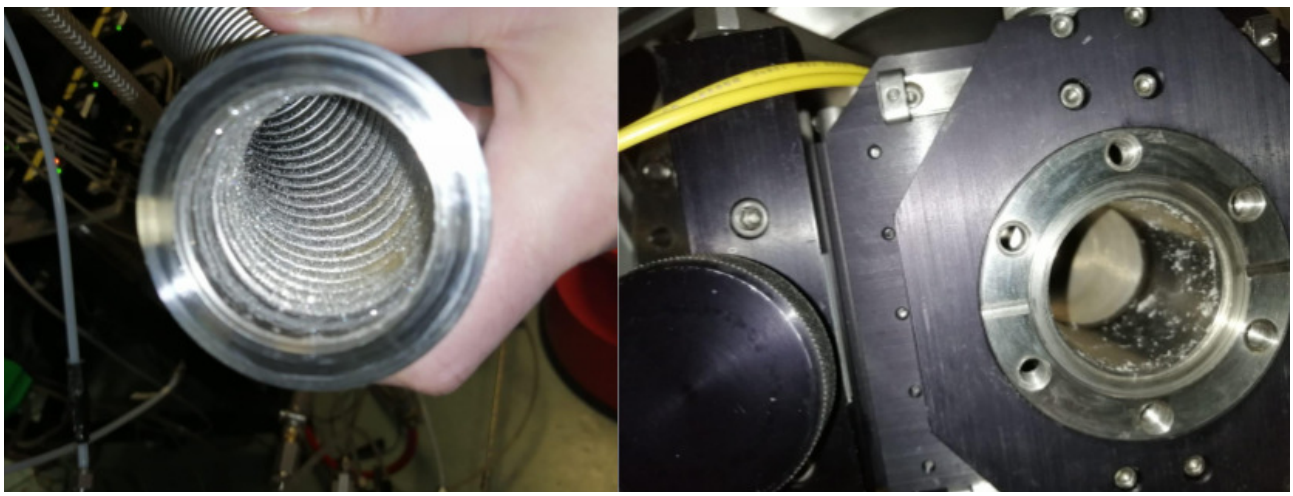

Figure 5.9: Two pictures of the salt contamination that was found in the chamber. The left hand picture is the pipe that was backing the main turbo pump on our chamber, and the right hand picture is the inside of the top xyz-manipulator in our chamber. Pictures taken by Oliver Alexander.

5.3.3.1 Contamination source

The question that arises from this is why did the differential pumping system fail to protect the optics from contamination? In order to answer this question we need to look at how likely it is that molecules could get through the differential pumping sections, and once they had got through how likely it is that they would contaminate the optics to such a degree.

The number of molecules impinging on a unit area of the wall of the vacuum chamber in unit time (J, assuming isotropic distribution) is [186].

$$J = n\frac{\bar{\nu}}{4}, \tag{5.10}$$

where n is the number density of molecules in the chamber and $\bar{\nu}$ is the expected value of the velocity of a molecule:

$$\bar{\nu} = \sqrt{\frac{8k_\mathrm{B}T}{\pi m}} \tag{5.11}$$

where k_B is the Boltzmann constant, T is the temperature and m is the mass of a molecule. Using the ideal

gas law (which should be fairly accurate under vacuum conditions) we can relate this to the pressure in the chamber as

$$J = \frac{p}{\sqrt{2\pi m k_{\mathrm{B}} T}}. \tag{5.12}$$

Therefore the number of molecules passing through an aperture of area A from an area of pressure p_1 to pressure p_2 in unit time is

$$\frac{dN}{dt} = (J_1 - J_2)A = \frac{A(p_1 - p_2)}{\sqrt{2\pi m k_{\mathrm{B}} T}}. \tag{5.13}$$

However, we have a pipe rather than an aperture, and some number of molecules that enter the pipe are going to bounce off the walls and exit back into the area they came from, which can be represented by a probability of transmission α. In general finding α is a non-trivial problem, but for cylindrical pipes a simple form has been proposed by Dushman [187] (which can be used to find the conduction equation in Chapter 4)

$$\alpha = \frac{1}{1 + \frac{3L}{4D}}. \tag{5.14}$$

where L is the length of the pipe and D is the diameter. We can therefore write the number of molecules passing through the differential pumping pipe per unit time as

$$\frac{dN}{dt} = \frac{\alpha A(p_1 - p_2)}{\sqrt{2\pi m k_{\mathrm{B}} T}}. \tag{5.15}$$

Using this equation, and values of 1×10^{-3} mbar and 1×10^{-7} mbar for the target chamber pressure and spectrometer chamber pressure respectively (the rough values recorded whilst the jet was running), we obtain a ballpark figure of around 3.3×10^{12} molecules per second entering the spectrometer chamber. Assuming no significant contribution from outgassing in the chambers, at equilibrium the overwhelming majority of these molecules will have come from the liquid jet.

However, just because the molecules are in the chamber does not mean that they will all settle on the sensitive optics. Condensation only happens when the partial pressure of the substance in the gas phase is greater than the equilibrium vapour pressure of the substance at the temperature of the surface that the substance is condensing on. This principle is already in operation in the spectrometer chamber — there is a cold finger in there, filled with liquid nitrogen. This is much colder than the chip of the X-ray CCD camera, so therefore any potential contaminants in the chamber should condense onto the cold finger preferentially, lowering the partial pressure of contaminant in the chamber to below the vapour pressure at the temperature of the chip.

This behaviour is expressed in the Hertz-Knudsen equation [188], giving the overall flux of molecules condensing onto a surface (J_{con}):

$$J_{\mathrm{con}} = \frac{p_{\mathrm{part}} - p_{\mathrm{vap}}}{\sqrt{2\pi m k_{\mathrm{B}} T}} \tag{5.16}$$

where p_{part} is the partial pressure of the gas and p_{vap} is the vapour pressure at temperature T. This equation is just derived from the above equations of flux, and doesn't take into account many microscopic factors such as instantaneous fluctuations in local pressure or temperature. Because of this it is often necessary to introduce a fudge factor in this equation to make its predictions match the experimentally observed results for the exact flux [188]. However, for our purpose the important observation is that if p_{vap} is greater than p_{part} then the net flux will be negative, *i.e.* away from the surface, and we should not encounter any problems with physical adsorption of molecules contaminating the surfaces.

The vapour pressure of water is non-trivially dependent on temperature. The toroidal mirror and diffraction grating should be at room temperature, as they should be at equilibrium with the surroundings. The lab is maintained at 20°C, and the vapour pressure of water at this temperature is known to be 23.388 mbar [189], so there should be no condensation on these optics. The CCD chip is cooled to −60°C. Many empirical formulae for the vapour pressure of water at various temperature have been proposed, and I used a recent one that was designed for the temperature range I was interested in (−100–0°C) [190]. This gave the vapour pressure of water at −60°C as 0.0108 mbar, which means the camera chip should not have experienced any significant condensation under these conditions.

There should therefore be no physical adsorption of the water onto the optics. The vapour pressure of NaCl at temperatures close to room temperature is essentially 0, so we would expect that if pure NaCl is adsorbed onto a surface then it should stay there. However, NaCl is hygroscopic (it has a tendency to absorb water from the surrounding atmosphere), with a deliquescence point (the relative humidity threshold at which this hygroscopic behaviour becomes active) relatively well defined at around 76% humidity [191]. As the majority of the particles in the chamber should be water, the relative humidity of the atmosphere in the chamber should actually be very high, and so we would expect that any adsorbed NaCl would be microsolvated, and therefore we might expect it to behave similarly to water, unless there is a stronger attraction than just Van der Waals forces — *i.e.* it is chemisorped onto the surface.

Chemisorption of NaCl would also explain the stubborn resistance of the contamination to water-based cleaning methods. Cl^- is known to chemisorp readily onto gold [192], which is the coating on our toroidal mirror and grating. At the time of writing we have not yet attempted to clean these optics using more harsh methods capable of removing the chemisorped Cl, but this is the next step. Hopefully this cleaning combined with the removal of the NaCl reservoir behind the turbo pumps will remove the contamination signal, allowing us to proceed with measurements once more.

5.4 Lessons learned and improvements for the future

Several lessons have been learned from these experiments. The first (and possibly most important) lesson this has taught us is the importance of theoretical support in planning experiments. The calculations presented for signal to noise in Aldrithiol were not particularly challenging, and could have been done beforehand to give us a better idea of the amount of flux we would need for this experiment. We did perform some rudimentary calculations using data from the online Henke database, but nothing as quantitative as the calculations above. As a result, when we discovered we did not have enough flux to perform the experiment we did not know if it would be feasible to increase the flux to the point where the experiment would become possible or not. As time was limited for this experimental run we felt forced into the alternative option, which in hindsight was probably the wrong choice, added a lot of experimental complexity and could have been avoided had we calculated the oscillator strengths and the likely flux that would be needed beforehand.

Another feature of experiments of this nature is the instability of the liquid jet. The freezing was a major issue, but a smaller yet in a way more significant issue was that the jet apparatus and the jet itself were not stable over long periods. This meant that longer integration times were not viable as a way of increasing the absolute size of the signal (although with a sufficiently small signal to noise ratio this may not have helped much), and the jet has to be monitored at all times making automated scans that would run overnight (for example) impractical. Whilst this will impact more on a time-resolved study (due to the much larger number of data points required), it is an issue that we need to resolve.

There are two parts to improving this, recirculation and stabilising the jet. Recirculation is discussed in Chapter 4, and will not be repeated here. Figure 5.10 shows a new design for a nozzle housing assembly that will hopefully be able to reduce the issues of jet stability. Both the nozzle and collector are housed inside a differential pumping jacket, with small holes to let the X-rays through. The fixed positions of the nozzle and collector should remove the issue of the jet missing the collector, and the higher pressure inside the jacket should mean that we can have a larger hole in the collector as the pressure differential between the reservoir and the jet should be less. The potential issue with this jacket is that the higher pressure inside it might increase the absorption from the solvent, but whether this increase is significant enough to reduce the feasibility of an XAS measurement remains to be seen.

The issues with the water jet freezing are now better understood, so aqueous targets are a feasible target type going forward. However the contamination issue indicates a much larger problem — if a solute or solvent that we want to use will bind strongly to optics and cause major contamination, how can we stop this? The differential pumping will not solve this issue, as the main source of gas in all of the connected chambers at equilibrium will be the liquid target, so there will still be a number of collisions between target molecules and optics no matter how good the differential pumping is. The pressure can be reduced sufficiently to avoid condensation, but it cannot stop a reaction that does not require the target molecule to be condensed in bulk

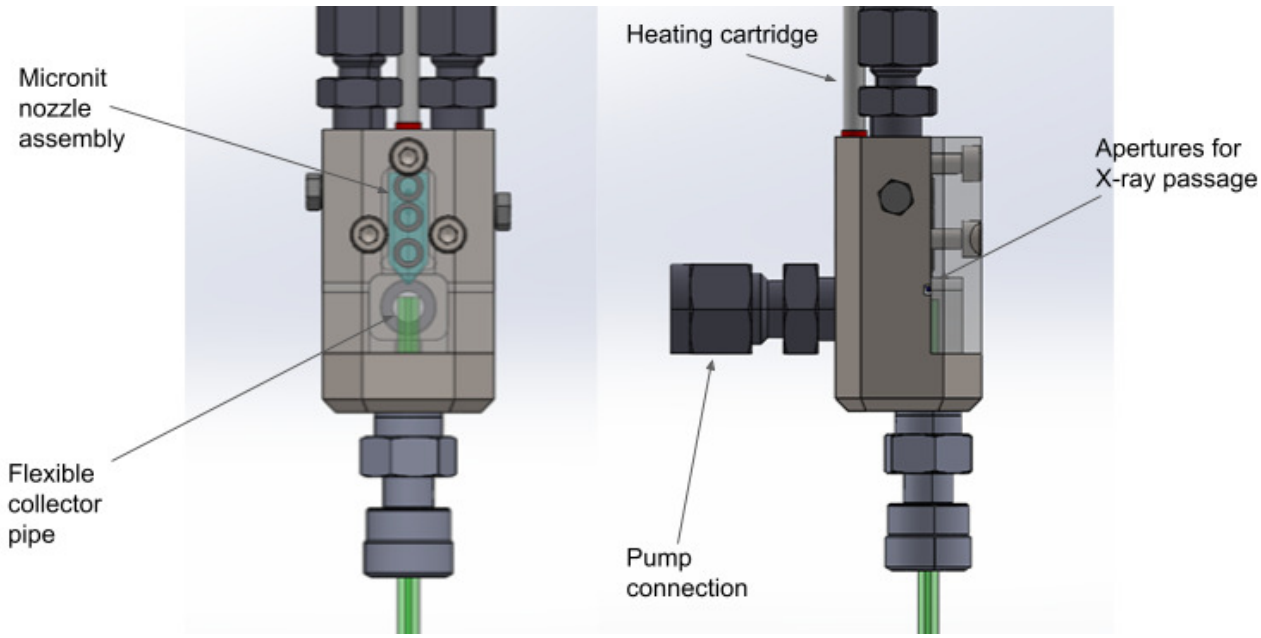

Figure 5.10: The design for a differential pumping jacket for the liquid jet. The Micronit nozzle assembly is held directly above the collector pipe, with just enough space to form a sheet before it flows into the collector pipe. The apertures are placed at the point where the center of the sheet should be. The collector tubing should be flexible to allow for the entire assembly to be aligned to the laser. Design drawn by Clement Ferchaud

on the optics.

One measure that could reduce this problem is the introduction of a buffer gas into an intermediate chamber. This would reduce the proportion of gas that enters into the optics chambers that is the target, further reducing the likelihood of a chemisorption event. However, the pressure of gas required might have the effect of raising the pressure in the beamline to an unacceptable level, and depending on how reactive the potential contaminant is the optics may still become contaminated anyway over a long time.

Thin liquid sheet targets in a vacuum chamber are an exciting new target type that could enable the probing of chemical reactions using soft X-ray spectroscopic techniques for the first time. However, as with all new techniques there are some implementation issues that needed to be worked out. I am hopeful that in the next few years we can start to use these targets to their full potential, to carry out time-resolved solution phase XAS experiments.

Chapter 6

High Harmonic generation from a liquid jet

6.1 Introduction

As has been indicated in Chapter 3, there has been an upswing in interest in condensed phase harmonics over the past decade, with the discovery of HHG from solid phase targets [80], and this is now a fairly well established field [35]. However, the majority of HHG results so far have been from (predominantly crystalline) solids, without much attention paid to the other common condensed phase — liquids. There have been some attempts at liquid phase HHG over the years, but until recently no one had observed any harmonics that were recognisably non-perturbative. In addition there was the belief that the density of the liquid, coupled with the lack of long range electronic structure would make HHG infeasible [74, 75].

An interesting result that speaks against this however, is that obtained by Ndabashimiye *et al.* [142]. They obtained high-order harmonics from solid argon and krypton, which both lack significant long range electronic structure due to their closed shell nature as noble gases and would have a density more of the order of a condensed phase target like a liquid. They found that they were able to generate harmonics from these rare gas solids, and that the spectrum had an interesting multiple-plateau structure which is explained by the presence of multiple different conduction bands in the solid, due to regularities in the structure.

However, in this chapter I will present results that show that HHG is not only possible in liquids, but occurs as readily as it does in a gas or solid. The initiation method is the same — strong field ionisation — but the question remains by what mechanism are these harmonics produced. I will present several different experimental runs that we performed in order to try and answer this question, and then I will tie all the different strands of evidence that we have together into a conclusion that will offer suggestions as to how HHG in a liquid is possible. This work is also currently in preparation for publication [81].

Concurrently Luu *et al.* have also performed some very similar measurements on a similar liquid target, with some similar results (although their results do differ from ours in certain key areas) [82]. Due to the great similarities between the two sets of experiments, at the end of the chapter I will contrast our results with those of Luu *et al.*, and discuss how the two datasets can be harmonised.

6.2 Harmonic Spectrum

Our headline result is that we have observed high harmonics from a liquid that extend all the way up to 50 eV, a similar range to that observed in the rare gas solids [142], as well as semiconductors [35]. However, there is no notable structure in the harmonic spectrum — just a single plateau and then a cut off. This makes sense as the liquid would not be expected to have the same kind of repeating unit that gave the rare gas solids their multiple-plateau structure. The cut off is also more of a gradual decline than would be anticipated, which is probably due to the scattering of the laser driven electrons by the liquid, which will be explained more below. We observe well separated harmonics, indicating that our generation is in the multi-cycle regime (although we do observe CEP dependence — see below). The well resolved harmonics aid in distinguishing harmonic orders.

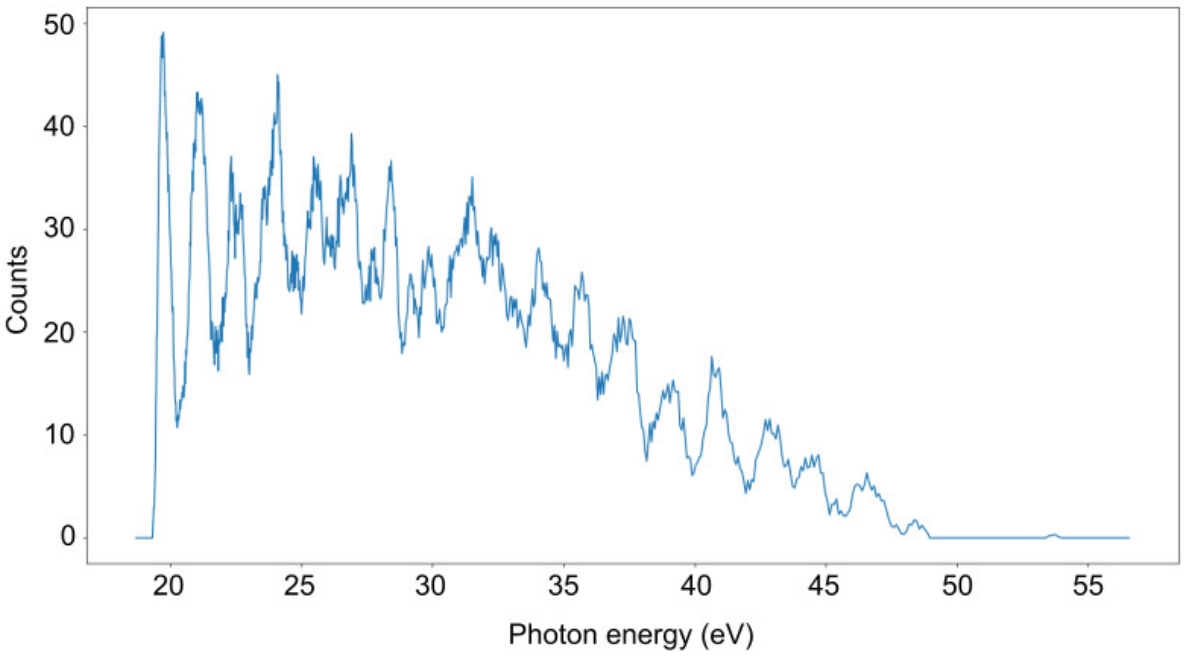

Figure 6.1: A harmonic spectrum generated from our isopropanol target. The harmonics stretch out to around 50 eV before cutting out. There is a slight plateau, before the harmonics decrease fairly monotonically. The driving laser was our 1.8 µm idler.

6.2.1 Calibration of the spectrometer

In order to determine the energy of our harmonic spectrum and to be sure of the value of 50 eV for our harmonic cutoff it was necessary to calibrate the spectrometer. This device had already been calibrated, but we were somewhat suspicious of this calibration as the grating had probably been moved since it was performed. We therefore undertook to calibrate the spectrometer using the procedure outlined in Chapter 4.

This calibration was performed using several 0.2 µm thick metal filters, which have clear windows of transmission with sharp and well known edges enabling us to pinpoint the pixels corresponding to these wavelengths. We used a tin filter, a zirconium filter and an aluminium filter. Tin and zirconium have small windows of transmission which cut off sharply at 52.1 nm (Sn) and 46 nm (Zr), and aluminium has a sharp rising edge at 82.6 nm. In addition to these sharp edges, Zr also has a less sharp rising edge that begins at 19 nm.

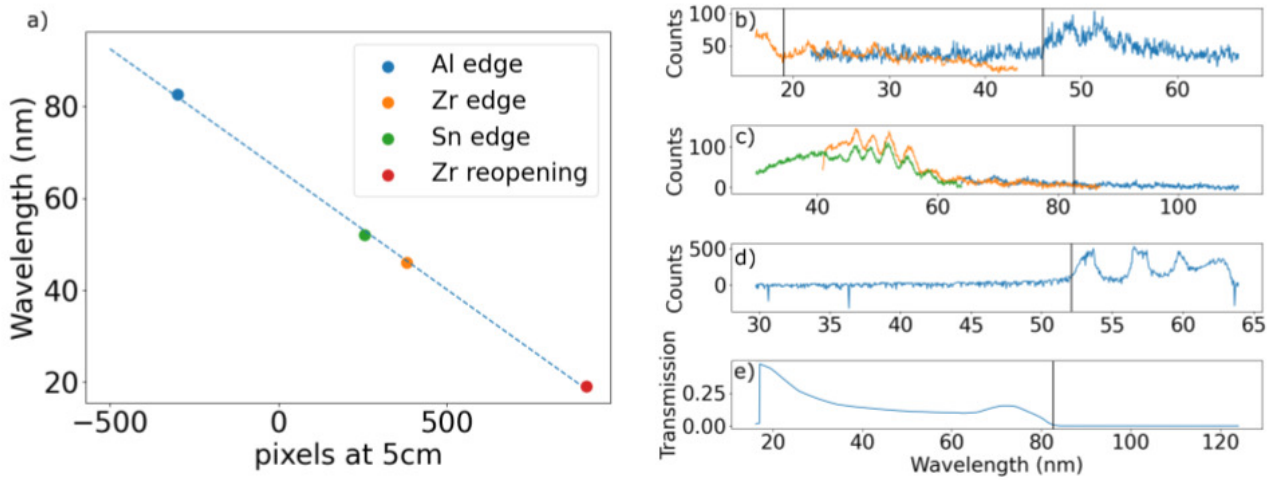

Figure 6.2: a) is a graph of the linear fit between the wavelengths of the edges and the position on the MCP. 0 pixels is the long wavelength side of the MCP in the position that the majority of the spectra were taken, and the pixel values of other positions are defined relative to this edge. b), c) and d) show Kr HHG spectra with the metal filters in, with b) Zr, c) Al and d) Sn. Differently coloured spectra are taken at different MCP positions, and the low energy Zr spectrum has been scaled up to be comparable to the high energy spectrum, as the latter was taken on a different day when the signal was better. The vertical lines indicate the positions of the filter edges. The two Zr points and the Sn point fits very well, but the sharp Al edge is not observed, rather a gradual increase in harmonic signal. This is due to oxidation of the Al filter, leading to a layer of Al_2O_3 on the surfaces. The transmission window of an oxidised Al filter (using the oxidation rates provided by Lebow) is shown in e), and from this a gradual increase in harmonic strength can be predicted.

Unfortunately, the aluminium filter had some oxidation on it, which meant that the sharp edge of the unoxidised metal was not a sharp edge at all, as Al_2O_3 has a different absorption spectrum in this region. According to Lebow (the manufacturer) an oxidised 0.2 nm thick Al filter will have around 10% oxidation, *i.e.* it will have an 0.02 nm thick layer of Al_2O_3 on the surface. The resulting absorption spectrum (using data from Henke *et al.* [42]) is shown in Figure 6.2. This meant that we were not able to assign the Al edge with much confidence to a pixel position, leaving us with only two points through which to fit a line. This was not ideal,

as any two points can be fitted with a line so in order to confirm our line was valid we needed a third point. This third point was the start of the second rising edge in Zr, which could be fixed with some more precision than the start of the Al transmission window. The pixel positions of these edges were then plotted against their known wavelengths, and a linear fit of pixel position against wavelength performed. The spectra used for the calibration are shown in Figure 6.2, with the calibrated wavelength on the x-axis.

6.2.2 Verification of liquid harmonics

In addition to the large cutoff energy the other part of our new finding is that these harmonics originate from the liquid itself. This is important, as if it could be shown that the harmonics were generated from the vapour surrounding the liquid then our claim would only be the much less significant one that we have generated harmonics in a gas phase molecule, which has been done before [36, 193]. To confirm that it is likely that the harmonics are being generated in the liquid and not in the surrounding vapour we carried out several tests. The first test was to try and generate harmonics from the isopropanol vapour given off by our jet. We rotated the sheet so that the flat side was parallel to the direction of the laser propagation and let the laser propagate very close to the surface, where the most vapour should be found, and saw no harmonic signal. We then turned the jet back into a perpendicular position to the laser and scanned horizontally across it in 100 µm steps. As can be seen in Figure 6.3, the harmonics drop off sharply at the edge of the jet (in fact they drop off slightly before, probably due to the round edges of the jet scattering the driving laser/harmonics), and no gas phase harmonics are observed beyond the edge.

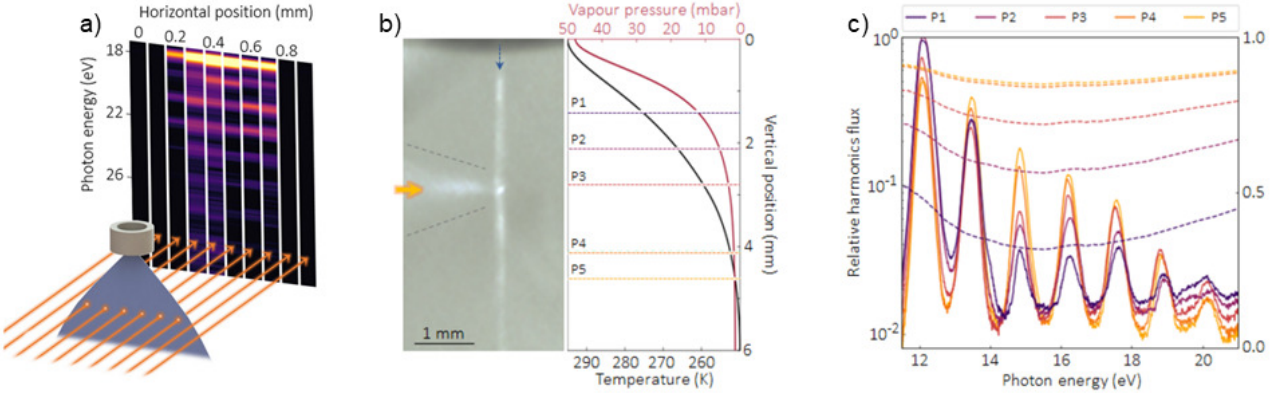

Figure 6.3: a) is a diagram showing the scan horizontally across the jet. Data was taken every 100 µm across the jet, which is shown on the screen. The jet and arrows are shown for illustration purposes only. b) is made of two parts, a side on view of the jet showing a plasma plume at position 3, and a graph showing the change in temperature of the jet and vapour pressure surrounding the jet as you go down the direction of propagation of the jet. Marked also are the five positions at which the harmonics in c) are generated. c) shows the harmonic spectrum (solid line) and isopropanol transmission through the vapour surrounding the jet (dashed line) for each position marked in b), showing that the harmonic spectra higher up show greater signs of absorption, but not of greater generation as one would expect if the harmonics were generated in the vapour. Figures created by Clement Ferchaud, for our paper in preparation [81]

The gas pressure around the jet is due to evaporation from the sheet into the vacuum. The jet cools rapidly due to evaporation as it propagates through the vacuum, which leads to a change of vapour pressure that the laser and the harmonics propagate through depending on how far down the jet the harmonics are generated [194]. As XUV radiation is readily absorbed by the isopropanol vapour, if the harmonics were to be produced in the liquid and then propagate through the gas we would see increased absorption (and therefore lower intensity harmonics) at points higher up the jet, whereas if they are produced by the gas we would expect the opposite, more harmonics with increased gas density.

Figure 6.3 shows the results of performing this experiment. At the higher positions (in the frame of the jet, so P1 and P2) we see a greater effect of absorption, as the difference between the higher order and lower order harmonics is greater at these positions. This is due to the greater absorption of the vapour at these positions, as shown by the dotted line. No significant enhancement of the harmonics is observed with the change of position, so it appears that if there are any harmonics generated in gas the contribution to the overall flux is small compared to absorption.

A small additional piece of evidence that is worth bringing up at this stage is the plasma plume. At relatively high intensities we observe a plasma plume where the laser meets the jet (shown in Figure 6.3). Under conditions where this plasma plume is observed, the presence and absence of the plasma plume corresponds exactly to the presence and absence of an extended harmonic cutoff, suggesting that harmonics are generated by the same process that is creating the plasma plume, *i.e.* the strong field ionisation of the liquid.

The other potential source of the harmonics is from the surface (rather than the bulk of the liquid). As our current belief (borne out by simulations that will be presented in our paper in preparation [81]) is that the harmonics that we observe are generated in the last $\approx 40\,\mathrm{nm}$ of the bulk due to the absorption length of the liquid, it is perhaps quite difficult to distinguish between this and a surface effect. As a part of the investigations above, we rotated the polarisation of the laser and found that there was a small gain in intensity when the laser is polarised horizontally in the lab frame and the jet is rotated almost parallel to the laser propagation direction. However this could just be due to a decrease in scattering probability as a large part of the free electron trajectory would then be in the vacuum. It certainly speaks against any effect due to electron transport along the surface of the liquid.

From the evidence shown above, we can be reasonably confident that the harmonics are produced from the interaction of the laser with the liquid itself, and not with the gas surrounding it. The next step is to establish how the laser field can generate these harmonics from a condensed and disordered medium.

6.3 Dependence of HHG upon laser intensity

Having established our ability to generate harmonics from the liquid sheet, the question of the mechanism by which the harmonics are generated arises. An identical mechanism to the gas phase seems unlikely, as the excursion distances involved of hundreds of nm are way beyond the mean free path in the liquid, so the electrons are likely to scatter away and never return to the parent ion. However, due to the absence of long range order in liquids a mechanism involving Bloch oscillations in a well defined conduction band does not seem very likely either, due to the lack of well defined electron conduction bands. One of the key diagnostics for which mechanism is occurring is observing the effects of varying the intensity of the harmonics generation field. In the gas phase three step model the cutoff varies linearly with the intensity of the field, whilst in solid harmonic generation it will vary linearly with the square root of the intensity, which is the electric field strength itself.

In order to perform this experiment we used two different methods of varying the intensity, a half waveplate and polariser setup and an iris, as laid out in Chapter 4. Initially the half waveplate and polariser combination was used, but we were concerned that this was stretching the pulse too much so we subsequently took data using only the iris. This gave us access to much higher intensities, which enabled us to observe the effects of crossing the intensity threshold above which strong ionisation disrupts the HHG process.

6.3.1 Determination of the laser intensity

In order to derive any meaningful conclusions from our intensity scans, we needed to know what the intensity of the generation field is at the location of harmonic generation. As it is not possible to measure this directly, we have to start with the beam in a position where we can measure the intensity, and then propagate the beam from there. In order to do this we assume the beam profile corresponds to the lowest order Gaussian mode for simplicity. Using the standard formula for focussing Gaussian beams with a lens, the beam radius at the focus (ω_0) can be written as

$$\omega_0 = \frac{f\lambda}{\omega_{\mathrm{r}}\pi}, \tag{6.1}$$

where ω_r is the radius of the beam filling the lens and f is the focal length of the lens. As the direction of the curvature of the wavefront of the focussing beam changes either side of the focus, the beam can be assumed to be a plane wave at the focal spot. The intensity of a plane wave is given by

$$I = \frac{c\epsilon_0}{2}|\mathcal{E}_0|^2, \tag{6.2}$$

where I is the intensity, c is the speed of light in the vacuum, ϵ_0 is the permittivity of the vacuum and $\mathcal{E}_0$ is the electric field amplitude. The power contained in the pulse is related to the electric field amplitude by [195]

$$\mathcal{E}_0 = A_0\sqrt{\frac{2P}{c\epsilon_0}} \tag{6.3}$$

where $P = E_\mathrm{p}/\tau_\mathrm{p}$ is the power of the pulse, with E_p the energy in the pulse and τ_p the pulse duration, and A_0 is a normalised amplitude function relating to the profile of the beam. For the lowest order Gaussian mode A_0 is given by

$$A_0 = \sqrt{\frac{2}{\pi}}\frac{1}{\omega_0}\mathrm{e}^{-r^2/\omega_0^2} \tag{6.4}$$

where r is the radius away from the center of the beam. The peak intensity of the beam at the focus will be in the center where $r = 0$ and so the exponential becomes 1. This leads to the following expression for the peak intensity of the beam at the focus:

$$I = \frac{2P}{\pi\omega_0^2} = \frac{2E_\mathrm{p}}{\tau_\mathrm{p}\pi\omega_0^2}. \tag{6.5}$$

However, due to the fact that we are attenuating the beam using an iris, the focused beam will not be a perfect Gaussian, but will have diffracted through the circular aperture. According to Fourier optics the beam profile at the focus is the Fourier transform of the beam profile at the far field [112]. For a beam with a Gaussian profile this is also a Gaussian which makes calculating the focus width simple, but for a circular beam of uniform intensity (such as would result if the iris aperture was uniformly illuminated) the Fourier transform is an Airy disk:

$$I(r, R) = I_0\left[\frac{2\mathrm{J}_1(\frac{\pi Dr}{\lambda R})}{\frac{\pi Dr}{\lambda R}}\right], \tag{6.6}$$

where D is the diameter of the aperture, λ is the wavelength, r is the radius as above, R is the distance from the aperture to the focus and J_1 is a Bessel function of the first kind [196]:

$$\mathrm{J}_n(x) = \frac{1}{\pi}\int_0^\pi \cos(n\tau - x\sin\tau)\,d\tau. \tag{6.7}$$

However, this form is fairly unwieldy to use. Fortunately, the center of the Airy disk is where the majority of the intensity is, and the subsequent rings are fairly shallow. We can then approximate the central profile with a Gaussian [197]

$$I(r) \approx I_0'\exp(\frac{-r^2}{2\sigma^2}) \tag{6.8}$$

where σ is the corrected RMS width of the disk center,

$$\sigma = 0.45\frac{\lambda f}{D}$$

where f is the focal length of the lens. Convolving this pattern with the Gaussian profile of the beam will give another Gaussian (another favourable property of the Gaussian distribution), with a corrected beam waist of

$$\omega_{\mathrm{g}} = 2\sqrt{(\omega_0)^2 + (0.45\lambda\frac{f}{D})^2}. \tag{6.9}$$

This then can be placed into equation 6.5 to give a corrected intensity

$$I = \frac{2E_{\mathrm{p}}}{\tau_{\mathrm{p}}\pi\omega_{\mathrm{g}}^2}. \tag{6.10}$$

The energy contained in the pulse was measured after the attenuation of the iris, but we also needed to take into account losses from windows and the sheet itself for an accurate value of E_{p}. From the point of measurement there is an uncoated CaF_2 lens, an uncoated fused sillica window into the chamber and the liquid sheet itself. The lens is 6.1 mm thick at the center, which is where we assumed our peak intensity is, with a transmission of 96.66%. The window is 3 mm thick, with a transmission of 98.06% (both values calculated from data provided by Thorlabs), and the jet is around 1.5 µm thick with a transmission of 99.97% (from data provided by the NIST [198]), all at a wavelength of 1.8 µm. This all adds up to an attenuation of 94.76% of the measured value of E_{p}.

In order to find the duration of the pulse we performed a second harmonic FROG measurement of the pulse before the lens in front of the chamber, the results of which are shown in Figure 6.4. From this measurement we determined the duration of the central pulse spike to be 13 fs before all the glass in it's path. Although there is a reasonable amount of glass between the point of measurement and the interaction point (the lens and the window), we assume that as we tune the wedges for optimal harmonic production the effects of these optics on the pulse leads only to a slight broadening to 15 fs or so. The other conclusion to note from Figure 6.4 is that only around 70% of the energy is contained in the central spike. Assuming that this is the only part of the driving pulse that is intense enough to produce harmonics, we therefore needed to reduce our measured energy further. After all of this, we can be confident that our estimated values for the intensity are a good apprioximation of the true value at the interaction region, and that we can base our conclusions on them.

6.3.2 Intensity scan using a half waveplate and a polariser

The first intensity scan we did was to use a half waveplate and polariser to reduce the intensity. This scan did not use an iris, so we could in theory use the uncorrected beam waist to determine the intensity. However this waist would be around 45 µm, which is much smaller than previous measurements of the focus spot in the chamber [194] that put it at around 100 µm. This could have been due to plasma defocusing either in or before the jet due to the vapour, or simply due to some non-ideality in our focusing. We therefore used the value of 100 µm as a minimum focus diameter. In addition, due to the extra glass in the beam path the pulse duration

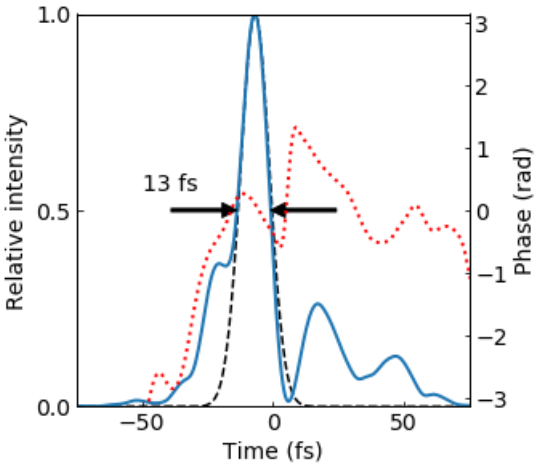

Figure 6.4: A second harmonic FROG trace, showing the retrieved temporal profile (blue line) and phase (red line). Around 70% of the intensity in the pulse is contained in the central spike, with FWHM of 13 fs. The rest of the power in the wings is assumed to not be enough to generate harmonics.

was stretched to around 33 fs. These effects decreased the peak intensity of the pulse dramatically, from around 90 $\mathrm{TWcm^{-2}}$ for the highest variable aperture scan (see Figure 6.7) to around 34 $\mathrm{TWcm^{-2}}$. As such we did not see any of the effects of the intensity being too high as we did in the variable aperture scan. The results of this scan are shown in Figure 6.5. Interestingly there appears to be a threshold intensity where the harmonics 'turn on', with several harmonics appearing at once at around 12 $\mathrm{TWcm^{-2}}$. This is somewhat similar to the phenomenon seen in Ndabishimiye *et al.* [142], where after a certain intensity threshold was passed suddenly a new range of harmonics became available to the system, as a new conduction band was able to be reached.

The power allowed through the half waveplate and polariser combination varies with the angle of the half waveplate nonlinearly, so in order to provide a mapping between the half waveplate position and the power allowed through as a proportion of the maximum power a calibration was performed, where the power was measured in several different half waveplate positions between the position of maximum transmission and the position of minimum transmission, and a polynomial fit was performed to these points. Using the coefficients of this fit we then determined the power relative to the maximum power that was allowed through at a given half waveplate position, and this could then be used in conjunction with the expected pulse length and the measured maximum power to work out the intensities.

6.3.3 Intensity scans using an iris

As we were concerned about the effects of stretching of the pulse by the half-wave plate on the intensity we also used an iris to aperture the beam in order to reduce the intensity without putting a lot of extra glass in the beam. As would be expected, this lead to a much higher peak intensity (70 $\mathrm{TWcm^{-2}}$ for a 15 fs pulse), and allowed us to observe the effects of increasing the intensity through the optimum value and over to the other

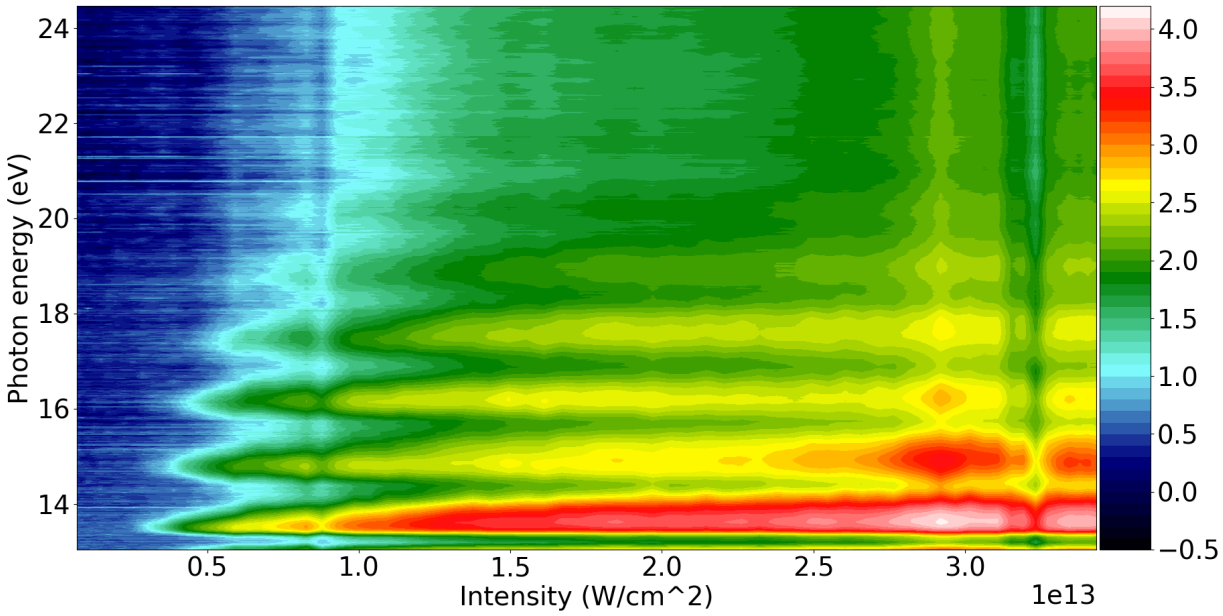

Figure 6.5: Intensity scan using the half waveplate and polariser. The clour scale shows the base 10 logarithm of the intensity values, in order to better show the variation of the cutoff. The MCP was at 3 cm when this data was taken, so the behaviour of the cutoff beyond around 15 TW/cm^2 cannot be determined. However, at this intensity it does seem like multiple harmonics 'switch on' at once, an effect that is not observed in the variable aperture scan.

side. The results of this scan are shown in Figure 6.6.

Above a certain threshold intensity we observed a significant distortion of the harmonics, which also coincided with their displacement in space and a reduction in the cutoff intensity. This occurs because (as described in Chapter 3) the plasma affects the refractive index as

$$n_{\text{plasma}}(\omega) = \sqrt{1 - \left(\frac{e^2 \rho_e}{\omega \epsilon_0 m_e}\right)^2}. \tag{6.11}$$

The larger ρ_e (the density of electrons) is, the further away from 1 the refractive index becomes and the greater the effect on the harmonics produced. The phasematching condition becomes less favourable, causing the loss of intensity and reduced cutoff, and the harmonics also distort due to this refractive index change.

The other point to make is that this shows that the cutoff is limited by the threshold at which the ionisation fraction becomes too much for the efficient phasematching of the HHG. In solid phase harmonics (even the rare gas ones) the cutoff is limited by the energies of the conduction bands that can be reached, and no plasma breakdown is observed [35, 142], although perhaps this is because the solid targets would be destroyed before they got to this point.

Interestingly, in this scan the harmonics do not appear to 'turn on', as they do in the half waveplate scan. The most likely explanation for this fact is that the half waveplate scan effectively takes a finely grained look

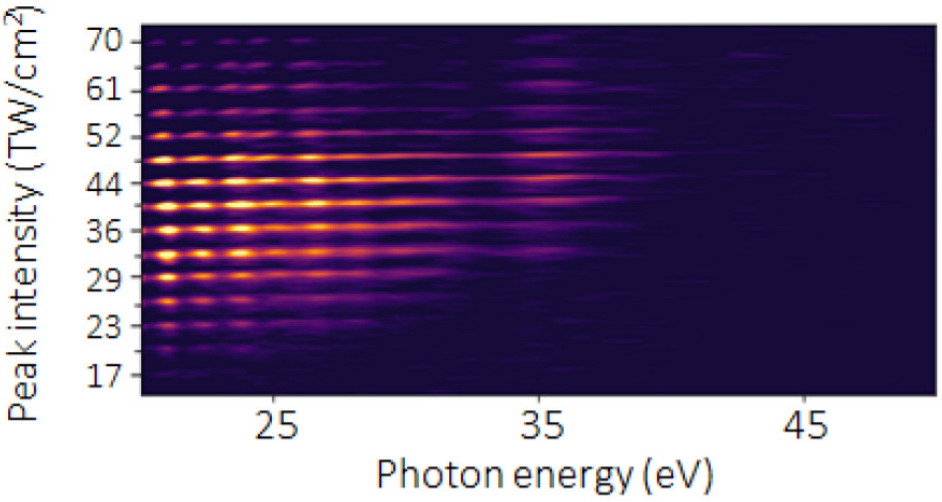

Figure 6.6: The intensity scan using the iris to control the intensity. Each horizontal line is a spectrum taken at the iris position corresponding to the intensity on the y-axis, rather than being interpolated together as has been done to the rest of the graphs in this chapter. The reason for doing this is to show how the spectrum changes with the intensity. Up to the 48 TW/cm^2 spectrum, the harmonics get more intense, and the cutoff extends linearly. No change of shape happens up to this point. Past that point however, the intensity starts to fade away and the shapes of the harmonics start to distort, whilst the cutoff no longer extends, and in fact starts to slide back. By the time 70 TW/cm^2 is reached the harmonics are very distorted, with only harmonics up to around 27 eV or so clearly visible and even those are very squashed. This is because past a certain threshold intensity the effect of the ionisation of the liquid on the refractive index becomes significant, which affects the phasematching leading to the lower cutoff and intensity, whilst also causing the distortion.

over the harmonic generation over a small range of relatively low intensities, whereas the iris scan cannot take this fine a step, and also looks over a much greater range of intensities. The iris scan might therefore be less sensitive to such effects than the half waveplate scan, as the intensity range between the harmonics appearing might be smaller than a step in the iris scan. The other possibility is that in the regime of either a shorter pulse or a higher intensity (or both) any band-like structure in the liquid is disrupted and the harmonic generation mechanism becomes more gas-like than that of the longer and less intense pulse of the half waveplate scan, which could be more solid-like.

6.3.4 Dependence of the harmonic cut-off energy upon laser intensity

Figure 6.7 shows the dependence of the cutoff on the intensity of the driving field. The two sets of data were both taken with the iris being used to control the intensity, as the half waveplate data does not extend high enough in intensity to show the overall trend, and the cutoff extends past the end of the MCP at the higher intensities. Both sets of data (taken on different days, with an estimated pulse duration of 15 fs for the yellow dots and 12 fs for the purple dots) show an increase in cutoff as the intensity increases towards 50 TW/cm^2,

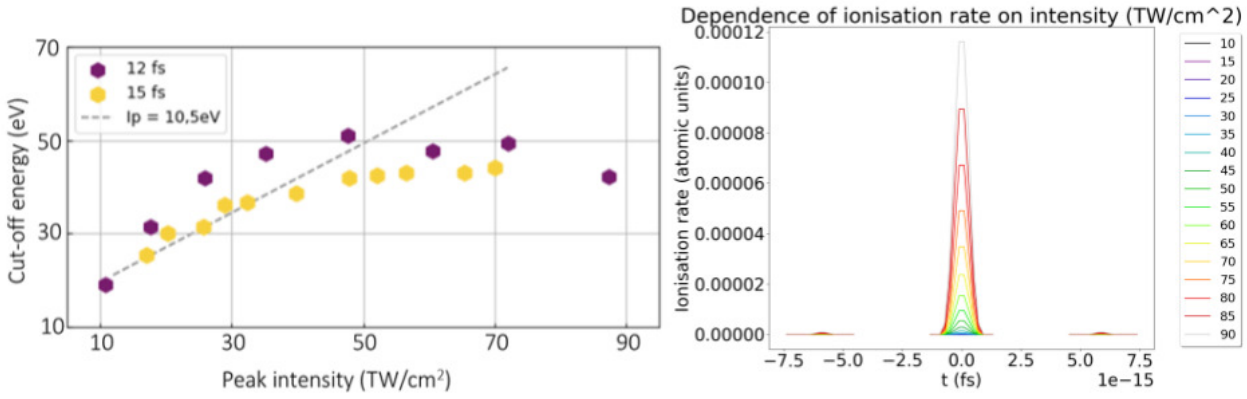

Figure 6.7: The left graph features two sets of data showing the change of cutoff with intensity, while the right graph shows the dependence of the ionisation rate on the intensity of the electric field. In the left figure the yellow dots are the cutoffs of the iris scan shown in Figure 6.6, and the purple dots are from a cutoff scan taken on the same day as the CEP scan, at which point the pulse was slightly better compressed. The dotted line shows the path that the cutoff would be expected to follow given an I_p of 10.5 eV. The right figure shows the ionisation rate calculated according to the equation in Chapter 3 for intensities from 10 to 90 TW/cm^2, an I_p of 10.12 eV and a 15 fs Gaussian pulse. From around 50 TW/cm^2 the ionisation rate shoots up (as might be expected due to the exponential dependence on the electric field). The left graph was made by Clement Ferchaud, for our paper in preparation [81]

and then a plateau afterwards. The yellow dots seem to follow this trend up to around 50 TW/cm^2, where the plasma breakdown starts to affect the harmonics that are generated. The purple dots also rise up to around 50 TW/cm^2, which seems to be the point at which the electron density becomes too much for HHG, and then plateau off. However the rise does not seem to follow a linear scaling, and in fact may be closer to a square root of the intensity dependence, which is what would be expected from a solid.

However, it must be said that these cutoffs were determined by counting the highest observable harmonic, and so the cutoff energies may be off by ± one harmonic separation. As such, the main conclusion that can be drawn from this analysis is that the threshold for plasma breakdown affecting the harmonic generation is consistently 50 TW/cm^2, which is also around the intensity at which the plasma plume appears. Some more light can be shed on this if we consider the ionisation rate of the isopropanol molecule. As set out in Chapter 3, we can use the ADK method with some modifications to work out the ionisation rate, an approach which has been shown to work for small molecules [199]. When we do so, we see that it is at around 50 TW/cm^2 that the ionisation rate starts growing noticeably, which provides a confirmation that the effects which switch on around that intensity really are due to the increase in the ionisation fraction.

6.4 Cut off variation by means of the CEP

One of the fundamental features of gas phase HHG is the attochirp. This is due to the electron recollision based mechanism of HHG as different energies are produced by electrons taking different excursion lengths into the continuum. These different lengths are due to the phase of the generating field at the time the electron is ionised into the continuum, and mean that different energy harmonics are produced at different times, leading to a fundamental chirping of the pulse, the attochirp. With a sufficiently short pulse, only a few attosecond bursts are created. As the pulse envelope is short, the optical cycles generating the harmonics have different peak intensities, and so photons of the same energy birthed from different cycles will have been born at a different time on the cycle, and thus will have a different phase. Or, to put it another way, as the CEP changes the phase profile of the attosecond burst is correspondingly altered. This leads to interference between these bursts, which leads to the harmonic structure, and therefore the harmonic structure can be changed by changing the CEP. Typically strong CEP effects can be observed for pulses shorter than two optical cycles.

One of the notable things about solid phase harmonics is that sometimes (depending on the solid), they do not exhibit this CEP dependence in their harmonics. This is explained by invoking an alternative mechanism to recombination, that of Bloch oscillations in the conduction band of the solid (see Chapter 3), which release their harmonics all at the same time, so there is no phase difference and therefore there should be no CEP dependence. In order to verify the recollision model of HHG in a liquid we performed a CEP scan of our harmonics to see if they exhibited this CEP dependence or not. As described in Chapter 4 we scanned the CEP by changing the thickness of a pair of fused silica wedges, and the resulting data is shown in Figure 6.8.

A CEP dependence can be clearly observed in this figure, with three separate regions of behaviour corresponding to the three regions identified in Rudawski *et al.* [200]. In this paper they show in gas phase harmonics with a few cycle pulse that the expected change of the relative phase of successive pulses ($\Delta\Phi$) of harmonics with CEP (ϕ_{cep}) is

$$\frac{\partial\Delta\Phi}{\partial\phi_{\text{cep}}} \propto \frac{\alpha_\Omega I_0}{\tau^2}, \tag{6.12}$$

where α_Ω is the photon energy, I_0 is the intensity of the driving field and τ is the duration of the generation pulse. When this fraction is small, so for long pulses, low intensities and low photon energies, there will be no change in phase and so no CEP dependence, as in region A. When this is not true, there will be a phase change between successive harmonic pulses, which leads to a CEP dependence in the position of the harmonic peaks, which will increase with increasing photon energy and intensity, and decreasing pulse duration. When the dependence is close to 1, then the phase difference changes as the CEP changes, leading to the nice 1:1 dependence in region C. Region B is the case when the dependence is non-integer, which leads to discontinuities in the CEP dependence.

This all therefore indicates that we are generating harmonics with a few cycle pulse (one cycle of a 1.8 µm

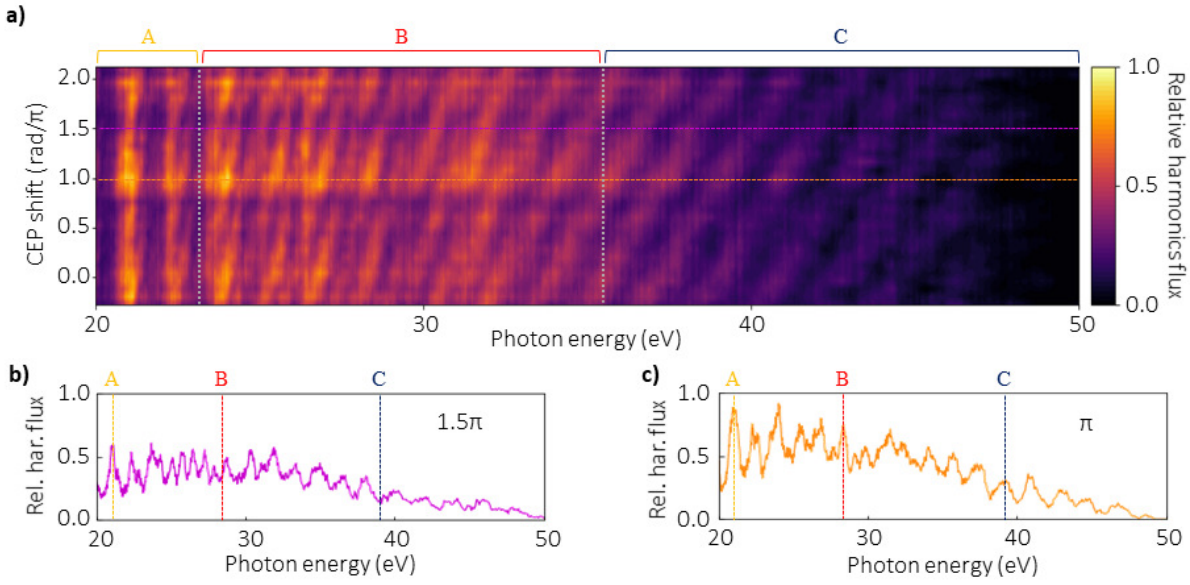

Figure 6.8: a) shows the variation of harmonics generated with the CEP of the generating field. Three sections of the spectrum can be clearly identified, corresponding to those in Rudawski *et al.* [200]: A is the low energy region where the phase of successive pulses is invariant with CEP, C is the high energy region where there is a 1:1 correspondance between the CEP and the phase of the pulses, so there is a nice diagonal interference pattern where odd and even harmonics turn on and off as the phase is scanned, and B is an intermediate region between the two. b) and c) are spectra taken at π and $1.5\,\pi$ radians, showing the lack of movement at the region marked A, compared with the region marked C where the harmonic disappears.

field is 6 fs long, so our generating pulse is 2.5 cycles long, which should be sort enough), and that the harmonics that we generate are generated with attochirp in exactly the same way as gas phase harmonics are. Whilst this does not rule out some kind of interband model as opposed to accelerating the electron in free space, it points very strongly towards a recollision model.

6.5 Cut off variation with ellipticity of the driving field

Another potentially important diagnostic test of the properties of this liquid phase HHG is the ellipticity dependence of the harmonics. A key difference between harmonics generated via an interband recollision mechanism and a propagation in the vacuum recollision mechanism is that the hole left behind by the excited electron in the interband mechanism can be delocalised across multiple sites. This means that the electron could recombine at several sites instead of just the parent site, and so an elliptical field will not have as drastic an effect on the efficiency of the harmonics as it does in gas phase HHG, where it kills the harmonics almost immediately [35] (see Chapter 3).

Although obviously highly liquid dependent, one could imagine a relatively short-chain alcohol like isopropanol might be able to form H-bond networks, and so when an electron is ionised from one of these molecules the charge could be delocalised over two or more molecules. Panman *et al.* [169] have suggested that such in-

stantaneous ordering could be relatively long lived (on the order of picoseconds for some molecules, in ethanol they predict it will be on the order of hundreds of femtoseconds). As the electron spends only a matter of femtoseconds being accelerated by the electric field such instantaneous ordering could conceivably have an effect.

In order to test this we used a quarter waveplate to change the ellipticity of the driving field (the ratio of the minor to the major axis). Due to the extra glass in the beam the driving pulse was stretched to around 19 fs, so we did not observe as high harmonics as in the CEP scan and they are a little more washed out. The results of the scan are shown in Figure 6.9.

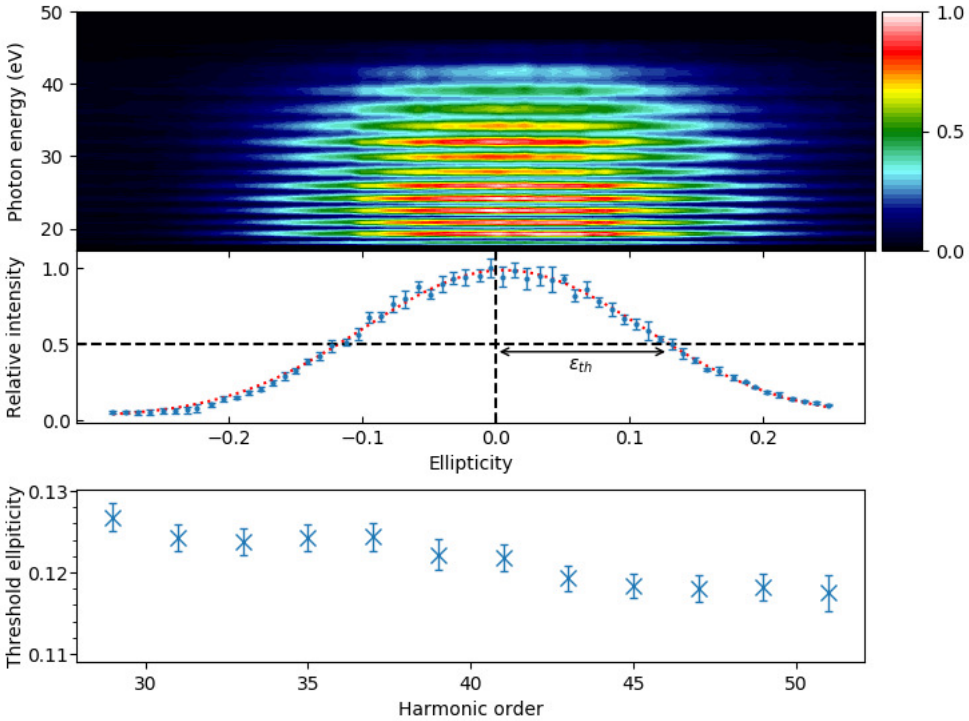

Figure 6.9: The ellipticity dependence of the harmonics generated from a liquid. The top panel shows the plots of the HHG spectra as the ellipticity is scanned through. These plots are made by averaging together four sets of ellipticity data. The middle panel shows the relative mean intensity at each value of ellipticity. The red curve behind them is a Gaussian that was fitted to the data, showing the FWHM of the ellipticity dependence. The third panel shows the threshold ellipticity (the ellipticity when the intensity passes 0.5 of the maximum) against harmonic order. There is a slight decrease in threshold ellipticity with increasing energy, which is what one would expect,as a shorter trajectory leads to less drift due to ellipticity, and thus the probability of recombination is higher.

The ellipticity dependence of these harmonics is roughly the same as produced by gas phase HHG (which is also similar to that produced by solid noble gases [142]) at 0.12–0.13 ellipticity FWHM, which indicates that such H-bond networks are not sufficiently long lived (or are insufficiently strong) in isopropanol to allow delocalisation of the electron hole. In solids with well defined band structures the ellipticity dependence is typically much greater than this. While this does not completely rule out a interband recollision mechanism

(as in rare gas solids an interband recollision mechanism is proposed whilst not having a delocalised hole left behind), we can at least say that there is no evidence from this data that there is an interband recollision type mechanism happening, and no evidence of any H-bond network allowing a delocalised hole.

6.6 Mechanism of HHG

Having assembled all of our evidence, we can make a conclusion about the likely mechanism of HHG from the liquid phase. From the half waveplate intensity scan we have a suggestion of an interband mechanism, as the harmonics appear to turn on at a threshold intensity, but this was not replicated in the variable aperture intensity scans, which we believe are superior due to a wider pulse energy span and a shorter pulse duration. The cutoff harmonic scales approximately linearly with the intensity of the driving field in the iris scan (although this is not conclusive), until a threshold is reached at which point the amount of ionisation overpowers the harmonic generation and destroys the phasematching. The CEP dependence of the harmonics demonstrates clear evidence of recollision and the ellipticity dependence suggests that the hole that the excited electron leaves behind during the HHG process is pretty much entirely localised in the molecule that was originally ionised.

From this we can conclude that the most likely mechanism is the 'gas like' recollision model, where the electron is tunnel ionised, accelerated in the continuum and then recollides with the parent ion to release a high energy photon. Whilst at low intensities it is possible that there are some effects from instantaneous local ordering leading to discrete energy levels that the electron can occupy, for the most part the liquid high harmonics generation source acts like a very dense gas.

However, this raises a big question, which I have avoided until now. This question is how can an electron propagate through the liquid without significant scattering due to the dense liquid environment?

6.6.1 Scattering

As the liquid does not possess long range order like a solid, and it is much more dense than a gas, the electron does not have conduction bands to be accelerated in without being scattered as it would in a solid, and it has many more particles that it can scatter off in its immediate vicinity than a gas. In fact the maximum excursion range for an electron in the three step model at 1.8 µm is 140 nm, which is orders of magnitude greater than the mean free path in most liquids — roughly the size of the liquid molecules themselves (a few angstroms).

As laid out in Chapter 3, the probability of the electron not scattering in the vacuum is

$$P_{\text{n.s.}} = \mathrm{e}^{-\rho_{\text{N}} \int_{t_0}^{t_{\max}} \sigma |v(t)|\, dt}, \tag{6.13}$$

where ρ_{N} is the number density of the medium, σ is the scattering cross section for the medium (in this case isopropanol), $v(t)$ is the velocity of the electron at time t, t_0 is the birth time of the electron into the continuum

and $t_{\max}$ is the time spent in the continuum by the electron before recollision. Two different scattering cross sections were used - a simple theoretical one based on an isopropanol molecule being a spherical potential well of diameter 4.9 Å [201] and one based on experimental values found from scattering low energy electrons off of isopropanol [74, 202, 203].

With the spherical well model the scattering cross section is simply πr^2, that is, independent of the velocity of the electron in the continuum. This makes the calculations easier, as the result of the integral is then just the distance travelled in the continuum. However in general scattering cross sections are not independent of the velocity of the electrons that are scattering, with higher energy electrons scattering less [204].

In order to take this into account I used the approach laid out in the supplementary information of Kurz *et al.* [204] for water. I fitted a straight line to a log/log plot of electron scattering data in propanol variants taken from [202] and [203]. One of these sources is for 1-propanol, not isopropanol (2-propanol). However, as is remarked upon in [202] itself, the data are very similar in the range where they overlap, and by including the 1-propanol data I have a dataset that extends to 100 eV meaning that all of the electron kinetic energies I am considering fall within the bounds of the data used to generate the fit. From this fit we obtain a power law relationship of

$$\sigma = \alpha E^{\gamma}, \tag{6.14}$$

where α and γ are the parameters of the fit and E is the kinetic energy of the electron — $\frac{1}{2}m_{\mathrm{e}}v^2$. We can combine the non time dependent parts together into a new constant,

$$\beta = \alpha \left(\frac{m_{\mathrm{e}}}{2}\right)^{\gamma}, \tag{6.15}$$

which then gives us a probability of not scattering of

$$P_{\mathrm{n.s.}} = \mathrm{e}^{-\rho_{\mathrm{N}}\beta \int_{t_0}^{t_{\max}} |v(t)^{2\gamma+1}|\, dt}. \tag{6.16}$$

The results from the two different approaches are shown in Figure 6.10. The cross sections from the published data show a much higher probability of scattering than those from the spherical quantum well, indicating that we should not see harmonics beyond around 10 eV as the scattering should quench them completely. Even the quantum spherical well model indicates we shouldn't see harmonics beyond around 20 eV. As we clearly do, the question arises — what are we missing from our model that allows the harmonics to be generated from this dense medium?

The answer is that this approach to using these cross sections (both the spherical well and measured) is too naive. These cross sections are found by scattering low energy electrons off a diffuse gas without a strong field in place to drive the electrons along and without any other potential sources of scattering in the vicinity.

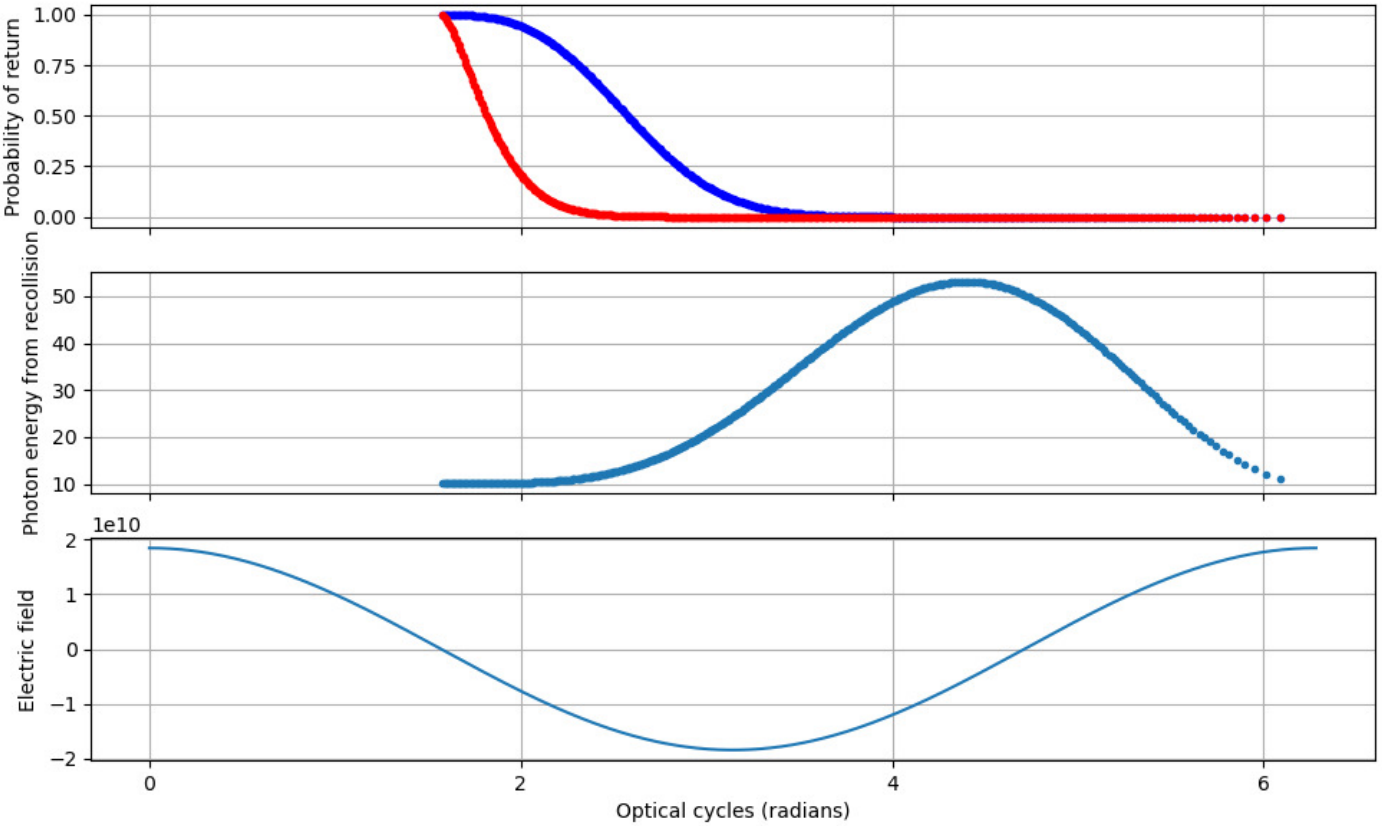

Figure 6.10: The first panel shows the probablity of return of an electron against the time of its return using the cross sections derived from the infinite spherical well (blue) and those interpolated from experimental data (red). The second panel shows the kinetic energy of the electrons at the time of recollision, and so therefore the energy of the photons released. The third panel shows the generating electric field. This simple approach predicts that no photons above 20 eV should be observed.

Neither of these assumptions is true in the case of liquid HHG, and this has a large effect on the probability of the free electron scattering.

This is highlighted by Boyle *et al.* [205] in discussing the difference between scattering cross sections between gaseous Xe and liquid Xe, and Michaud *et al.* [206], discussing the difference between gaseous water and ice electron scattering. In the gas phase the electron will only encounter one atom at a time, so the scattering effect can be found by just considering single atom scattering effects (as above). However, in the condensed phase the quantum mechanical nature of the electron have to be taken into account — it can act like a wave, scattering off many centers at once and interfering with itself in the process [205]. In addition, due to the proximity of the neighbouring molecules to the molecules the electron is scattering off, there is a screening effect on the scattering, which decreases the magnitude of the effect compared to the gas phase case where the scattering center is unaffected by its nearest neighbours [206]. This all means that the magnitude of the scattering effect

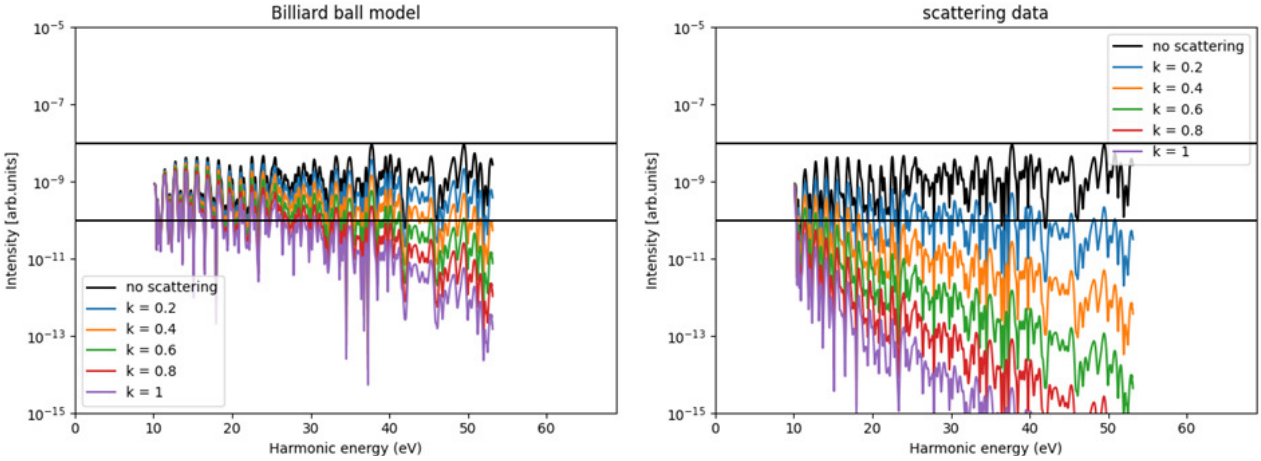

Figure 6.11: Harmonics calculated using the saddle point approximation to the single atom response in Lewenstein *et al.* [31] for a 2.5 cycle 1800 nm pulse (which gives a duration of around 15 fs). The black lines in each figure are the calculated harmonics, and then the coloured lines show these harmonics multiplied by the probability of not scattering calculated using our two approaches. The k in the legend is a scaling factor that the calculated crossection is multiplied by, and the black lines show a gap of two orders of magnitude, which is the dynamic range of our detector.

in condensed phase should be several times smaller than would be predicted using gas phase cross-sections and a classical model of the electron trajectory.

Boyle *et al.* suggest that the liquid cross-sections (Σ) can be obtained from gas phase cross sections (σ) by

$$\Sigma(v) = \mathrm{k}(v)\sigma(v), \tag{6.17}$$

where v is the velocity of the electron, and

$$\mathrm{k}(v) = \frac{\tilde{\nu}(v)}{\nu(v)}. \tag{6.18}$$

$\nu(v)$ and $\tilde{\nu}(v)$ are the collision frequencies in the gas and the liquid respectively. The calculation of these collision frequencies involves expanding in terms of Legendre polynomials and the calculation of the structure factors of isopropanol, and as such are beyond the scope of this thesis. However we can estimate the value of $\mathrm{k}(v)$ from comparing our measured harmonics to the predicted ones.

Figure 6.11 shows the results for this comparison. It shows simulated harmonics from a single atom calculated using the strong field approximation (see Chapter 3), multiplied by the probability of not scattering for various different values of $\mathrm{k}(v)\sigma(v)$, calculated with the two models above. As might be expected from the results shown in Figure 6.10 the billiard ball model is somewhat more optimistic than the fitted scattering data. However, taking the observed cutoff to be the point by which the spectrum dips below two orders of magnitude, the value of k for both models is around 0.4 for billliard ball and around 0.2 for the fitted data. Given the inherent error in judging something like this, these two results are not too dissimilar. The fitted scattering data are more likely to be reflective of the actual scattering properties of the isopropanol molecule, so from this we can conclude

that a k of around 0.2 is likely to be roughly right.

This conclusion can be compared to various papers that have been published on this subject. In Boyle *et al.* [205] for the highest energy electrons they looked at (which only went up to 10 eV) the liquid phase crossections are around 5 times smaller than the gas phase ones. In Michaud *et al.* [206] (which went up to 100 eV) the crossections are between 2% and 80% of those in the gas, with the gap decreasing for higher energy electrons (50 eV is somewhere between 8 and 10 times smaller). Our observed spectra therefore make sense, as a scattering effect is observed with roughly the same magnitude as that calculated from electron scattering experiments in solids, liquids and gases. Obviously this is quite crude, as k should depend on v rather than being constant, but for a simple check to see if we are on the right track it seems to show that we are.

6.6.2 Plasma distortion and comparison to Luu *et al.*

Above a threshold of $50\,\mathrm{TW/cm^2}$ the harmonics produced by the jet start to distort due to the plasma produced by ionisation of the liquid sheet target. The jet thus acts as we would expect a dielectric material to act — with no noticeable damage effects until a threshold is crossed, at which point there is a close to 100% certainty that the material will break down [207]. We observe this easily in the experiment by the presence or absence of the plasma plume, as under the conditions in which we observed the highest harmonic cutoff we generated no visible plume at all from the interaction between our sheet and the laser, and the presence of the plume coincides with the distortion of the harmonics.

The damage threshold of isopropanol is not readily available, but that of water is, and has been measured to be $25\,\mathrm{TW/cm^2}$ for a 140 fs pulse at 800 nm [208]. Assuming (as in Luu *et al.*) that isopropanol resembles water in terms of its bulk electronic structure then this would seem to be around half as high as we expect. However, the damage threshold of a dielectric material actually increases as the pulse length decreases and the wavelength increases [207, 209–211], so our figure of $50\,\mathrm{TW/cm^2}$ for a pulse of shorter duration and longer wavelength is not unexpected.

Luu *et al.* do not give an intensity figure for their generating field, but they do give an electric field strength, which is between 0.9 and $1.3\,\mathrm{V\AA^{-1}}$. Converting to intensity using Equation 6.2 gives intensities of between 107 and $224\,\mathrm{TW/cm^2}$, which are both much above the damage threshold that we find for isopropanol. As a result we would expect their jet to experience plasma breakdown during the generation process, and this is what we find — they report a much bigger plume than we do, which extends out of both sides of the jet. This points to the interaction region being totally destroyed by the laser pulse, such that it becomes entirely plasma and the plasma is then ejected out of both sides equally. This also explains the other two main differences to our spectra — the cutoff and the distorted look of the harmonics. Both of these are due to the large ionisation fraction changing the refractive index, throwing the phasematching off (which reduces flux and cutoff) and causing the distortion of the XUV pulses themselves.

6.7 Conclusion

We have created high harmonics from a thin liquid sheet target consisting of isopropanol. The harmonics show CEP and ellipticity dependence similar to harmonics produced by a gas. The intensity dependence of the harmonics is indeterminate, but it is likely to be the same as that produced by a gas. Our cutoff is roughly twice that of the other reported study of HHG in liquids (by Luu *et al.*), most likely due to the lower intensities used in our experiment, which avoids ionisation clamping and so allows for the best phasematching conditions.

We propose a recollision based model for understanding the generation of these harmonics, as it seems unlikely that any long range electronic structure is present in liquids to allow any other type of mechanism, and this is confirmed by our experimental results which are as would be expected for a recollision model. There was a concern that scattering might stop the generation of harmonics due to the fact that the mean free path in a liquid is much smaller than the necessary excursion depth of the electron, but analogies from other experimental and theoretical work on electron scattering show that at this density we cannot just look at the probability of a collision along the path using a gas phase scattering cross section, but instead need to scale the scattering cross sections by a factor to take into account the fact that the electron is interacting with many molecules at once.

Chapter 7

Conclusion

Over the course of this thesis I have covered the development of a thin liquid sheet target that is suitable to be deployed in a vacuum chamber. I have also covered the use of this thin liquid target as an HHG source, and the progress so far towards the use of this target as a source for solution-phase X-ray spectroscopy. Here I will summarise my conclusions and ideas for future research.

The main subject of this thesis has been the development and use of the liquid sheet jet as a target in vacuum. This has some significant advantages over the other major design of colliding jets, as the sheet produced is alignment free and easily reproducible, whilst also producing a larger target with fewer surface waves than the colliding jets geometry. This target was shown to be usable as an HHG source and also stable enough to be run in vacuum for several hours at a pressure of around 10^{-3} mbar, low enough to allow XUV light transmission. The HHG produced does not appear to have any special properties, but a liquid sheet target as a source could potentially be useful in a situation where a solid HHG source would need replacing due to laser damage. No meaningful solution phase spectra have yet been collected when the sheet was used as an X-ray absorption spectroscopy target, but the X-ray transmission spectrum of an ethanol sheet proves absorption spectrum measurements are possible using this target.

Maintaining a liquid sheet in vacuum is, somewhat understandably, a significant challenge. However, the challenge arises not so much for the integrity of the vacuum (a problem that can be solved by improved pumping of the vapour) but from the stability of the flat sheet. The sheet would often change size slowly as the local pressures, both that forcing the liquid through the nozzle and the ambient pressure around the sheet, were not constant. This was especially a problem with the gas driven system, where the pressure differential between the atmosphere and the backing pressure could change quite significantly over time. Since the nozzle and the collector orifice are not fixed in relation to each other the system occasionally misaligns. The above quirks did not prevent measurements, but rather made it harder to collect them. A fixed system, as presented in Chapter 5, could solve the misalignment problem, but an automated system would be needed to run the experiments

without supervision. Such a system would be required to adjust the flow rate to keep the sheet size constant and to be able to shut everything off in the case of a major issue.

The 3D printing of the nozzles was very effective for the initial rapid prototyping of the nozzle design, but once the design was finalised the downsides of this manufacturing process became obvious. The nozzles are fragile, especially during printing, and as a consequence had a high failure rate (around one in ten were usable). Fortunately successful nozzles wore out at an acceptable rate. In the future we would like to machine or cast the nozzles out of a more durable material such as ceramic, hopefully avoiding the difficulties we had with manufacture.

As a result of the issues we had with the printed nozzles, for the X-ray absorption experiments we switched to one of the gas dynamic virtual nozzles invented by Koralek *et al.* [55]. As these can be bought off the shelf (from Micronit GmbH) we were spared having to create our own nozzles, and our technical workshop were able to create a housing for these chips. They are cheap and reliable (£1240 for 10 at time of writing), but do not produce as large or as smooth a sheet as our fan spray nozzle.

Moving forward, the possibilities of thin liquid sheets in vacuum are only starting to be explored. They are not particularly efficient HHG sources, but the self healing nature means that liquid sheets could replace solid targets where target replacement is an issue but a low cut off is not. The potential for liquid sheets in time-resolved soft X-ray spectroscopy is only just starting to be explored. We hope to carry out some experiments to expand our knowledge of the ultrafast behaviour of chemically and biologically relevant molecules in the solution phase within the next few years .

The existence of HHG from the liquid in and of itself was surprising, especially the extent of the cutoff. Such extended HHG was not expected because the mean free path is far too short to allow free space acceleration of the electrons, and liquids as disordered dielectric materials do not have any conduction bands into which an electron can be promoted and accelerated. The observation of harmonics therefore indicated that our initial thoughts about the likely mechanism of HHG in a liquid were wrong, and required a more nuanced explanation. At this initial stage it appears that the electron will scatter less than predicted due to correlation effects within the condensed phase.

There is a lot more scope for investigation of HHG in liquids. The difference (or lack thereof) in the harmonics produced in very polar liquids (like water) and very non-polar ones (like cyclohexane) would be able to tell us some more about how this form of HHG takes place. Comparing the results to amorphous solids would also be interesting — is there any difference between the two forms of matter on the timescale of HHG? Ultimately the differences observed (or lack thereof) would function as an HHG spectroscopy experiment on the behaviour of disordered condensed phase materials when exposed to a strong laser field.

The time-resolved X-ray absorption spectroscopy experiment proved much more challenging than expected — the lack of flux, issues with the target and then finally severe contamination problems meant that little

meaningful data was collected. However, once the contamination issue is resolved, we will be in good stead to make progress towards solution phase measurements.

The use of the thin liquid sheet target is not confined to our lab. Although facility experiments have many orders of magnitude more flux than our HHG beamline, the transmission of common solvents in the soft X-ray range is sufficiently low that a thin target will be of use in soft X-ray transient absorption experiments at these facilities as well. In addition the low thickness and optical flatness means that they potentially enable more ambitious solution phase experiments at these facilities for which the flux is not necessarily very high, such as a two pulse soft X-ray pump-probe scheme or an impulsive Raman type experiment using a very broad bandwidth X-ray pulse to simultaneously excite a liquid molecule to a high lying electronic state and de-excite it to a lower state.

Thin liquid sheet jets of this type will surely have a long and productive future enabling soft X-ray experiments on liquid and solution phase targets, and hopefully the findings presented in this thesis will be of use to experimenters wishing to build and operate thin sheet jets of their own.

Bibliography

[1] Jamie Durrani. The Holy Grails of chemistry. *Chem. World*, pages 24–41, oct 2020.

[2] John C. Polanyi and Ahmed H. Zewail. Direct Observation of the Transition State. *Acc. Chem. Res.*, 28(3):119–132, 1995.

[3] Ahmed H Zewail. Femtochemistry: Ultrafast Dynamics of the Chemical Bond. *J. Phys. Chem. A*, 104(24):5660–5694, 2000.

[4] Henry Eyring. The activated complex in chemical reactions. *J. Chem. Phys.*, 3(2):63–71, 1935.

[5] M. G. Evans and M. Polanyi. Some applications of the transition state method to the calculation of reaction velocities, especially in solution. *Trans. Faraday Soc.*, 31:875–894, 1935.

[6] Nils Huse, Hana Cho, Kiryong Hong, Lindsey Jamula, Frank M F De Groot, Tae Kyu Kim, James K. McCusker, and Robert W. Schoenlein. Femtosecond soft X-ray spectroscopy of solvated transition-metal complexes: Deciphering the interplay of electronic and structural dynamics. *J. Phys. Chem. Lett.*, 2(8):880–884, 2011.

[7] Nils Huse, Haidan Wen, Dennis Nordlund, Erzsi Szilagyi, Dan Daranciang, Timothy A. Miller, Anders Nilsson, Robert W. Schoenlein, and Aaron M. Lindenberg. Probing the hydrogen-bond network of water via time-resolved soft X-ray spectroscopy. *Phys. Chem. Chem. Phys.*, 11(20):3951–3957, 2009.

[8] Andrew R. Attar, Aditi Bhattacherjee, and Stephen R. Leone. Direct Observation of the Transition-State Region in the Photodissociation of CH3I by Femtosecond Extreme Ultraviolet Transient Absorption Spectroscopy. *J. Phys. Chem. Lett.*, 6(24):5072–5077, 2015.

[9] Kirsten Schnorr, Aditi Bhattacherjee, Katherine J. Oosterbaan, Mickaël G. Delcey, Zheyue Yang, Tian Xue, Andrew R. Attar, Adam S. Chatterley, Martin Head-Gordon, Stephen R. Leone, and Oliver Gessner. Tracing the 267 nm-Induced Radical Formation in Dimethyl Disulfide Using Time-Resolved X-ray Absorption Spectroscopy. *J. Phys. Chem. Lett.*, 10(6):1382–1387, 2019.

[10] Stephen R Leone, C. William Mccurdy, Joachim Burgdörfer, Lorenz S Cederbaum, Zenghu Chang, Nirit Dudovich, Johannes Feist, Chris H Greene, Misha Ivanov, Reinhard Kienberger, Ursula Keller, Matthias F Kling, Zhi Heng Loh, Thomas Pfeifer, Adrian N Pfeiffer, Robin Santra, Kenneth Schafer, Albert Stolow, Uwe Thumm, and Marc J.J. Vrakking. What will it take to observe processes in 'real time'? *Nat. Photonics*, 8(3):162–166, 2014.

[11] Joseph Duris, Siqi Li, Taran Driver, Elio G. Champenois, James P. MacArthur, Alberto A. Lutman, Zhen Zhang, Philipp Rosenberger, Jeff W. Aldrich, Ryan Coffee, Giacomo Coslovich, Franz Josef Decker, James M. Glownia, Gregor Hartmann, Wolfram Helml, Andrei Kamalov, Jonas Knurr, Jacek Krzywinski, Ming Fu Lin, Jon P. Marangos, Megan Nantel, Adi Natan, Jordan T. O'Neal, Niranjan Shivaram, Peter Walter, Anna Li Wang, James J. Welch, Thomas J.A. Wolf, Joseph Z. Xu, Matthias F. Kling, Philip H. Bucksbaum, Alexander Zholents, Zhirong Huang, James P. Cryan, and Agostino Marinelli. Tunable isolated attosecond X-ray pulses with gigawatt peak power from a free-electron laser. *Nat. Photonics*, 14(1):30–36, 2020.

[12] Hubert Jean-Ruel, Meng Gao, Michal A. Kochman, Cheng Lu, Lai Chung Liu, Ryan R. Cooney, Carole A. Morrison, and R. J. Dwayne Miller. Ring-closing reaction in diarylethene captured by femtosecond electron crystallography. *J. Phys. Chem. B*, 117(49):15894–15902, 2013.

[13] Vern L. Schramm. Enzymatic Transition States and Drug Design. *Chem. Rev.*, 118(22):11194–11258, 2018.

[14] A S Johnson, L Miseikis, D A Wood, D R Austin, C Brahms, S Jarosch, C S Struber, P Ye, and J P Marangos. Measurement of sulfur L2,3 and carbon K edge XANES in a polythiophene film using a high harmonic supercontinuum. *Struct. Dyn.*, 3(062603), 2016.

[15] Fivos Perakis and Christian Gutt. Towards molecular movies with X-ray photon correlation spectroscopy. *Phys. Chem. Chem. Phys.*, 22(35):19443–19453, 2020.

[16] Kirill Prozument, Joshua H. Baraban, P. Bryan Changala, G. Barratt Park, Rachel G. Shaver, John S. Muenter, Stephen J. Klippenstein, Vladimir Y. Chernyak, and Robert W. Field. Photodissociation transition states characterized by chirped pulse millimeter wave spectroscopy. *Proc. Natl. Acad. Sci. U. S. A.*, 117(1):146–151, 2020.

[17] Giovanni Batignani, Carino Ferrante, and Tullio Scopigno. Accessing Excited State Molecular Vibrations by Femtosecond Stimulated Raman Spectroscopy. *J. Phys. Chem. Lett.*, 11(18):7805–7813, 2020.

[18] Zheng Li, Sandeep Gyawali, Anatoly A. Ischenko, Stuart Hayes, and R. J.Dwayne Miller. Mapping Atomic Motions with Electrons: Toward the Quantum Limit to Imaging Chemistry. *ACS Photonics*, 2020.

[19] Kelvin Anggara, Lydie Leung, Matthew J. Timm, Zhixin Hu, and John C. Polanyi. Approaching the forbidden fruit of reaction dynamics: Aiming reagent at selected impact parameters. *Sci. Adv.*, 4(10):1–7, 2018.

[20] E Goulielmakis, M Schultze, M Hofstetter, V S Yakovlev, J Gagnon, M Uiberacker, A L Aquila, E M Gullikson, D T Attwood, R Kienberger, F Krausz, and U Kleineberg. Single-cycle nonlinear optics. *Science*, 320(5883):1614–7, 2008.

[21] Douglas Garratt. *Time Resolved X-ray Absorption Spectroscopy in Poly (3-hexylthiophene)*. PhD thesis, Imperial College London, 2020.

[22] Simon P. Neville, Vitali Averbukh, Serguei Patchkovskii, Marco Ruberti, Renjie Yun, Majed Chergui, Albert Stolow, and Michael S. Schuurman. Beyond structure: Ultrafast X-ray absorption spectroscopy as a probe of non-adiabatic wavepacket dynamics. *Faraday Discuss.*, 194:117–145, 2016.

[23] Frank De Groot. High-resolution X-ray emission and X-ray absorption spectroscopy. *Chem. Rev.*, 101(6):1779–1808, 2001.

[24] J J Rehr and R C Albers. Theoretical approaches to x-ray absorption fine structure. *Rev. Mod. Phys.*, 72(3):621–654, 2000.

[25] Joachim Stohr. *NEXAFS Spectroscopy*. Springer-Verlag, Berlin, 1 edition, 1992.

[26] Pierre Agostini and Louis F. DiMauro. The physics of attosecond light pulses. *Reports Prog. Phys.*, 67(6):813–855, 2004.

[27] Douglas Garratt. P3HT time resolved study. *Prep.*

[28] Anne L'Huillier and Phillipe Balcou. High-order harmonic generation in rare gases with a 1 ps 1053 nm laser. *Phys. Rev. Lett.*, 70(6):774–777, 1993.

[29] P. B. Corkum. Plasma perspective on strong field multiphoton ionization. *Phys. Rev. Lett.*, 71(13):1994–1997, 1993.

[30] K. C. Kulander, K. J. Schafer, and J. L. Krause. Dynamics of Short-Pulse Excitation, Ionization and Harmonic Conversion. In *Super-Intense Laser-Atom Phys.*, pages 95–110. 1993.

[31] M. Lewenstein, M. Yu Balcou, P.and Ivanov, Anne L'Huillier, and P. B. Corkum. Theory of high-harmonic generation by low-frequency laser fields. *Phys. Rev. A*, 49(3):2117–2132, 1994.

[32] P B Corkum, N H Burnett, and M Y Ivanov. Subfemtosecond pulses. *Opt. Lett.*, 19(22):1870, 1994.

[33] Thomas Gaumnitz, Arohi Jain, Yoann Pertot, Martin Huppert, Inga Jordan, Fernando Ardana-Lamas, and Hans Jakob Wörner. Streaking of 43 attosecond soft X-ray pulses generated by a passively CEP-stable mid-infrared driver. *Opt. Express*, 25(22):27506, 2017.

[34] ELI Beamlines Strategic Development Plan. (November 2017).

[35] Shambhu Ghimire and David A. Reis. High-harmonic generation from solids. *Nat. Phys.*, 15(1):10–16, 2019.

[36] C. Vozzi, R. Torres, M. Negro, L. Brugnera, T. Siegel, C. Altucci, R. Velotta, F. Frassetto, L. Poletto, P. Villoresi, S. De Silvestri, S. Stagira, and J. P. Marangos. High harmonic generation spectroscopy of hydrocarbons. *Appl. Phys. Lett.*, 97(24):1–4, 2010.

[37] H. G. Muller. Reconstruction of attosecond harmonic beating by interference of two-photon transitions. *Appl. Phys. B Lasers Opt.*, 74(SUPPL.):17–21, 2002.

[38] Álvaro Jiménez-Galán, Rui E. F. Silva, and Misha Ivanov. Reconstruction of attosecond beating by interference of two-photon transitions in solids. *arXiv*, 2009.02214, 2020.

[39] Ferenc Krausz and Misha Ivanov. Attosecond physics. *Rev. Mod. Phys.*, 2009.

[40] Allan S Johnson, Dane R Austin, David A Wood, Christian Brahms, Andrew Gregory, Konstantin B Holzner, Sebastian Jarosch, Esben W Larsen, Susan Parker, Christian S Strüber, Peng Ye, John W.G. Tisch, and Jon P Marangos. High-flux soft x-ray harmonic generation from ionization-shaped few-cycle laser pulses. *Sci. Adv.*, 4(5):eaar3761, 2018.

[41] G W C Kaye, T H Laby, and National Physical Laboratory. Tables of physical and chemical constants. volume 1, pages 10–14. 2016.

[42] B.L. Henke, E.M. Gullikson, and J.C. Davis. X-Ray Interactions: Photoabsorption, Scattering, Transmission, and Reflection at E = 50–30,000 eV, Z = 1–92. *At. Data Nucl. Data Tables*, 54(2):181–342, 1993.

[43] Inga Jordan, Martin Huppert, Dominik Rattenbacher, Michael Peper, Denis Jelovina, Conaill Perry, Aaron Von Conta, Axel Schild, Hans Jakob Wörner, and Monte Carlo. Attosecond spectroscopy of liquid water. *Science (80-.).*, 369(August):974–979, 2020.

[44] H. Iglev, A. Trifonov, A. Thaller, I. Buchvarov, T. Fiebig, and A. Laubereau. Photoionization dynamics of an aqueous iodide solution: The temperature dependence. *Chem. Phys. Lett.*, 403(1-3):198–204, 2005.

[45] Toshinori Suzuki. Time-resolved photoelectron spectroscopy of non-adiabatic electronic dynamics in gas and liquid phases. *Int. Rev. Phys. Chem.*, 31(2):265–318, 2012.

[46] Gaurav Kumar, Anirban Roy, Ryan S. McMullen, Shanmukh Kutagulla, and Stephen E. Bradforth. The influence of aqueous solvent on the electronic structure and non-adiabatic dynamics of indole explored by liquid-jet photoelectron spectroscopy. *Faraday Discuss.*, 212:359–381, 2018.

[47] Jamie W. Riley, Bingxing Wang, Michael A. Parkes, and Helen H. Fielding. Design and characterization of a recirculating liquid-microjet photoelectron spectrometer for multiphoton ultraviolet photoelectron spectroscopy. *Rev. Sci. Instrum.*, 90(8), 2019.

[48] Richard Neutze, Remco Wouts, Simone Techert, Jan Davidsson, Menhard Kocsis, Adam Kirrander, Friedrich Schotte, and Michael Wulff. Visualizing photochemical dynamics in solution through picosecond X-ray scattering. *Phys. Rev. Lett.*, 87(19):3–6, 2001.

[49] Ivan V. Tomov, Dmitri A. Oulianov, Peilin Chen, and Peter M. Rentzepis. Ultrafast time-resolved transient structures of solids and liquids studied by means of X-ray diffraction and EXAFS. *J. Phys. Chem. B*, 103(34):7081–7091, 1999.

[50] Jeongho Kim, Kyung Hwan Kim, Key Young Oang, Jae Hyuk Lee, Kiryong Hong, Hana Cho, Nils Huse, Robert W Schoenlein, Tae Kyu Kim, and Hyotcherl Ihee. Tracking reaction dynamics in solution by pump–probe X-ray absorption spectroscopy and X-ray liquidography (solution scattering). *Chem. Commun.*, 52(19):3734–3749, 2016.

[51] C. Mueller, M. Harb, J. R. Dwyer, and R. J Dwayne Miller. Nanofluidic cells with controlled pathlength and liquid flow for rapid, high-resolution in situ imaging with electrons. *J. Phys. Chem. Lett.*, 4(14):2339–2347, 2013.

[52] A. Charvat, E. Lugovoj, M. Faubel, and B. Abel. New design for a time-of-flight mass spectrometer with a liquid beam laser desorption ion source for the analysis of biomolecules. *Rev. Sci. Instrum.*, 75(5 PART 1):1209–1218, 2004.

[53] D. P. DePonte, U. Weierstall, K. Schmidt, J. Warner, D. Starodub, J. C.H. Spence, and R. B. Doak. Gas dynamic virtual nozzle for generation of microscopic droplet streams. *J. Phys. D. Appl. Phys.*, 41(19), 2008.

[54] Maria Ekimova, Wilson Quevedo, Manfred Faubel, Philippe Wernet, and Erik T J Nibbering. A liquid flatjet system for solution phase soft-x-ray spectroscopy. *Struct. Dyn.*, 2(5):054301, 2015.

[55] Jake D. Koralek, Jongjin B. Kim, Petr Brůža, Chandra B. Curry, Zhijiang Chen, Hans A. Bechtel, Amy A. Cordones, Philipp Sperling, Sven Toleikis, Jan F. Kern, Stefan P. Moeller, Siegfried H. Glenzer, and Daniel P. DePonte. Generation and characterization of ultrathin free-flowing liquid sheets. *Nat. Commun.*, 9(1), 2018.

[56] N. Dombrowski, D. Hasson, and D. E. Ward. Some aspects of liquid flow through fan spray nozzles. *Chem. Eng. Sci.*, 12(1):35–50, 1960.

[57] M. Faubel, S. Schlemmer, and J. P. Toennies. A molecular beam study of the evaporation of water from a liquid jet. *Zeitschrift für Phys. D Atoms, Mol. Clust.*, 10(2-3):269–277, 1988.

[58] Alfonso M. Gañán-Calvo. Generation of steady liquid microthreads and micron-sized monodisperse sprays in gas streams. *Phys. Rev. Lett.*, 80(2):285–288, 1998.

[59] Akira Watanabe, Hiroshi Saito, Yuzo Ishida, Masafumi Nakamoto, and Tatsuo Yajima. A new nozzle producing ultrathin liquid sheets for femtosecond pulse dye lasers. *Opt. Commun.*, 71(5):301–304, 1989.

[60] Geoffrey Taylor. Formation of thin flat sheets of water. *Proc. R. Soc. A Math. Phys. Eng. Sci.*, 259, 1960.

[61] Jens Eggers and Emmanuel Villermaux. Physics of liquid jets. *Reports Prog. Phys.*, 71(3):036601, 2008.

[62] Christophe Clanet and Emmanuel Villermaux. Life of a smooth liquid sheet. *J. Fluid Mech.*, 462:307–340, 2002.

[63] N. Dombrowski and G. Munday. Spray drying. In *Biochem. Biol. Eng. Sci.*, volume 2, pages 209–320. 1968.

[64] John W.M. Bush and Alexander E. Hasha. On the collision of laminar jets: Fluid chains and fishbones. *J. Fluid Mech.*, 511:285–310, 2004.

[65] N. Bremond and E. Villermaux. Atomization by jet impact. *J. Fluid Mech.*, 549:273–306, 2006.

[66] E. A. Ibrahim and A. J. Przekwas. Impinging jets atomization. *Phys. Fluids A*, 3(12):2981–2987, 1991.

[67] David Hasson and Ralph E. Peck. Thickness distribution in a sheet formed by impinging jets. *AIChE J.*, 10(5):752–754, 1964.

[68] Y. J. Choo and B. S. Kang. The velocity distribution of the liquid sheet formed by two low-speed impinging jets. *Phys. Fluids*, 14(2):622–627, 2002.

[69] G. Taylor. The Dynamics of Thin Sheets of Fluid. I. *Proc. R. Soc. A Math. Phys. Eng. Sci.*, 253(1274):289–295, 1959.

[70] Geoffrey Taylor. The dynamics of thin sheets of fluid II. Waves on fluid sheets. *Proc. R. Soc. A Math. Phys. Eng. Sci.*, 253:296–312, 1959.

[71] Geoffrey Taylor. Dynamics of thin sheets of fluid III. Disintegration of fluid sheets. *Proc. R. Soc. A Math. Phys. Eng. Sci.*, 253:313–321, 1959.

[72] P. K. Runge and R. Rosenberg. Unconfined flowing-dye films for CW dye lasers. *IEEE J. Quantum Electron.*, 8(12):910–911, 1972.

[73] Gediminas Galinis, Jergus Strucka, Jonathan C.T. Barnard, Avi Braun, Roland A. Smith, and Jon P. Marangos. Micrometer-thickness liquid sheet jets flowing in vacuum. *Rev. Sci. Instrum.*, 88(8), 2017.

[74] Heiko G Kurz, Martin Kretschmar, Thomas Binhammer, Tamas Nagy, Detlev Ristau, Manfred Lein, Uwe Morgner, and Milutin Kovacev. Revealing the microscopic real-space excursion of a laser-driven electron. *Phys. Rev. X*, 6(3):1–7, 2016.

[75] Anthony D DiChiara, Emily Sistrunk, Terry A Miller, Pierre Agostini, and Louis F DiMauro. An investigation of harmonic generation in liquid media with a mid-infrared laser. *Opt. Express*, 17(23):20959–20965, 2009.

[76] A. Flettner, T. Pfeifer, D. Walter, C. Winterfeldt, C. Spielmann, and G. Gerber. High-harmonic generation and plasma radiation from water microdroplets. In *Appl. Phys. B Lasers Opt.*, volume 77, pages 747–751, 2003.

[77] Heiko G. Kurz, Daniel S. Steingrube, Detlev Ristau, Manfred Lein, Uwe Morgner, and Milutin Kovačev. High-order-harmonic generation from dense water microdroplets. *Phys. Rev. A - At. Mol. Opt. Phys.*, 87(6), 2013.

[78] B. Kim, B. Ahn, D. Lee, J. Kim, and D. Kim. Optimization of laser parameters for the maximum efficiency in the generation of water-window radiation using a liquid nitrogen jet. *Appl. Phys. Lett.*, 88(14), 2006.

[79] P Heissler, E Lugovoy, R Horlein, L Waldecker, J Wenz, M Heigoldt, K Khrennikov, S Karsch, F Krausz, B Abel, and G D Tsakiris. Using the third state of matter: high harmonic generation from liquid targets. *New J. Phys*, 16, 2014.

[80] Shambhu Ghimire, Anthony D. Dichiara, Emily Sistrunk, Pierre Agostini, Louis F. Dimauro, and David A. Reis. Observation of high-order harmonic generation in a bulk crystal. *Nat. Phys.*, 7(2):138–141, 2011.

[81] Oliver Alexander, Jonathan C T Barnard, Esben W Larsen, Timur Avni, Sebastian Jarosch, Clement Ferchaud, Andrew Gregory, Susan Parker, Gediminas Galinis, Alexandra Tofful, Douglas Garratt, Mary Matthews, and Jon P Marangos. The mechanism of high harmonic generation in liquid alcohol. *(Submitted to Nat. Physics).*

[82] Tran Trung Luu, Zhong Yin, Arohi Jain, Thomas Gaumnitz, Yoann Pertot, Jun Ma, and Hans Jakob Wörner. Extreme–ultraviolet high–harmonic generation in liquids. *Nat. Commun.*, 9(1):3723, 2018.

[83] Aghiad Ghazal, Josiane P. Lafleur, Kell Mortensen, Jörg P. Kutter, Lise Arleth, and Grethe V. Jensen. Recent advances in X-ray compatible microfluidics for applications in soft materials and life sciences. *Lab Chip*, 16(22):4263–4295, 2016.

[84] Christophe Clanet. Waterbells and Liquid Sheets. *Annu. Rev. Fluid Mech.*, 2007.

[85] J D Naber and R D Reitz. Modeling engine spray/wall impingement. In *SAE Tech. Pap.*, 1988.

[86] N Dombrowski and R P Fraser. A photographic investigation into the disintegration of liquid sheets. *Math. Phys. Sci.*, 247(924):101–130, 1954.

[87] N. Dombrowski and P. C. Hooper. A study of the sprays formed by impinging jets in laminar and turbulent flow. *J. Fluid Mech.*, 18(3):392–400, 1964.

[88] Y. J. Choo and S. B. Kang. Parametric study on impinging-jet liquid sheet thickness distribution using an interferometric method. *Exp. Fluids*, 31(1):56–62, 2001.

[89] W. E. Ranz. Some experiments on the dynamics of liquid films. *J. Appl. Phys.*, 30(12):1950–1955, 1959.

[90] K. Dexter Miller. Distribution of spray from impinging liquid jets [14]. *J. Appl. Phys.*, 31(6):1132–1133, 1960.

[91] J. Y. Choo and S. B. Kang. The effect of jet velocity profile on the characteristics of thickness and velocity of the liquid sheet formed by two impinging jets. *Phys. Fluids*, 19(11), 2007.

[92] Byunghang Ha, Daniel P. Deponte, and Juan G. Santiago. Device design and flow scaling for liquid sheet jets. *Phys. Rev. Fluids*, 3(11), 2018.

[93] C. J. Clark and N. Dombrowski. The dynamics of the rim of a fan spray sheet. *Chem. Eng. Sci.*, 26(11):1949–1952, 1971.

[94] C. J. Clark and N. Dombrowski. On the formation of drops from the rims of fan spray sheets. *J. Aerosol Sci.*, 3(3), 1972.

[95] Daniel Weihs. Stability of thin, radially moving liquid sheets. *J. Fluid Mech.*, 87(2):289–298, 1978.

[96] J. C.P. Huang. The break-up of axisymmetric liquid sheets. *J. Fluid Mech.*, 43(2):305–319, 1970.

[97] T. Funada and D. D. Joseph. Viscous potential flow analysis of Kelvin–Helmholtz instability in a channel. *J. Fluid Mech.*, 445:263–283, 2001.

[98] E. Villermaux and C. Clanet. Life of a flapping liquid sheet. *J. Fluid Mech.*, 462:341–363, 2002.

[99] L. Rayleigh. On the Tension of Recently Formed Liquid Surfaces. *Proc. R. Soc. London*, 47:281–287, 1889.

[100] Lord Rayleigh. On the Instability of Jets. *Proc. London Math. Soc.*, 14(June 1873):4–13, 1878.

[101] Sungjune Jung, Stephen D. Hoath, Graham D. Martin, and Ian M. Hutchings. Atomization patterns produced by the oblique collision of two Newtonian liquid jets. *Phys. Fluids*, 22(4):1–10, 2010.

[102] Nanoscribe;. Nanoscribe Photonic Professional (GT) User Manual Revision: 496. (July):46,85–86,, 2015.

[103] J. P. Letouzey and S. O. Sari. Continuous pulse train dye laser using an open flowing passive absorber. *Appl. Phys. Lett.*, 23(6):311–313, 1973.

[104] Jason Mark Cole. *Diagnosis and Application of Laser Wakefield Accelerators*. PhD thesis, Imperial College Lodon, 2016.

[105] Christopher E Heil and David F Walnut. Continuous and Discrete wavelet transforms. *SIAM Rev.*, 31(4):628–666, 1989.

[106] Richard M. Goldstein, Howard A. Zebker, and Charles L. Werner. Satellite radar interferometry: Two dimensional phase unwrapping. *Radio Sci.*, 23(4):713–720, 1988.

[107] Milton van Dyke and Frank M. White. An Album of Fluid Motion. *J. Fluids Eng.*, 104(4):542–543, 1982.

[108] Geoffrey New. *Nonlinear Optics*. 2011.

[109] Dane Austin. *High-resolution Interferometric Diagnostics for Ultrashort Pulses*. PhD thesis, Oxford University, 2010.

[110] Andrius Baltuska, Th. Udem, M. Uiberacker, M. Hentschel, E. Goulielmakis, Ch. Gohle, R. Holzwarth, V. S. Yakovlev, A. Scrinzi, T. W. Hansch, and Ferenc Krausz. Attosecond control of electronic processes by intense light fields. *Nature*, 421(6923):611–615, 2003.

[111] R Paschotta. Chromatic dispersion. In *RP Photonics Encycl. <https//www.rp-photonics.com/>*.

[112] Eugene Hecht. *Optics, Global Edition*. Pearson Education Limited, 2016.

[113] R. Paschotta. waveplates. In *RP Photonics Encycl. <https//www.rp-photonics.com/>*.

[114] John Kerr. XL. A new relation between electricity and light: Dielectrified media birefringent . *London, Edinburgh, Dublin Philos. Mag. J. Sci.*, 50(332):337–348, 1875.

[115] F. Pockels. Ueber den Einfluss des elektrostatischen Feldes auf das optische Verhalten piezoelektrischer Krystalle. Mit 14 Textfiguren. *Abhandlungen der Gesellschaft der Wissenschaften Göttingen, Math. Klasse*, 39:1–204, 1894.

[116] T. H. Maiman. Stimulated optical radiation in ruby. *Nature*, 187:493–494, 1960.

[117] P. A. Franken, A. E. Hill, C. W. Peters, and G. Weinreich. Generation of optical harmonics. *Phys. Rev. Lett.*, 7(4):118–119, 1961.

[118] Andrius Baltuška, Takao Fuji, Takayoshi Kobayashi, Andrius Baltuska, Takao Fuji, and Takayoshi Kobayashi. Controlling the carrier-envelope phase of ultrashort light pulses with optical parametric amplifiers. *Phys. Rev. Lett.*, 88(13):133901, 2002.

[119] A. Ducasse, C. Rullière, and B. Couillaud. Methods for the generation of ultrashort laser pulses: mode-locking. In *Femtosecond Laser Pulses*, chapter 3, pages 57–87. Springer, 2 edition, 2005.

[120] Kasra Amini, Jens Biegert, Francesca Calegari, Alexis Chacón, Marcelo F. Ciappina, Alexandre Dauphin, Dmitry K. Efimov, Carla Figueira de Morisson Faria, Krzysztof Giergiel, Piotr Gniewek, Alexandra S. Landsman, Michał Lesiuk, Michał Mandrysz, Andrew S. Maxwell, Robert Moszyński, Lisa Ortmann, Jose Antonio Pérez-Hernández, Antonio Picón, Emilio Pisanty, Jakub Prauzner-Bechcicki, Krzysztof Sacha, Noslen Suárez, Amelle Zaïr, Jakub Zakrzewski, and Maciej Lewenstein. Symphony on strong field approximation. *Rep. Prog. Phys.*, 82(11):116001, 2019.

[121] Enikoe Seres, Jozsef Seres, and Christian Spielmann. X-ray absorption spectroscopy in the keV range with laser generated high harmonic radiation. *Appl. Phys. Lett.*, 89(18):1–4, 2006.

[122] L. V. Keldysh. Ionization in the field of a strong electromagnetic wave. *J. Exp. Theor. Phys.*, 47(20, 5):1307–1314, 1964.

[123] Gennady L. Yudin and Misha Yu Ivanov. Nonadiabatic tunnel ionization: Looking inside a laser cycle. *Phys. Rev. A - At. Mol. Opt. Phys.*, 64(1):4, 2001.

[124] M. V. Ammosov, N. B. Delone, and V. P. Krainov. Tunnel ionisation of complex atoms and of atomic ions in an alternating electromagnetic field. *Zh. Eksp. Teor. Fiz.*, 91.

[125] A. Perelomov, V. Popov, and M. Terent'ev. Ionization of atoms in an alternating electric field. *Sov. J. Exp. Theor. Phys.*, 23(5):924, 1966.

[126] E. Clementi, D. L. Raimondi, and W. P. Reinhardt. Atomic screening constants from SCF functions. II. Atoms with 37 to 86 electrons. *J. Chem. Phys.*, 47(4):1300–1307, 1967.

[127] K. Schiessl, K. L. Ishikawa, E. Persson, and J. Burgdörfer. Wavelength dependence of high-harmonic generation from ultrashort pulses. *J. Mod. Opt.*, 55(16):2617–2630, 2008.

[128] Kenichi L. Ishikawa, Eiji J. Takahashi, and Katsumi Midorikawa. Wavelength dependence of high-order harmonic generation with independently controlled ionization and ponderomotive energy. *Phys. Rev. A - At. Mol. Opt. Phys.*, 80(1):1–4, 2009.

[129] Yunhui Wang, Chao Yu, Qi Shi, Yadong Zhang, Xu Cao, Shicheng Jiang, and Ruifeng Lu. Reexamination of wavelength scaling of harmonic yield in intense midinfrared fields. *Phys. Rev. A - At. Mol. Opt. Phys.*, 89(2):2–9, 2014.

[130] Anh Thu Le, Hui Wei, Cheng Jin, and C. D. Lin. Strong-field approximation and its extension for high-order harmonic generation with mid-infrared lasers. *J. Phys. B At. Mol. Opt. Phys.*, 49(5), 2016.

[131] Esben Witting Larsen. *Attosecond Sources and Interferometers — Developments and Applications*. PhD thesis, Lund University, 2016.

[132] Philippe Balcou, Pascal Salières, Anne L'Huillier, and Maciej Lewenstein. Generalized phase-matching conditions for high harmonics: The role of field-gradient forces. *Phys. Rev. A - At. Mol. Opt. Phys.*, 55(4):3204–3210, 1997.

[133] H. Dachraoui, T. Auguste, A. Helmstedt, P. Bartz, M. Michelswirth, N. Mueller, W. Pfeiffer, P. Salieres, and U. Heinzmann. Interplay between absorption, dispersion and refraction in high-order harmonic generation. *J. Phys. B At. Mol. Opt. Phys.*, 42(17), 2009.

[134] Thomas Pfeifer. *Adaptive Control of Coherent Soft X-Rays*. PhD thesis, Julius-Maximilians-Universitat Wurzburg, 2004.

[135] A. Penzkofer H.J. Lehmeier, W. Leupacher. Nonresonant third order hyperpolarizability. *Opt. Commun.*, 56(1):67–72, 1985.

[136] E Constant, D Garzella, P Breger, E Mével, Ch Dorrer, C Le Blanc, F Salin, and P Agostini. Optimizing high harmonic generation in absorbing gases: Model and experiment. *Phys. Rev. Lett.*, 82(8):1668–1671, 1999.

[137] Shambhu Ghimire, Georges Ndabashimiye, Anthony D. Dichiara, Emily Sistrunk, Mark I. Stockman, Pierre Agostini, Louis F. Dimauro, and David A. Reis. Strong-field and attosecond physics in solids. *J. Phys. B At. Mol. Opt. Phys.*, 47(20), 2014.

[138] Edyta N. Osika, Alexis Chacón, Lisa Ortmann, Noslen Suárez, Jose Antonio Pérez-Hernández, Bartlomiej Szafran, Marcelo F. Ciappina, Fernando Sols, Alexandra S. Landsman, and Maciej Lewenstein. Wannier-Bloch approach to localization in high-harmonics generation in solids. *Phys. Rev. X*, 7(2):1–14, 2017.

[139] G. Vampa, C. R. McDonald, G. Orlando, D. D. Klug, P. B. Corkum, and T. Brabec. Theoretical analysis of high-harmonic generation in solids. *Phys. Rev. Lett.*, 113(7):1–5, 2014.

[140] G. Vampa, C. R. McDonald, G. Orlando, P. B. Corkum, and T. Brabec. Semiclassical analysis of high harmonic generation in bulk crystals. *Phys. Rev. B - Condens. Matter Mater. Phys.*, 91(6), 2015.

[141] M. Garg, M. Zhan, T. T. Luu, H. Lakhotia, T. Klostermann, A. Guggenmos, and E. Goulielmakis. Multi-petahertz electronic metrology. *Nature*, 538(7625):359–363, 2016.

[142] Georges Ndabashimiye, Shambhu Ghimire, Mengxi Wu, Dana A. Browne, Kenneth J. Schafer, Mette B. Gaarde, and David A. Reis. Solid-state harmonics beyond the atomic limit. *Nature*, 534(7608):520–523, 2016.

[143] Dean M. DeLongchamp, Eric K. Lin, and Daniel A. Fischer. Organic semiconductor structure and chemistry from near-edge X-ray absorption fine structure (NEXAFS) spectroscopy. *Org. Field-Effect Transistors IV*, 5940(August 2005):59400A, 2005.

[144] J. J. Rehr and A. L. Ankudinov. Progress in the theory and interpretation of XANES. *Coord. Chem. Rev.*, 249(1-2):131–140, 2005.

[145] Jacob W. Smith and Richard J. Saykally. Soft X-ray absorption spectroscopy of liquids and solutions. *Chem. Rev.*, 117(23):13909–13934, 2017.

[146] D. Strickland and G. Mourou. Compression of ampified chirped optical pulses. *Opt. Commun.*, 55(6):447–449, 1985.

[147] P. Maine, D. Strickland, P. Bado, M. Pessot, and G. Mourou. Generation of ultrahigh peak power pulses by chirped pulse amplification. *IEEE J. Quantum Electron.*, 1988.

[148] R L Fork, W J Tomlinson, C V Shank, C. Hirlimann, and R Yen. Femtosecond white-light continuum pulses. *Opt. Lett.*, 8(1):1, 1983.

[149] Giulio Cerullo and Sandro De Silvestri. Ultrafast optical parametric amplifiers. *Rev. Sci. Instrum.*, 74(1 I):1–18, 2003.

[150] R Baumgartner and R Byer. Optical parametric amplification. *Quantum Electron. IEEE J.*, 15(6):432–444, 1979.

[151] Th. Udem, J. Reichert, R. Holzwarth, and T. W. Hänsch. Accurate measurement of large optical frequency differences with a mode-locked laser. *Opt. Lett.*, 24(13):881, 1999.

[152] J. Reichert, R. Holzwarth, Th Udem, and T. W. Hänsch. Measuring the frequency of light with mode-locked lasers. *Opt. Commun.*, 172(1):59–68, 1999.

[153] A. Apolonski, A. Poppe, G. Tempea, Ch Spielmann, Th Udem, R. Holzwarth, T. W. Hänsch, and F. Krausz. Controlling the phase evolution of few-cycle light pulses. *Phys. Rev. Lett.*, 85(4):740–743, 2000.

[154] R Kienberger, E. Goulielmakis, M. Uiberacker, A. Baltuska, V S Yakovlev, F. Bammer, A. Scrinzi, Th. Westerwalbesloh, U Kleineberg, U. Heinzmann, M. Drescher, and F. Krausz. Atomic transient recorder. *Nature*, 427:817–821, 2004.

[155] D. Anderson and M. Lisak. Nonlinear asymmetric self-phase modulation and self-steepening of pulses in long optical waveguides. *Phys. Rev. A*, 27(3):1393–1398, 1983.

[156] Bruno E. Schmidt, Pierre Béjot, Mathieu Giguère, Andrew D. Shiner, Carlos Trallero-Herrero, Éric Bisson, Kasparian Jérôme, Jean Pierre Wolf, David M. Villeneuve, Jean Claude Kieffer, Paul B. Corkum, and François Légaré. Compression of 1.8 μm laser pulses to sub two optical cycles with bulk material. *Appl. Phys. Lett.*, 96(12):0–3, 2010.

[157] Rick Trebino, Kenneth W. DeLong, David N. Fittinghoff, John N. Sweetser, Marco A. Krumbügel, Bruce A. Richman, and Daniel J. Kane. Measuring ultrashort laser pulses in the time-frequency domain using frequency-resolved optical gating. *Rev. Sci. Instrum.*, 68(9):3277–3295, 1997.

[158] Chris Iaconis, Matthew E. Anderson, and Ian A. Walmsley. Spectral phase interferometry for direct electric field reconstruction of ultrashort optical pulses. *Springer Ser. Chem. Phys.*, 63(10):103–105, 1998.

[159] D. J. Kane and R. Trebino. Single-shot measurement of the intensity and phase of a femtosecond pulse. *Springer Ser. Chem. Phys.*, 18(55):138–139, 1993.

[160] Mitsuo Takeda, Hideki Ina, and Seiji Kobayashi. Fourier-transform method of fringe-pattern analysis for computer-based topography and inteferometry. *J. Opt. Soc. Am.*, 72(1):156–160, 1982.

[161] Toshiaki Kita, Tatsuo Harada, N. Nakano, and H. Kuroda. Mechanically ruled aberration-corrected concave gratings for a flat-field grazing-incidence spectrograph. *Appl. Opt.*, 28(2):418–420, 1983.

[162] Noboru Nakano, Hiroto Kuroda, Toshiaki Kita, and Tatsuo Harada. Development of a flat-field grazing-incidence XUV spectrometer and its application in picosecond XUV spectroscopy. *Appl. Opt.*, 23(14):2386, 1984.

[163] Tatsuo Harada, Kaoru Takahashi, Hideo Sakuma, and Andrzej Osyczka. Optimum design of a grazing-incidence flat-field spectrograph with a spherical varied-line-space grating. *Appl. Opt.*, 38(13):2743, 1999.

[164] Dominic A Ventura and John G Nikelly. Pulse Dampening System for High Pressure Liquid Chromatography. *Anal. Chem.*, 50(7):1017, 1978.

[165] Leybold. Fundametnals of Leybold Vacuum Technology. (199).

[166] E. Lorek, E. W. Larsen, C. M. Heyl, S. Carlström, D. Paleček, D. Zigmantas, and J. Mauritsson. High-order harmonic generation using a high-repetition-rate turnkey laser. *Rev. Sci. Instrum.*, 2014.

[167] R. Zürl and H. Graener. High-harmonic generation of mid-IR pulses in simple liquids. *Appl. Phys. B Lasers Opt.*, 66(2):213–216, 1998.

[168] Felicity McGrath, Allan S. Johnson, Dane R. Austin, Peter Hawkins, David Wood, Lukas Miseikis, Emma R. Simpson, Marta Castillejo, Ricardo Torres, Susan Parker, Thomas Siegel, and Jon P. Marangos. An apparatus for quantitative high-harmonic generation spectroscopy in molecular vapours. *Rev. Sci. Instrum.*, 88(10), 2017.

[169] M. R. Panman, D. J. Shaw, B. Ensing, and S. Woutersen. Local orientational order in liquids revealed by resonant vibrational energy transfer. *Phys. Rev. Lett.*, 113(20):1–5, 2014.

[170] I. H. Malitson. Interspecimen comparison of the refractive index of fused silica. *J. Opt. Soc. Am.*, 55(10):1205, 1965.

[171] Allan Stewart Johnson. *Long and Short Wavelength Optical Sources for Attosecond Science*. PhD thesis, Imperial College London, 2017.

[172] Lin X. Chen and Xiaoyi Zhang. Photochemical processes revealed by X-ray transient absorption spectroscopy. *J. Phys. Chem. Lett.*, 4(22):4000–4013, nov 2013.

[173] R. W. Schoenlein, S. Chattopadhyay, H. H.W. Chong, T. E. Glover, P. A. Heimann, C. V. Shank, A. A. Zholents, and M. S. Zolotorev. Generation of femtosecond pulses of synchrotron radiation. *Science (80-.).*, 287(5461):2237–2240, 2000.

[174] Majed Chergui and Ahmed H. Zewail. Electron and X-ray methods of ultrafast structural dynamics: Advances and applications. *ChemPhysChem*, 10(1):28–43, 2009.

[175] Malte Christian Brahms. *Few-Femtosecond Deep-UV Pulses for Transient-Absorption Experiments*. PhD thesis, Imperial College London, 2018.

[176] Miguel Ochmann, Abid Hussain, Inga Von Ahnen, Amy A. Cordones, Kiryong Hong, Jae Hyuk Lee, Rory Ma, Katrin Adamczyk, Tae Kyu Kim, Robert W. Schoenlein, Oriol Vendrell, and Nils Huse. UV-photochemistry of the disulfide bond: evolution of early photoproducts from picosecond X-ray absorption spectroscopy at the sulfur K-edge. *J. Am. Chem. Soc.*, 140(21):6554–6561, 2018.

[177] Tyler J. Bechtel and Eranthie Weerapana. From structure to redox: The diverse functional roles of disulfides and implications in disease. *Proteomics*, 17(6), 2017.

[178] Maria Teresa Neves-Petersen, Zygmunt Gryczynski, Joseph Lakowicz, Peter Fojan, Shona Pedersen, Evamaria Petersen, and Steffen Bjørn Petersen. High probability of disrupting a disulphide bridge mediated by an endogenous excited tryptophan residue. *Protein Sci.*, 11(3):588–600, 2009.

[179] Manho Lim, S. Gnanakaran, and Robin M. Hochstrasser. Charge shifting in the ultrafast photoreactions of ClO^- in water . *J. Chem. Phys.*, 106(9):3485–3493, mar 1997.

[180] Robert C. Hilborn. Einstein coefficients, cross sections, f values, dipole moments, and all that. *Am. J. Phys.*, 50:982–986, 1982.

[181] Timothy R. Gosnell. *Fundamentals of Spectroscopy and Laser Physics*. Cambridge University Press, 2002.

[182] Farideh Jalilehvand. Sulfur: Not a "silent" element any more. *Chem. Soc. Rev.*, 35(12):1256–1268, 2006.

[183] P. A. Lee, P. H. Citrin, P. Eisenberger, and B. M. Kincaid. Extended X-ray absorption fine structure its strengths and limitations as a structural tool. *Rev. Mod. Phys.*, 53(4):769–806, 1981.

[184] Martin Chaplin. *Water Structure and Science*.

[185] M. Kasrai, M. E. Fleet, G. M. Bancroft, K. H. Tan, and J. M. Chen. X-ray-absorption near-edge structure of alkali halides: The interatomic-distance correlation. *Phys. Rev. B*, 43(2):1763–1772, 1991.

[186] R D Present. *Kinetic theory of gases*. McGraw-Hill, 1st edition, 1958.

[187] Saul Dushman. Production and measurement of high vacuum. *Gen. Electr. Rev.*, 25, 1922.

[188] Robert Hołyst, Marek Litniewski, and Daniel Jakubczyk. A molecular dynamics test of the Hertz-Knudsen equation for evaporating liquids. *Soft Matter*, 11(36):7201–7206, 2015.

[189] David R Lide. *CRC Handbook of Chemistry and Physics*. CRC Press, 85 edition, 2004.

[190] Jianhua Huang. A simple accurate formula for calculating saturation vapor pressure of water and ice. *J. Appl. Meteorol. Climatol.*, 57(6):1265–1272, 2018.

[191] Matthew E. Wise, Trudi A. Semeniuk, Roelof Bruintjes, Scot T. Martin, Lynn M. Russell, and Peter R. Buseck. Hygroscopic behavior of NaCl-bearing natural aerosol particles using environmental transmission electron microscopy. *J. Geophys. Res. Atmos.*, 112(10):1–12, 2007.

[192] N. D. Spencer and R. M. Lambert. Chlorine chemisorption and surface chloride formation on Au(111). *Surf. Sci.*, 107(1):237–248, 1981.

[193] Jon P. Marangos, Sarah Baker, Nathaniel Kajumba, Joseph S. Robinson, John W.G. Tisch, and Ricardo Torres. Dynamic imaging of molecules using high order harmonic generation. *Phys. Chem. Chem. Phys.*, 10(1):35–48, 2008.

[194] Sebastian Jarosch. *High harmonic generation driven by few-cycle infrared fields in gases and liquids*. PhD thesis, Imperial College London, 2019.

[195] Enrique J Galvez. *Gaussian Beams*. Colgate University, 2014.

[196] Nico M. Temme. *Special Functions: An introduction to the classical functions of mathematical physics*. New York: Wiley, 2nd print edition, 1996.

[197] Bo Zhang and Josiane Zerubia. Point-Spread Function Models. *Appl. Opt.*, 46(10):1819–1829, 2007.

[198] Coblentz Society Inc. *Evaluated Infrared Reference Spectra*. 2021-01-13. Webb. Chim. NIST, Base Données Stand. référence NIST numéro 69, Eds. P.J. Linstrom W.G. Mallard, Natl. Inst. Stand. Technol. Gaithersbg. MD, 20899.

[199] C. Cornaggia and Ph Hering. Nonsequential double ionization of small molecules induced by a femtosecond laser field. *Phys. Rev. A - At. Mol. Opt. Phys.*, 62(2):13, 2000.

[200] Piotr Rudawski, Anne Harth, Chen Guo, Eleonora Lorek, Miguel Miranda, Christoph M. Heyl, Esben W. Larsen, Jan Ahrens, Oliver Prochnow, Thomas Binhammer, Uwe Morgner, Johan Mauritsson, Anne L'Huillier, and Cord L. Arnold. Carrier-envelope phase dependent high-order harmonic generation with a high-repetition rate OPCPA-system. *Eur. Phys. J. D*, 69(3), 2015.

[201] Zhuonan Song, Yi Huang, Weiwei L. Xu, Lei Wang, Yu Bao, Shiguang Li, and Miao Yu. Continuously adjustable, molecular-sieving gate on 5A zeolite for distinguishing small organic molecules by size. *Sci. Rep.*, 5(August):1–7, 2015.

[202] M. H.F. Bettega, C. Winstead, V. McKoy, A. Jo, A. Gauf, J. Tanner, L. R. Hargreaves, and M. A. Khakoo. Collisions of low-energy electrons with isopropanol. *Phys. Rev. A - At. Mol. Opt. Phys.*, 84(4):1–5, 2011.

[203] M. A. Khakoo, J. Muse, H. Silva, M. C.A. Lopes, C. Winstead, V. McKoy, E. M. De Oliveira, R. F. Da Costa, M. T.Do N. Varella, M. H.F. Bettega, and M. A.P. Lima. Elastic scattering of slow electrons by n -propanol and n -butanol. *Phys. Rev. A - At. Mol. Opt. Phys.*, 78(6):1–8, 2008.

[204] Heiko G. Kurz, Martin Kretschmar, Thomas Binhammer, Tamas Nagy, Detlev Ristau, Manfred Lein, Uwe Morgner, and Milutin Kovacev. Revealing the microscopic real-space excursion of a laser-driven electron. *Phys. Rev. X*, 6(3):1–12, 2016.

[205] G. J. Boyle, R. P. McEachran, D. G. Cocks, M. J. Brunger, S. J. Buckman, S. Dujko, and R. D. White. Ab initio electron scattering cross-sections and transport in liquid xenon. *J. Phys. D. Appl. Phys.*, 49(35), 2016.

[206] M. Michaud, A. Wen, and L. Sanche. Cross sections for low-energy (1-100 eV) electron elastic and inelastic scattering in amorphous ice. *Radiat. Res.*, 159(1):3–22, 2003.

[207] N. Sanner, O. Ueéza, B. Chimier, M. Sentis, P. Lassonde, F. Légaré, and J. C. Kieffer. Toward determinism in surface damaging of dielectrics using few-cycle laser pulses. *Appl. Phys. Lett.*, 96(7):1–4, 2010.

[208] Beop Min Kim, Aleksey M. Komashko, Alexander M. Rubenchik, Michael D. Feit, Steffen Reidt, Luiz B. Da Silva, and Jürgen Eichler. Interferometric analysis of ultrashort pulse laser-induced pressure waves in water. *J. Appl. Phys.*, 94(1):709–715, 2003.

[209] M. Mero, J. Liu, W. Rudolph, D. Ristau, and K. Starke. Scaling laws of femtosecond laser pulse induced breakdown in oxide films. *Phys. Rev. B - Condens. Matter Mater. Phys.*, 71(11):1–7, 2005.

[210] L. Gallais, D. B. Douti, M. Commandré, G. Batavičite, E. Pupka, M. Ščiuka, L. Smalakys, V. Sirutkaitis, and A. Melninkaitis. Wavelength dependence of femtosecond laser-induced damage threshold of optical materials. *J. Appl. Phys.*, 117(22), 2015.

[211] E. G. Gamaly, A. V. Rode, B. Luther-Davies, and V. T. Tikhonchuk. Ablation of solids by femtosecond lasers: Ablation mechanism and ablation thresholds for metals and dielectrics. *Phys. Plasmas*, 9(3):949, 2002.

Appendix

Republishing permissions

AIP PUBLISHING LICENSE
TERMS AND CONDITIONS

May 07, 2021

This Agreement between Mr. Jonathan Barnard ("You") and AIP Publishing ("AIP Publishing") consists of your license details and the terms and conditions provided by AIP Publishing and Copyright Clearance Center.

License Number	5061990599326
License date	May 04, 2021
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Review of Scientific Instruments
Licensed Content Title	New design for a time-of-flight mass spectrometer with a liquid beam laser desorption ion source for the analysis of biomolecules
Licensed Content Author	A. Charvat, E. Lugovoj, M. Faubel, et al
Licensed Content Date	May 1, 2004
Licensed Content Volume	75
Licensed Content Issue	5
Type of Use	Thesis/Dissertation
Requestor type	Student

Format	Print and electronic
Portion	Figure/Table
Number of figures/tables	1
Title	X-ray generation from andspectroscopy of a thin liquid sheet
Institution name	Imperial College London
Expected presentation date	Jun 2021
Portions	Panel a of Figure 2
Requestor Location	Mr. Jonathan Barnard [redacted]
Total	0.00 GBP

Terms and Conditions

AIP Publishing -- Terms and Conditions: Permissions Uses

AIP Publishing hereby grants to you the non-exclusive right and license to use and/or distribute the Material according to the use specified in your order, on a one-time basis, for the specified term, with a maximum distribution equal to the number that you have ordered. Any links or other content accompanying the Material are not the subject of this license.

1. You agree to include the following copyright and permission notice with the reproduction of the Material:"Reprinted from [FULL CITATION], with the permission of AIP Publishing." For an article, the credit line and permission notice must be printed on the first page of the article or book chapter. For photographs, covers, or tables, the notice may appear with the Material, in a footnote, or in the reference list.
2. If you have licensed reuse of a figure, photograph, cover, or table, it is your responsibility to ensure that the material is original to AIP Publishing and does not

contain the copyright of another entity, and that the copyright notice of the figure, photograph, cover, or table does not indicate that it was reprinted by AIP Publishing, with permission, from another source. Under no circumstances does AIP Publishing purport or intend to grant permission to reuse material to which it does not hold appropriate rights.
You may not alter or modify the Material in any manner. You may translate the Material into another language only if you have licensed translation rights. You may not use the Material for promotional purposes.

3. The foregoing license shall not take effect unless and until AIP Publishing or its agent, Copyright Clearance Center, receives the Payment in accordance with Copyright Clearance Center Billing and Payment Terms and Conditions, which are incorporated herein by reference.
4. AIP Publishing or Copyright Clearance Center may, within two business days of granting this license, revoke the license for any reason whatsoever, with a full refund payable to you. Should you violate the terms of this license at any time, AIP Publishing, or Copyright Clearance Center may revoke the license with no refund to you. Notice of such revocation will be made using the contact information provided by you. Failure to receive such notice will not nullify the revocation.
5. AIP Publishing makes no representations or warranties with respect to the Material. You agree to indemnify and hold harmless AIP Publishing, and their officers, directors, employees or agents from and against any and all claims arising out of your use of the Material other than as specifically authorized herein.
6. The permission granted herein is personal to you and is not transferable or assignable without the prior written permission of AIP Publishing. This license may not be amended except in a writing signed by the party to be charged.
7. If purchase orders, acknowledgments or check endorsements are issued on any forms containing terms and conditions which are inconsistent with these provisions, such inconsistent terms and conditions shall be of no force and effect. This document, including the CCC Billing and Payment Terms and Conditions, shall be the entire agreement between the parties relating to the subject matter hereof.

This Agreement shall be governed by and construed in accordance with the laws of the State of New York. Both parties hereby submit to the jurisdiction of the courts of New York County for purposes of resolving any disputes that may arise hereunder.

V1.2

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.

Copyright Clearance Center Marketplace™

IOP Publishing, Ltd - License Terms and Conditions

This is a License Agreement between Mr. Jonathan Barnard ("You") and IOP Publishing, Ltd ("Publisher") provided by Copyright Clearance Center ("CCC"). The license consists of your order details, the terms and conditions provided by IOP Publishing, Ltd, and the CCC terms and conditions.

All payments must be made in full to CCC.

Order Date	04-May-2021	**Type of Use**	Republish in a thesis/dissertation
Order license ID	1116448-1		
ISSN	0022-3727	**Publisher**	IOP Publishing
		Portion	Image/photo/illustration

LICENSED CONTENT

Publication Title	Journal of Physics D : Applied Physics	**Country**	United Kingdom of Great Britain and Northern Ireland
Author/Editor	Institute of Physics (Great Britain), Institute of Physics and the Physical Society.	**Rightsholder**	IOP Publishing, Ltd
		Publication Type	Journal
Date	01/01/1970		
Language	English		

REQUEST DETAILS

Portion Type	Image/photo/illustration	**Distribution**	U.K. and Commonwealth (excluding Canada)
Number of images / photos / illustrations	1	**Translation**	Original language of publication
Format (select all that apply)	Print, Electronic	**Copies for the disabled?**	No
Who will republish the content?	Academic institution	**Minor editing privileges?**	No
Duration of Use	Life of current edition	**Incidental promotional use?**	No
Lifetime Unit Quantity	Up to 499	**Currency**	GBP
Rights Requested	Main product		

NEW WORK DETAILS

Title	X-ray generation from andspectroscopy of a thin liquid sheet	**Institution name**	Imperial College London
		Expected presentation date	2021-06-30
Instructor name	Professor Jonathan Marangos		

ADDITIONAL DETAILS

Order reference number	N/A	**The requesting person / organization to appear on the license**	Mr. Jonathan Barnard

REUSE CONTENT DETAILS

Title, description or numeric reference of the portion(s)	Figure 2	Title of the article/chapter the portion is from	Gas dynamic virtual nozzle for generation of microscopic droplet streams
Editor of portion(s)	N/A	Author of portion(s)	DePonte, D. P. Weierstall, U. Schmidt, K. Warner, J. Starodub, D. Spence, J. C.H. Doak, R. B.
Volume of serial or monograph	N/A	Issue, if republishing an article from a serial	N/A
Page or page range of portion	5	Publication date of portion	2008-09-18

PUBLISHER TERMS AND CONDITIONS

These special terms and conditions are in addition to the standard terms and conditions for CCC's Republication Service and, together with those standard terms and conditions, govern the use of the Works. As the User you will make all reasonable efforts to contact the author(s) of the article which the Work is to be reused from, to seek consent for your intended use. Contacting one author who is acting expressly as authorised agent for their co-author(s) is acceptable. User will reproduce the following wording prominently alongside the Work: the source of the Work, including author, article title, title of journal, volume number, issue number (if relevant), page range (or first page if this is the only information available) and date of first publication; and a link back to the article (via DOI); and if practicable, and IN ALL CASES for new works published under any of the Creative Commons licences, the words "© IOP Publishing. Reproduced with permission. All rights reserved" Without the express permission of the author(s) and the Rightsholder of the article from which the Work is to be reused, User shall not use it in any way which, in the opinion of the Rightsholder, could: (i) distort or alter the author(s)' original intention(s) and meaning; (ii) be prejudicial to the honour or reputation of the author(s); and/or (iii) imply endorsement by the author(s) and/or the Rightsholder. This licence does not apply to any article which is credited to another source and which does not have the copyright line '© IOP Publishing Ltd'. User must check the copyright line of the article from which the Work is to be reused to check that IOP Publishing Ltd has all the necessary rights to be able to grant permission. User is solely responsible for identifying and obtaining separate licences and permissions from the copyright owner for reuse of any such third party material/figures which the Rightsholder is not the copyright owner of. The Rightsholder shall not reimburse any fees which User pays for a republication license for such third party content. This licence does not apply to any material/figure which is credited to another source in the Rightsholder's publication or has been obtained from a third party. User must check the Version of Record of the article from which the Work is to be reused, to check whether any of the material in the Work is third party material. Third party citations and/or copyright notices and/or permissions statements may not be included in any other version of the article from which the Work is to be reused and so cannot be relied upon by the User. User is solely responsible for identifying and obtaining separate licences and permissions from the copyright owner for reuse of any such third party material/figures where the Rightsholder is not the copyright owner. The Rightsholder shall not reimburse any fees which User pays for a republication license for such third party content. User and CCC acknowledge that the Rightsholder may, from time to time, make changes or additions to these special terms and conditions without express notification, provided that these shall not apply to permissions already secured and paid for by User prior to such change or addition. User acknowledges that the Rightsholder (which includes companies within its group and third parties for whom it publishes its titles) may make use of personal data collected through the service in the course of their business. If User is the author of the Work, User may automatically have the right to reuse it under the rights granted back when User transferred the copyright in the article to the Rightsholder. User should check the copyright form and the relevant author rights policy to check whether permission is required. If User is the author of the Work and does require permission for proposed reuse of the Work, User should select 'Author of requested content' as the Requestor Type. The Rightsholder shall not reimburse any fees which User pays for a republication license. If User is the author of the article which User wishes to reuse in User's thesis or dissertation, the republication licence covers the right to include the Version of Record of the article, provided it is not then shared or deposited online. User must include citation details. Where User wishes to share their thesis or dissertation online, they should remove the Version of Record before uploading it. User may include a Preprint or the Accepted Manuscript (after the embargo period) in the online version of the thesis or dissertation, provided they do so in accordance with the Rightsholder's policies on sharing Preprints or Accepted Manuscripts. User may need to obtain separate permission for any third party content included within the article. User must check this with the copyright owner of such third party content. Any online or commercial use of User's thesis or dissertation containing the article, including publication via ProQuest, would need to be expressly notified in writing to the Rightsholder at the time of request and would require separate written permission from the Rightsholder. As well as CCC, the Rightsholder shall have the right to bring any legal action that it deems necessary to enforce its rights should it consider that the Work infringes those rights in any way. For content reuse requests that qualify for permission under the STM Permissions Guidelines, which may be updated from time to time, the STM Permissions Guidelines supplement the terms and conditions contained in this license.

CCC Republication Terms and Conditions

1. Description of Service; Defined Terms. This Republication License enables the User to obtain licenses for republication of one or more copyrighted works as described in detail on the relevant Order Confirmation (the "Work(s)"). Copyright Clearance Center, Inc. ("CCC") grants licenses through the Service on behalf of the rightsholder identified on the Order Confirmation (the "Rightsholder"). "Republication", as used herein, generally means the inclusion of a Work, in whole or in part, in a new work or works, also as described on the Order Confirmation. "User", as used herein, means the person or entity making such republication.

2. The terms set forth in the relevant Order Confirmation, and any terms set by the Rightsholder with respect to a particular Work, govern the terms of use of Works in connection with the Service. By using the Service, the person transacting for a republication license on behalf of the User represents and warrants that he/she/it (a) has been duly authorized by the User to accept, and hereby does accept, all such terms and conditions on behalf of User, and (b) shall inform User of all such terms and conditions. In the event such person is a "freelancer" or other third party independent of User and CCC, such party shall be deemed jointly a "User" for purposes of these terms and conditions. In any event, User shall be deemed to have accepted and agreed to all such terms and conditions if User republishes the Work in any fashion.

3. Scope of License; Limitations and Obligations.

    3.1. All Works and all rights therein, including copyright rights, remain the sole and exclusive property of the Rightsholder. The license created by the exchange of an Order Confirmation (and/or any invoice) and payment by User of the full amount set forth on that document includes only those rights expressly set forth in the Order Confirmation and in these terms and conditions, and conveys no other rights in the Work(s) to User. All rights not expressly granted are hereby reserved.

    3.2. General Payment Terms: You may pay by credit card or through an account with us payable at the end of the month. If you and we agree that you may establish a standing account with CCC, then the following terms apply: Remit Payment to: Copyright Clearance Center, 29118 Network Place, Chicago, IL 60673-1291. Payments Due: Invoices are payable upon their delivery to you (or upon our notice to you that they are available to you for downloading). After 30 days, outstanding amounts will be subject to a service charge of 1-1/2% per month or, if less, the maximum rate allowed by applicable law. Unless otherwise specifically set forth in the Order Confirmation or in a separate written agreement signed by CCC, invoices are due and payable on "net 30" terms. While User may exercise the rights licensed immediately upon issuance of the Order Confirmation, the license is automatically revoked and is null and void, as if it had never been issued, if complete payment for the license is not received on a timely basis either from User directly or through a payment agent, such as a credit card company.

    3.3. Unless otherwise provided in the Order Confirmation, any grant of rights to User (i) is "one-time" (including the editions and product family specified in the license), (ii) is non-exclusive and non-transferable and (iii) is subject to any and all limitations and restrictions (such as, but not limited to, limitations on duration of use or circulation) included in the Order Confirmation or invoice and/or in these terms and conditions. Upon completion of the licensed use, User shall either secure a new permission for further use of the Work(s) or immediately cease any new use of the Work(s) and shall render inaccessible (such as by deleting or by removing or severing links or other locators) any further copies of the Work (except for copies printed on paper in accordance with this license and still in User's stock at the end of such period).

    3.4. In the event that the material for which a republication license is sought includes third party materials (such as photographs, illustrations, graphs, inserts and similar materials) which are identified in such material as having been used by permission, User is responsible for identifying, and seeking separate licenses (under this Service or otherwise) for, any of such third party materials; without a separate license, such third party materials may not be used.

    3.5. Use of proper copyright notice for a Work is required as a condition of any license granted under the Service. Unless otherwise provided in the Order Confirmation, a proper copyright notice will read substantially as follows: "Republished with permission of [Rightsholder's name], from [Work's title, author, volume, edition number and year of copyright]; permission conveyed through Copyright Clearance Center, Inc. " Such notice must be provided in a reasonably legible font size and must be placed either immediately adjacent to the Work as used (for example, as part of a by-line or footnote but not as a separate electronic link) or in the place where substantially all other credits or notices for the new work containing the republished Work are located. Failure to include the required notice results in loss to the Rightsholder and CCC, and the User shall be liable to pay liquidated damages for each such failure equal to twice the use fee specified in the Order Confirmation, in addition to the use fee itself and any other fees and charges specified.

3.6. User may only make alterations to the Work if and as expressly set forth in the Order Confirmation. No Work may be used in any way that is defamatory, violates the rights of third parties (including such third parties' rights of copyright, privacy, publicity, or other tangible or intangible property), or is otherwise illegal, sexually explicit or obscene. In addition, User may not conjoin a Work with any other material that may result in damage to the reputation of the Rightsholder. User agrees to inform CCC if it becomes aware of any infringement of any rights in a Work and to cooperate with any reasonable request of CCC or the Rightsholder in connection therewith.

4. Indemnity. User hereby indemnifies and agrees to defend the Rightsholder and CCC, and their respective employees and directors, against all claims, liability, damages, costs and expenses, including legal fees and expenses, arising out of any use of a Work beyond the scope of the rights granted herein, or any use of a Work which has been altered in any unauthorized way by User, including claims of defamation or infringement of rights of copyright, publicity, privacy or other tangible or intangible property.

5. Limitation of Liability. UNDER NO CIRCUMSTANCES WILL CCC OR THE RIGHTSHOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL OR INCIDENTAL DAMAGES (INCLUDING WITHOUT LIMITATION DAMAGES FOR LOSS OF BUSINESS PROFITS OR INFORMATION, OR FOR BUSINESS INTERRUPTION) ARISING OUT OF THE USE OR INABILITY TO USE A WORK, EVEN IF ONE OF THEM HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. In any event, the total liability of the Rightsholder and CCC (including their respective employees and directors) shall not exceed the total amount actually paid by User for this license. User assumes full liability for the actions and omissions of its principals, employees, agents, affiliates, successors and assigns.

6. Limited Warranties. THE WORK(S) AND RIGHT(S) ARE PROVIDED "AS IS". CCC HAS THE RIGHT TO GRANT TO USER THE RIGHTS GRANTED IN THE ORDER CONFIRMATION DOCUMENT. CCC AND THE RIGHTSHOLDER DISCLAIM ALL OTHER WARRANTIES RELATING TO THE WORK(S) AND RIGHT(S), EITHER EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. ADDITIONAL RIGHTS MAY BE REQUIRED TO USE ILLUSTRATIONS, GRAPHS, PHOTOGRAPHS, ABSTRACTS, INSERTS OR OTHER PORTIONS OF THE WORK (AS OPPOSED TO THE ENTIRE WORK) IN A MANNER CONTEMPLATED BY USER; USER UNDERSTANDS AND AGREES THAT NEITHER CCC NOR THE RIGHTSHOLDER MAY HAVE SUCH ADDITIONAL RIGHTS TO GRANT.

7. Effect of Breach. Any failure by User to pay any amount when due, or any use by User of a Work beyond the scope of the license set forth in the Order Confirmation and/or these terms and conditions, shall be a material breach of the license created by the Order Confirmation and these terms and conditions. Any breach not cured within 30 days of written notice thereof shall result in immediate termination of such license without further notice. Any unauthorized (but licensable) use of a Work that is terminated immediately upon notice thereof may be liquidated by payment of the Rightsholder's ordinary license price therefor; any unauthorized (and unlicensable) use that is not terminated immediately for any reason (including, for example, because materials containing the Work cannot reasonably be recalled) will be subject to all remedies available at law or in equity, but in no event to a payment of less than three times the Rightsholder's ordinary license price for the most closely analogous licensable use plus Rightsholder's and/or CCC's costs and expenses incurred in collecting such payment.

8. Miscellaneous.

8.1. User acknowledges that CCC may, from time to time, make changes or additions to the Service or to these terms and conditions, and CCC reserves the right to send notice to the User by electronic mail or otherwise for the purposes of notifying User of such changes or additions; provided that any such changes or additions shall not apply to permissions already secured and paid for.

8.2. Use of User-related information collected through the Service is governed by CCC's privacy policy, available online here:https://marketplace.copyright.com/rs-ui-web/mp/privacy-policy

8.3. The licensing transaction described in the Order Confirmation is personal to User. Therefore, User may not assign or transfer to any other person (whether a natural person or an organization of any kind) the license created by the Order Confirmation and these terms and conditions or any rights granted hereunder; provided, however, that User may assign such license in its entirety on written notice to CCC in the event of a transfer of all or substantially all of User's rights in the new material which includes the Work(s) licensed under this Service.

8.4. No amendment or waiver of any terms is binding unless set forth in writing and signed by the parties. The Rightsholder and CCC hereby object to any terms contained in any writing prepared by the User or its

principals, employees, agents or affiliates and purporting to govern or otherwise relate to the licensing transaction described in the Order Confirmation, which terms are in any way inconsistent with any terms set forth in the Order Confirmation and/or in these terms and conditions or CCC's standard operating procedures, whether such writing is prepared prior to, simultaneously with or subsequent to the Order Confirmation, and whether such writing appears on a copy of the Order Confirmation or in a separate instrument.

8.5. The licensing transaction described in the Order Confirmation document shall be governed by and construed under the law of the State of New York, USA, without regard to the principles thereof of conflicts of law. Any case, controversy, suit, action, or proceeding arising out of, in connection with, or related to such licensing transaction shall be brought, at CCC's sole discretion, in any federal or state court located in the County of New York, State of New York, USA, or in any federal or state court whose geographical jurisdiction covers the location of the Rightsholder set forth in the Order Confirmation. The parties expressly submit to the personal jurisdiction and venue of each such federal or state court.If you have any comments or questions about the Service or Copyright Clearance Center, please contact us at 978-750-8400 or send an e-mail to support@copyright.com.

v 1.1

ELSEVIER LICENSE
TERMS AND CONDITIONS

May 07, 2021

This Agreement between Mr. Jonathan Barnard ("You") and Elsevier ("Elsevier") consists of your license details and the terms and conditions provided by Elsevier and Copyright Clearance Center.

License Number	5062080486313
License date	May 04, 2021
Licensed Content Publisher	Elsevier
Licensed Content Publication	Chemical Engineering Science
Licensed Content Title	Some aspects of liquid flow through fan spray nozzles
Licensed Content Author	N. Dombrowski,D. Hasson,D.E. Ward
Licensed Content Date	Mar 1, 1960
Licensed Content Volume	12
Licensed Content Issue	1
Licensed Content Pages	16
Start Page	35
End Page	50
Type of Use	reuse in a thesis/dissertation

Portion	figures/tables/illustrations
Number of figures/tables/illustrations	1
Format	both print and electronic
Are you the author of this Elsevier article?	No
Will you be translating?	No
Title	X-ray generation from andspectroscopy of a thin liquid sheet
Institution name	Imperial College London
Expected presentation date	Jun 2021
Portions	Panel a from Figure 1
Requestor Location	Mr. Jonathan Barnard [redacted] Attn: Mr. Jonathan Barnard
Publisher Tax ID	GB 494 6272 12
Total	0.00 GBP

Terms and Conditions

INTRODUCTION

1. The publisher for this copyrighted material is Elsevier. By clicking "accept" in connection with completing this licensing transaction, you agree that the following terms and conditions apply to this transaction (along with the Billing and Payment terms and conditions established by Copyright Clearance Center, Inc. ("CCC"), at the time that you opened your Rightslink account and that are available at any time at

http://myaccount.copyright.com).

GENERAL TERMS

2. Elsevier hereby grants you permission to reproduce the aforementioned material subject to the terms and conditions indicated.

3. Acknowledgement: If any part of the material to be used (for example, figures) has appeared in our publication with credit or acknowledgement to another source, permission must also be sought from that source. If such permission is not obtained then that material may not be included in your publication/copies. Suitable acknowledgement to the source must be made, either as a footnote or in a reference list at the end of your publication, as follows:

"Reprinted from Publication title, Vol /edition number, Author(s), Title of article / title of chapter, Pages No., Copyright (Year), with permission from Elsevier [OR APPLICABLE SOCIETY COPYRIGHT OWNER]." Also Lancet special credit - "Reprinted from The Lancet, Vol. number, Author(s), Title of article, Pages No., Copyright (Year), with permission from Elsevier."

4. Reproduction of this material is confined to the purpose and/or media for which permission is hereby given.

5. Altering/Modifying Material: Not Permitted. However figures and illustrations may be altered/adapted minimally to serve your work. Any other abbreviations, additions, deletions and/or any other alterations shall be made only with prior written authorization of Elsevier Ltd. (Please contact Elsevier's permissions helpdesk here). No modifications can be made to any Lancet figures/tables and they must be reproduced in full.

6. If the permission fee for the requested use of our material is waived in this instance, please be advised that your future requests for Elsevier materials may attract a fee.

7. Reservation of Rights: Publisher reserves all rights not specifically granted in the combination of (i) the license details provided by you and accepted in the course of this licensing transaction, (ii) these terms and conditions and (iii) CCC's Billing and Payment terms and conditions.

8. License Contingent Upon Payment: While you may exercise the rights licensed immediately upon issuance of the license at the end of the licensing process for the transaction, provided that you have disclosed complete and accurate details of your proposed use, no license is finally effective unless and until full payment is received from you (either by publisher or by CCC) as provided in CCC's Billing and Payment terms and conditions. If full payment is not received on a timely basis, then any license preliminarily granted shall be deemed automatically revoked and shall be void as if never granted. Further, in the event that you breach any of these terms and conditions or any of CCC's Billing and Payment terms and conditions, the license is automatically revoked and shall be void as if never granted. Use of materials as described in a revoked license, as well as any use of the materials beyond the scope of an unrevoked license, may constitute copyright infringement and publisher reserves the right to take any and all action to protect its copyright in the materials.

9. Warranties: Publisher makes no representations or warranties with respect to the licensed material.

10. Indemnity: You hereby indemnify and agree to hold harmless publisher and CCC, and their respective officers, directors, employees and agents, from and against any and all claims arising out of your use of the licensed material other than as specifically authorized pursuant to this license.

11. No Transfer of License: This license is personal to you and may not be sublicensed, assigned, or transferred by you to any other person without publisher's written permission.

12. No Amendment Except in Writing: This license may not be amended except in a writing signed by both parties (or, in the case of publisher, by CCC on publisher's behalf).

13. Objection to Contrary Terms: Publisher hereby objects to any terms contained in any purchase order, acknowledgment, check endorsement or other writing prepared by you, which terms are inconsistent with these terms and conditions or CCC's Billing and Payment terms and conditions. These terms and conditions, together with CCC's Billing and Payment terms and conditions (which are incorporated herein), comprise the entire agreement between you and publisher (and CCC) concerning this licensing transaction. In the event of any conflict between your obligations established by these terms and conditions and those established by CCC's Billing and Payment terms and conditions, these terms and conditions shall control.

14. Revocation: Elsevier or Copyright Clearance Center may deny the permissions described in this License at their sole discretion, for any reason or no reason, with a full refund payable to you. Notice of such denial will be made using the contact information provided by you. Failure to receive such notice will not alter or invalidate the denial. In no event will Elsevier or Copyright Clearance Center be responsible or liable for any costs, expenses or damage incurred by you as a result of a denial of your permission request, other than a refund of the amount(s) paid by you to Elsevier and/or Copyright Clearance Center for denied permissions.

LIMITED LICENSE

The following terms and conditions apply only to specific license types:

15. **Translation**: This permission is granted for non-exclusive world **English** rights only unless your license was granted for translation rights. If you licensed translation rights you may only translate this content into the languages you requested. A professional translator must perform all translations and reproduce the content word for word preserving the integrity of the article.

16. **Posting licensed content on any Website**: The following terms and conditions apply as follows: Licensing material from an Elsevier journal: All content posted to the web site must maintain the copyright information line on the bottom of each image; A hyper-text must be included to the Homepage of the journal from which you are licensing at http://www.sciencedirect.com/science/journal/xxxxx or the Elsevier homepage for books at http://www.elsevier.com; Central Storage: This license does not include permission for a scanned version of the material to be stored in a central repository such as that provided by Heron/XanEdu.

Licensing material from an Elsevier book: A hyper-text link must be included to the Elsevier homepage at http://www.elsevier.com . All content posted to the web site must maintain the copyright information line on the bottom of each image.

Posting licensed content on Electronic reserve: In addition to the above the following clauses are applicable: The web site must be password-protected and made available only to bona fide students registered on a relevant course. This permission is granted for 1 year only. You may obtain a new license for future website posting.

17. **For journal authors:** the following clauses are applicable in addition to the above:

Preprints:

A preprint is an author's own write-up of research results and analysis, it has not been peer-reviewed, nor has it had any other value added to it by a publisher (such as formatting, copyright, technical enhancement etc.).

Authors can share their preprints anywhere at any time. Preprints should not be added to or enhanced in any way in order to appear more like, or to substitute for, the final versions of articles however authors can update their preprints on arXiv or RePEc with their Accepted Author Manuscript (see below).

If accepted for publication, we encourage authors to link from the preprint to their formal publication via its DOI. Millions of researchers have access to the formal publications on ScienceDirect, and so links will help users to find, access, cite and use the best available version. Please note that Cell Press, The Lancet and some society-owned have different preprint policies. Information on these policies is available on the journal homepage.

Accepted Author Manuscripts: An accepted author manuscript is the manuscript of an article that has been accepted for publication and which typically includes author-incorporated changes suggested during submission, peer review and editor-author communications.

Authors can share their accepted author manuscript:

- immediately
    - via their non-commercial person homepage or blog
    - by updating a preprint in arXiv or RePEc with the accepted manuscript
    - via their research institute or institutional repository for internal institutional uses or as part of an invitation-only research collaboration work-group
    - directly by providing copies to their students or to research collaborators for their personal use
    - for private scholarly sharing as part of an invitation-only work group on commercial sites with which Elsevier has an agreement
- After the embargo period
    - via non-commercial hosting platforms such as their institutional repository
    - via commercial sites with which Elsevier has an agreement

In all cases accepted manuscripts should:

- link to the formal publication via its DOI
- bear a CC-BY-NC-ND license - this is easy to do
- if aggregated with other manuscripts, for example in a repository or other site, be shared in alignment with our hosting policy not be added to or enhanced in any way to appear more like, or to substitute for, the published journal article.

Published journal article (JPA): A published journal article (PJA) is the definitive final record

of published research that appears or will appear in the journal and embodies all value-adding publishing activities including peer review co-ordination, copy-editing, formatting, (if relevant) pagination and online enrichment.

Policies for sharing publishing journal articles differ for subscription and gold open access articles:

Subscription Articles: If you are an author, please share a link to your article rather than the full-text. Millions of researchers have access to the formal publications on ScienceDirect, and so links will help your users to find, access, cite, and use the best available version.

Theses and dissertations which contain embedded PJAs as part of the formal submission can be posted publicly by the awarding institution with DOI links back to the formal publications on ScienceDirect.

If you are affiliated with a library that subscribes to ScienceDirect you have additional private sharing rights for others' research accessed under that agreement. This includes use for classroom teaching and internal training at the institution (including use in course packs and courseware programs), and inclusion of the article for grant funding purposes.

Gold Open Access Articles: May be shared according to the author-selected end-user license and should contain a CrossMark logo, the end user license, and a DOI link to the formal publication on ScienceDirect.

Please refer to Elsevier's posting policy for further information.

18. **For book authors** the following clauses are applicable in addition to the above: Authors are permitted to place a brief summary of their work online only. You are not allowed to download and post the published electronic version of your chapter, nor may you scan the printed edition to create an electronic version. **Posting to a repository:** Authors are permitted to post a summary of their chapter only in their institution's repository.

19. **Thesis/ Dissertation**: If your license is for use in a thesis/dissertation your thesis may be submitted to your institution in either print or electronic form. Should your thesis be published commercially, please reapply for permission. These requirements include permission for the Library and Archives of Canada to supply single copies, on demand, of the complete thesis and include permission for Proquest/UMI to supply single copies, on demand, of the complete thesis. Should your thesis be published commercially, please reapply for permission. Theses and dissertations which contain embedded PJAs as part of the formal submission can be posted publicly by the awarding institution with DOI links back to the formal publications on ScienceDirect.

Elsevier Open Access Terms and Conditions

You can publish open access with Elsevier in hundreds of open access journals or in nearly 2000 established subscription journals that support open access publishing. Permitted third party re-use of these open access articles is defined by the author's choice of Creative Commons user license. See our open access license policy for more information.

Terms & Conditions applicable to all Open Access articles published with Elsevier:

Any reuse of the article must not represent the author as endorsing the adaptation of the

article nor should the article be modified in such a way as to damage the author's honour or reputation. If any changes have been made, such changes must be clearly indicated.

The author(s) must be appropriately credited and we ask that you include the end user license and a DOI link to the formal publication on ScienceDirect.

If any part of the material to be used (for example, figures) has appeared in our publication with credit or acknowledgement to another source it is the responsibility of the user to ensure their reuse complies with the terms and conditions determined by the rights holder.

Additional Terms & Conditions applicable to each Creative Commons user license:

CC BY: The CC-BY license allows users to copy, to create extracts, abstracts and new works from the Article, to alter and revise the Article and to make commercial use of the Article (including reuse and/or resale of the Article by commercial entities), provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. The full details of the license are available at http://creativecommons.org/licenses/by/4.0.

CC BY NC SA: The CC BY-NC-SA license allows users to copy, to create extracts, abstracts and new works from the Article, to alter and revise the Article, provided this is not done for commercial purposes, and that the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, indicates if changes were made and the licensor is not represented as endorsing the use made of the work. Further, any new works must be made available on the same conditions. The full details of the license are available at http://creativecommons.org/licenses/by-nc-sa/4.0.

CC BY NC ND: The CC BY-NC-ND license allows users to copy and distribute the Article, provided this is not done for commercial purposes and further does not permit distribution of the Article if it is changed or edited in any way, and provided the user gives appropriate credit (with a link to the formal publication through the relevant DOI), provides a link to the license, and that the licensor is not represented as endorsing the use made of the work. The full details of the license are available at http://creativecommons.org/licenses/by-nc-nd/4.0. Any commercial reuse of Open Access articles published with a CC BY NC SA or CC BY NC ND license requires permission from Elsevier and will be subject to a fee.

Commercial reuse includes:

- Associating advertising with the full text of the Article
- Charging fees for document delivery or access
- Article aggregation
- Systematic distribution via e-mail lists or share buttons

Posting or linking by commercial companies for use by customers of those companies.

20. **Other Conditions**:

v1.10

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.

AIP PUBLISHING LICENSE
TERMS AND CONDITIONS

May 07, 2021

This Agreement between Mr. Jonathan Barnard ("You") and AIP Publishing ("AIP Publishing") consists of your license details and the terms and conditions provided by AIP Publishing and Copyright Clearance Center.

License Number	5062080821526
License date	May 04, 2021
Licensed Content Publisher	AIP Publishing
Licensed Content Publication	Review of Scientific Instruments
Licensed Content Title	Micrometer-thickness liquid sheet jets flowing in vacuum
Licensed Content Author	Gediminas Galinis, Jergus Strucka, Jonathan C. T. Barnard, et al
Licensed Content Date	Aug 1, 2017
Licensed Content Volume	88
Licensed Content Issue	8
Type of Use	Thesis/Dissertation
Requestor type	Author (original article)
Format	Print and electronic
Portion	Figure/Table

Number of figures/tables	3
Title	X-ray generation from andspectroscopy of a thin liquid sheet
Institution name	Imperial College London
Expected presentation date	Jun 2021
Portions	Figures 1, 6 and 7
Requestor Location	Mr. Jonathan Barnard [redacted] Attn: Mr. Jonathan Barnard
Total	0.00 GBP

Terms and Conditions

AIP Publishing -- Terms and Conditions: Permissions Uses

AIP Publishing hereby grants to you the non-exclusive right and license to use and/or distribute the Material according to the use specified in your order, on a one-time basis, for the specified term, with a maximum distribution equal to the number that you have ordered. Any links or other content accompanying the Material are not the subject of this license.

1. You agree to include the following copyright and permission notice with the reproduction of the Material:"Reprinted from [FULL CITATION], with the permission of AIP Publishing." For an article, the credit line and permission notice must be printed on the first page of the article or book chapter. For photographs, covers, or tables, the notice may appear with the Material, in a footnote, or in the reference list.
2. If you have licensed reuse of a figure, photograph, cover, or table, it is your responsibility to ensure that the material is original to AIP Publishing and does not contain the copyright of another entity, and that the copyright notice of the figure, photograph, cover, or table does not indicate that it was reprinted by AIP Publishing, with permission, from another source. Under no circumstances does AIP Publishing purport or intend to grant permission to reuse material to which it does not hold appropriate rights.
   You may not alter or modify the Material in any manner. You may translate the Material into another language only if you have licensed translation rights. You may

not use the Material for promotional purposes.

3. The foregoing license shall not take effect unless and until AIP Publishing or its agent, Copyright Clearance Center, receives the Payment in accordance with Copyright Clearance Center Billing and Payment Terms and Conditions, which are incorporated herein by reference.
4. AIP Publishing or Copyright Clearance Center may, within two business days of granting this license, revoke the license for any reason whatsoever, with a full refund payable to you. Should you violate the terms of this license at any time, AIP Publishing, or Copyright Clearance Center may revoke the license with no refund to you. Notice of such revocation will be made using the contact information provided by you. Failure to receive such notice will not nullify the revocation.
5. AIP Publishing makes no representations or warranties with respect to the Material. You agree to indemnify and hold harmless AIP Publishing, and their officers, directors, employees or agents from and against any and all claims arising out of your use of the Material other than as specifically authorized herein.
6. The permission granted herein is personal to you and is not transferable or assignable without the prior written permission of AIP Publishing. This license may not be amended except in a writing signed by the party to be charged.
7. If purchase orders, acknowledgments or check endorsements are issued on any forms containing terms and conditions which are inconsistent with these provisions, such inconsistent terms and conditions shall be of no force and effect. This document, including the CCC Billing and Payment Terms and Conditions, shall be the entire agreement between the parties relating to the subject matter hereof.

This Agreement shall be governed by and construed in accordance with the laws of the State of New York. Both parties hereby submit to the jurisdiction of the courts of New York County for purposes of resolving any disputes that may arise hereunder.

V1.2

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.

SPRINGER NATURE LICENSE
TERMS AND CONDITIONS

May 07, 2021

This Agreement between Mr. Jonathan Barnard ("You") and Springer Nature ("Springer Nature") consists of your license details and the terms and conditions provided by Springer Nature and Copyright Clearance Center.

License Number	5062081075386
License date	May 04, 2021
Licensed Content Publisher	Springer Nature
Licensed Content Publication	Springer eBook
Licensed Content Title	Methods for the Generation of Ultrashort Laser Pulses: Mode-Locking
Licensed Content Author	A. Ducasse, C. Rullière, B. Couillaud
Licensed Content Date	Jan 1, 2005
Type of Use	Thesis/Dissertation
Requestor type	academic/university or research institute
Format	print and electronic
Portion	figures/tables/illustrations
Number of figures/tables /illustrations	1

Will you be translating?	no
Circulation/distribution	1 - 29
Author of this Springer Nature content	no
Title	X-ray generation from andspectroscopy of a thin liquid sheet
Institution name	Imperial College London
Expected presentation date	Jun 2021
Portions	Figure describing Kerr lens modelocking in Chapter 3
Requestor Location	Mr. Jonathan Barnard [redacted] Attn: Mr. Jonathan Barnard
Total	0.00 GBP

Terms and Conditions

Springer Nature Customer Service Centre GmbH
Terms and Conditions

This agreement sets out the terms and conditions of the licence (the **Licence**) between you and **Springer Nature Customer Service Centre GmbH** (the **Licensor**). By clicking 'accept' and completing the transaction for the material (**Licensed Material**), you also confirm your acceptance of these terms and conditions.

1. Grant of License

1. 1. The Licensor grants you a personal, non-exclusive, non-transferable, world-wide licence to reproduce the Licensed Material for the purpose specified in your order only. Licences are granted for the specific use requested in the order and for no other use, subject to the conditions below.

1. 2. The Licensor warrants that it has, to the best of its knowledge, the rights to license reuse of the Licensed Material. However, you should ensure that the material you are requesting is original to the Licensor and does not carry the copyright of another entity (as credited in the published version).

1. 3. If the credit line on any part of the material you have requested indicates that it was reprinted or adapted with permission from another source, then you should also seek permission from that source to reuse the material.

2. Scope of Licence

2. 1. You may only use the Licensed Content in the manner and to the extent permitted by these Ts&Cs and any applicable laws.

2. 2. A separate licence may be required for any additional use of the Licensed Material, e.g. where a licence has been purchased for print only use, separate permission must be obtained for electronic re-use. Similarly, a licence is only valid in the language selected and does not apply for editions in other languages unless additional translation rights have been granted separately in the licence. Any content owned by third parties are expressly excluded from the licence.

2. 3. Similarly, rights for additional components such as custom editions and derivatives require additional permission and may be subject to an additional fee. Please apply to Journalpermissions@springernature.com/bookpermissions@springernature.com for these rights.

2. 4. Where permission has been granted **free of charge** for material in print, permission may also be granted for any electronic version of that work, provided that the material is incidental to your work as a whole and that the electronic version is essentially equivalent to, or substitutes for, the print version.

2. 5. An alternative scope of licence may apply to signatories of the STM Permissions Guidelines, as amended from time to time.

3. Duration of Licence

3. 1. A licence for is valid from the date of purchase ('Licence Date') at the end of the relevant period in the below table:

Scope of Licence	Duration of Licence
Post on a website	12 months
Presentations	12 months
Books and journals	Lifetime of the edition in the language purchased

4. Acknowledgement

4. 1. The Licensor's permission must be acknowledged next to the Licenced Material in print. In electronic form, this acknowledgement must be visible at the same time as the

figures/tables/illustrations or abstract, and must be hyperlinked to the journal/book's homepage. Our required acknowledgement format is in the Appendix below.

5. Restrictions on use

5. 1. Use of the Licensed Material may be permitted for incidental promotional use and minor editing privileges e.g. minor adaptations of single figures, changes of format, colour and/or style where the adaptation is credited as set out in Appendix 1 below. Any other changes including but not limited to, cropping, adapting, omitting material that affect the meaning, intention or moral rights of the author are strictly prohibited.

5. 2. You must not use any Licensed Material as part of any design or trademark.

5. 3. Licensed Material may be used in Open Access Publications (OAP) before publication by Springer Nature, but any Licensed Material must be removed from OAP sites prior to final publication.

6. Ownership of Rights

6. 1. Licensed Material remains the property of either Licensor or the relevant third party and any rights not explicitly granted herein are expressly reserved.

7. Warranty

IN NO EVENT SHALL LICENSOR BE LIABLE TO YOU OR ANY OTHER PARTY OR ANY OTHER PERSON OR FOR ANY SPECIAL, CONSEQUENTIAL, INCIDENTAL OR INDIRECT DAMAGES, HOWEVER CAUSED, ARISING OUT OF OR IN CONNECTION WITH THE DOWNLOADING, VIEWING OR USE OF THE MATERIALS REGARDLESS OF THE FORM OF ACTION, WHETHER FOR BREACH OF CONTRACT, BREACH OF WARRANTY, TORT, NEGLIGENCE, INFRINGEMENT OR OTHERWISE (INCLUDING, WITHOUT LIMITATION, DAMAGES BASED ON LOSS OF PROFITS, DATA, FILES, USE, BUSINESS OPPORTUNITY OR CLAIMS OF THIRD PARTIES), AND WHETHER OR NOT THE PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION SHALL APPLY NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY PROVIDED HEREIN.

8. Limitations

8. 1. ***BOOKS ONLY:*** Where **'reuse in a dissertation/ thesis'** has been selected the following terms apply: Print rights of the final author's accepted manuscript (for clarity, NOT the published version) for up to 100 copies, electronic rights for use only on a personal website or institutional repository as defined by the Sherpa guideline (www.sherpa.ac.uk/romeo/).

8. 2. For content reuse requests that qualify for permission under the STM Permissions Guidelines, which may be updated from time to time, the STM Permissions Guidelines supersede the terms and conditions contained in this licence.

9. Termination and Cancellation

9. 1. Licences will expire after the period shown in Clause 3 (above).

9. 2. Licensee reserves the right to terminate the Licence in the event that payment is not received in full or if there has been a breach of this agreement by you.

Appendix 1 — Acknowledgements:

For Journal Content:
Reprinted by permission from [**the Licensor**]: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication)

For **Advance Online Publication papers:**
Reprinted by permission from [**the Licensor**]: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication), advance online publication, day month year (doi: 10.1038/sj.[JOURNAL ACRONYM].)

For Adaptations/Translations:
Adapted/Translated by permission from [**the Licensor**]: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication)

Note: For any republication from the British Journal of Cancer, the following credit line style applies:

Reprinted/adapted/translated by permission from [**the Licensor**]: on behalf of Cancer Research UK: : [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT]** (year of publication)

For **Advance Online Publication** papers:
Reprinted by permission from The [**the Licensor**]: on behalf of Cancer Research UK: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication), advance online publication, day month year (doi: 10.1038/sj.[JOURNAL ACRONYM])

For Book content:
Reprinted/adapted by permission from [**the Licensor**]: [**Book Publisher** (e.g. Palgrave Macmillan, Springer etc) [**Book Title**] by [**Book author**(s)] [**COPYRIGHT**] (year of publication)

Other Conditions:

Version 1.3

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.

American Physical Society
Reuse and Permissions License

04-May-2021

This license agreement between the American Physical Society ("APS") and Jonathan Barnard ("You") consists of your license details and the terms and conditions provided by the American Physical Society and SciPris.

Licensed Content Information

License Number:	**RNP/21/MAY/039617**
License date:	04-May-2021
DOI:	10.1103/PhysRevLett.82.1668
Title:	Optimizing High Harmonic Generation in Absorbing Gases: Model and Experiment
Author:	E. Constant et al.
Publication:	Physical Review Letters
Publisher:	American Physical Society
Cost:	USD $ 0.00

Request Details

Does your reuse require significant modifications: No

Specify intended distribution locations:	UK & Commonwealth (excluding Canada)
Reuse Category:	Reuse in a thesis/dissertation
Requestor Type:	Student
Items for Reuse:	Figures/Tables
Number of Figure/Tables:	1
Figure/Tables Details:	Figure 1
Format for Reuse:	Print and Electronic
Total number of print copies:	Up to 1000

Information about New Publication:

University/Publisher:	Imperial College London
Title of dissertation/thesis:	X-ray generation from and spectroscopy of a thin liquid sheet
Author(s):	Jonathan C T Barnard
Expected completion date:	Jun. 2021

License Requestor Information

Name:	Jonathan Barnard
Affiliation:	Individual
Email Id:	[redacted]
Country:	United Kingdom

American Physical Society
Reuse and Permissions License

TERMS AND CONDITIONS

The American Physical Society (APS) is pleased to grant the Requestor of this license a non-exclusive, non-transferable permission, limited to Print and Electronic format, provided all criteria outlined below are followed.

1. You must also obtain permission from at least one of the lead authors for each separate work, if you haven't done so already. The author's name and affiliation can be found on the first page of the published Article.
2. For electronic format permissions, Requestor agrees to provide a hyperlink from the reprinted APS material using the source material's DOI on the web page where the work appears. The hyperlink should use the standard DOI resolution URL, http://dx.doi.org/{DOI}. The hyperlink may be embedded in the copyright credit line.
3. For print format permissions, Requestor agrees to print the required copyright credit line on the first page where the material appears: "Reprinted (abstract/excerpt/figure) with permission from [(FULL REFERENCE CITATION) as follows: Author's Names, APS Journal Title, Volume Number, Page Number and Year of Publication.] Copyright (YEAR) by the American Physical Society."
4. Permission granted in this license is for a one-time use and does not include permission for any future editions, updates, databases, formats or other matters. Permission must be sought for any additional use.
5. Use of the material does not and must not imply any endorsement by APS.
6. APS does not imply, purport or intend to grant permission to reuse materials to which it does not hold copyright. It is the requestor's sole responsibility to ensure the licensed material is original to APS and does not contain the copyright of another entity, and that the copyright notice of the figure, photograph, cover or table does not indicate it was reprinted by APS with permission from another source.
7. The permission granted herein is personal to the Requestor for the use specified and is not transferable or assignable without express written permission of APS. This license may not be amended except in writing by APS.
8. You may not alter, edit or modify the material in any manner.
9. You may translate the materials only when translation rights have been granted.
10. APS is not responsible for any errors or omissions due to translation.
11. You may not use the material for promotional, sales, advertising or marketing purposes.
12. The foregoing license shall not take effect unless and until APS or its agent, Aptara, receives payment in full in accordance with Aptara Billing and Payment Terms and Conditions, which are incorporated herein by reference.
13. Should the terms of this license be violated at any time, APS or Aptara may revoke the license with no refund to you and seek relief to the fullest extent of the laws of the USA. Official written notice will be made using the contact information provided with the permission request. Failure to receive such notice will not nullify revocation of the permission.
14. APS reserves all rights not specifically granted herein.
15. This document, including the Aptara Billing and Payment Terms and Conditions, shall be the entire agreement between the parties relating to the subject matter hereof.

SPRINGER NATURE LICENSE
TERMS AND CONDITIONS

May 07, 2021

This Agreement between Mr. Jonathan Barnard ("You") and Springer Nature ("Springer Nature") consists of your license details and the terms and conditions provided by Springer Nature and Copyright Clearance Center.

License Number	5062090071761
License date	May 04, 2021
Licensed Content Publisher	Springer Nature
Licensed Content Publication	Nature Physics
Licensed Content Title	High-harmonic generation from solids
Licensed Content Author	Shambhu Ghimire et al
Licensed Content Date	Nov 19, 2018
Type of Use	Thesis/Dissertation
Requestor type	academic/university or research institute
Format	print and electronic
Portion	figures/tables/illustrations
Number of figures/tables /illustrations	1
High-res required	no

Will you be translating?	no
Circulation/distribution	1 - 29
Author of this Springer Nature content	no
Title	X-ray generation from andspectroscopy of a thin liquid sheet
Institution name	Imperial College London
Expected presentation date	Jun 2021
Portions	Figure 2
Requestor Location	Mr. Jonathan Barnard [redacted] Attn: Mr. Jonathan Barnard
Total	0.00 GBP

Terms and Conditions

Springer Nature Customer Service Centre GmbH
Terms and Conditions

This agreement sets out the terms and conditions of the licence (the **Licence**) between you and **Springer Nature Customer Service Centre GmbH** (the **Licensor**). By clicking 'accept' and completing the transaction for the material (**Licensed Material**), you also confirm your acceptance of these terms and conditions.

1. Grant of License

1. 1. The Licensor grants you a personal, non-exclusive, non-transferable, world-wide licence to reproduce the Licensed Material for the purpose specified in your order only. Licences are granted for the specific use requested in the order and for no other use, subject to the conditions below.

1. 2. The Licensor warrants that it has, to the best of its knowledge, the rights to license reuse of the Licensed Material. However, you should ensure that the material you are requesting is original to the Licensor and does not carry the copyright of another entity (as credited in the published version).

1. 3. If the credit line on any part of the material you have requested indicates that it was reprinted or adapted with permission from another source, then you should also seek permission from that source to reuse the material.

2. Scope of Licence

2. 1. You may only use the Licensed Content in the manner and to the extent permitted by these Ts&Cs and any applicable laws.

2. 2. A separate licence may be required for any additional use of the Licensed Material, e.g. where a licence has been purchased for print only use, separate permission must be obtained for electronic re-use. Similarly, a licence is only valid in the language selected and does not apply for editions in other languages unless additional translation rights have been granted separately in the licence. Any content owned by third parties are expressly excluded from the licence.

2. 3. Similarly, rights for additional components such as custom editions and derivatives require additional permission and may be subject to an additional fee. Please apply to Journalpermissions@springernature.com/bookpermissions@springernature.com for these rights.

2. 4. Where permission has been granted **free of charge** for material in print, permission may also be granted for any electronic version of that work, provided that the material is incidental to your work as a whole and that the electronic version is essentially equivalent to, or substitutes for, the print version.

2. 5. An alternative scope of licence may apply to signatories of the STM Permissions Guidelines, as amended from time to time.

3. Duration of Licence

3. 1. A licence for is valid from the date of purchase ('Licence Date') at the end of the relevant period in the below table:

Scope of Licence	Duration of Licence
Post on a website	12 months
Presentations	12 months
Books and journals	Lifetime of the edition in the language purchased

4. Acknowledgement

4. 1. The Licensor's permission must be acknowledged next to the Licenced Material in print. In electronic form, this acknowledgement must be visible at the same time as the

figures/tables/illustrations or abstract, and must be hyperlinked to the journal/book's homepage. Our required acknowledgement format is in the Appendix below.

5. Restrictions on use

5. 1. Use of the Licensed Material may be permitted for incidental promotional use and minor editing privileges e.g. minor adaptations of single figures, changes of format, colour and/or style where the adaptation is credited as set out in Appendix 1 below. Any other changes including but not limited to, cropping, adapting, omitting material that affect the meaning, intention or moral rights of the author are strictly prohibited.

5. 2. You must not use any Licensed Material as part of any design or trademark.

5. 3. Licensed Material may be used in Open Access Publications (OAP) before publication by Springer Nature, but any Licensed Material must be removed from OAP sites prior to final publication.

6. Ownership of Rights

6. 1. Licensed Material remains the property of either Licensor or the relevant third party and any rights not explicitly granted herein are expressly reserved.

7. Warranty

IN NO EVENT SHALL LICENSOR BE LIABLE TO YOU OR ANY OTHER PARTY OR ANY OTHER PERSON OR FOR ANY SPECIAL, CONSEQUENTIAL, INCIDENTAL OR INDIRECT DAMAGES, HOWEVER CAUSED, ARISING OUT OF OR IN CONNECTION WITH THE DOWNLOADING, VIEWING OR USE OF THE MATERIALS REGARDLESS OF THE FORM OF ACTION, WHETHER FOR BREACH OF CONTRACT, BREACH OF WARRANTY, TORT, NEGLIGENCE, INFRINGEMENT OR OTHERWISE (INCLUDING, WITHOUT LIMITATION, DAMAGES BASED ON LOSS OF PROFITS, DATA, FILES, USE, BUSINESS OPPORTUNITY OR CLAIMS OF THIRD PARTIES), AND WHETHER OR NOT THE PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION SHALL APPLY NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY PROVIDED HEREIN.

8. Limitations

8. 1. ***BOOKS ONLY:*** Where **'reuse in a dissertation/ thesis'** has been selected the following terms apply: Print rights of the final author's accepted manuscript (for clarity, NOT the published version) for up to 100 copies, electronic rights for use only on a personal website or institutional repository as defined by the Sherpa guideline (www.sherpa.ac.uk/romeo/).

8. 2. For content reuse requests that qualify for permission under the STM Permissions Guidelines, which may be updated from time to time, the STM Permissions Guidelines supersede the terms and conditions contained in this licence.

9. Termination and Cancellation

9. 1. Licences will expire after the period shown in Clause 3 (above).

9. 2. Licensee reserves the right to terminate the Licence in the event that payment is not received in full or if there has been a breach of this agreement by you.

Appendix 1 — Acknowledgements:

For Journal Content:
Reprinted by permission from [**the Licensor**]: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication)

For **Advance Online Publication papers:**
Reprinted by permission from [**the Licensor**]: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication), advance online publication, day month year (doi: 10.1038/sj.[JOURNAL ACRONYM].)

For Adaptations/Translations:
Adapted/Translated by permission from [**the Licensor**]: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication)

Note: For any republication from the British Journal of Cancer, the following credit line style applies:

Reprinted/adapted/translated by permission from [**the Licensor**]: on behalf of Cancer Research UK: : [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT]** (year of publication)

For **Advance Online Publication** papers:
Reprinted by permission from The [**the Licensor**]: on behalf of Cancer Research UK: [**Journal Publisher** (e.g. Nature/Springer/Palgrave)] [**JOURNAL NAME**] [**REFERENCE CITATION** (Article name, Author(s) Name), [**COPYRIGHT**] (year of publication), advance online publication, day month year (doi: 10.1038/sj.[JOURNAL ACRONYM])

For Book content:
Reprinted/adapted by permission from [**the Licensor**]: [**Book Publisher** (e.g. Palgrave Macmillan, Springer etc) [**Book Title**] by [**Book author**(s)] [**COPYRIGHT**] (year of publication)

Other Conditions:

Version 1.3

Questions? customercare@copyright.com or +1-855-239-3415 (toll free in the US) or +1-978-646-2777.
