UNIVERSITY OF SOUTHAMPTON

Faculty of Engineering, Science and Mathematics

School of Physics and Astronomy

DEVELOPMENT AND SPATIO-SPECTRAL MAPPING OF A CAPILLARY HIGH HARMONIC SOURCE

by

Matthew Praeger

Thesis for the degree of Doctor of Philosophy

March 2008

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMETICS
SCHOOL OF PHYSICS AND ASTRONOMY

Doctor of Philosophy

DEVELOPMENT AND SPATIO-SPECTRAL MAPPING OF A CAPILLARY HIGH HARMONIC SOURCE

by Matthew Praeger

This work describes the development and operation of a capillary based High Harmonic Generation (HHG) system. Using this system a coherent beam of soft x-rays is generated, studied and applied. A series of experiments was then undertaken in order to deepen our knowledge of the HHG process and to optimise the performance of the source. Notable contributions made to the field are: A novel laser mode quality measuring device. (Laser mode quality strongly affects the efficiency of the capillary launch). A study of the spectral output of the system as a function of gas pressure, laser power, and laser spectral phase. An analysis technique for recovering spatially-resolved spectral information about a beam by studying the Fresnel diffraction pattern produced at an array of apertures. A study of pulse compression using cascaded quadratic nonlinearity for spectral broadening.

Contents

Abstract i

List of Figures v

Declaration of Authorship xiii

Abbreviations xv

Physical Constants xviii

Symbols xix

1 Introduction 1

2 Apparatus 5
2.1 Laser System . . . 6
2.1.1 Overview and Schematic . . . 6
2.1.2 Kapteyne & Murnane Labs Ti:sapphire . . . 7
2.1.3 Spitfire CPA . . . 9
2.1.4 Dazzler . . . 14
2.1.5 Laser Upgrade . . . 15
2.1.5.1 Tsunami . . . 15
2.1.5.2 Spitfire Pro . . . 17
2.2 Laser Diagnostics . . . 19
2.2.1 Autocorrelator . . . 19
2.2.2 FROG . . . 24
2.2.3 SPIDER . . . 27
2.2.4 M^2 Meter . . . 30
2.3 X-ray Control . . . 38
2.3.1 Gas system . . . 38
2.3.1.1 Vacuum Pumping . . . 38
2.3.1.2 Capillary and Mounting Schemes . . . 39
2.3.1.3 Pressure Control . . . 43
2.3.2 Filtering . . . 47
2.3.2.1 Physics of filters . . . 47
2.3.2.2 Practical implementation of filters . . . 51
2.3.3 X-ray optics . . . 52

2.3.3.1 Grazing Incidence Reflection . . . 52
2.3.3.2 Multilayer Mirrors . . . 59
2.3.3.3 Zone Plates . . . 61
2.4 X-ray detection . . . 63
2.4.1 Photodiode . . . 63
2.4.2 Andor CCD . . . 65
2.4.3 Spectrometer . . . 67

3 Capillary High Harmonic Generation 70
3.1 Theory of HHG . . . 70
3.1.1 Three Step Model . . . 71
3.1.2 Beyond the three step model . . . 72
3.1.3 Regimes of electric field strength . . . 73
3.1.4 Phase matching . . . 75
3.1.5 Quasi phase matching . . . 79
3.1.6 Motivation for use of capillary geometry. . . . 80
3.2 Launch . . . 81
3.2.1 Capillary Modes . . . 81
3.2.2 Focusing . . . 83
3.2.3 Laser M^2 . . . 85
3.3 Pressure and Power dependence . . . 88
3.3.1 Pressure Dependence . . . 92
3.3.2 Power Dependence . . . 96
3.4 Dispersion Dependence . . . 98

4 X-ray diffraction 104
4.1 Overview of Diffraction . . . 105
4.2 Fresnel and Fraunhofer Diffraction . . . 111
4.2.1 Fraunhofer Diffraction . . . 111
4.2.2 Fresnel Diffraction . . . 115
4.2.2.1 Fresnel diffraction at an edge . . . 119
4.2.2.2 Fresnel diffraction through a slit . . . 122
4.2.2.3 Fresnel diffraction by a rectangular aperture . . . 124
4.3 Spatio-Spectral Reconstruction . . . 129
4.3.1 Experimental Setup . . . 129
4.3.2 Initial Findings . . . 131
4.3.3 Least squares fitting . . . 136
4.3.4 Results . . . 142
4.3.5 Maximum entropy . . . 153

5 Pulse Compression 161
5.1 Capillary Compressor and Filamentation . . . 162
5.2 Cascaded $\chi^{(2)}$ compression . . . 162
5.3 Experiment . . . 163
5.4 Results and analysis . . . 166
5.5 Conclusions . . . 168

6 Conclusions **169**
6.1 Paths to higher energy x-rays . . . 172
6.2 Paths to higher flux . . . 178

A Capillary Alignment **180**

B Matlab Code **184**
B.1 SEM Image Analysis . . . 184
B.2 Diffraction Simulation . . . 189
B.3 Least Squares Fitting Routine . . . 192
B.3.1 Input data processing . . . 192

Bibliography **197**

List of Figures

2.1 Internal layout of the KML Ti:sapphire laser cavity with the infrared path illustrated. CM1 and CM2 are the curved dichroic mirrors. C is the Ti:sapphire crystal. L is the lens. FM is the fold mirror. P1 and P2 are the prisms. HR is the high reflector. OC is the output coupler. 8

2.2 Schematic of a prism pair. Red light travels a geometrically shorter path and propagates faster in the prism (normal dispersion) however, the red light has to travel through much more material in the second prism which can produce anomalous dispersion. 9

2.3 Schematic representation of the CPA laser system. Where M is the Millennia Vs, D is the Dazzler, E is the Evolution, QR is the Quanta Ray, X is the X-path amplifier, LA is the linear amplifier and GC is the grating compressor. 10

2.4 Annotated photograph of the Spitfire CPA system. The regen cavity runs along the front of the box and cannot be seen in this picture. D is the Dazzler. In the stretcher: R is the vertical retroflector, G is the grating, CM1 and CM2 are the curved mirrors and BWD is the bandwidth protection device. X is the crystal for the X path amplifier. PC3 is the third Pockels cell and LA is the crystal for the linear amplifier. GC is the compressor for the 1 kHz line. 10

2.5 This diagram (taken from reference [22],) shows how a chirped pulse can be compressed using the Dazzler crystal. 14

2.6 Optical path through a Positive Light femtosecond Tsunami Ti:sapphire oscillator. This figure was extracted from the Tsunami Brochure[24]. . . . 16

2.7 Photograph of Spitfire Pro CPA laser system, extracted from the spitfire pro datasheet[25]. 17

2.8 The timewarp E750 autocorrelator[27]. 20

2.9 Interferometric autocorrelation of a 36 fs laser pulse. 20

2.10 Intensity autocorrelation of an ideal Gaussian 36 fs laser pulse. 21

2.11 Schematic of the PulseCheck autocorrelator[28]. 22

2.12 Photograph of the SSA autocorrelator. 23

2.13 Diagram showing two replica fs pulses crossing in a nonlinear crystal, this was extracted from Trebino's GRENOUILLE tutorial[29]. The autocorrelation can be imaged onto a CCD array and thus collected for a single laser shot. 24

2.14 (a) Simplified diagram showing the main components of the GRENOUILLE in order to illustrate the concept, taken from Trebino's GRENOUILLE tutorial[29]. (b) Example frog trace for a 45 fs laser pulse, delay axis horizontal, spectral axis vertical. 25

2.15 Schematic showing the internal configuration of the GRENOUILLE[29]. . 26

2.16 Schematic showing the optical path inside the Avoca SP-10 SPIDER - reproduced from the users manual. 28
2.17 Illustration of Gaussian beam propagation with the addition of M^2. Upper plot: beam waist is kept the same, in the far field $M^2 = 2$ beam has an M^2 larger radius than $M^2 = 1$ beam. Lower plot: $M^2 = 2$ beam and its hypothetical underlying Gaussian. Hypothetical beam has $M^2 = 1$ and is everywhere M times smaller than the 'real' beam. 31
2.18 Diagram of M^2 meter. 33
2.19 Nonlinear fit (red line) using equation 2.8 to some example data (black markers) calculated for an $M^2 = 1.2$ beam at 780 nm with a waist radius of 25 μm. The blue lines show how varying M^2 by just 0.2 affects the fit. . 36
2.20 M^2 meter software output. 37
2.21 Average transmission for 20–50 eV x-rays through 1 m of Ar gas. 39
2.22 Composite ($\times$ 10) microscope image showing the cross section of a laser-drilled hole (looking down the length of the capillary). This hole has a square cross section (produced by translating the laser focus along a square path). The hole was first drilled from one side, then the capillary was rotated 180° and a second hole was drilled. The cross-section view was achieved by cleaving the capillary at the hole position. 41
2.23 Schematic of the capillary in its T-piece holder (top). Sketch showing the approximate pressure profile that the holes in the capillary wall produce (below). 42
2.24 Photograph of a simple T-piece with the capillary installed; in this case the seal between the capillary and the T-piece is made using two o-rings. The o-rings used here are in fact x-rings; they have an "x" shaped cross section. This provides two sealing surfaces which prevents twisting and improves performance compared with traditional o-rings. 43
2.25 Photograph of the capillary mounting block with a capillary fitted. The bolts provide enough force to compress the o-rings and make a seal without exerting undue stress on the capillary. The custom o-ring is the smaller central one; gas is introduced into this section through the fitting at the bottom of the block. The block is made compatible with the existing T-piece holding vacuum system with the addition of metal tubes glued to either side. 44
2.26 Exploded view taken from the design drawings for the capillary mounting block. In this view the shape and function of the block is a little easier to follow. 44
2.27 Left: Schematic of the pressure control system. Right: Legend of components. The pressure controller used was a Tescom ER3000. 45
2.28 The upper plot shows the real (blue) and imaginary (green) parts of the refractive index of aluminium over a broad range of the electromagnetic spectrum (note the logarithmic scales). The lower plot shows the reflectance of aluminium over the same spectral range. (Reflectance figures are for unpolarised light at normal incidence). Data taken from Webber[34]. 49

2.29 The upper plot shows the real(blue) and imaginary (green) parts of the refractive index of aluminium in the soft x-ray region of the spectrum. The lower plot shows the transmittance (blue) and reflectance (green) of aluminium over the same spectral range. (Reflectance figures are for unpolarised light at normal incidence and transmittance is calculated for a 200 nm thick aluminium foil). This data was taken from CXRO[31]. . . . 50
2.30 Left: Cutaway diagram showing the internal structure of the filtering unit. Right: View of the whole unit; the optical path runs along the horizontal axis of the central cross, filters slide in the vertical axis. 52
2.31 Angular reflectivity at... Upper Left: A vacuum-SiO_2 interface at 633 nm. Upper Right: A vacuum-SiO_2 interface at 25 nm. Lower Left: An SiO_2-vacuum interface at 633 nm. Lower Right: An SiO_2-vacuum interface at 25 nm. 54
2.32 Angular reflectivity at a vacuum to silicon interface - S and P polarisations are both shown. The upper plot shows the reflectivity curves for 780 nm whilst the lower plot is calculated for 25 nm. 55
2.33 3D model of a Kirkpatrick-Baez mirror system (KB-SR1) produced by Xradia - this figure was extracted from their description brochure[36]. . . . 57
2.34 XMM-Newton employs a two step optical system consisting of a set of parabolic shells followed by a set of hyperbolic shells. This image is taken from ESA's XMM-Newton users handbook (v2.6)[37]. 57
2.35 Composite image of a tapered capillary used for x-ray focusing - original images produced using an SEM. 58
2.36 Left: Ray trace of the taper profile extracted from the SEM image in figure 2.35. Centre: Ray trace of an improved taper profile - extracted from optical microscope images. Right: Ray trace of a truncated parabolic taper profile. Note: the aspect ratio of these plots has been deliberately distorted to make the data easier to see. 59
2.37 A partial reflection occurs at each interface within a multilayer mirror (typical optics consist of between 30 and 100 layers). By carefully designing the layer indices and thicknesses the reflections from each interface can be made to constructively interfere - enhancing reflected power. Note that a 180° phase change occurs on reflection when an increase in refractive index occurs across the interface. 60
2.38 SEM image of an intensity mask Fresnel zone plate - image taken from[42]. 61
2.39 A plot of quantum efficiency as a function of illumination energy for the Andor camera. Note that the model used for this project was back illuminated and is represented by the solid line. This data and image are reproduced from the Andor DX434 specification sheet[46]. 65
2.40 A plot of dark current as a function of detector temperature. This data is reproduced from the Andor DX434 specification sheet[46]. 67
2.41 Grazing incidence spectrometer: 1, Entrance slit. 2, Grating. 3, MCP. 4, Lens system. 5, CCD camera. 6, Turbo pump. 7, Vacuum gauge. . . . 68

3.1 Electric field line diagrams and radial intensity profile for the lowest order electrical hybrid modes of a hollow dielectric waveguide. Taken from reference [71]. 82
3.2 Calculated coupling efficiency for the first three EH_{1m} modes as a function of input spot size. 84

3.3 Photodiode signal as a function of pressure. Four different data sets are shown; each recorded at a very slightly different alignment. This highlights the variation in the output of the HHG capillary. High order polynomial fits have been added as a guide for the eye. (The apparent DC current offset is likely to be caused by laser light leaking through the aluminium filter). 90
3.4 Photodiode signal as a function of pressure. The thin blue lines represent the experimental data from two different pressure scans. The thicker dashed lines are high order polynomial fits to this data. The red dotted line is the average of the data sets shown in figure 3.3. 91
3.5 Image plot showing a series of spectra plotted against argon pressure. . . 93
3.6 Series of x-ray CCD images showing the spatial intensity distribution as a function of pressure. Pressures are (from left to right) 10, 60 and 200 mbars. 94
3.7 X-ray CCD camera image, recorded at 80 mbar. This shows an intense central core surrounded by a weaker more divergent spot. The central core is believed to be primarily due to one particular harmonic that is well phase matched. 96
3.8 X-ray spectra (horizontal axis) recorded against input laser power (vertical axis) The effect is similar to that of varying pressure as shown in figure 3.5. 97
3.9 Top left: Real part of the electric field of a 30 fs FWHM, Gaussian pulse, centred at 780 nm. Top right: After adding 300 fs^2 of GDD the pulse is increased to $\sim$ 41 fs FWHM. Lower left and right: These plots show the variation in instantaneous wavelength with time. Note that the leading edge of the pulse (negative times) contains the longer wavelength 'red' components since normal dispersion (positive GDD) was added. 99
3.10 ADK based calculation of the ionisation fraction for two different pulse durations. Other laser parameters are: centre wavelength 780 nm, repetition rate 1 kHz, average power 650 mW, spot size 51 μm as required for launch into a 150 μm bore capillary. This leads to peak intensities of 2.8×10^{14} and 2×10^{14} $W.cm^{-2}$ for the 30 fs and 41 fs pulses respectively. The 4 % ionisation fraction is reached at approximately -8 fs and -4 fs for the 30 fs and 41 fs pulses respectively. 100
3.11 Image plot showing the x-ray spectra (horizontal) as a function of quadratic spectral phase (vertical). 101
3.12 Image plot showing the x-ray spectra (horizontal) as a function of cubic spectral phase (vertical). 101
3.13 Image plot showing x-ray spectrum as a function of quadratic spectral phase. In this data set average laser power is adjusted to maintain a constant peak power. As a result ionisation is increased for large quadratic spectral phase offsets 102

4.1 (Left: Finite element mesh for a 2D simulation of interference between light emitted from two waveguides. The boundary conditions are set so that the end of each waveguide is a plane wave source. Right: The intensity map obtained by solving for these boundary conditions. This is actually the output of a demonstration file included with the Comsol Multiphysics package. 106

4.2 Illustration of the Huygens-Fresnel principle applied to diffraction through a slit. The spherical waves re-radiated from every point along the wavefront combine to approximate the diffracted light. 108
4.3 Diagram showing the geometry of simple aperture diffraction problem. . . 109
4.4 Diagram showing the sign convention used for the angles (n, r), (n, s) and δ. This figure also demonstrates why the centre of the aperture (at O) must be close to the axis between the source P_{-1} and the measurement point P_1. Note that this figure shows a 2D example where P_{-1} and P_1 are in the same y-plane; typically this is not the case and x_1 would for example be replaced with $\sqrt{x_1^2 + y_1^2}$. 112
4.5 Comparison between the phase produced by a point light source as calculated under different approximations. Left column: Point light source is on the optical axis. Right column: Point source is off axis. Top row: No approximation - distances calculated with exact geometry. Middle row: Fresnel approximation. Bottom row: Fraunhofer approximation. 117
4.6 Diagram showing the geometry used in calculating the Fresnel diffraction pattern produced by an infinite straight edge. 119
4.7 Left: Plot showing the behaviour of the Fresnel $\mathscr{C}(u)$ function (blue-solid) and the Fresnel $\mathscr{S}(u)$ function (red-dashed). These are both plotted against q here since u is a function of q. Right: Plot of the spiral of Cornu created by plotting $(\mathscr{C}(u) + i\mathscr{S}(u))$ in the form of an Argand diagram. 120
4.8 Image plot showing the calculated Fresnel diffraction pattern produced by the edge of a semi-infinite plane. The conditions used here are: $\lambda =$ 25 nm, $r\prime$ and $s\prime$ are both approximately 0.5 m and q takes the range $\pm$ 1 mm. 122
4.9 The diffracted fields of two spatially separated infinite edges may be subtracted (according to Babinet's principle) in order to simulate in infinite slit. Image and line plots both show the real part of the E-field amplitude as a function of position q. Left and middle show the contributions from each edge. Right shows the amplitude resulting from the subtraction. . . 123
4.10 The calculated E-field (left) is multiplied by its complex conjugate to recover the intensity distribution (right). This simulation is calculated for a 300 μm wide slit diffracting 25 nm radiation. The source-slit and slit-screen distances are both $\sim$ 0.5 m. Note that the plots on the left differ from those in the right hand column of figure 4.9. This is because here the absolute value of the E-field is plotted whereas previously the real part was plotted. 124
4.11 Calculated intensity profiles for a selection of square apertures (aperture width is written in the title above each plot). The simulated wavelength was 25 nm, the source-slit and slit-screen distances are both $\sim$ 0.5 m. The axis scales represent translation on the detector screen in mm. 127
4.12 Diffraction pattern produced by 25 nm light passing through a 150 μm square aperture. Left: plotted on a linear scale and displayed in monochrome. Right: plotted on a logarithmic intensity scale using a false colour scale (dark red is most intense, dark blue is least intense, the markers on the colorbar give the scale in dB). 128
4.13 Diagram showing the setup used to create Fresnel diffraction patterns using a wire mesh as a diffraction target. 130

4.14 Example of a measured diffraction pattern (340 μm square apertures separated by 38 μm bars). The insets show the variation in the diffraction pattern produced by different apertures. The scale bar represents 4 mm in the CCD plane. 132
4.15 Typical output from the SEM image analysis software described in appendix B.1. This image shows analysis of the grid geometry of a Ni wire mesh diffraction target. Errors in aperture size are approximately 1%. . . . 133
4.16 Calculated diffraction patterns for $\lambda = 41$ nm. Pattern A has a 340 μm square aperture whilst pattern B is 345 μm square. 133
4.17 Plot showing the wavelength selectivity for the 19^{th} (dashed) and 29^{th} (solid) harmonics, in terms of pixel error (% full scale). The red lines indicate the harmonics that are typically generated by the HHG system. 134
4.18 Image on left shows the calculated diffraction pattern for an isolated aperture illuminated by the 27^{th} harmonic. The right hand image shows calculated diffraction pattern for an array of identical apertures under the same illumination conditions. The average pixel error between the two boxed regions is $\sim$ 7.5 % of the full scale. 135
4.19 Diffraction patterns produced by an aperture illuminated by a discrete spectral line at 33 nm (left); and by a 1 nm FWHM spectral intensity distribution centred at 33 nm (right). The average pixel error between the two patterns is $\sim$ 3 % of the full scale. 138
4.20 Comparison between raw input data and processed data for two different apertures in the same CCD image. Note that the processed data has been symmetrised and the background signal has been removed from the dark area surrounding the diffraction pattern. Despite the processing it is clear that the underlying shape of the diffraction pattern has been maintained. 139
4.21 Graphical representation of the least squares fitting routine. Each vertical column in the main block contains the simulated diffraction data calculated for a different harmonic. (The 2D 60 $\times$ 60 grid of simulated pixels has simply been wrapped to produce a 1 $\times$ 3600 column). The measured diffraction pattern (60 $\times$ 60 pixels) is also re-sampled into a column vector format and is shown on the right. The row vector at the bottom of the figure shows the coefficients for each harmonic that are required to obtain the best fit to the measured data. 141
4.22 Comparison between spectrum measured with a grazing incidence spectrometer (blue line plot) and spectral coefficients recovered from the diffraction pattern (red stem plot). Argon pressures in the HHG system ranges from 10 to 100 mbar as labeled. 143
4.23 Left: Examples of the retrieved spatio-spectral data. Right: Rotated CCD image from which the data is recovered. HHG gas pressure set to 10 mbar 145
4.24 Left: Examples of the retrieved spatio-spectral data. Right: Rotated CCD image from which the data is recovered. HHG gas pressure set to 20 mbar 146
4.25 Left: Examples of the retrieved spatio-spectral data. Right: Rotated CCD image from which the data is recovered. HHG gas pressure set to 40 mbar 147

4.26 Left: Examples of the retrieved spatio-spectral data. Right: Rotated CCD image from which the data is recovered. HHG gas pressure set to 60 mbar . . . 147
4.27 Left: Examples of the retrieved spatio-spectral data. Right: Rotated CCD image from which the data is recovered. HHG gas pressure set to 80 mbar . . . 148
4.28 Left: Examples of the retrieved spatio-spectral data. Right: Rotated CCD image from which the data is recovered. HHG gas pressure set to 100 mbar . . . 148
4.29 Measurements made using an imaging x-ray spectrometer. Different spectra are found to arrive at different parts of the slit. . . . 150
4.30 Spectral intensity map sequences recorded at a several different pressures. Upper row was extracted from a CCD image taken at 10 mbar whilst the image used for the lower row was recorded at 100 mbar. . . . 152
4.31 Spectral intensity map sequences extracted from four different CCD images. Each sequence contains panes that appear to have an annular intensity profile. . . . 153
4.32 Test of the maximum entropy analysis algorithm. The red lines on the lower plot show the coefficient values that were used to create the input data diffraction pattern. The blue line on the lower plot shows the spectral coefficients that were recovered by the maximum entropy algorithm. Multiplying the simulated data by these coefficients and then summing produced the fitted data diffraction pattern. For this simple test pattern the fit of the recovered spectrum to the input spectrum is good. . . . 156
4.33 Layout of this figure is the same as that of 4.32. The agreement with the input data is also good for this simple test pattern. . . . 156
4.34 Fitting accuracy for a more complicated pattern is slightly reduced. . . . 157
4.35 This figure shows the recovered spectrum that is obtained when a single spectral coefficient is input into the algorithm. . . . 157
4.36 When attempting to fit to a simulated harmonic spectrum composed of narrow spectral peaks the maximum entropy algorithm interprets the peaks as noise. Rather than reconstructing the harmonic peak structure the algorithm returns the spectral envelope of the x-ray emission. . . . 158
4.37 Here the maximum entropy method is applied to some measured diffraction patterns. The agreement between measured and reconstructed diffraction patterns is qualitatively rather good. There is some residual error although this is slightly exaggerated by the colormap. . . . 159
4.38 The dip in the peak of this fitted curve suggests that the maximum entropy method is on the verge of being able to resolve individual harmonic peaks from the measured diffraction data. . . . 160

5.1 Results of a Gaussian beam focusing simulation. This was used to find suitable focusing conditions. . . . 164
5.2 Upper plot: Second harmonic signal as a function of BBO crystal angle. Lower plot: Laser spectrum. . . . 165
5.3 Image plot showing the spectrum of the second harmonic at the crystal exit as a function of crystal tuning angle. . . . 165

5.4 Output pulse duration (intensity FWHM) is plotted as a function of input laser power both for positive and negative detuning of 1.5° from phase matching. . . . 166
5.5 Images A and B show the measured FROG trace recorded at ∼ 12 mW for ± 1.5° detuning from the phase matched angle. . . . 167
5.6 Sequence of FROG traces showing the evolution of the pulse shape with crystal tuning angle. Crystal angle (relative to face normal) is marked above each trace. The horizontal axis for each trace is the delay axis. The vertical axis of each trace represents wavelength. . . . 167

B.1 A: Plot of data measured using a scanning electron microscope. B: The Sobel method is used to detect the vertical edges in image A. . . . 187
B.2 Left: X,Y plot of all data points determined to be on edges. Middle: Plot of Y data values in ascending order. Right: Plot the differential of the ascending Y data values, the spikes clearly separate the data into groups corresponding to each edge. . . . 187
B.3 X,Y scatter plot of the data contained in one group. The red lines show the best fit to the data after the itterative loop, the parameter values are shown in the title. Notice that data is transposed whilst the fit is performed. 188
B.4 Example output with vertical lines fitted. Average values and errors are quoted in the title these are now scaled to units of microns. . . . 188
B.5 Example output of CornuVariable.m. The axis units are simply array indices (the output size was set to 1000). . . . 191
B.6 Example output when the optional plotting commands are executed - this is for visual checking that the aperture spacing values have been set correctly. . . . 195
B.7 Comparison between raw input data and processed data for two different apertures. Note that the processed data has been symmetrised and the background signal has been removed from the dark area surrounding the diffraction pattern. . . . 196

Declaration of Authorship

I, Matthew Praeger, declare that this thesis entitled, 'DEVELOPMENT AND SPATIO-SPECTRAL MAPPING OF A CAPILLARY HIGH HARMONIC SOURCE' and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated
- Where I have consulted the published work of others, this is always clearly attributed
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work
- I have acknowledged all main sources of help
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself
- Parts of this work have been published as: [1] and [2]

Signed: _____

Date: _____

"Procrastination is the thief of time."

Edward Young (1683 - 1765)

Abbreviations

ADK	**A**mmosov **D**elone **K**rainov (Ionisation theory)
AOPDF	**A**cousto **O**ptic **P**rogrammable **D**ispersive **F**ilter
APE	**A**ngewandte **P**hysik & **E**lektronik
BBO	**B**eta **B**arium B**o**rate (BaB_2O_4)
BNC	**B**ayonet **N**eill-**C**oncelman connector
BWD	**B**and**W**idth protection **D**iode
C	Ti:sapphire **C**rystal
CCD	**C**harge **C**oupled **D**evice
CF	**C**on**F**lat
CM#	**C**urved **M**irror (Numbered)
CPA	**C**hirped **P**ulse **A**mplification
CW	**C**ontinuous **W**ave
CXRO	**C**entre for **X**-**R**ay **O**ptics
D	**D**azzler (AOPDF)
DFG	**D**ifference **F**requency **G**eneration
E	**E**volution Laser
EDX	**E**nergy **D**ispersive **X**-ray spectrometer
EH	**E**lectrical **H**ybrid
ESA	**E**uropean **S**pace **A**gency
ET	**ET**alon
FEL	**F**ree **E**lectron **L**aser
FEM	**F**inite **E**lement **M**ethod
FF	**F**undamental **F**ield
FFT	**F**ast **F**ourier **T**ransform
FLASH	**F**ree electron **LAS**er in **H**amburg

FM	**F**old **M**irror
FROG	**F**requency **R**esolved **O**ptical **G**ating device
FWHM	**F**ull **W**idth **H**alf **M**aximium
G	**G**rating
GC	**G**rating **C**ompressor
GaAsP	**G**allium **A**rsenide **P**hosphide
GDD	**G**roup **D**elay **D**ispersion
GPIB	**G**eneral **P**urpose **I**nterface **B**us (IEEE-488)
GRENOUILLE	**GR**ating **E**liminated **N**o-nonsense **O**bservation of ... **U**ltrafast **I**ncident **L**aser **L**ight **E**-fields
HHG	**H**igh **H**armonic **G**eneration
HR	**H**igh **R**eflector
HV	**H**igh **V**acuum
IR	**I**nfra **R**ed
ISO−K	**I**nternational **S**tandards **O**rganisation style **K** flange
KF	**K**lein **F**lange
KML	**K**apteyne and **M**urnane **L**abs
L	**L**ens
LA	**L**inear **A**mplifier
LASER	**L**ight **A**mplification by **S**timulated **E**mission of **R**adiation
LCLS	**L**inac **C**oherent **L**ight **S**ource
LED	**L**ight **E**mitting **D**iode
LP	**L**inear **P**olarised
M	**M**illennia Vs Laser
MCP	**M**ulti **C**hannel **P**late
Nd:YVO_4	**N**eo**d**ymium **Y**ttrium **V**andate
Nd:YAG	**N**eo**d**ymium **Y**ttrium **A**luminium **G**arnate
OC	**O**utput **C**oupler
OTB	**O**ver **T**he **B**arrier
P#	**P**rism (Numbered)
PC#	**P**ockels **C**ell (Numbered)
PID	**P**roportional **I**ntegral **D**ifferential
QE	**Q**uantum **E**fficiency

QPM	**Q**uasi **P**hase **M**atching
QR	**Q**uanta **R**ay Laser
R	Vertical **R**etroflector
RMS	**R**oot **M**ean **S**quared
RS485	**R**ecommended **S**tandard 485 (serial comms)
SEM	**S**canning **E**lectron **M**icroscope
SH	**S**econd **H**armonic
SHG	**S**econd **H**armonic **G**eneration
SPIDER	**S**pectral **P**hase **I**nterferometry for **D**irect ...
	E-field **R**econstruction
SPM	**S**elf **P**hase **M**odulation
SSA	**S**ingle **S**hot **A**utocorrelator
TEM	**T**ransverse **E**lectro-**M**agnetic
Ti:sapphire	**Ti**tanium doped **sapphire** (Ti:Al_2O_3)
TOD	**T**hird **O**rder **D**ispersion
UHV	**U**ltra **H**igh **V**acuum
USB	**U**niversal **S**erial **B**us
UV	**U**ltra **V**iolet
VUV	**V**acuum **U**ltra **V**iolet
X	**X**-path amplifier
XMM-Newton	**X**-ray **M**ulti **M**irror **M**ission ...
	named in honor of Sir Isaac **Newton**
XUV	e**X**treme **U**ltra **V**iolet

Physical Constants

Constant Name	Symbol		Constant Value
Speed of Light	c	=	$2.997\ 924\ 58 \times 10^{8}$ ms^{-1}
Vacuum permittivity	ϵ_0	=	$8.854\ 187\ 817 \times 10^{-12}$ $\mathrm{F\ m}^{-1}$
Electron mass	m_e	=	$9.109\ 381\ 88 \times 10^{-31}$ kg
Electron charge	q_e	=	$1.602\ 176\ 46 \times 10^{-19}$ C

Symbols

Section 2.1

symbol	name	unit
P_{avg}	Average laser power	W
P_{peak}	Peak laser power	W
$\Delta\tau$	Pulse duration	s
f_{rep}	Laser repetition rate	Hz

Section 2.2

T_{ac}	Time calibration factor	
$N_{fringes}$	Number of fringes	
λ	Wavelength	m
T_{scope}	Oscilloscope time window	s
S_S	Measured Spider signal intensity	W cm^{-2} Hz^{-1}
ω	Angular frequency	rad s^{-1}
Ω	Frequency shear	rad s-1
$\vec{E}$	Electric field	V m^{-1}
ψ	Spectral phase	rad
τ	Time delay	s
S_R	Spider reference signal intensity	W cm^{-2} Hz^{-1}
M^2	Beam propagation factor	
I	Laser intensity	W cm^{-2}
ω_{0R}	Measured beam waist ($1/e^2$ radius)	m
ω_R	Measured beam radius ($1/e^2$)	m

ω_0	Beam waist of underlying Gaussian	m
θ_R	Measured far field divergence	rad
θ	Divergence of underlying Gaussian	rad
z	Propagation axis position	m
x and y	Transverse axis position	m
M_x^2 and M_y^2	Beam propagation factors in x and y axes	
x_0	First moment	W m^{-3}
σ_x^2	Second moment	W m^{-4}

Section 2.3

n	Complex refractive index	
α	Real part of complex refractive index	
β	Imaginary part of complex refractive index	
σ	Electrical conductivity	Sm^{-1}
ω	Angular frequency	rads^{-1}
τ	Mean free path	s
N	Number density of electrons per unit volume	m^{-3}
r	Reflectivity	
θ_i	Angle of incidence	rad
θ_r	Angle of reflection	rad
r_s	Reflectivity for S polarised light	
r_p	Reflectivity for P polarised light	
$n_\#$	Refractive index of the $\#^{th}$ medium	
θ_c	Critical angle	rad

Section 3.1

E_{max}	Cutoff energy	J or eV
I_p	Ionisation potential	J or eV
U_p	Ponderomotive potential	J or eV
E	Electric field amplitude	V m^{-1}
χ	Electric susceptibility	m A^{-1}

P	Electric polarisation	C m^{-2}
$\chi^{(1)}$	Linear susceptibility	
$\chi^{(2...)}$	Non-linear susceptibility coefficients	
γ	Keldysh parameter	rad J s A^{-1}
q	Harmonic order	
k_q	k-vector of the q^{th} harmonic	m^{-1}
k_f	k-vector of the fundamental	m^{-1}
Δk	Phase mismatch	rad
l_c	Coherence length	m
l_a	Absorption length	m
σ	Ionisation cross-section	m^2
ρ	Gas density	kg m^{-3}
P	Gas pressure in atmospheres	Bar
η	Ionisation fraction	
$\delta(\lambda)$	Linear dispersion of the neutral gas	s m^{-1}
n_2	Nonlinear refractive index of the neutral gas	
N_{atm}	Number of gas molecules per atmosphere	
r_e	Classical electron radius	m
a_0	Bohr radius	m
α	Fine structure constant	
u_{11}	First root of the J_0 Bessel function	
a	Capillary radius	m
η_c	Critical ionisation fraction	
λ_L	Laser wavelength	m
$\Delta\delta(\lambda_L)$	Difference in refractive index ...	
	between the harmonics and the laser	

Section 3.2

λ_c	Cutoff wavelength	m
l and m	Azimuthal and radial mode indices	
$V_c(lm)$	Cutoff normalised frequency for the mode (lm)	
n_i and n_e	Refractive index of internal end external materials	

u_m	The m^{th} root of the J_0 Bessel function	
r	Radial coordinate	m
E_0	On-axis Electric field amplitude	V m^{-1}
$E(r)$	Electric field amplitude as a function of radius	V m^{-1}
ω	Beam width ($1/e^2$ radius)	m
f	Focal length	m
D	Beam diameter ($1/e^2$ diameter)	m

Section 3.3

E_p	Photon energy	J
N_p	Number of photons	
I	Current	A

Section 3.4

φ	Spectral phase	rad
φ_0	Absolute phase	rad
φ_1	Delay	s
φ_2	Quadratic spectral phase, (GDD)	s^2, (s^2 m^{-1})
φ_3	Third order dispersion	s^3

Section 4.1

E_x E_y E_z	Orthoganal components of vector E-field	V m^{-1}
U	Magnitude of the vector E-field (scalar)	V m^{-1}
P_{-1}	Position of the point light source (x_{-1}, y_{-1}, z_{-1})	
P_0	Position in the aperture plane (x_0, y_0, z_0)	
P_1	Position on detector screen (x_1, y_1, z_1)	
A	Amplitude scaling factor	V m^{-1}
r	Distance from P_{-1} to P_0 (source to aperture)	m
s	Distance from P_0 to P_1 (aperture to screen)	m
O	Position of the coordinate system origin	
$r\prime$	Distance from P_{-1} to O (source to origin)	m
$s\prime$	Distance from O to P_1 (origin to screen)	m
n	Unit vector, normal to the aperture	

(n, r)	Accute angle between r and n	rad
(n, s)	Obtuse angle between s and n	rad

Section 4.2

F	Fresnel number	
a	Characteristic size (radius)	m
L	Object to measurement plane distance	m
δ	Angle between n and $P_{-1}P_1$	rad
Q	Intersection of $P_{-1}P_1$ and the aperture plane	
ρ	Substitution as defined in equation 4.6	
x	Substitution as defined in equation 4.8	
G	The aperture function	
k	Wavevector	m^{-1}
k_x and k_y	Orthoganal wavevector components	m^{-1}
$\mathscr{F}$	Denotes the Fourier transform	
f	Substitution as defined in equation 4.15	
B	Constant terms from equation 4.14	
C	Real terms from equation 4.14	
S	Imaginary terms from equation 4.14	
b	Constant scaling factor given by equation 4.20	
u and v	Variables of integration defined in 4.19	
$d/2$	Displacement of the edge from the origin	m
w	Value of u that corresponds to the edge	
$\mathscr{S}$ and $\mathscr{C}$	Fresnel Sine and Cosine integrals 4.22	
I	Intensity distribution on the screen	$\mathrm{W\ cm}^{-2}$
q	Position on the screen	
w_1 and w_2	Values of u that correspond to the slit edges	
U_0	Field amplitude at P_1 without screen	$\mathrm{V\ m}^{-1}$
I_0	Intensity normalisation constant	$\mathrm{W\ cm}^{-2}$

Section 4.3

P_m	Measured diffraction pattern	
i and j	The apertures indices	

x and y	Coordinates within each aperture	m
h	Odd integer denoting harmonic order	
P_s	Simulated diffraction pattern	
I	Harmonic intensity coefficients	W cm^{-2}
N_s	Shot noise	Counts
S	Measured signal in electron counts	Counts

For the friends and family who helped me get through this...

Chapter 1

Introduction

Throughout the history of science the development of new sources of electromagnetic radiation has driven innovation and increased our understanding of the physical processes that make up the world around us. Each new source opens further regions of the electromagnetic spectrum to scientific study and allows us to probe our environment at different energy levels.

Naturally the visible light from the sun (readily detected by the human eye) was the first part of the electromagnetic spectrum to be studied in detail. Simple lenses, made from polished rock crystals, and dating back as far as 1400 BC have been discovered[3]. These early lenses are thought to have been used for magnification and for focusing the suns light in order to ignite fires.

Our understanding of visible light progressed gradually over the subsequent centuries with, for example, Euclid's work on geometrical optics[4] (ca. 300BC). It was not until much later (the start of the 17^{th} century) that more complex optical devices such as the telescope and microscope were invented. An interesting historical review of this development has been written by Herzberger[5].

It was, however, the study of the nature of colour and of refraction by prisms that led Newton[6] (1642–1727) to coin the term 'spectrum' for the spread of colours he observed. Huygens (1629–1695) later proposed the wave-like nature of light and the concept of wavelength appeared.

By burning a variety of salts using a spirit flame and observing these through an aperture and prism, in 1752 Melvill (1726–1753) became the first to observe a line spectrum[7] and so the field of spectroscopy was born. This was aided by the development of gas discharge and incandescent bulbs which occurred at around the same time.

Huygens' concept of light as a wave was elegantly applied to diffraction by Fresnel (1788–1827) (see chapter 4). This theory of optics was given physical interpretation and a firm mathematical foundation in 1864 by Maxwell[8] (1831–1879) who consolidated optics with work on electromagnetic induction carried out by Faraday (1791–1867). Maxwell's work brought with it our modern concept of the electromagnetic spectrum and initiated a period of rapid development of new sources.

In 1887 the invention of the spark-gap transmitter made Hertz (1857–1894) the first to generate and then detect electromagnetic radiation outside of the visible region. Hertz produced radiation in what would now be termed the microwave and radio wave regions of the spectrum - his demonstration was sufficient to prove that Maxwell's theory of electromagnetic waves was correct. Later the vacuum-tube and the work of Tesla (1856–1943) allowed these radio waves to be developed as a means of wireless communication.

In 1895, shortly after the invention of radio, Roentgen (1845–1923) discovered x-rays whilst working with a cathode ray tube[9]. High energy (or hard) x-rays have a wavelength so short that their probability of interaction with matter is low unless the matter is very dense. This property permits x-rays to be used for transmission imaging through objects that are opaque to visible light. In medical x-ray imaging the x-rays pass through the softer tissues of the body with little attenuation; in contrast the denser bone attenuates the x-rays producing a shadow-like image. When the wavelength of x-rays is comparable to the separation of atoms in a crystal lattice this leads to the strong diffraction that makes x-ray crystallography possible. Using this technique it is possible to establish the composition of a crystallised molecular or protein sample down to the atomic level.

In 1890 Villard (1860–1934) discovered a still more energetic form of electromagnetic radiation; gamma-rays, emitted from radioactive uranium and radium[10]. This was part of the work that led to our modern understanding of radioactivity and the structure of the atom.

In 1893 Schumann (1841–1913) who was using hydrogen and mercury lamps substituted quartz prisms for fluorite ones and performed his experiment under vacuum[11]. In so doing, Schumann discovered the vacuum ultraviolet (VUV) region of the spectrum (sometimes called the Schumann region) the study of which had previously been precluded by material absorption. This later allowed Lyman to discover the spectral lines of hydrogen.

Over the years a great variety of sources of electromagnetic radiation have been produced; these sources combined mean that we can produce radiation in almost every part of the electromagnetic spectrum. However, the region between x-rays and the VUV,

referred to as the soft x-ray region, has persistently eluded study since it corresponds to a wealth of atomic resonances which causes very high absorption in all materials[12] and therefore a lack of suitable optics. Recent developments in x-ray optics (see section 2.3.3 and in high harmonic generation (HHG) of soft x-rays (see chapter 3) has renewed interest in this region of the spectrum.

Since the invention of the laser[13] (Light Amplification by Stimulated Emission of Radiation), the intensity with which light can be generated has steadily increased, now reaching levels that were previously unimaginable[14]. These high intensity sources opened up a new regime in the interaction between light and matter and gave rise to the field of nonlinear optics (see section 3.1.3).

High intensity (photons per second per unit area) is typically achieved using a pulsed source. This leads to the related challenge of producing the shortest possible pulse duration (and in so doing gaining the highest possible temporal resolution). Trains of optical pulses with individual pulse durations in the range of $\sim$ 100 attoseconds are produced by HHG sources such as that demonstrated in this work. Others have carried out work with the aim of isolating individual pulses from these trains in order to obtain single attosecond bursts[15].

Short pulse duration is necessarily accompanied by large spectral bandwidth. Possibly the most promising new source in this respect is the free electron laser (FEL). Depending upon the optimisation, a FEL may operate anywhere from the terahertz to the soft-x-ray region. The current record for shortest wavelength generated is held by the free electron laser at Hamburg (FLASH) with light being generated down to 6.5 nm though this is expected to be exceeded when LCLS comes on-line later this year[16]. Just as was the case with earlier storage ring based synchrotrons, the repulsion between electrons imposes limits on the operation of the FEL. In the beam of a FEL this repulsion causes dispersion of the electron pulse which has a negative effect on coherence and limits the optical pulse duration to be greater than a few $10's$ of femtoseconds.

Throughout the course of this research the primary focus has been the development of all aspects of a capillary based HHG system. This system converts light from an intense femtosecond laser system into a coherent beam of soft x-rays. In doing this the photon energy is increased from 1.6 eV to $\sim$ 50 eV (780 nm to 25 nm).

A long-term aspiration of this project is to perform x-ray diffraction experiments, for example to determine the shape of individual protein molecules. (Similar experiments have since been performed on FEL sources except with radiation-hard samples and exposure times of the order of 500 seconds[17, 18, 19]). Whilst HHG (as employed in this work) is unable to produce the high-energy radiation that is obtained by synchrotrons and

FELs; the wavelengths it generates are sufficiently short to resolve useful shape information from such molecules. Aside from the obvious benefits of size, cost and convenience that HHG offers relative to electron accelerator sources, HHG also offers shorter pulse durations. Short pulse durations may prove to be essential when imaging non-radiation hard samples (as predicted by Neutze[20]) since the energy of the illuminating light may be sufficient to destroy the sample before a long exposure can be completed. Short pulse durations might also be exploited in a pump probe geometry in order to study the shape of a single protein molecule as it undergoes conformational change.

Although the aim stated above provides motivation for this project, the major focus of the work presented here remains on understanding the physics of the generation system itself. The immediate aims of this work are therefore to develop a capillary based HHG source, to achieve reliable operation and to further our understanding of the HHG process. The items chosen for inclusion here reflect the focus on source development and consequently chapter 2 begins the thesis with an explanation of some of the apparatus used in this work.

Chapter 2

Apparatus

In this chapter the apparatus and equipment that have been used as part of this project are described. This chapter incorporates discussion of theory concepts (relevant to the project and to later chapters) by using the underlying physics of the apparatus as an example.

Section 2.3 contains some discussion of the optical properties of materials (particularly how they differ for x-ray and visible light) whilst the main topic of the section is the optics and equipment that are used to manipulate the x-ray beam.

The topic of section 2.4 is the x-ray detectors used in this project, however, there is also discussion of how equipment and experiment design can affect the signal-to-noise ratio on measurements.

This chapter begins with a discussion of the laser source in section 2.1 and in section 2.2 the diagnostics systems that are used to maintain it.

2.1 Laser System

In this section the operation of the laser systems used will be described (towards the end of this project the laser system was upgraded). Subsection 2.1.1 will first give a brief overview of the concepts underlying this type of laser system. The subsequent subsections will then discuss the specifics of the original system in more depth. The final subsection 2.1.5 will detail the differences between the original and new systems, pointing out where design improvements have been made.

2.1.1 Overview and Schematic

In order to carry out extreme nonlinear optical experiments such as HHG a laser source capable of producing very large peak intensities is required. In practice this means that a pulsed laser system must be used:

For a laser with a given average output power (P_{avg}), the shorter the pulse duration ($\Delta\tau$) and the lower the repetition rate (f_{rep}) the higher the peak power (P_{peak}) will be:

$$P_{peak} \; \alpha \; \frac{P_{avg}}{f_{rep} \; \Delta\tau} \tag{2.1}$$

Typically mode-locking is used to produce a train of pulses from the laser cavity. Essentially the cavity is modified in such a way as to make light propagating in a short pulse receive less loss than the continuous wave (CW) light. If this is achieved then the laser can operate stably in a mode-locked state. All that is required is a little random noise (could be simply flicking a mirror mount) in order to get the laser to jump from CW into mode-locked operation - a train of pulses will be automatically generated. The repetition rate of such a mode-locked laser system is determined by the cavity round trip time. Typically the path length inside the cavity cannot be longer than a few metres and so oscillator repetition rates are usually approximately 100 MHz.

In order to generate the shortest possible pulses a large gain bandwidth is required; titanium doped sapphire (Ti:sapphire) currently offers the largest gain bandwidth ($\sim$ 200 nm FWHM) and in theory could support pulses as short as 3 fs. Ti:sapphire also has excellent thermal properties making it the gain medium of choice in most ultrafast applications.

Ti:sapphire based oscillators can easily achieve sub 50 fs pulses at a repetition rate of 100 MHz; because of this duty cycle such a laser will have a peak power approximately 5 orders of magnitude higher than a CW laser delivering the same time-averaged power. Supposing the average power is $\sim$ 1 W, the peak power will be in the region of 2×10^5 W

whilst the pulse energy will be approximately 10 nJ. Unfortunately for many experiments this will be insufficient and so further amplification is required.

As the laser is amplified, steps must be taken to avoid material damage and nonlinearity: Firstly Pockels cell gating can be used to reduce the repetition rate of the laser for example from $\sim$ 100 MHz down to 1 kHz. For a given peak power this has the effect of reducing the average power passing through the optics and hence the heating effect. Importantly reducing repetition rate to 1 kHz also makes it possible to use a Q-switched pump laser. The second and most crucial step is to temporally stretch the pulse prior to amplification (achieved by adding large amounts of quadratic spectral phase to the pulse using a grating stretcher). This greatly reduces the peak power which remained unchanged by the step down in repetition rate. The reduction in peak power is required in order to avoid nonlinearity as the beam propagates through the laser. Any nonlinearity that the beam experiences will add high order spectral phase to the pulse which will not be properly corrected by the compressor. A subsidiary benefit of the stretched pulse is that it overlaps temporally with more of the pump pulse - increasing pumping efficiency. The short pulse duration is reclaimed after amplification by removing the appropriate dispersion using a grating compressor. Because of the use of dispersion to lengthen the pulse this type of amplifier is referred to as a chirped pulse amplifier (CPA).

2.1.2 Kapteyne & Murnane Labs Ti:sapphire

In this project a Kapteyne & Murnane Labs Ti:sapphire laser (KML) was used to produce a train of pulses at 96 MHz each with a pulse energy in the nJ range. The pump energy for this laser was provided by a Spectra Physics Millennia Vs; which is a frequency-doubled neodymium-yttrium-vanadate laser (Nd:YVO_4) producing up to 5 W of CW radiation at 532 nm. Selected pulses from the KML beam are used to seed the CPA system described in subsection 2.1.3.

Figure 2.1 shows the internal configuration of the KML:

The optical path is as follows: The lens focuses the pump laser through the back of the first curved, dichroic mirror (CM1) and into the Ti:sapphire crystal. The pump beam then passes through the second curved, dichroic mirror (CM2) and onto a beam block having deposited some of its energy into generating a population inversion in the Ti:sapphire crystal. CM2 reflects and collimates the infrared light emitted from the crystal with high efficiency, sending the beam to the output coupler. The cavity round trip can be thought of as starting at the output coupler. The portion of the beam that is back reflected from the output coupler is focused by CM2 back through the crystal (where the beam gains more energy) before being collimated by CM1. The beam then

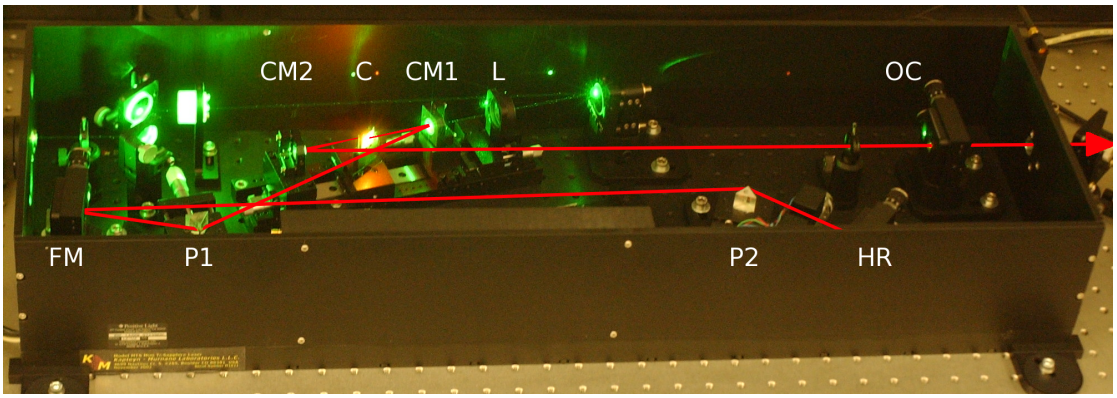

FIGURE 2.1: Internal layout of the KML Ti:sapphire laser cavity with the infrared path illustrated. CM1 and CM2 are the curved dichroic mirrors. C is the Ti:sapphire crystal. L is the lens. FM is the fold mirror. P1 and P2 are the prisms. HR is the high reflector. OC is the output coupler.

passes through a prism pair (via the fold mirror) before arriving at the high reflector. At this point the beam is reflected back along the same path until it completes the cavity round trip by arriving at the output coupler. The output coupler is of course a partial reflector and so part of the beam passes through, providing the output of the laser. The CW output power of the laser can be optimised by making small adjustments to the alignment of the high reflector and output coupler.

In order to change to mode-locked operation the Kerr lens effect is used. CM2 is translated a few mm in towards the crystal (reducing the CW output power considerably); this optimises the cavity for intense pulses since they will undergo self focusing inside the crystal and will thus diverge more quickly. It may also be necessary to translate the crystal slightly toward CM1. Light in short pulses should now see less loss than CW light as it passes through the crystal and the curved mirrors.

The role of the prism pair is intra-cavity dispersion control. In order for the cavity to support short pulses the net group velocity dispersion of a cavity round-trip must be as close to zero as possible. If this is not the case then the dispersion will tend to broaden the pulses in time; leading to longer output pulses, higher loss and possibly a rapid collapse back into CW operation. The prism pair can be arranged to produce anomalous second order dispersion in order to negate the normal dispersion introduced by the prisms themselves and the other optics in the cavity, see figure 2.2:

The first prism diffracts each spectral component of the beam through a slightly different angle - the second prism (inserted from the opposite direction) collimates the beam so that all spectral components run parallel but spatially separated (spatial chirp). At this point (between the second prism and the high reflector) the beam can be spectrally

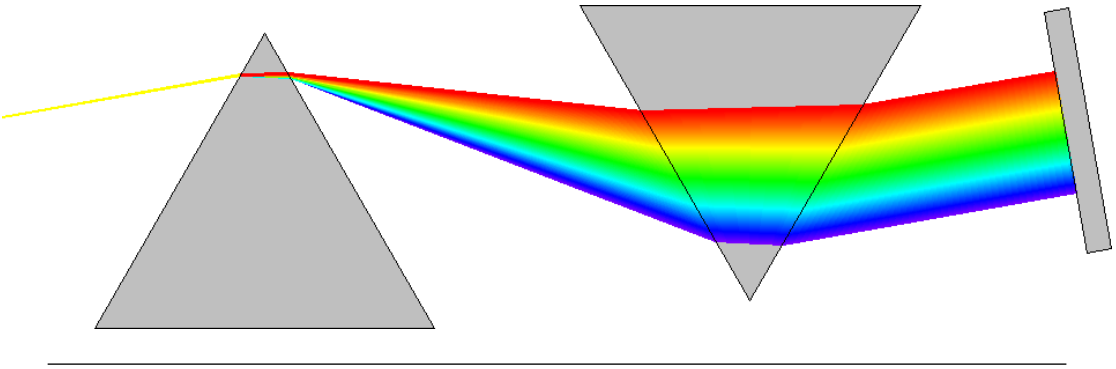

FIGURE 2.2: Schematic of a prism pair. Red light travels a geometrically shorter path and propagates faster in the prism (normal dispersion) however, the red light has to travel through much more material in the second prism which can produce anomalous dispersion.

filtered using movable shutter blades making it possible to adjust the wavelength range of the laser (in fact these were typically kept in the fully open position to ensure the maximum possible bandwidth). The high reflector sends the beam back along the same path through the prisms, recombining all of the frequency components spatially but with a modified spectral phase due to their difference in journey times. The precise amount of dispersion may be optimised by adjusting the prism insertion whilst watching the mode-locked output spectrum.

Once intra-cavity loss and dispersion are optimised, light in short pulses sees higher gain from the Ti:sapphire crystal than CW light. Once in this situation gently nudging the second prism usually provides enough of an impulse to initiate mode-locking. The laser output power should be observed to jump up (perhaps doubling) once mode-locking has been initiated; the spectrum will also broaden considerably indicating the presence of temporally short pulses at some point inside the cavity. Note: the exact duration of the pulses as they leave the output coupler is not necessarily optimised (only the net cavity dispersion is optimised). The pulses may therefore be rather longer than their bandwidth would suggest unless external dispersion compensation is employed. In this experiment the duration of these pulses is irrelevant since the pulse is to be stretched to an even longer duration as soon as it is passed into the CPA.

2.1.3 Spitfire CPA

A schematic overview (figure 2.3) and an annotated photograph (figure 2.4) help to illustrate the optical path through the Positive Light Spitfire CPA system.

The seed beam enters the Spitfire towards the top left of the image and is immediately passed through a Faraday isolator. This prevents any back reflections that may occur

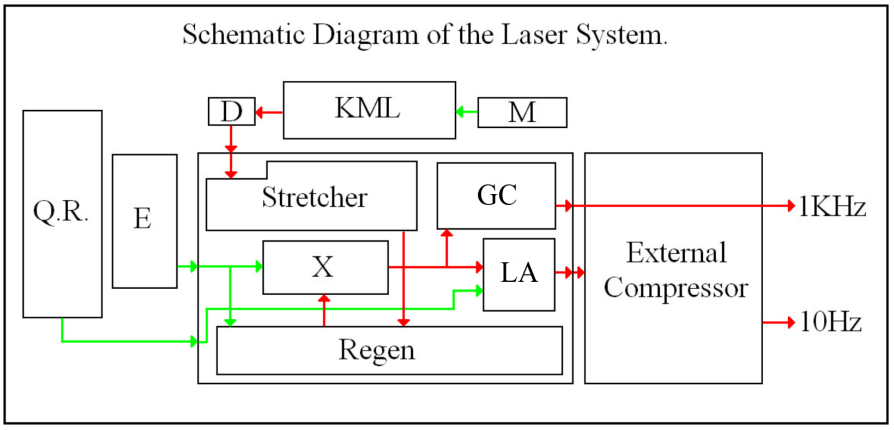

FIGURE 2.3: Schematic representation of the CPA laser system. Where M is the Millennia Vs, D is the Dazzler, E is the Evolution, QR is the Quanta Ray, X is the X-path amplifier, LA is the linear amplifier and GC is the grating compressor.

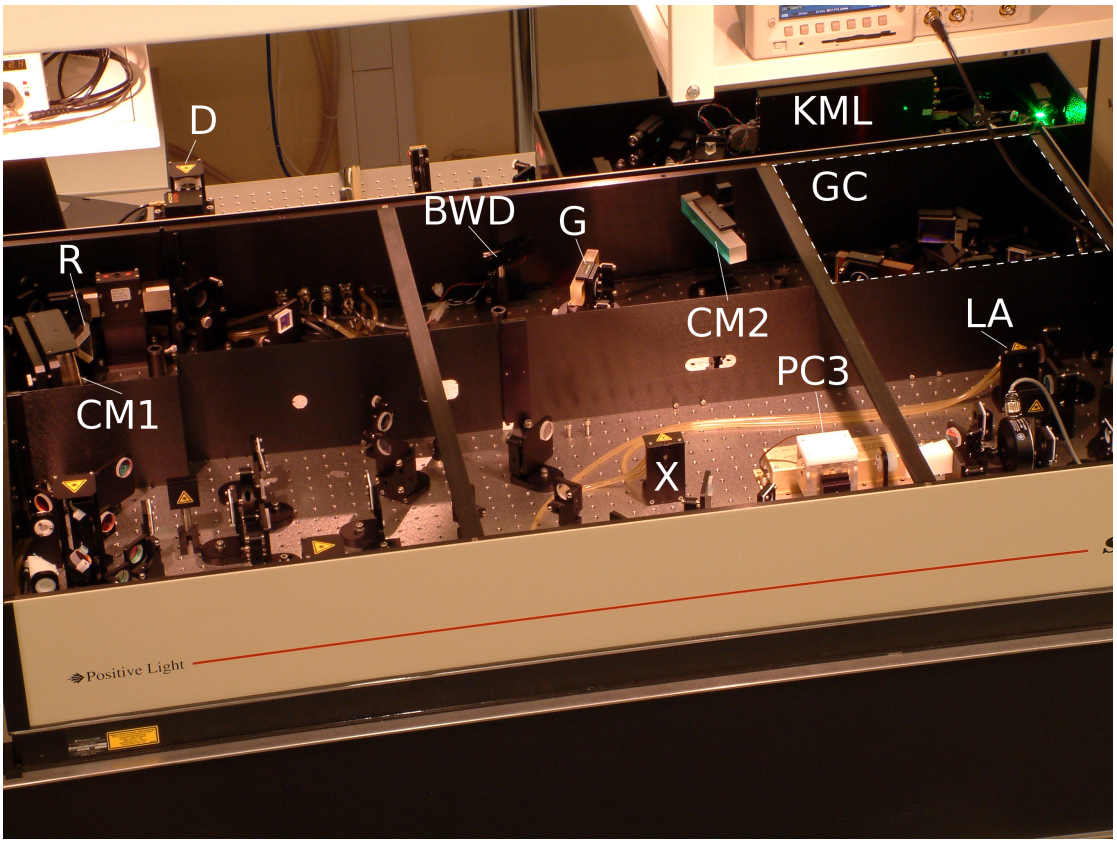

FIGURE 2.4: Annotated photograph of the Spitfire CPA system. The regen cavity runs along the front of the box and cannot be seen in this picture. D is the Dazzler. In the stretcher: R is the vertical retroflector, G is the grating, CM1 and CM2 are the curved mirrors and BWD is the bandwidth protection device. X is the crystal for the X path amplifier. PC3 is the third Pockels cell and LA is the crystal for the linear amplifier. GC is the compressor for the 1 kHz line.

in the Spitfire from upsetting the stable operation of the seed laser. The seed beam is then reflected from two steering mirrors before being launched into the stretcher:

The stretcher used is an all-reflective design; material dispersion of transmissive optics introduces unwanted cubic and quartic spectral phase to the pulse (which would not be corrected properly by the compressor). The stretcher is a two level design with a vertical retroflector (R) to produce a step in height and to fold the optical path. There are a total of four reflections - all from the same grating (G). Between bounces 1 and 2 and bounces 3 and 4 the beam is reflected from the large gold concave mirror (CM1) to the dielectric mirror (CM2) and then back to CM1. Between bounces 2 and 3 the spatially chirped beam is sent to the retroflector (R). After bounce 4 the step in height introduced by the retroflector allows the beam to be picked off by a small mirror.

The reflections from the curved mirrors CM1 and CM2 serve as a telescope in order to make the optical distance between the image of grating bounce 1 and grating bounce 2 effectively negative (same applies to bounces 3 and 4). Martinez realized that this effective negative optical distance between gratings could be used to create a stretcher the dispersion of which was exactly the inverse of a compressor made using the same grating [21]. In fact the gratings used in the spitfire stretcher and compressor are deliberately not identical; this is to compensate for dispersion that occurs inside the amplifier.

The second diffracted order from grating (G) is very weak but is sufficient to provide a signal on the bandwidth protection diodes (BWD); these sense any drop in seed laser bandwidth and stop the Pockels cells from switching the seed laser into the regen cavity if bandwidth is too low. This prevents possible damage to the compressor grating which would be caused by the drop in laser bandwidth and the consequent increase in spectral intensity.

The stretcher takes an input pulse duration of approximately 50 fs and adds sufficient positive quadratic chirp to extend the pulse to several ps. A theoretical upper limit on the stretch ratio is determined by the rep rate of the seed laser; i.e. if pulses are stretched up 1/rep rate then spectral components from adjacent pulses will begin to overlap in time. There are many other practical factors that limit the highest possible stretch ratio and in any case it is not desirable to stretch the pulse to a longer duration than the pump laser pulse.

The $\sim$ 96 MHz chain of pulses from the stretcher is horizontally polarised but this is rotated to vertical by a periscope with a 90° twist. The stretched pulse is now launched into the regenerative amplifier cavity by reflecting the beam from the Brewster cut face

of the Ti:sapphire rod. The cavity is defined by two high reflecting mirrors and contains two Pockels cells and two polarising beam splitters.

Once inside the cavity the seed beam passes through the first Pockels cell and then hits the first polarising beam splitter. Most of the seed pulses are dumped out of the cavity by the beam splitter, however, timing electronics fire the Pockels cell at 1 kHz in order to select individual seed pulses for amplification[†]. The now horizontally polarised pulse is trapped in the cavity where it is amplified progressively each time it passes through the amplifier. The build up of amplification can be observed by using an oscilloscope to display the signal from a fast photodiode positioned to accept the leak from one of the high reflectors.

Once $\sim$ 7 or 8 round trips have been made the amplifier begins to go into saturation; pulse energy stops increasing and additional round trips just add unnecessary dispersion. At this point the second Pockels cell is fired which rotates the amplified pulse to vertical polarisation and allows the second beamsplitter to extract the pulse from the regen cavity. Immediately after the regenerative amplifier the pulse energy is approximately 1.2 mJ.

After the regenerative amplifier there are two possible paths the light can take, switching between these is controlled by a third Pockels cell. Light passes through the third Pockels cell, through a half wave-plate, then arrives at a beam splitter. If the Pockels cell is not fired then the vertical polarisation is rotated to horizontal and passes through the beam splitter. If the Pockels cell is fired then the polarisation is rotated from vertical to horizontal and then returned to vertical by the wave-plate - the beam is therefore reflected from the beam splitter.

The direct path through the third beamsplitter leads the beam into a two-pass, x-path, external amplifier. This second Ti:sapphire crystal may be pumped at 1 kHz using part of the evolution pump laser[‡]. A beam block can be translated into the path of the external amplifier pump if additional amplification is not required; however, the beam still makes its x-path through the external amplifier Ti:sapphire crystal. Assuming the external amplifier is pumped the pulse energy is increased to $\sim$ 2.2 mJ.

Whether the external amplifier is pumped or not the beam is then sent to the compressor in order to remove the positive quadratic chirp (group delay dispersion, GDD) introduced by the stretcher. Before being sent to the compressor the beam is expanded

[†]There is a photodiode just outside the oscillator, the signal from this is frequency divided and passed to the timing electronics as is the evolution pump laser trigger out signal. The Pockels cell drive signals are synchronised with respect to both the pump and seed laser trigger signals.

[‡]The evolution pump laser generates 20 W of average power; this is split using a half-wave plate and polarising beamsplitter to provide $\sim$ 8 W to the regenerative amplifier and $\sim$ 12 W to the external amplifier.

with a telescope; this is to prevent damage to the compressor grating due to high power densities.

Like the stretcher the compressor is all-reflective and consists of just one grating in a two level arrangement with four passes. The beam enters the compressor at a low level and is dispersed by the grating. The spatially chirped beam propagates for a short distance and then hits the horizontal retroflector which returns the beam to the grating at the same angle it left. (The distance propagated on this arm determines the amount of negative GDD added to the beam). Upon its second reflection from the grating the beam is spatially chirped but collimated so that all spectral components run parallel; the vertical retroflector steps the beam up in height and folds the optical path, returning the beam to the grating. The return pass through the compressor undoes the spatial chirp that has been introduced to the beam and after the final bounce from the grating the beam passes over the top of the mirror that launched the beam into the compressor. Once optimised the pulse duration at the output should be approximately 35 fs. A fundamental limit on the shortest possible output pulse is defined by the spectral bandwidth of the amplifier; this necessitates the use of large optics where the beam is spatially dispersed. In practice the lowest possible pulse duration is further limited by imperfect cancellation of higher order dispersion.

If Pockels cell three is fired then the beam bypasses the external amplifier and compressor. Instead the beam makes a linear double pass through the third Ti:sapphire crystal. Pump energy for this linear amplifier is provided by a Spectra Physics Quanta-Ray. This laser is a Q-switched flashlamp pumped Nd:YAG laser which is then frequency doubled to 533 nm and produces almost 1 J pulse energy in a $\sim$ 10 ns pulse. The Quanta-Ray, and therefore the third Pockels cell, can only be operated at 10 Hz. The outcome is a 10 Hz beam at 780 nm with pulse energies upwards of 20 mJ.

In order to recompress this high energy 10 Hz beam the beam must be expanded considerably to avoid damage to optics. A second compressor is used; this works on the same principles as the compressor integrated into the amplifier box but the optics are considerably larger.

During work carried out for this project the regenerative amplifier was almost always used without the additional amplification stages. We found that the degradation of mode quality introduced by the x-path amplifier negated the increase in pulse energy it offered. The lower repetition rate and propensity to damage optical components made the 10 Hz line less practical to use in this HHG experiment and so this was almost never used.

2.1.4 Dazzler

The Dazzler[22] is an example of an acousto-optic programmable dispersive filter (AOPDF). This is an optical device that enables arbitrary control of laser pulse spectral phase by modifying the second and higher order dispersion in a section of the beam path. In the laser system used for this project the Dazzler is installed between the oscillator and the amplifier[†] as shown in figure 2.3.

The ability to arbitrarily control the spectral phase of our laser pulses makes some very interesting experiments possible and it also allows us to compensate for any dispersion that is inadvertently incurred by passing through optics etc.

A more mathematical explanation of how the Dazzler works is available in reference [23]. However the general premise is as follows:

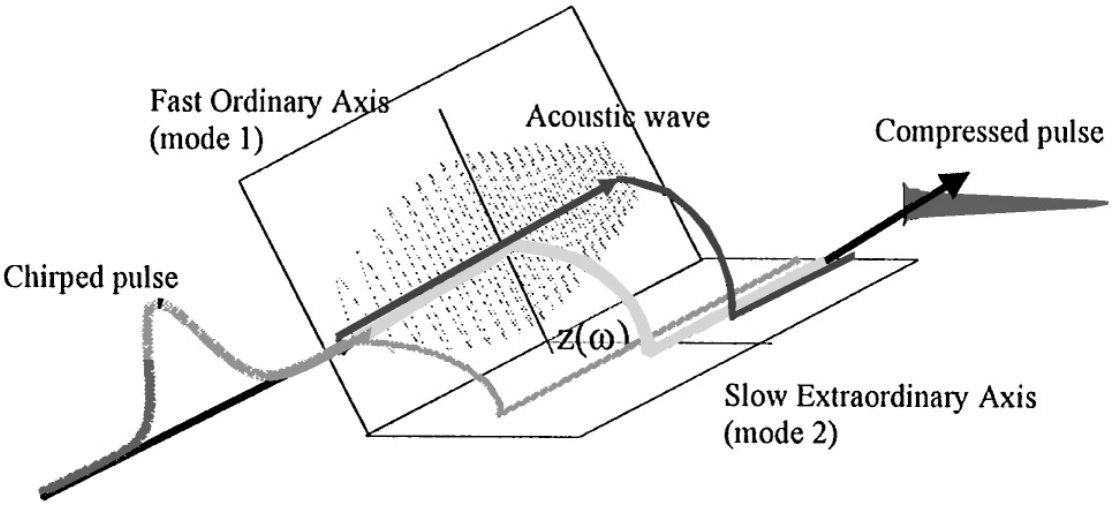

FIGURE 2.5: This diagram (taken from reference [22],) shows how a chirped pulse can be compressed using the Dazzler crystal.

A short laser pulse is launched into the Dazzler crystal in mode 1 and a radio frequency acoustic wave propagates through the crystal along the same axis. The acoustic wave modulates the crystal lattice spacing to create an acoustic grating with a spatial frequency proportional to the frequency of the acoustic wave. This acoustic grating can diffract light out of mode 1 and into mode 2 where it sees a higher refractive index and therefore propagates more slowly. The grating only diffracts one particular optical frequency (wavelength) efficiently for any given spatial frequency. It is therefore possible, by controlling the shape and timing of the acoustic wave, to diffract different wavelengths into mode 2 at different depths through the crystal. In this way the total time

[†]Dazzler crystals are now available with increased clear aperture, this allows them to be installed after the amplifier with the beam expanded so that the damage threshold of the crystal is not exceeded.

for each wavelength to propagate through the crystal can be controlled (see figure 2.5). Using this effect the dispersion (spectral phase) of the mode 2 beam can be arbitrarily controlled.

The largest possible delay between spectral components is determined by the length of the Dazzler crystal and the difference between the refractive indices of the two modes. The crystal used in this project is capable of a maximum programmable delay of the order of a few picoseconds.

The amplitude of the acoustic wave also allows control of how much light is diffracted into mode 2 so arbitrary filtering of spectral intensity is also provided. Note that it is not possible to create additional spectral content; only to change relative spectral intensities within the spectral envelope of the input pulse.

The output beams from modes 1 and 2 are spatially and angularly separated making it easy to select only the mode 2 beam to be sent to the experiment. Unfortunately the maximum throughput into mode 2 is approximately 50 % and this is further reduced to 25 % towards the edges of the Dazzler's 200 nm spectral bandwidth. Using a Dazzler therefore introduces a considerable loss of laser power, however, in the case of a CPA laser system this is of little importance since very little power is required to seed the amplifier.

2.1.5 Laser Upgrade

The KML oscillator described in section 2.1.2 was replaced with a Spectra Physics Tsunami described below 2.1.5.1. The Positive Light Spitfire CPA was replaced with a new 3 mJ Spitfire Pro system, the differences between this and the old system are outlined in section 2.1.5.2. The Millenia and Evolution pump lasers were retained and continue to be used.

2.1.5.1 Tsunami

Figure 2.6 shows the optical path through the Positive Light femtosecond Tsunami Ti:sapphire oscillator. Rather than being mounted on a horizontal aluminium breadboard (as was the case with the KML) the components of the Tsunami are bolted to a large cylindrical invar beam. This should improve laser stability in particularly preventing thermal heating from increasing cavity length or disturbing alignment.

The plane of the prisms has rotated from horizontal to vertical and so too must the pump laser polarisation; this was easily achieved by removing a half-wave plate that had been used with the KML.

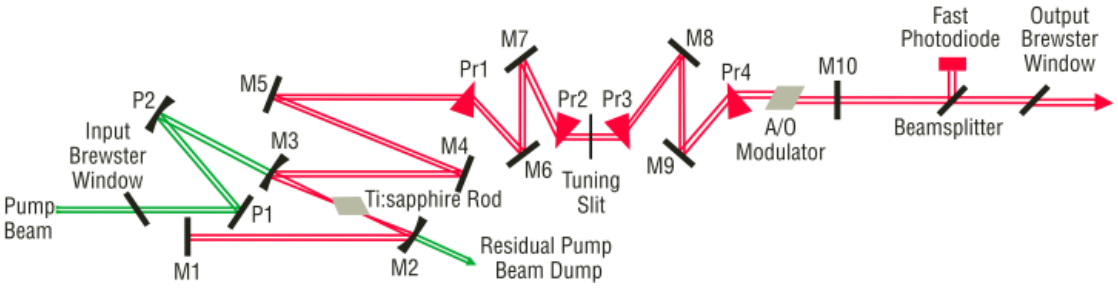

FIGURE 2.6: Optical path through a Positive Light femtosecond Tsunami Ti:sapphire oscillator. This figure was extracted from the Tsunami Brochure[24].

One of the main differences between the KML and Tsunami cavities is that the Tsunami incorporates an acousto-optic modulator in order to initiate and maintain mode locking; in the KML this was achieved by simply translating one of the prisms.

The second main difference is partly a result of the first - the KML used just 2 prisms whilst the tsunami contains 4. The extra dispersion compensation is required to counteract the chirp added by the acousto-optic modulator. This also allows the Tsunami to output a temporally short (< 35fs) pulse, whereas the KML had the bandwidth to support this but outputted a chirped pulse.

Additional coarse tuning of the cavity dispersion is provided by chirped mirrors which can be installed in the M6 to M9 positions as required. Fine tuning is provided by an external actuator which simultaneously controls insertion of prisms PR2 and PR3. Although a little extra complexity is added this offers the prospect of more accurate dispersion control.

In terms of specification the Tsunami and KML are roughly equivalent. The KML operated at $\sim$ 96 MHz whilst the Tsunami has a slightly longer cavity and operates at nominally 80 MHz. Output power varies greatly depending upon laser alignment, however estimating figures for average output power the KML usually produced $\sim$ 450 mW whilst the Tsunami averages closer to 350 mW. Spectral bandwidth also varies considerably however the KML usually produced $\sim$ 60 nm FWHM whilst the Tsunami is able to regularly make more like 80 nm FWHM.

One of the nicest features of the Tsunami is that adjustments of the high-reflector, output coupler, prism insertion and tuning slit are all possible with the laser housing

closed. With its Brewster cut windows the housing is almost completely sealed (gas purge is possible) helping to eliminate air currents and keep optics clean.

2.1.5.2 Spitfire Pro

In the new Spitfire CPA system a lot of the unused complexity has been removed. For example there is no 10 Hz line so the Quanta-Ray and the external compressor are no longer needed (this considerably reduces the footprint of the system, see figure 2.7). There is also no x-path amplifier so the Ti:sapphire crystal count is down from three to one. This reduced optical complexity has a positive effect on both laser stability and mode quality.

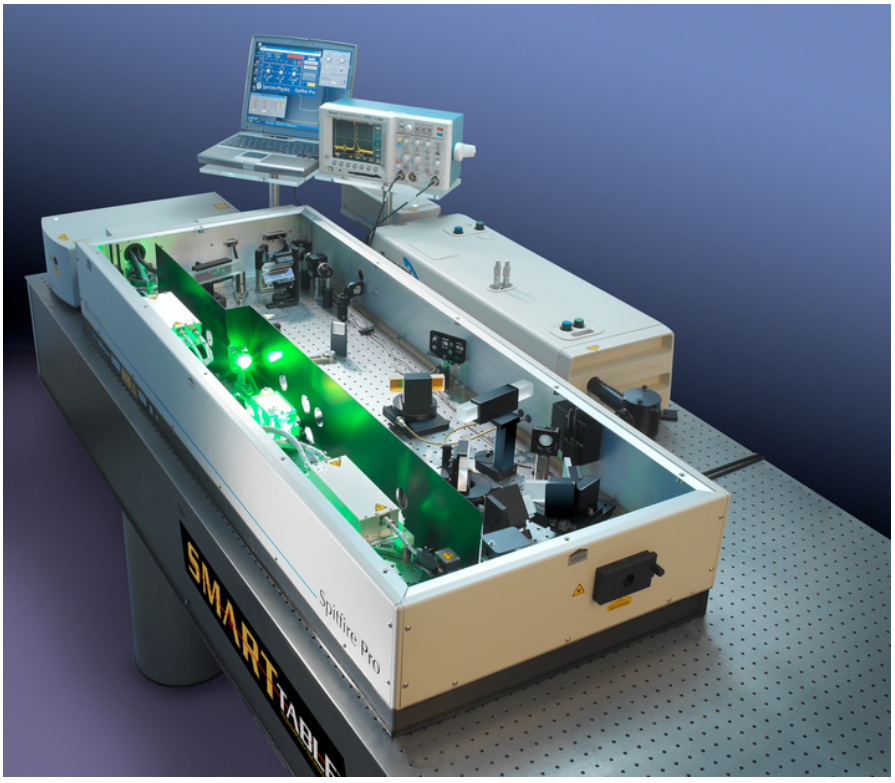

FIGURE 2.7: Photograph of Spitfire Pro CPA laser system, extracted from the spitfire pro datasheet[25].

Despite having just one Ti:sapphire crystal technological advancements have made it possible to produce a 3 mJ output pulse direct from the regen cavity (in fact the same system is also available in a 5 mJ version). The main advances that make this possible are discussed below:

The Evolution pump power is split 50:50 and used to simultaneously pump the Ti:sapphire crystal from both ends. This allows higher total pump energy to be used as it effectively divides the heat load between the two ends of the rod; lowering the chance of fracture.

Thermo-electric cooling of the crystal allows still higher pump power to be used without causing damage to the crystal[†]. This also reduces the radial heat gradient produced by the pump laser and the thermal lensing this causes. Unchecked thermal lensing will degrade the mode quality of the laser; this is usually countered by introducing spatial filtering irises into the cavity. However, spatial filtering in this way introduces additional loss to the cavity and reduces overall efficiency. Reducing the thermal lensing by cooling the crystal therefore improves both mode quality and amplifier efficiency.

Another significant step forward is that only one Pockels cell is used inside the cavity, this reduces the intra-cavity dispersion and loss. Whereas before separate Pockels cells were used to switch the pulse in and out of the cavity now the same Pockels cell is kept on to hold the pulse in the cavity, once the required number of round trips has been made the Pockels cell is switched off and the pulse is allowed out of the cavity.

The design of the stretcher and compressor seems to have remained unchanged.

[†]In order to provide efficient cooling and to prevent the danger of condensation forming on its facets; the Ti:sapphire crystal is held in a sealed container with Brewster angle input and output windows. A gentle flow of dried, filtered, air is passed through this container to prevent condensation and to ensure that dust particles are not deposited on the crystal.

2.2 Laser Diagnostics

HHG is a high-order nonlinear effect, it is therefore extremely sensitive to any fluctuation of the input laser properties - particularly peak power. Whilst average laser power may be tested very simply using a laser power meter (this is often a very helpful diagnostic) in a pulsed laser system it is also necessary to monitor the spectral, temporal and spatial properties of the laser very carefully. Spectrometers are also invaluable tools in the lab; however, both of the aforementioned instruments are very well known. In the following sections I will describe some of the specialised measurement techniques used in this project.

2.2.1 Autocorrelator

During this project three different autocorrelators have been used to measure the laser pulse duration at various stages through the laser system - they are described in the following paragraphs. The underlying concept of all of these autocorrelators is the same, however, there are slight differences in the execution which means that each operate optimally for slightly different input laser beams. An excellent discussion of the operation and the failings of autocorrelators can be found in Trebino's book [26].

Lightspeed Technologies Timewarp E750

This autocorrelator is used to measure the pulse duration exiting from the Ti-Sapphire oscillator that seeds the amplifier system. A Michelson interferometer type geometry is used with a 50:50 beamsplitter to produce two identical replica beams. The beamsplitter is very thin to keep dispersion to a minimum, however, this factor sets the lower limit of measurable pulse duration to be 15 fs for this device. At the end of each arm is a retroflector, the first of these is held fixed (although there is an adjuster to ensure the nominal path length of the two arms can be set the same). The second retroflector is mounted on a loudspeaker cone which is driven at 18 Hz in order to modulate the path length of this arm. The two replica beams are recombined by the beamsplitter and are then focused onto a GaAsP LED by a microscope objective. As the path length of the second arm is modulated successive replica pairs in the laser pulse train arrive at the detector with a different relative delay. At the two extremes of the loudspeaker travel the two replica pulses should be completely separated in time. (The throw of the loudspeaker therefore sets the upper limit of the measurable pulse duration; it is 500 fs for this device). The bandgap of the GaAsP LED is wider than the energy of a single photon from the oscillator and so two photon absorption is required in order to register

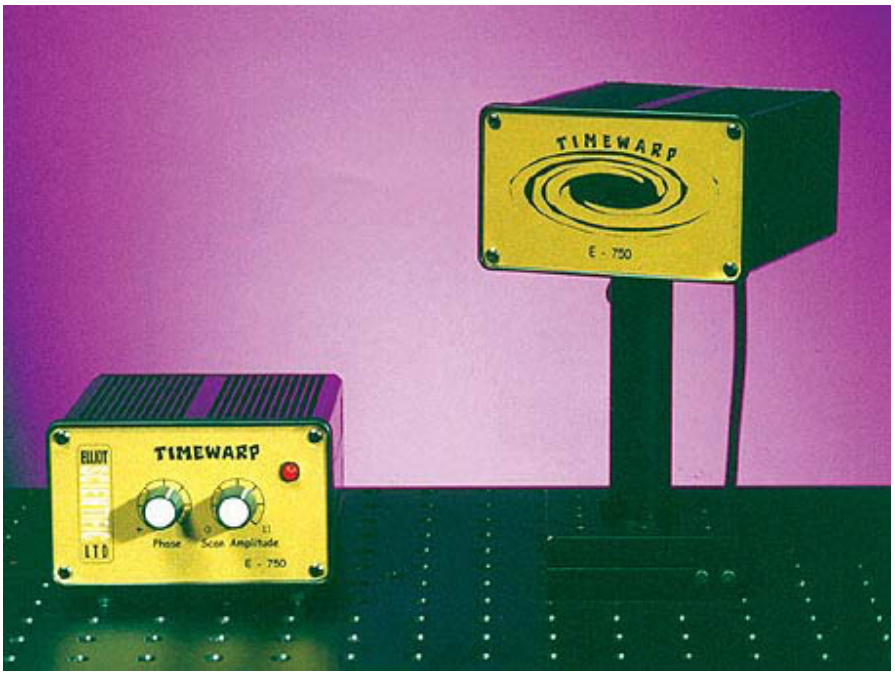

FIGURE 2.8: The timewarp E750 autocorrelator[27].

a signal. (The band structure of the diode therefore sets the spectral range of this device to be $700 - 1250$ nm). An oscilloscope is used to display the output signal of the diode (the electronics also provide a trigger signal). Since the time constant of the LED is short compared to the mirror scan rate, the result is a field autocorrelation as shown in figure 2.9.

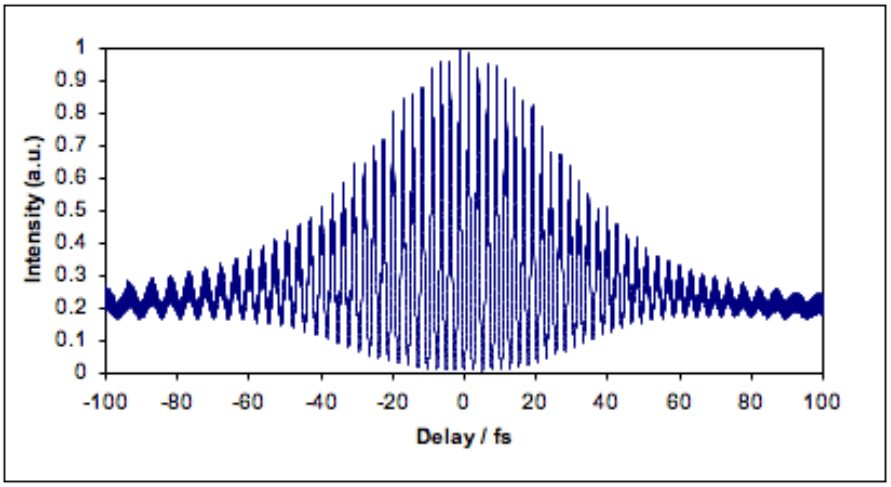

FIGURE 2.9: Interferometric autocorrelation of a 36 fs laser pulse.

The fringes, caused as the carrier waves of the pulse replicas step in and out of phase

with each other, can be used to calibrate the time axis of the oscilloscope (assuming the wavelength is known). I.e. the time calibration factor T_{ac} is calculated by counting the number of fringes $N_{fringes}$ occurring in a given time window T_{scope} on the oscilloscope:

$$T_{ac} = \frac{N_{fringes}\ \lambda}{c\ T_{scope}} \tag{2.2}$$

Where as usual λ and c are the central wavelength and group velocity of the laser being measured.

This autocorrelator also has a switch which adds a resistance in series with the LED. This increases the LED's time constant causing it to output a time averaged signal proportional to the square of the laser intensity. This produces an intensity autocorrelation such as that shown in figure 2.10.

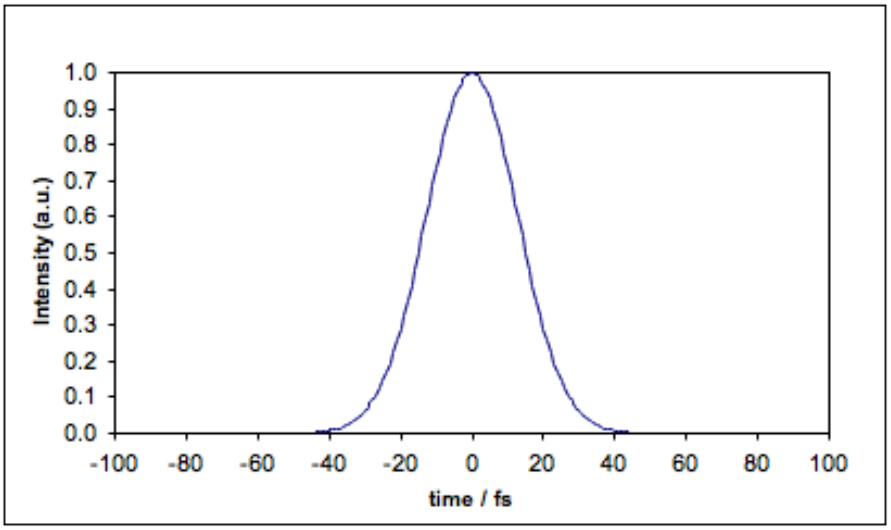

FIGURE 2.10: Intensity autocorrelation of an ideal Gaussian 36 fs laser pulse.

Once in this form the FWHM trace width (typically ms) can be read off using the oscilloscopes measurement tools. This can be converted into an autocorrelation width (typically fs) using the calibration factor calculated from the interference fringes.

The final step in the measurement is to multiply the autocorrelation width by the appropriate deconvolution factor in order to obtain a pulse duration in time. Unfortunately the temporal intensity profile is almost never known and must therefore be assumed. The factor for a pulse with a Gaussian time profile is 0.707, however experimenters are often tempted to assume a sech^2 pulse shape without justification simply because the factor of 0.648 indicates a shorter pulse duration.

The temporal resolution of this autocorrelator is quoted as $\sim$ 2 fs, however, this autocorrelator is only able to measure lasers with a repetition rate higher than 10 MHz, meaning it is suitable for the Ti-Sapphire oscillator (80 MHz) but cannot be used on the amplified laser (1 kHz). This limit is determined by the scan rate of the mirror and the time constant of the LED since the pulses must be frequent enough to produce a smooth time averaged intensity signal at all positions throughout the mirror scan.

APE Pulse Check 50

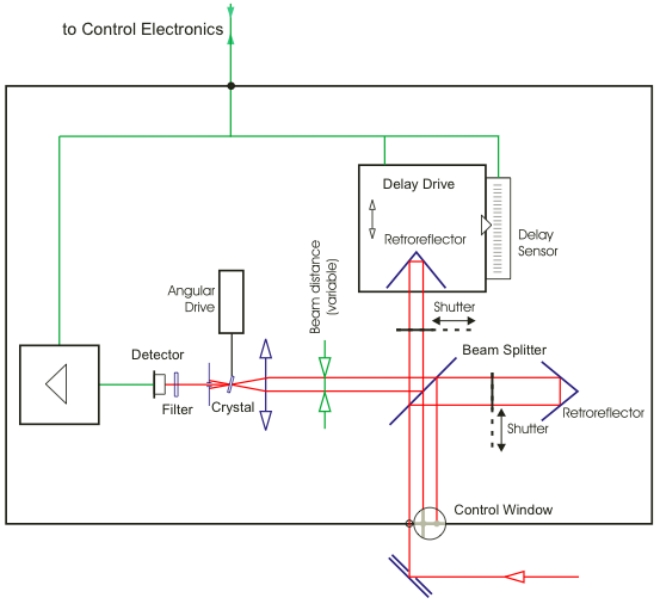

FIGURE 2.11: Schematic of the PulseCheck autocorrelator[28].

An APE pulse check autocorrelator was also used during this project. A schematic of this device is shown in figure 2.11. This device operates in a very similar fashion to the Timewarp E750 described above with two main alterations:

In this device the GaAsP LED is replaced by a second harmonic generation (SHG) crystal; the effect is the same, it produces a signal proportional to the square of the laser intensity. The use of an SHG crystal adds expense and complexity but it does offer some advantages. The crystal angle control allows phase matching conditions to be met over a wide wavelength range for any given crystal ($700 - 1100$ nm for the device used). Furthermore different nonlinear crystals may be inserted to further extend the operating range (in this case the filter, intended to block the fundamental, may also need to be replaced). Since the nonlinear response of the LED is no longer required this has also

been substituted for a far more sensitive detector. In the model used here the crystal and beamsplitter dispersion set the measurable pulse duration, lower limit to be $\sim$ 50 fs.

The second alteration is that the loudspeaker has been replaced with a spring loaded linear delay drive. This allows the scan range to be varied depending upon the pulse duration to be measured. Since the temporal resolution is quoted to be $\sim$ 1% of the scan range - reducing the scan range effectively zooms in on shorter pulses improving the temporal resolution. The full travel of the delay drive sets the upper limit of measurable pulse duration to be $\sim$ 500 fs. Additionally the delay drive is encoded and sends its signal to the control electronics, this eliminates the need to calibrate the time axis of the autocorrelation trace.

Importantly this device incorporates a triggered mode which allows it to operate on the 1 kHz amplified laser beam in both field and intensity autocorrelation modes. Since it is limited to pulse durations $<$ 50 fs this device can only be employed in locations where the pulse is expected to be chirped.

Positive Light Single Shot Autocorrelator

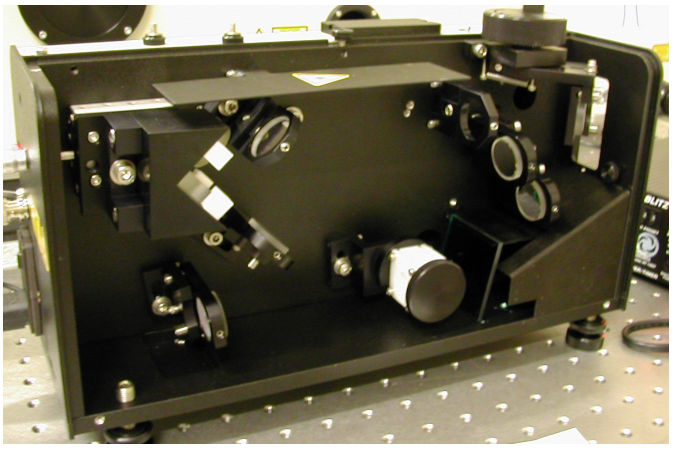

FIGURE 2.12: Photograph of the SSA autocorrelator.

This is the last autocorrelator used on this project, its internal configuration (see figure 2.12) at first appears rather different to the previous two but is functionally almost identical to the APE PulseCheck. The major difference with this autocorrelator is the fact that the two replica beams cross at an angle in the second harmonic generation crystal (see figure 2.13. Since the pulse wavefronts are tilted in opposite directions - different relative delays between the two pulses are encoded spatially across the crystal. A filter eliminates the fundamental allowing a CCD array to image the second harmonic signal; this gives the intensity autocorrelation directly from a single laser shot.

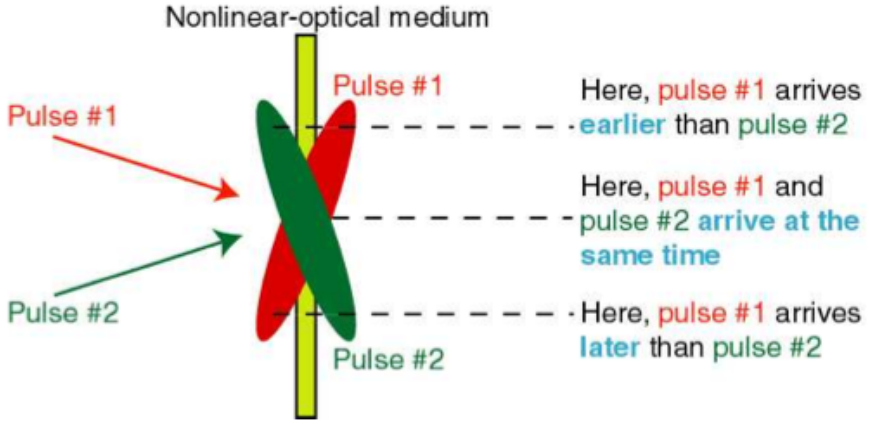

FIGURE 2.13: Diagram showing two replica fs pulses crossing in a nonlinear crystal, this was extracted from Trebino's GRENOUILLE tutorial[29]. The autocorrelation can be imaged onto a CCD array and thus collected for a single laser shot.

Note that a telescope is included in order to expand the central part of the beam; this helps to ensure uniform spatial intensity which would otherwise bias the intensity autocorrelation.

This autocorrelator can operate in an externally triggered mode allowing the 1kHz rep-rate amplifier beam to be monitored. A thin nonlinear crystal is used giving a wide angle tuning range (530−1600 nm) and allows this device to measure pulse durations down to ∼ 30 fs. The pulse duration upper limit of ∼ 500 fs is imposed by the maximum geometric wavefront tilt; however, this upper limit can be extended by using a grating to increase the wavefront tilt - in this case pulses of up to 5 ps duration may be measured. These specifications make this autocorrelator particularly suited to monitoring the output of the amplifier system.

2.2.2 FROG

Throughout this section the grenouille chapter from Trebino's book[26] was used for reference.

In this project a Swamp Optics 8:20 Frequency-Resolved Optical Gating device (FROG) was used. Such a device combines both temporal and spectral measurements the results of which are passed into an iterative algorithm which is able to retrieve the laser's complex E-field. This information can be outputted from the control software in a variety of ways; i.e. as a spectrum and autocorrelation, as spectral intensity and spectral phase, as temporal intensity and temporal phase, or of course as a simple FWHM pulse duration.

The particular FROG device used here bears a somewhat dubious acronym; it is of the GRENOUILLE type (GRating Eliminated No-nonsense Observation of Ultrafast Incident Light Laser E-fields). Figure 2.14(a) illustrates the underlying concept of the GRENOUILLE. The same input laser beam undergoes different optical operations in the horizontal and vertical planes. In the horizontal plane (top view) a Fresnel bi-prism causes the two halves of the beam to cross over inside the SHG crystal. The second harmonic light is then imaged onto a CCD camera. This is effectively the same setup used in the SSA (see section 2.2.1) and so produces an intensity autocorrelation across the horizontal axis of the camera. In the vertical plane a cylindrical lens (curved mirror in the actual device) is used to focus the laser beam onto the relatively thick (1.65 mm) SHG crystal. The thickness of the crystal means that the peak of the phase-matching angle tuning curve will be very narrow; however since the laser wavefront is curved when it reaches the crystal an angular spread of laser k-vectors is present. The outcome of this is that the crystal produces an SHG signal that as a function of angle is proportional to the spectral autoconvolution of the input laser pulse. This is imaged onto the same camera that collects the autocorrelation signal and thus produces a 2D representation of the temporal and spectral profiles of the input laser pulse (an SHG FROG trace) see figure 2.14(b).

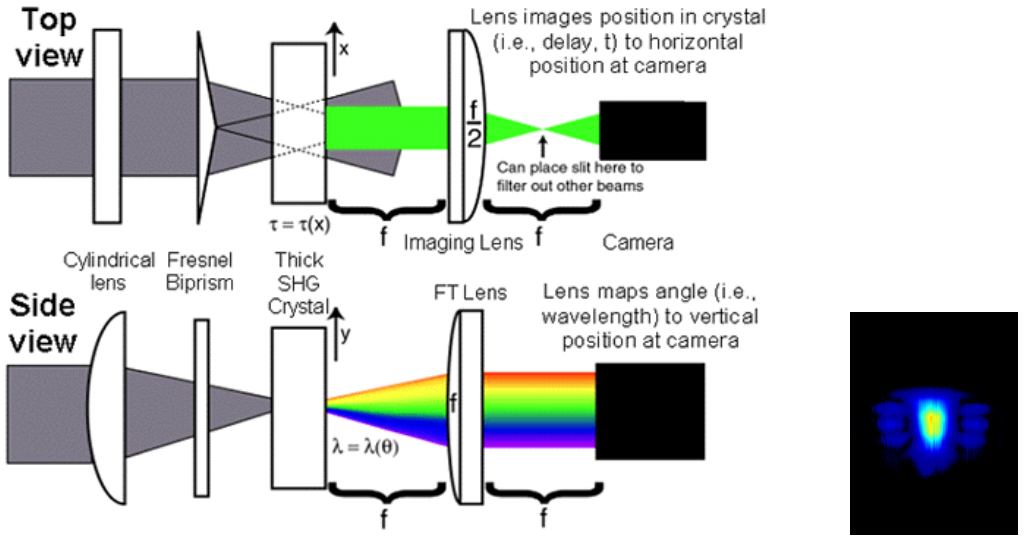

FIGURE 2.14: (a) Simplified diagram showing the main components of the GRENOUILLE in order to illustrate the concept, taken from Trebino's GRENOUILLE tutorial[29]. (b) Example frog trace for a 45 fs laser pulse, delay axis horizontal, spectral axis vertical.

It is clear that a compromise has to be made between a very thick crystal (for better spectral resolution) and a very thin crystal (for less dispersion and therefore the ability to measure shorter pulses). Since the dispersion of the beta barium borate (BBO) SHG

crystal changes with wavelength so too does the measurable pulse duration range. At 800 nm the 8-20 GRENOUILLE can measure $20 - 200$ fs pulses whilst at 1100 nm the range is $15 - 150$ nm.

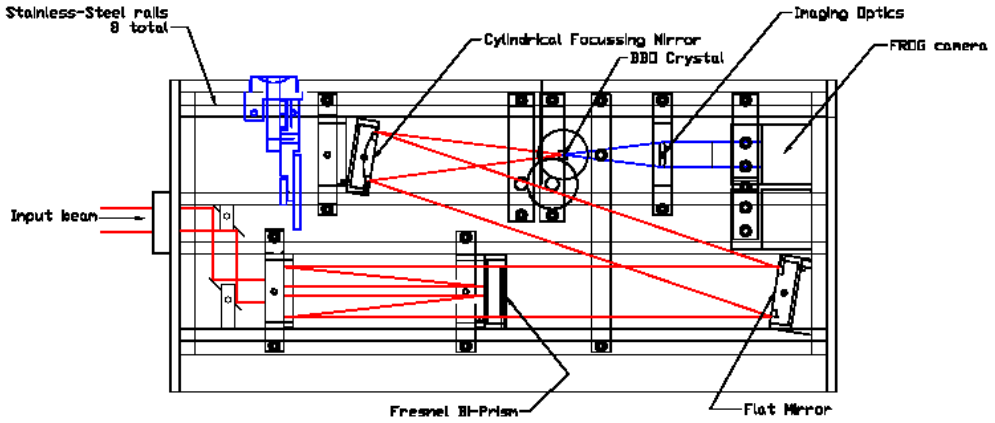

FIGURE 2.15: Schematic showing the internal configuration of the GRENOUILLE[29].

It is also clear that the GRENOUILLE relies heavily upon the beam being spatially symmetric with a flat intensity profile. If this is not the case the two beams created by the bi-prism will not be true replicas and errors will be produced in the autocorrelation. Similarly the spectral autoconvolution can be biased by spatial intensity variations on the input laser. It is for this reason that the GRENOUILLE incorporates a Cassegrain telescope to expand the central region of the beam, see figure 2.15.

Assuming a flat spatial intensity profile the GRENOUILLE trace should always be symmetric in the time axis, just as the single shot intensity autocorrelation must be symmetric. This leads to an ambiguity in the time axis associated with the retrieved pulse, which in turn leads to ambiguity in the sign of the spectral phase. This is easily solved by inserting some medium known to have normal dispersion (i.e. a block of glass) into the beam whilst observing the effect on the measured pulse duration.

In order to provide the best possible data to the reconstruction algorithm the intensity of the signal must be carefully controlled in order to fill the dynamic range of the camera without saturation. The background light level must also be minimised and subtracted. One particular requirement of the reconstruction algorithm is that the entire FROG trace must be recorded. In order to achieve this the measured FROG trace should appear completely surrounded by a dark, background level signal (this assumption is incorporated into the reconstruction algorithm). If the FROG trace extends too close to the edge of the camera or is clipped in any way then the results of the reconstruction will be inaccurate.

Since the GRENOUILLE uses an iterative pulse retrieval algorithm it is possible for this to become stuck in local minima. In order to check against this a visual comparison between the measured and retrieved frog traces should be made. If the two traces appear rather different then the algorithm may have located a local minima. In order to escape the local minima and collapse to a more accurate reconstruction it may be necessary to reset the FROG algorithm.

Numerically the FROG error is defined simply as the RMS difference between the measured and retrieved FROG traces; this can provide misleading results. For example if the camera background level is incorrectly adjusted the measured frog trace may appear surrounded by a uniformly dark background and low intensity FROG signal may be clipped. The reported FROG error in this situation would be very low but the reconstruction would be less accurate than that of a FROG trace with a correctly adjusted background level.

The GRENOUILLE is the most frequently used pulse duration diagnostic in the lab thanks to its simple alignment procedure and the speed with which results can be recorded.

2.2.3 SPIDER

Spectral Phase Interferometry for Direct Electric-field Reconstruction (SPIDER) is an alternative self referencing technique for measuring the spectral phase of an ultrafast laser pulse. A major advantage of this technique is that the phase-retrieval algorithm is not iterative (such as the one used by FROG). In addition, since the algorithm consists mainly of fast Fourier Transforms it can complete very rapidly.

Figure 2.16 shows the optical arrangement inside the Del Mar Photonics Avoca SP-10 SPIDER that is used on this project. The optical path is as follows:

When the input laser pulse arrives at the thin etalon (ET) the beam is split into two. Firstly the light reflected from the two interfaces of (ET) produces a double pulsed beam consisting of two replica pulses with a time delay set by the thickness of (ET). The second beam consists of the pulses that pass straight through (ET); these pulses are then stretched in time by adding large amounts of second order dispersion in a grating stretcher. The two beams are then aligned so that they run parallel (but at different heights) before hitting a curved mirror. The curved mirror focuses the two beams so that they cross at an angle to each other in the SHG crystal. Because the stretched pulse is highly chirped each frequency component occurs at a different time within the pulse. The two time delayed replica pulses will therefore each interfere with

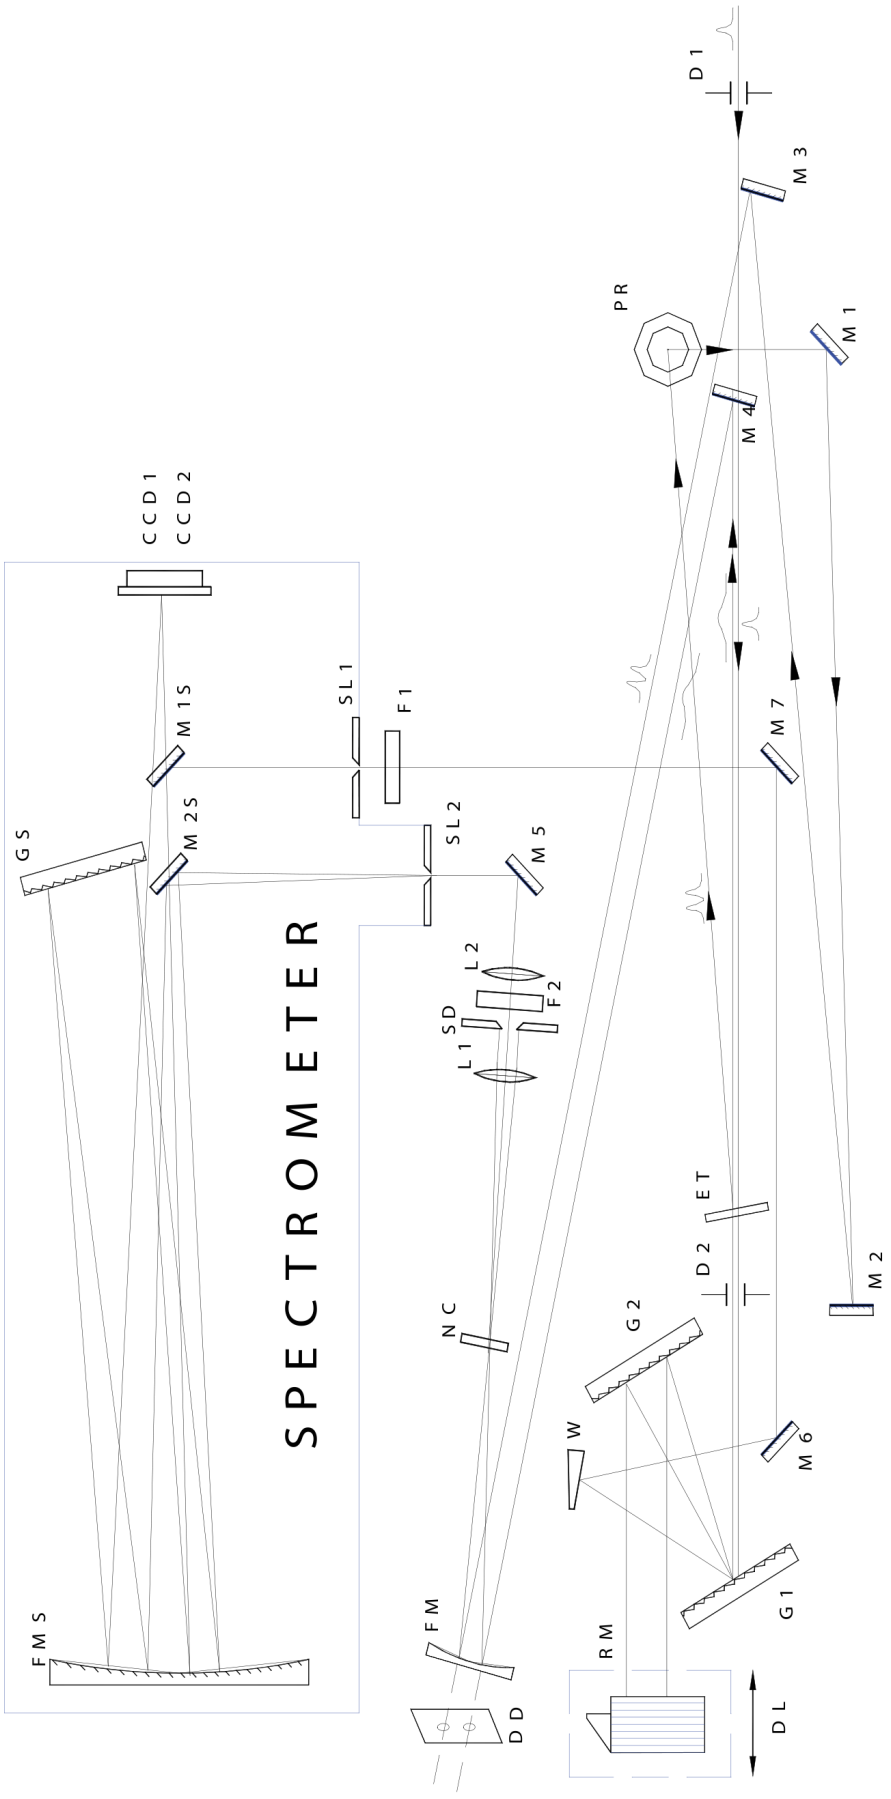

FIGURE 2.16: Schematic showing the optical path inside the Avoca SP-10 SPIDER - reproduced from the users manual.

a different frequency component of the stretched pulse. Second harmonic generation in the crystal produces two up-converted pulse replicas (still delayed in time) but now with a frequency shear. (Note that the stretched pulse should be sufficiently highly chirped that its instantaneous frequency is approximately constant over the duration of each pulse replica). A filter immediately behind the crystal blocks the infrared beams but allows the second harmonic light to pass. A vertical translating slit is then used to send either the reference signal or the spider signal to the spectrometer: If the slit is on-axis with the centre of the two crossing beams then the spider signal is recorded on the spectrometer (the interference between the two beams). The signal measured ($S_S(\omega)$) is related to the input laser pulse by expression 2.3:

$$S_S(\omega) = |\vec{E}(\omega)|^2 + |\vec{E}(\omega+\Omega)|^2 + 2|\vec{E}(\omega)||\vec{E}(\omega+\Omega)|\cos[\psi(\omega+\Omega) - \psi(\omega) + \omega\tau] \quad (2.3)$$

Where $\vec{E}(\omega)$ and $\vec{E}(\omega+\Omega)$ are the electric field of the up-shifted pulse replicas - having a frequency shift of Ω relative to each other. Similarly in the third term $\psi(\omega)$ and $\psi(\omega+\Omega)$ represent the spectral phases of the two up-shifted, frequency-sheared pulses. The third term also includes $\omega\tau$ which represents a linear phase change due to the time delay τ between the two pulse replicas.

If the slit is raised then the second harmonic light generated solely by the double pulsed beam is recorded (this is the reference signal). In this case the two pulse replicas are not spectrally sheared (since the stretched pulse is not present) and so the expression for the interferogram becomes:

$$S_R(\omega) = 2|\vec{E}(\omega)|^2(1 + \cos[\omega\tau]) \quad (2.4)$$

Since there is no frequency shear between the two pulses the unknown spectral phase terms drop out of the equation. The fringe pattern produced now depends only on $\omega\tau$.

Taking the Fourier transform of the spider and reference interferograms converts them into the time domain where the signals can be filtered by selecting only the component centred at $+\tau$. The inverse Fourier transform then converts back to the frequency domain. For the spider signal this yields a signal related to the phase terms in equation 2.3. In the case of the reference beam the linear phase term $\omega\tau$ is recovered; this can then be subtracted from the filtered spider signal leaving $\psi(\omega+\Omega) - \psi(\omega)$. This is the relative phase difference between successive pairs of frequency components separated by Ω. The spectral phase can be recovered by concatenating these differences.

The outcome of this algorithm is several spectral phase data sets each sampled with a point spacing of Ω. (The number of data sets is determined by the number of pixels that traverse each Ω sampling step). Unfortunately the phase relationship between each

data set is not known and must be assumed constant or interpolated using some fitting routine. Although the SPIDER is therefore able to output spectral phase data at the resolution of the spectrometer features smaller than the true sampling frequency Ω may simply be the result of errors in recombining the data sets and consequently should be dealt with cautiously.

The device also contains another spectrometer to provide an independent measure of the spectral intensity at each of the sample positions.

The optical path is slightly more complex than that of the GRENOUILLE although day to day alignment is still trivial.

2.2.4 M^2 Meter

The first few paragraphs below explain why laser mode profile is so important and what the M^2 property is. Some existing schemes for measuring M^2 will be briefly described before a more in-depth discussion of the novel M^2 meter that was designed and made as part of this project. Siegman's tutorial[30] was used for reference throughout this section.

In any experiment that depends upon peak laser intensity the mode quality of the input laser beam is very important. Any deviation from a smooth, radially symmetric Gaussian intensity profile and the beam propagation will not be in agreement with Gaussian beam theory; this typically results in a focal spot that is larger than would otherwise be predicted and hence lower peak intensity. Even quite subtle changes in the laser mode, for example clipping of the beam as it passes through an aperture or reflection from a scratched mirror, can be enough to significantly increase focal spot size and/or to disturb the intensity profile at focus.

HHG experiments are not only highly sensitive to peak laser intensity (generation being approximately proportional to I^5) in this experiment we are also attempting to launch into a specific mode of the generation capillary (see section 3 for more details). This places extremely high tolerances on the phase profile as well as the intensity profile produced at focus. In fact capillary throughput can be observed to fall from 95% to less than 10% as a result of laser mode disruption caused by mirror burn damage. As the input laser spot size is increased, damage to the input facet of the generation capillary is also accelerated; further reducing coupling efficiency. The necessity for a quick and easy method of accurately monitoring our input laser beam quality is therefore apparent.

The quality of a laser beam can be quantitatively expressed in an M^2 value; sometimes also referred to as the propagation constant of the beam. This M^2 value has a real

physical interpretation; it is the factor by which the divergence of the beam is increased as a result of the degradation of the intensity profile. I.e. if two beams have the same wavelength and spot size but one has an M^2 of 1 and the other has an M^2 of 2 once into the far field the radius of the $M^2 = 2$ beam will have increased by twice as much as that of the $M^2 = 1$ beam. This situation is illustrated in the upper plot of figure 2.17 and can be expressed as follows:

$$M^2 = \frac{\omega_{0R}.\theta_R}{\omega_0.\theta} \tag{2.5}$$

Where ω_{0R} is the measured beam waist radius and θ_R is the measured far field divergence ($\omega_R(z)/z$ for $z \gg$ the Rayleigh range). ω_0 and θ are the beam waist and far field divergence for a hypothetical perfect Gaussian beam whose waist is M times smaller than but in the same location as that of the measured beam. This hypothetical beam is illustrated in the lower plot of figure 2.17 - the radius of the underlying Gaussian is everywhere M times less than the 'real' beam.

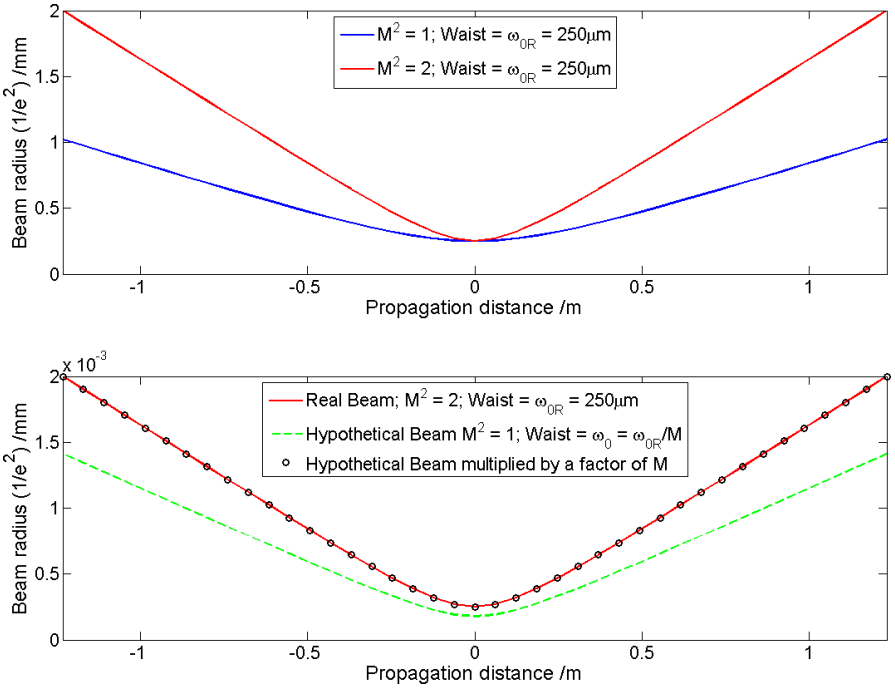

FIGURE 2.17: Illustration of Gaussian beam propagation with the addition of M^2. Upper plot: beam waist is kept the same, in the far field $M^2 = 2$ beam has an M^2 larger radius than $M^2 = 1$ beam. Lower plot: $M^2 = 2$ beam and its hypothetical underlying Gaussian. Hypothetical beam has $M^2 = 1$ and is everywhere M times smaller than the 'real' beam.

Since it is not possible to directly measure hypothetical objects, M^2 values are usually calculated by fitting beam propagation equations to measured beam profiles either side

of a waist. This approach is taken in several M^2 meters that are now commercially available.

The main difference between currently available M^2 meters is in the method by which the beam profile is measured at each sample position along the beam path.

- Photodiode behind translating slit: A measurement of the laser beam width is made by analysing the photodiode signal as the slit is translated across the laser beam. Often two orthogonal slits are used and these may be mounted either on a linear translating probe or on a rotating drum. In this method the width resolution is limited mainly by the width of the slit. Unfortunately this method is currently only suitable for CW or very high rep rate lasers because at lower rep rates the beam intensity will modulate whilst the slit is being translated. A possible solution to this would be to sync the slit translation mechanism (with a slight offset) to a frequency divided laser trigger signal. This would, however, significantly increase the complexity of the device and as yet this approach has not been taken on a commercial product.
- CCD camera beam profiler: In this type of design width resolution is limited by the CCD pixel size and is typically lower than that of a translating slit device. In order to get an accurate measure of beam width the beam must span several CCD pixels. This means that comparatively longer focal length lenses must be used to produce the beam waist thus increasing the amount of bench space taken up by the device.

In either case this beam profile measurement should be performed at several positions relative to a fixed focusing lens - sampling particularly frequently within the Rayleigh range. Translating the focusing lens introduces errors unless the beam is perfectly collimated so typically either the beam profiler itself or a retroflector is translated (either manually or on an automated stage) to produce the required change in path length.

At the time this work was carried out none of the commercially available M^2 meters claimed to be able to cope with kHz rep rate lasers. We therefore began designing our own system.

The novel M^2 meter described here uses a fixed focusing lens and a tilted etalon to create a series of spots on a CCD camera. Each spot samples the beam intensity profile at a different z-position through the waist, see figure 2.18. Using the CCD camera approach has the advantage of performing a time average over the integration time of each camera frame. Since no moving parts are required all problems due to the 1 KHz laser repetition rate are eliminated. Additionally, since all sampling positions are

measured simultaneously the technique is potentially single shot. An online computer algorithm analyses the data from the CCD camera, performs a fit to the measured beam profile, and is able to calculate the M_x^2 and M_y^2 values for the beam. (In the case of non-radially symmetric beams it is common to quote M^2 values for two orthogonal directions).

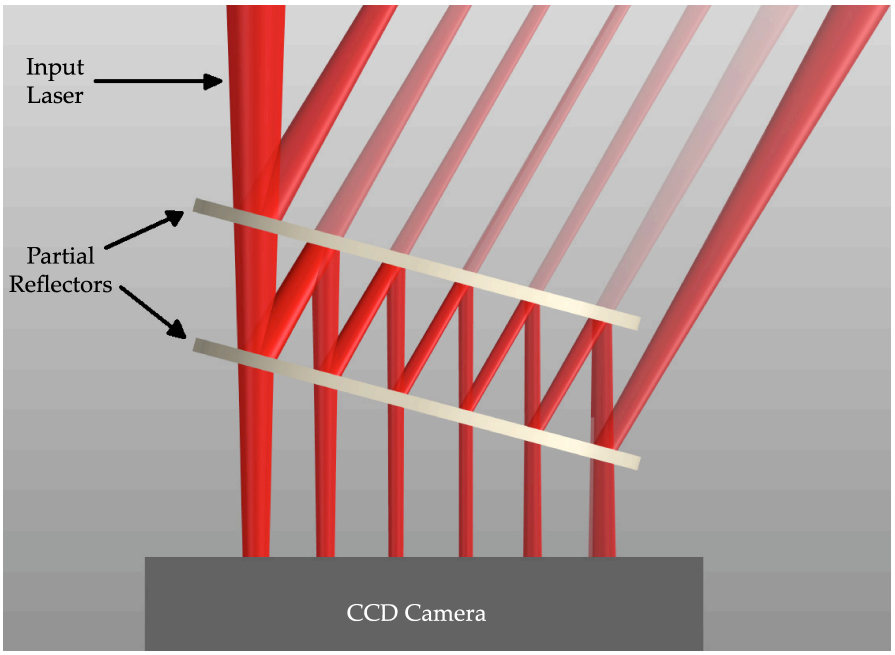

FIGURE 2.18: Diagram of M^2 meter.

Selecting appropriate reflectivity for the two partial reflectors that make up the etalon is critical. High reflectivities limit the amount of light transmitted to the camera but they also minimise the drop in power for each subsequent bounce. After carrying out computer simulations we opted for a 99% reflector as the first plate and a 95% reflector for the second plate. Using the 95% reflector reduced the amount of power that was required in the beam pick off and still left sufficient signal for an accurate measurement in the final bounce.

Lens selection is also very important and has an impact upon how many samples fit onto the camera - too short a focal length and it will be impossible to accurately resolve the beam at focus - too long a focal length and the beam profiles will become so large that they overlap. A focal length of 0.5 m was found to offer good performance.

The number of beam profile samples that can fit across the camera has an impact upon the accuracy of the M^2 measurement. A camera with native resolution of $1628 \times 1236 \times 10$ bit was used. It was found that fitting five beam profiles across the camera offered a good compromise This was sufficient to obtain good M^2 values, permitted real time

operation of the computer algorithm and allowed beams with larger M^2 values to be measured without causing adjacent beam profiles to overlap.

Alignment of the M^2 meter is relatively straight forward. A lens is positioned in the beam pick off so as to produce a focus inside the M^2 meter. There are then four adjustments that must be made: The separation of the two partial reflectors should if possible be set to ensure that the path length over which the beam profile samples are taken encompasses the Rayleigh range for the beam. (This separation must be entered into the computer fitting algorithm). The M^2 meter should then be translated laterally until the first beam profile is positioned at the edge of the CCD sensor. The angle of the two partial reflectors should be adjusted until the first five beam profiles are evenly spaced across the sensor. The z position of the M^2 meter should then be set so that the beam waist occurs at the position of the third beam profile.

The even separation and horizontal alignment of the profiles can be used as a visual indicator that the two partial reflectors are parallel. The back reflection from the first reflector plate can be used to accurately measure the angle of the reflector plates to the input beam. (This information is also required by the computer fitting algorithm).

If accurate measurements of the beam width are to be made, it is imperative that the camera does not saturate for any of the beam profiles. Saturated beam profiles typically result in measured beam widths that are larger than the actual beam. In fact the exact measured width would depend upon the degree of saturation. It would be impossible to accurately recover an M^2 value from data that has been corrupted in this manner.

In the following section I will describe the algorithm employed by the M^2 meter software:

In order to be able to calculate M^2 rigorously it is very important to measure the width of each intensity profile correctly. Ideally second moment width measurements should be made. Assuming an arbitrary intensity profile $I(x, y)$, the first moment x_0 (centre of mass) may be calculated as follows:

$$x_0 = \frac{\int_\infty^\infty x \int_\infty^\infty I(x,y)\ dy\ dx}{\int_\infty^\infty I(x,y)\ dx\ dy} \tag{2.6}$$

This can then be used to calculate the second moment (equivalent to the square of the variance) of the beam:

$$\sigma_x^2 = \frac{\int_\infty^\infty (x - x_0)^2 I(x,y)\ dx\ dy}{\int_\infty^\infty I(x,y)\ dx\ dy} \tag{2.7}$$

This definition for the width of the beam will work equally well for any arbitrarily shaped beam. Unfortunately it is rather problematic to implement in a practical system. The integrals in equations 2.6 and 2.7 although very easy to calculate numerically are particularly sensitive to background noise. Any stray back reflection or change in ambient light conditions is likely to drastically alter the measured width.

As a more stable backup a routine has also been included that fits Gaussians in x any y to each of the intensity profiles. The $1/e^2$ width of each of these Gaussians is then used in the beam profile fitting algorithm. Provided the laser beam profiles to be measured are approximately Gaussian and that the same definition is employed at all points throughout the algorithm this does not result in an appreciable error.

The beam profile fitting algorithm simply uses a nonlinear least squares fit to the Gaussian beam propagation equation to determine the best fit value for M^2. The propagation equation used is shown below:

$$\omega_R(z) = \omega_{0R}\left[1 + \left(\frac{z\lambda M^2}{\pi\omega_{0R}^2}\right)^2\right]^{1/2} \tag{2.8}$$

Figure 2.19 shows an example of a fit made using this equation. The black markers show calculated beam radii for an $M^2 = 1.2$ beam with $\omega_{0R} = 25$ μm and $\lambda = 780$ nm. The red line shows a nonlinear fit of equation 2.8 to the black markers. In the absence of noise the fit repeats the simulated data exactly. The blue lines above and below the red one show the effect of adding or subtracting just 0.2 from the M^2 value. The considerable beam profile deviation that results from a small change in M^2 value demonstrates that this fitting technique is potentially an accurate means of measuring M^2.

The nonlinear fitting algorithm employed in figure 2.19 was incorporated into the M^2 meter software which was developed as part of this project. In this software the fit is applied directly to the x and y beam widths extracted from the CCD image. A screenshot of the M^2 meter software is shown in figure 2.20; the x and y measured profiles can be seen on the right. A simple "block beam" background subtraction feature was added to the software to improve the accuracy of the width measurements. The live CCD image is displayed on the screen in order to assist in alignment. Saturation indicators for each beam profile were also added; these make it very easy to check against saturation and to select an appropriate level of beam attenuation. The outputted M_x^2 and M_y^2 values are calculated as a rolling average (from a user variable number of CCD frames) so as to smooth out the effect of any frame-to-frame noise.

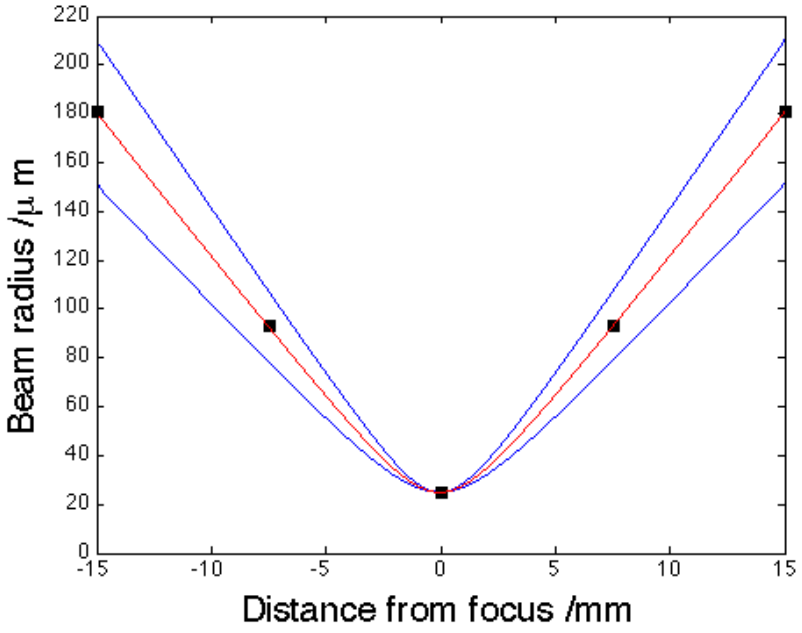

FIGURE 2.19: Nonlinear fit (red line) using equation 2.8 to some example data (black markers) calculated for an $M^2 = 1.2$ beam at 780 nm with a waist radius of 25 μm. The blue lines show how varying M^2 by just 0.2 affects the fit.

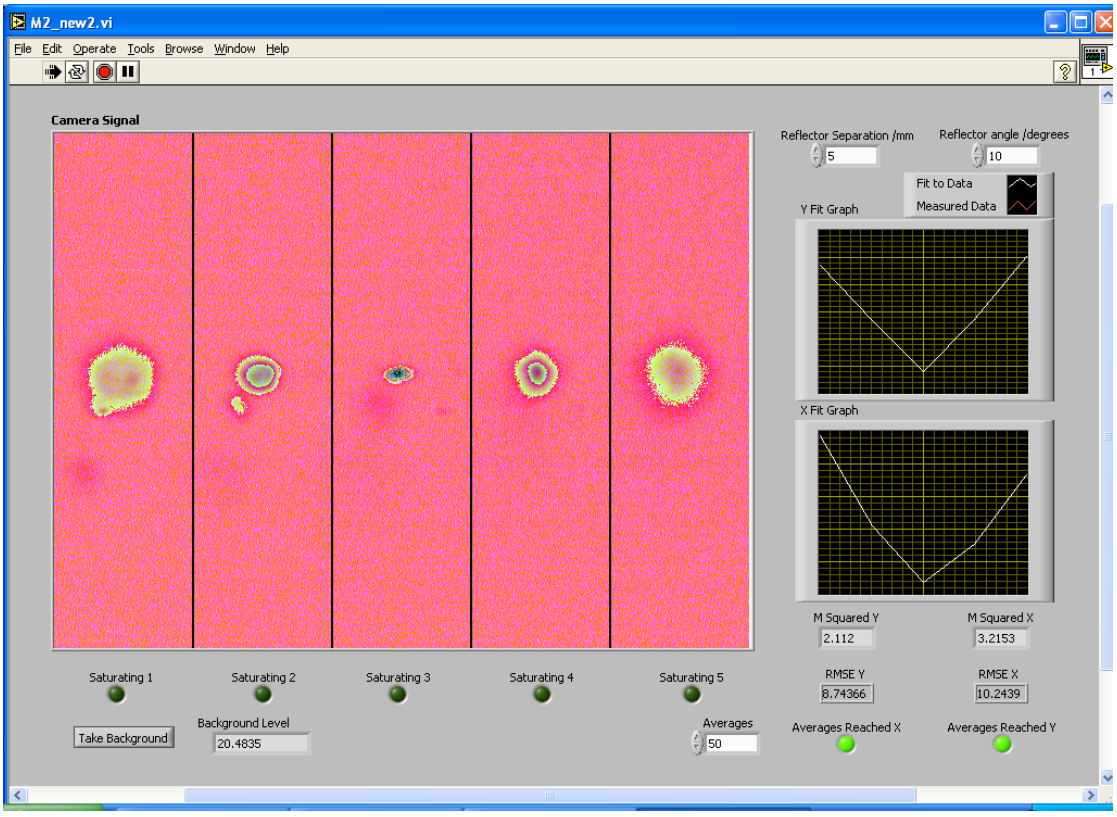

FIGURE 2.20: M^2 meter software output.

2.3 X-ray Control

Once the HHG process has produced x-rays, apparatus that allows these x-rays to be manipulated is required. This section provides details of some of the techniques for x-ray control that have been employed as part of this project.

2.3.1 Gas system

The x-rays currently produced by the HHG process are referred to as soft x-rays or extreme ultra-violet light. Their energy (typically $< \sim$50 eV) whilst above that of visible light ($\sim$ 2–3 eV) is relatively low compared to harder x-rays produced by a synchrotron ($\sim$ 100–10000 eV). HHG therefore occupies a niche energy range which is difficult to produce by other means. A particular feature of this energy range is that in many materials it corresponds to transitions between inner electronic states and the continuum. This leads to very strong absorption and makes it necessary to carry out these experiments under vacuum conditions. During the course of this project many different vacuum systems have been devised and used. Rather than describing each system in detail this section will give an overview of the important design considerations.

2.3.1.1 Vacuum Pumping

The first major consideration is the level of vacuum that is required. Obviously the fewer gas molecules that remain in the chamber (outside of the interaction region) the lower the x-ray absorption will be. There will also be a subsidiary benefit that focusing of the infrared light will not be disturbed by thermal lensing if there is no medium to be heated. Vacuum systems that are capable of attaining pressures below 1×10^{-8} mbar are designated ultra high vacuum (UHV). In order to reach these pressures it is often necessary to bake the chamber to aid the removal of residual gas molecules. These systems usually rely on conflat flanges (CF) where a hard knife edge is pressed into a copper gasket to form an effective vacuum seal. Whilst effective, these systems can be rather difficult to work with and require replacement gaskets each time a flange is opened and re-sealed.

There comes a point where the practical difficulties and expense involved with obtaining low pressures outweigh the advantage that can be gained. Using argon transmission data obtained from the Centre for X-Ray Optics (CXRO)[31] figure 2.21 shows the average transmission through a path length of 1 m in a spectral window between 20 and 50 eV as a function of gas pressure.

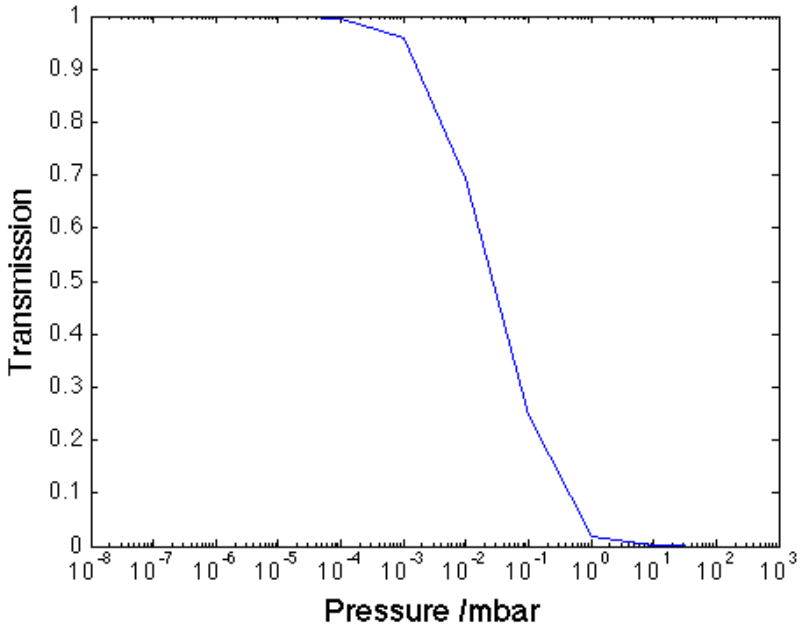

FIGURE 2.21: Average transmission for 20–50 eV x-rays through 1 m of Ar gas.

From figure 2.21 it is clear that (at least for argon) there is little advantage gained by obtaining pressures lower than 10^{-5} mbar. This is convenient as it permits the use of o-ring seals and other high vacuum (HV) components which tend to be far easier to use and cheaper than their UHV counterparts.

All of the vacuum systems constructed as part of this project have been primarily based on o-ring sealed chambers (either KF or ISO-K flanges) pumped by turbo pumps which are in turn backed by rotary or piston based roughing pumps. Typically this type of system is able to reach between 10^{-6} and 10^{-7} mbar. Where used, oil pumps were fitted with fore-line traps and exhaust gasses were piped out of the lab - this reduces the possibility of oil particles settling on optics and the subsequent damage this would cause.

2.3.1.2 Capillary and Mounting Schemes

Another major consideration in the gas system is the ability to maintain a steep pressure gradient between the active region (which must contain gas at ~ 100 mbar) and the rest of the chamber which must be kept as low as possible. Ultimately the gradient is determined by the pumping capacity on the chamber and the flow rate from the active region into the chamber. Keeping the flow rate from the active region as low as possible puts less load on the pumps and allows the chamber to reach a lower ultimate pressure,

thus making the pressure gradient from the active region into the chamber steeper. Similarly increasing pumping capacity causes the chamber to reach a lower ultimate pressure for a given leak rate from the active region.

In this work a (150 μm diameter) capillary waveguide is used to achieve phase matching between the laser pulse and the generated x-rays (see section 3). A narrow bore capillary (because of it's low conductance) also provides an efficient means of obtaining a pressure differential between the active region and the rest of the vacuum chamber.

The means of incorporating the capillary into the vacuum chamber is one of the most challenging aspects of designing a capillary HHG vacuum system. Consequently several different designs have been conceived and tested as part of this project. Whilst displaying a broad range in execution these designs remain similar in concept and so the discussion below will focus on some of the simplest designs.

In order to inject gas into the capillary, holes (typically $\sim$ 300 μm in diameter) were drilled through the side walls of the capillary. Initially this drilling was done mechanically using an ultrasonic mill; later these holes were laser machined [32]. Figure 2.22 shows a cross section of a laser drilled hole; this microscope image was obtained by cleaving the capillary at the position where the hole was drilled.

In a typical system two holes with a separation of a few centimeters are drilled in the side walls of the capillary. The two ends of the capillary are connected to the vacuum chamber. Delivering gas through the holes in the side of the capillary defines a central high-pressure region. The low conductance of the end sections of the capillary restricts the flow of gas out into the vacuum chamber without impeding the laser or x-rays at all. The pressure profile is therefore expected to be as shown in the lower portion of figure 2.23.

A glass T-piece provides a means of mounting and delivering gas to the capillary and also provides additional re-enforcement (see figure 2.23). The primary function of the T-piece is to deliver gas to the holes in the side wall of the capillary whilst preventing it from flowing directly into the chamber (i.e. to ensure that the gas flows through the capillary).

Depending on the implementation the T-piece can itself form part of the outer wall of the vacuum chamber. In this configuration the glass walls of the T-piece allow the capillary to be seen which can provide helpful prompts during alignment. The seal between the outside of the capillary and the inside of the T-piece was initially made by gluing the capillary into narrow sections built into the T-piece (see figure 2.23). These narrower sections were produced by the skilled work of the university's glass blowers who were

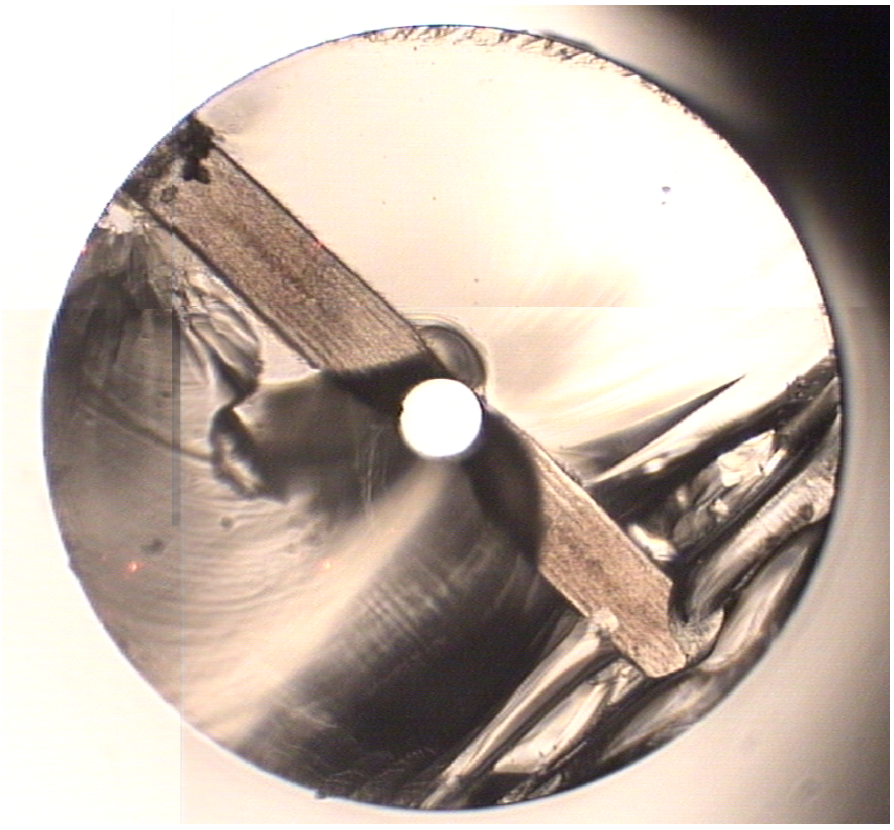

FIGURE 2.22: Composite (× 10) microscope image showing the cross section of a laser-drilled hole (looking down the length of the capillary). This hole has a square cross section (produced by translating the laser focus along a square path). The hole was first drilled from one side, then the capillary was rotated 180° and a second hole was drilled. The cross-section view was achieved by cleaving the capillary at the hole position.

able to bond sections of tube with differing inner diameters together. Unfortunately this design meant that each T-piece required a considerable amount of work to be manufactured. In addition the bonding of several sections introduced an inherent likelihood that the capillary holding tube would not be perfectly straight. Furthermore since the capillary was glued into the T-piece (with UV cure epoxy) it was not possible to re-use the T-piece once the capillary had been damaged by the laser.

In order to tackle the problems of the 'glued in' design attempts were made to use o-ring seals between the capillary and holder (figure 2.24 shows an example of this). In this design the capillary holding tube is all one section (and should therefore be intrinsically quite straight). The gas input tube was again bonded in place by the glass blowers, however, this represented an almost trivial task in comparison to the manufacture of the 'glued in' T-piece design. Different o-ring designs were tested and it was found that x-rings worked very well. These are similar to conventional o-rings except that the ring material has an 'x' shaped cross section. This provides better sealing because it offers

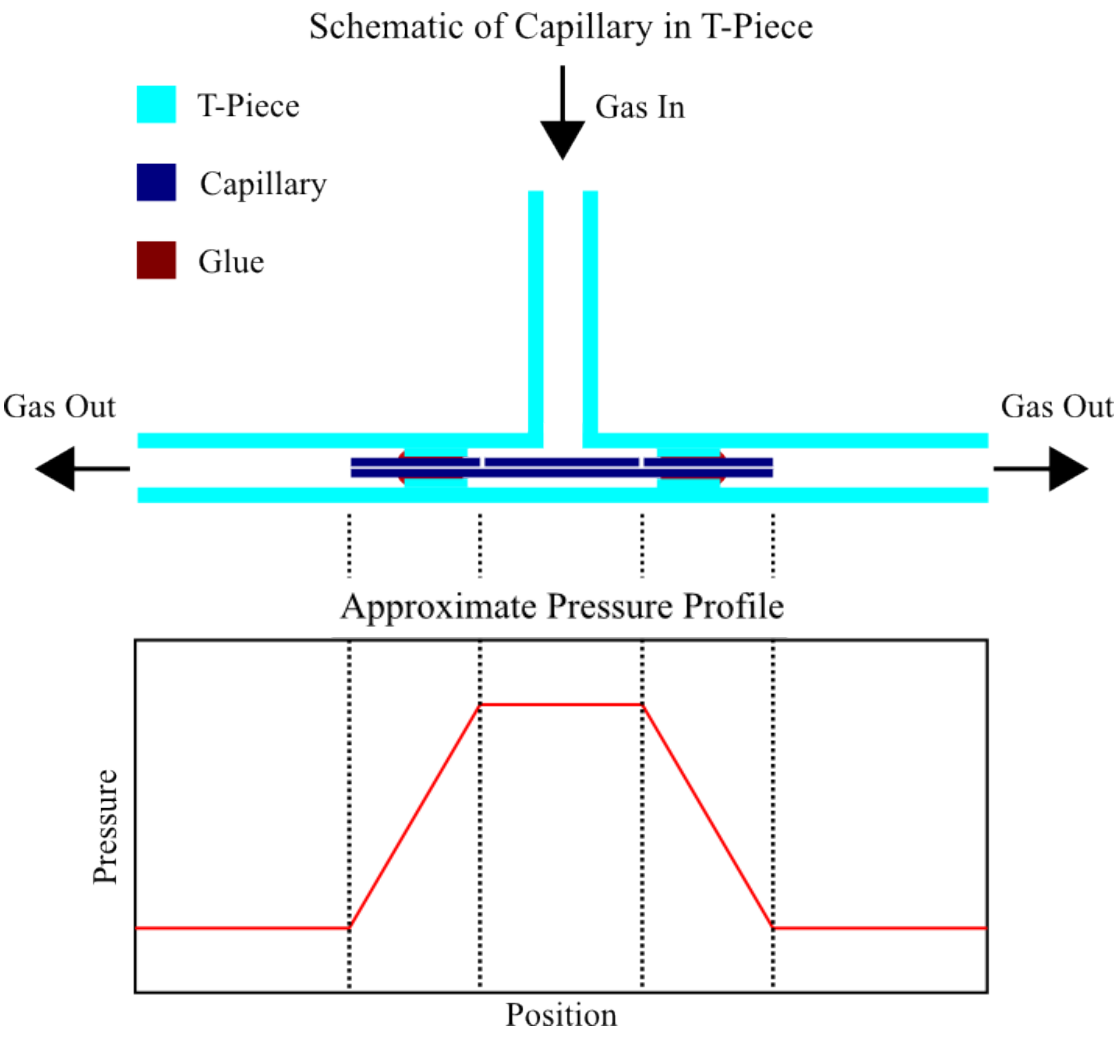

FIGURE 2.23: Schematic of the capillary in its T-piece holder (top). Sketch showing the approximate pressure profile that the holes in the capillary wall produce (below).

two sealing surfaces as opposed to the one surface offered by a traditional o-ring. In addition the shape of the x-ring helps to prevent twisting which is a common cause of o-ring seal failure. Unfortunately with this design whilst it was possible to replace damaged capillaries the operation was rather difficult. During insertion the capillary was placed under considerable stresses and several drilled capillaries (particularly those with many drill holes) were broken during insertion.

In order to make the replacement of capillaries a quick and simple operation (something that is required if different capillary designs are to be tested) an alternative capillary holding device was designed (see figure 2.25). Rather than a capillary, this design uses a cylindrical channel drilled through a block. The block is able to separate into two parts in order to facilitate capillary replacement. O-rings are used to provide a vacuum seal between the two parts of the block and between the gas input section and the vacuum

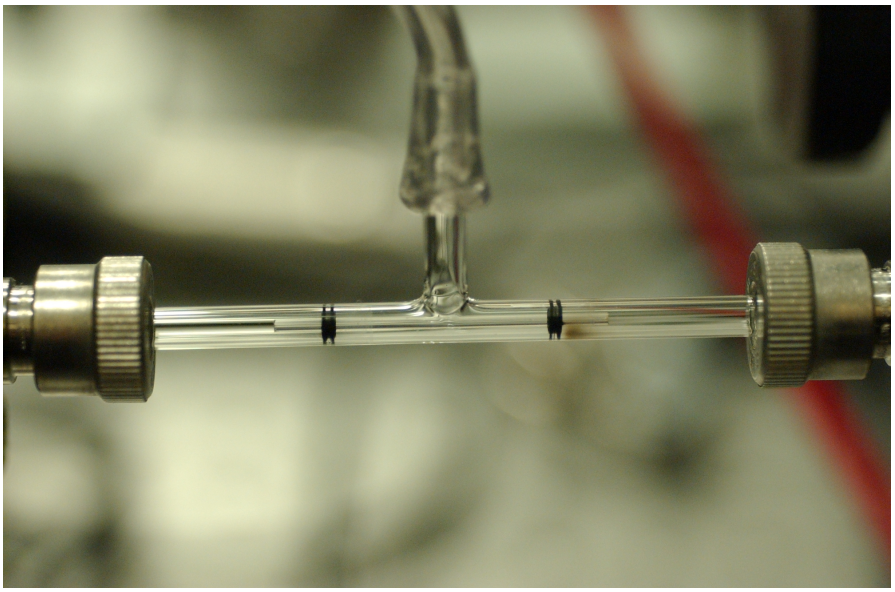

FIGURE 2.24: Photograph of a simple T-piece with the capillary installed; in this case the seal between the capillary and the T-piece is made using two o-rings. The o-rings used here are in fact x-rings; they have an "x" shaped cross section. This provides two sealing surfaces which prevents twisting and improves performance compared with traditional o-rings.

system. A custom o-ring was produced by gluing two x-rings (that fitted the capillary and the bore of the channel) to a standard 2 cm diameter 1.5 mm thickness o-ring. This custom o-ring effectively defines a high pressure gas region inside the block, with the only means of escape being through the capillary. Part of the block was machined from perspex so as to allow observations whilst in use and to assist with alignment. Figure 2.26 shows an exploded view of the block taken from the design drawings; in this view the structure of the block is a little easier to see.

2.3.1.3 Pressure Control

The ability to accurately control the gas pressure delivered to the generation capillary is of significant importance to this experiment. An electronic pressure controller (typically applied in industrial process applications) was used to achieve this level of control. The system consists of two solenoid valves, a pressure transmitter and some control electronics (see the schematic in figure 2.27. The input solenoid valve is supplied with the test gas via a step down pressure regulator. The output solenoid valve is connected to a vacuum backing pump (capable of achieving $\sim 10^{-3}$ mbar). The central part of the system (to which the pressure transmitter and the gas input arm of the T-piece are both attached) forms a pressure controlled reservoir.

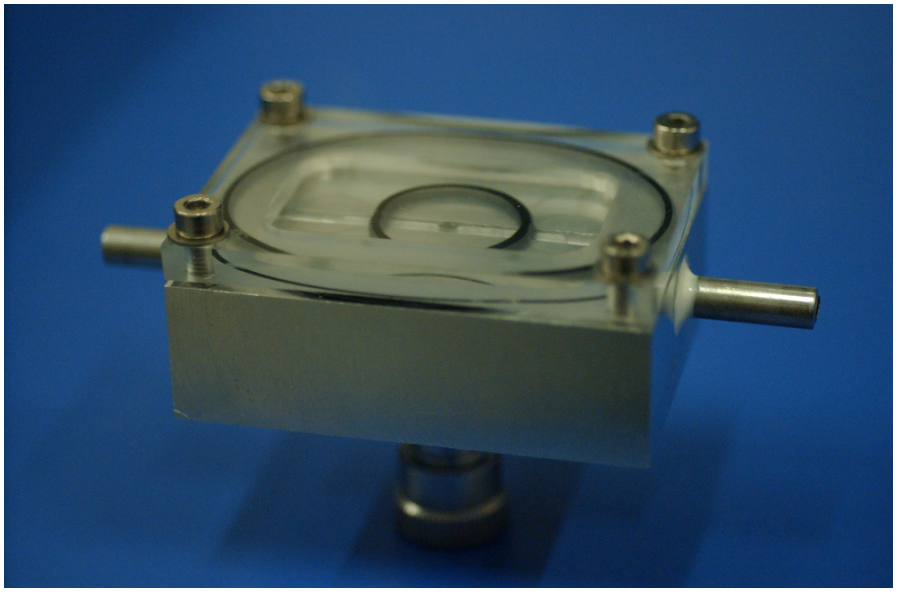

FIGURE 2.25: Photograph of the capillary mounting block with a capillary fitted. The bolts provide enough force to compress the o-rings and make a seal without exerting undue stress on the capillary. The custom o-ring is the smaller central one; gas is introduced into this section through the fitting at the bottom of the block. The block is made compatible with the existing T-piece holding vacuum system with the addition of metal tubes glued to either side.

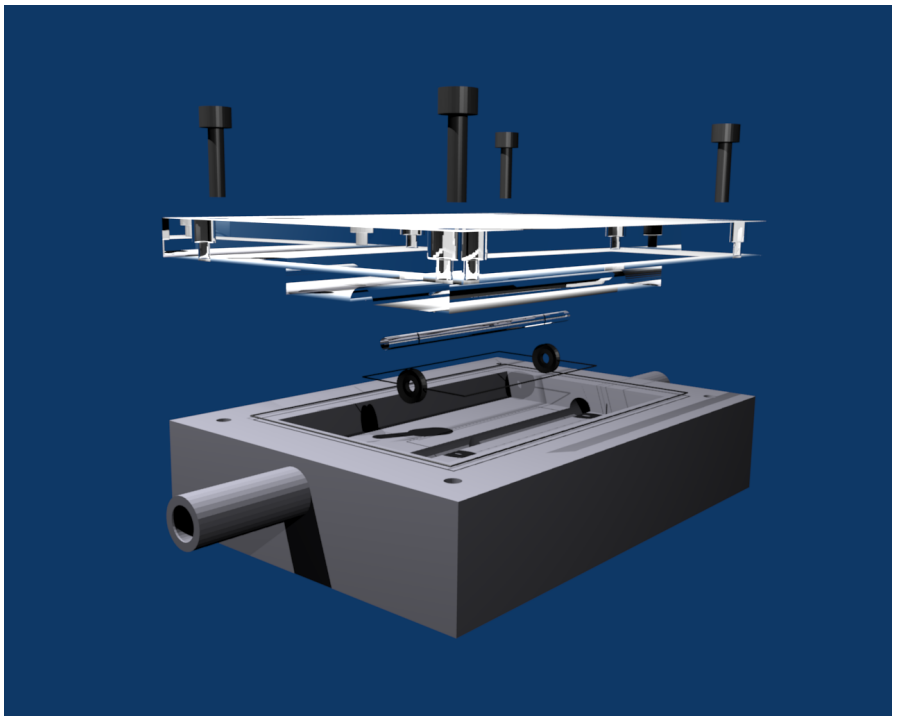

FIGURE 2.26: Exploded view taken from the design drawings for the capillary mounting block. In this view the shape and function of the block is a little easier to follow.

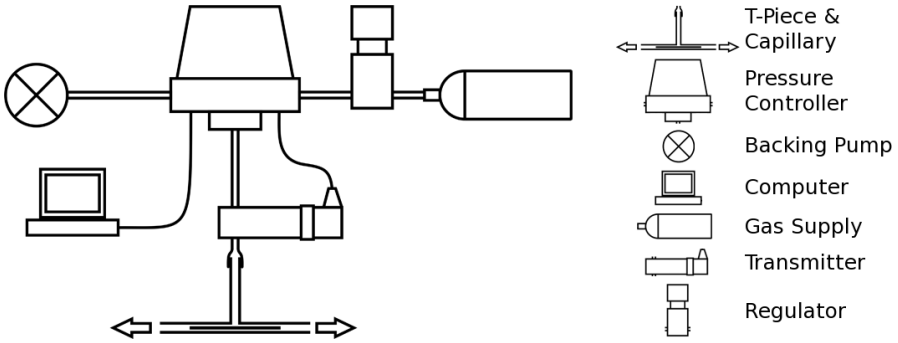

FIGURE 2.27: Left: Schematic of the pressure control system. Right: Legend of components. The pressure controller used was a Tescom ER3000.

The pressure control device contains Proportional Integral Differential (PID) electronic circuits which allow it to automatically operate the solenoid valves in such a way as to maintain the signal from the pressure transmitter as close as possible to the current setpoint value. This enables the system to accurately maintain a specific gas pressure for an arbitrarily long period of time, however, it also makes it possible to program complex pressure profiles as a function of time. This facilitates automated data taking and also allows experiments with, for example, burst or ramp pressure profiles.

Although the pressure controller does contain an internal pressure sensor; an external pressure transducer is used for feedback to the PID electronics. Using an external transducer allows the pressure measurement to be made as close as possible to the location of the capillary (improving accuracy). Inevitably there is a slight pressure drop between the transducer location and the inside of the capillary (mostly occurring at the drill holes). For this reason the absolute pressure calibration of the system is expected to be slightly out - this fact must be appreciated when comparing theoretical and experimental work.

A secondary benefit of using external pressure transmitters is that they can be readily exchanged. Pressure transmitters typically provide an analogue current signal in the range of $4 - 20$ ma (voltage transmitters are also available). With (for example) 4 ma corresponding to the lower limit of the pressure range and 20 ma signal given at the upper limit of the pressure range. The accuracy of the transmitter is therefore typically quoted as a percentage of their full pressure range. Furthermore the pressure controller then digitises this analogue signal into a fixed number of steps over the full signal range. Because of these facts the accuracy of the pressure reading can be significantly affected by the choice of transmitter pressure range - selecting too large a pressure range will significantly reduce accuracy.

Whilst the PID electronics can be tuned to a certain extent there are also several aspects of the system that can also be adjusted in order to optimise its response. Rapid system response is desirable as it allows experiments to be carried out more quickly and it permits the use of sharper temporal profiles, however, there is usually a trade off between response time and system stability.

The input pressure to the control system is best set to slightly above atmospheric pressure. This allows a large proportional term to be used (improving response time) without inducing too much overrun. Keeping this pressure above atmospheric helps to prevent air seeping into the system.

The volume of the regulated part of the system also affects the response time. Initially a buffer volume was incorporated into the system in an attempt to smooth its response; ultimately this was found to be unnecessary and system response time was greatly improved once it was removed.

Initially the system was operated without the backing pump on the output solenoid. This worked to a degree but the system took a long time to reduce pressure after an overshoot. (The whole volume of the system had to be pumped out through the very low conductance capillary). This resulted in a sawtooth pressure profile against time. The addition of a vacuum pump to reduce the pressure at the output solenoid allows the system pressure to drop more quickly thus reducing the average error in the pressure value significantly.

The backing pump connected to the output solenoid valve therefore sets the lower limit on the pressure that can be effectively controlled. Although a turbo pump could achieve a lower pressure, a backing pump is suitable in this application since it can cope with sudden increases in pressure and is capable of very high flow rates. The ultimate pressure of the entire vacuum system is not affected by the use of a backing pump here since (when pumping down) this is isolated from the rest of the system by the output solenoid valve.

The flow rate through the capillary also has a significant effect on the functioning of the pressure control system. It may therefore be necessary, particularly when testing capillaries with multiple large holes, to adjust the PID values accordingly. In fact the pumping capacity of the vacuum pumps used in the main chamber restricts capillary design slightly by placing an upper limit on acceptable capillary conductance. Taking the most extreme case as an example; for a very free flowing capillary even if the pressure at the input solenoid is drastically increased it may not be possible to reach the setpoint pressure without overwhelming (and possibly damaging) the turbo pumps.

2.3.2 Filtering

This subsection is divided into a general discussion of the physics of filtering followed by a subsection which is more concerned with the practical implementation of filters in our experiment.

2.3.2.1 Physics of filters

The optical properties of all materials vary as a function of the frequency with which they are irradiated. When using filters this fact is exploited in order to attenuate specific parts of the electromagnetic spectrum relative to others present in the experiment.

In this project the main use for filtering is to separate the IR laser radiation from the UV and soft x-rays produced by HHG. Aluminium has proved to be a useful filter material in this field because it attenuates the IR considerably whilst allowing a significant part of the soft x-ray beam through to the experiment. The extinction ratio that can be achieved between x-ray and IR is absolutely critical to maintaining a good signal-to-noise ratio and high accuracy in our measurements.

Insight into the operation of a filter can be gained by considering its complex refractive index n. This is typically separated into its real and imaginary parts ($n = \alpha - i\beta$). For all physical materials both α and β are functions of the illumination frequency. When selecting a transmission filter material it is the β value (often referred to as the extinction coefficient) as a function of illumination frequency that is of prime importance. The higher its value the stronger the material's absorption at that frequency. (Negative values of β can be used to describe media with gain). The property α is inversely related to the phase velocity inside the material; higher values of α imply slower phase-front propagation within the material. The ratio of phase velocities either side of an interface determines the angle through which the transmitted light is refracted - this in turn affects the proportion of energy that is reflected.

The interaction between electromagnetic waves and metals such as aluminium can be modelled (to first order) by considering only the motion of the conduction electrons[33]. This leads to equation 2.9 for the refractive index of a metal:

$$n^2 = 1 + \frac{\sigma/\epsilon_0}{i\omega(1 + i\omega\tau)} \tag{2.9}$$

Where ω is the angular frequency of the wave, σ is the electrical conductivity of the metal, ϵ_0 is the vacuum permittivity and τ is the average time between collisions in the electron's path:

$$\tau = \frac{m_e \sigma}{N q_e^2} \tag{2.10}$$

Here N is the number density of free electrons per unit volume - each with charge q_e and mass m_e.

As equation 2.9 tends towards the low frequency limit it is possible to neglect the real term[†] and terms with ω^2 dependence - the equation can then be written in the following form:

$$n^2 = -i\frac{\sigma}{\omega \epsilon_0} \tag{2.11}$$

By using the identity $\sqrt{-i} = \frac{1-i}{\sqrt{2}}$ equation 2.11 becomes:

$$n = \sqrt{\sigma / 2\epsilon_0 \omega}(1 - i) \tag{2.12}$$

The implication of this equation is that for low frequencies the real and imaginary parts of the complex refractive index are both large. The high β value will mean that low frequency electromagnetic waves are rapidly attenuated in a metal. The combined effect of high α and β values will be to produce a high optical impedance and therefore a large impedance mismatch with air (or vacuum). This impedance mismatch makes the metal highly reflecting. Conversely for high frequencies equation 2.9 becomes:

$$n^2 = 1 - \frac{\sigma}{\epsilon_0 \omega^2 \tau} \tag{2.13}$$

At high frequencies the refractive index therefore becomes real and slightly less than 1. The very low β values mean there is little attenuation and the reduced impedance mismatch causes a drastic reduction in reflectance.

The behaviour described by equations 2.12 and 2.13 is borne out by experimental measurement of the refractive index. Notice the sharp transition from high to very low reflectance that occurs at approximately 200 nm (see figure 2.28):

Aside from blocking the IR aluminium displays some interesting features in the spectral range of interest to this project (from $\sim$ 50 down to a few nm). Figure 2.29 shows

[†]For typical metal conductivities and at low frequencies the imaginary term dominates: i.e. the conductivity σ of aluminium $\approx$ 37.8 $\times 10^6$ S m^{-1} and taking $\lambda = 1$ μm gives $\omega \approx 1.9 \times 10^{15}$ rad s^{-1}. If the real term is retained in equation 2.11 we find $n^2 \approx 1 - 1124\, i$ or $n = 23.7 - 23.7i$. Neglecting the real term we obtain (to three significant figures) an identical value of n.

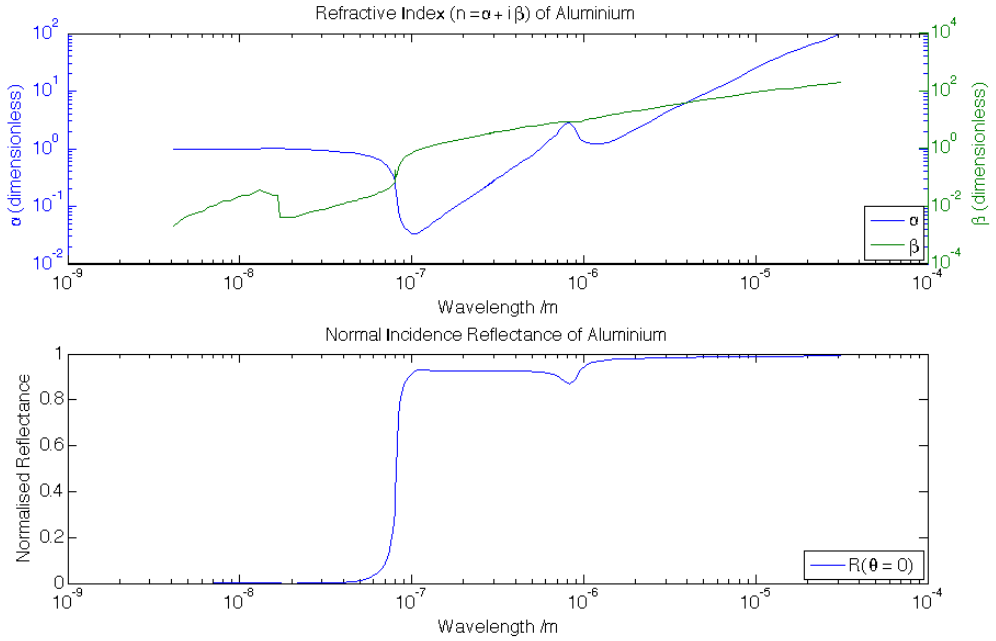

FIGURE 2.28: The upper plot shows the real (blue) and imaginary (green) parts of the refractive index of aluminium over a broad range of the electromagnetic spectrum (note the logarithmic scales). The lower plot shows the reflectance of aluminium over the same spectral range. (Reflectance figures are for unpolarised light at normal incidence). Data taken from Webber[34].

the refractive index data over this spectral range. The real and imaginary parts of the refractive index for aluminium are shown in the upper plot; the lower plot displays the transmittance and reflectance data for the same spectral range.

The peak in absorption (β) between 5 and 17 nm is due to the fact that these photon energies correspond to electronic transitions within the aluminium atom. In fact this particular peak is mainly due to transitions between aluminium L shell electrons and the continuum[†]. The two functions α and β are not completely independent since they are subject to a Kramers-Kronig relation. The effect of this relation is illustrated by the fact that the sudden step in β at $\sim$ 17 nm gives rise to a peak in α at the same wavelength. (Also note that the absorption leads to anomalous dispersion between 7 and 17 nm). The increase in β has a straightforward effect on the transmittance of the aluminium - effectively blocking any signal between 5 and 17 nm. Consequently a different type of filter would be required in order to carry out experiments in this region of the spectrum.

The reflectance curve in figure 2.29 is for a normal incidence, input beam. As can be seen by the axis scale, under these conditions the reflectance is very low and has a negligible effect when compared to the transmittance. However, reflectance is considerably altered

[†](L_{III} edge @ 75 eV (16.5 nm), L_{II} edge @ 76 eV (16.3 nm) and L_I edge @ 87 eV (14.2 nm).

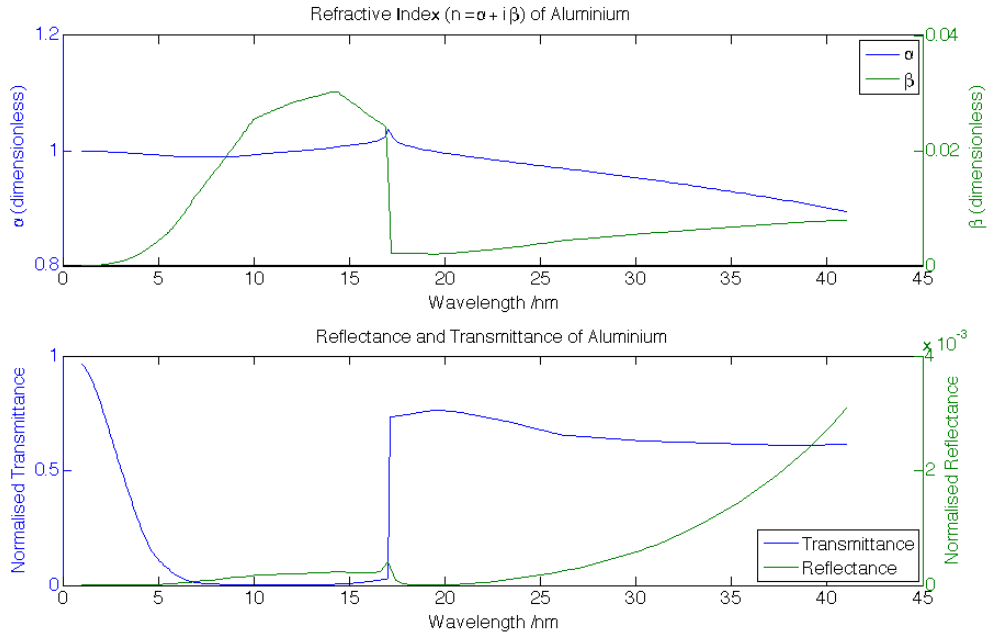

FIGURE 2.29: The upper plot shows the real(blue) and imaginary (green) parts of the refractive index of aluminium in the soft x-ray region of the spectrum. The lower plot shows the transmittance (blue) and reflectance (green) of aluminium over the same spectral range. (Reflectance figures are for unpolarised light at normal incidence and transmittance is calculated for a 200 nm thick aluminium foil). This data was taken from CXRO[31].

by factors such as the surface flatness, the angle at which radiation is incident on the filter, the polarisation angle relative to the surface and the thickness of the filter. It is for example possible for a grazing incidence x-ray beam to undergo total external reflection (see subsection 2.3.3) at a metal surface. This is because x-rays see a drop in index when moving from vacuum into metal just as light sees a drop in index as it moves from the glass of a prism into the air.

In order to operate effectively as a filter the aluminium must be thick enough to block the IR without being so thick that it severely attenuates the XUV. Since most of the attenuation of the IR is due to reflection at the first interface the filter can be very thin which improves XUV transmission. The lower limit for filter thickness is in theory set by leaking due to quantum mechanical tunneling. Since the IR is approximately five orders of magnitude more intense than the XUV, several skin depths[†] are required to attenuate the IR sufficiently. In practice the thickness is limited by the mechanical strength of the aluminium (if it is required to be self supporting) and by fabrication limitations (very thin filters tend to suffer from pinholes that allow the IR to pass).

[†]The skin depth for aluminium at 780 nm is approximately 4 nm.

Filters of approximately 200 nm (sometimes composed of two 100 nm thick layers) have proved the staple for this project.

2.3.2.2 Practical implementation of filters

As discussed in the previous section transmission filters for use in the XUV typically have to be very thin ($\sim$ a few hundred nanometers) to counter the high material absorption (transmission improves further into the x-ray region). This means that the filters are very delicate and must be handled carefully.

In order for the aluminium filters to operate effectively a vacuum compatible mounting system is required. For convenience of use the filtering system is required to allow filters to be changed without disturbing the vacuum inside the chamber. This was achieved by mounting the filters on two horizontal sliders. Each slider has five positions (though one is typically kept open to allow a through path whilst pumping). Each position can hold up to two filters - this allows many combinations of different filters to be used without opening the vacuum chamber.

One of the most important features of the filter holder is that it should incorporate a baffle that prevents scattered laser light from arriving at the detector. (Because the laser light is several orders of magnitude more intense than the x-rays even a very small percentage leak of scattered light could be enough to swamp the x-ray signal). At the same time, because the filters are so thin they are not mechanically strong enough to support a large pressure differential. For this reason there must be a gas path from one side of the filter to the other in order to dissipate any pressure imbalance that may occur.

Figure 2.30 shows a schematic view of the filter unit designed for this project. The design makes use of two translation feedthroughs which allow the filter carriers to be slid across the beam. These translation feedthroughs consist of a metal rod passing through two o-ring seals. The seals are sufficient to prevent gas from leaking into the chamber but are loose enough to allow linear translation and rotation of the rod. The design was based on a standard KF 40 cross in order to maintain compatibility with the existing vacuum system.

The apertures either side of the filter unit effectively block the vast majority of scattered light whilst allowing gas to bypass the filter preventing filter damage due to pressure build-up. The aperture blocks also engage with the filter carriers preventing them from rotating in the cross.

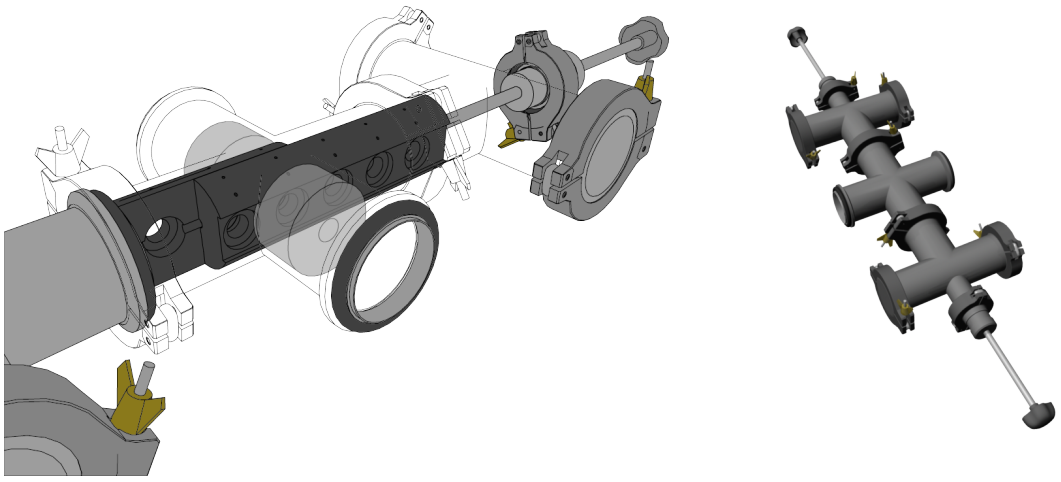

FIGURE 2.30: Left: Cutaway diagram showing the internal structure of the filtering unit. Right: View of the whole unit; the optical path runs along the horizontal axis of the central cross, filters slide in the vertical axis.

An XUV sensitive photodiode (see section 2.4.1) was installed in one of the filter positions to allow rapid measurements of XUV intensity to be made. The wiring for the photodiode is passed along a hole drilled through the carrier then soldered to an electrical feedthrough (BNC) where the signal can be measured.

2.3.3 X-ray optics

Distinct from detection devices (discussed in section 2.4) and from filters (described in section 2.3.2) it is useful to be able to further manipulate the soft x-ray beam after it has been generated. To this end x-ray optics may be utilised. This subsection will give a very brief explanation of those x-ray optics that have been used or considered for use on this project.

2.3.3.1 Grazing Incidence Reflection

Because of material absorption in the soft x-ray region of the spectrum it is impossible to use transmission optics such as the glass lenses that are typically used in the visible. Similarly, traditional near normal incidence mirrors are ineffective because of their high absorption and low reflectivity. However, at grazing incidence, reflectivity is considerably improved.

The Fresnel reflection equations describe how the reflectivity (r) of an interface changes with the angle of incidence (θ_i). Note that the reflectivity behaves differently for S and P polarisations and so two separate equations are required:

$$r_s = -\frac{\sin(\theta_i - \theta_t)}{\sin(\theta_i + \theta_t)} \tag{2.14}$$

and

$$r_p = \frac{\tan(\theta_i - \theta_t)}{\tan(\theta_i + \theta_t)} \tag{2.15}$$

Note that r_s and r_p represent the E-field amplitude of the reflected S and P waves. These equations can therefore be applied as a simple multiplication factor to the input E-field. In order to convert to reflected intensities one simply takes the squared modulus or multiplies by the complex conjugate in the usual way.

The refractive indices of the two materials are brought into the calculation through Snell's law which is used to calculate the angle of transmission (θ_t):

$$\theta_t = \arcsin\left(\frac{n_1}{n_2}\sin(\theta_i)\right) \tag{2.16}$$

Where n_1 and n_2 are the refractive indices of the input and output sides of the interface respectively.

If the refractive index $n_1 > n_2$ (i.e. the light is propagating into an optically less dense medium) then there exists a critical angle (θ_c) beyond which total reflection will occur.

$$\theta_c = \arcsin\left(\frac{n_2}{n_1}\right) \tag{2.17}$$

By substituting 2.17 into 2.16 we see that when $\theta_i = \theta_c$ the angle of transmission (θ_t) is 90∘ (i.e. along the interface). Furthermore when θ_i exceeds θ_c the arcsin term in 2.16 no longer has a real solution and so θ_t becomes complex. This in turn leads to complex values for the transmitted and reflected intensities.

Above the critical angle the complex amplitude means that the transmitted wave becomes evanescent (its propagation direction is along the interface and the field diminishes exponentially with distance into the second medium). Providing this evanescent wave is not coupled into another nearby medium (by quantum mechanical tunneling) then there will be no net transfer of energy across the interface.

In order to conserve energy, the incident wave therefore undergoes total reflection and the intensity of the reflected wave is equal to that of the incident wave. Note that this only holds true given the assumptions that the media are not absorbing and that the

interface is perfectly smooth. In a non-ideal case scattering and absorption losses will tend to smooth the shoulder of the step in reflectivity[35]. Figure 2.31 shows the angular reflectivity calculated for a vacuum - SiO_2 interface at 633 and at 25 nm - the dotted lines illustrate the approximate effect of absorption and scattering.

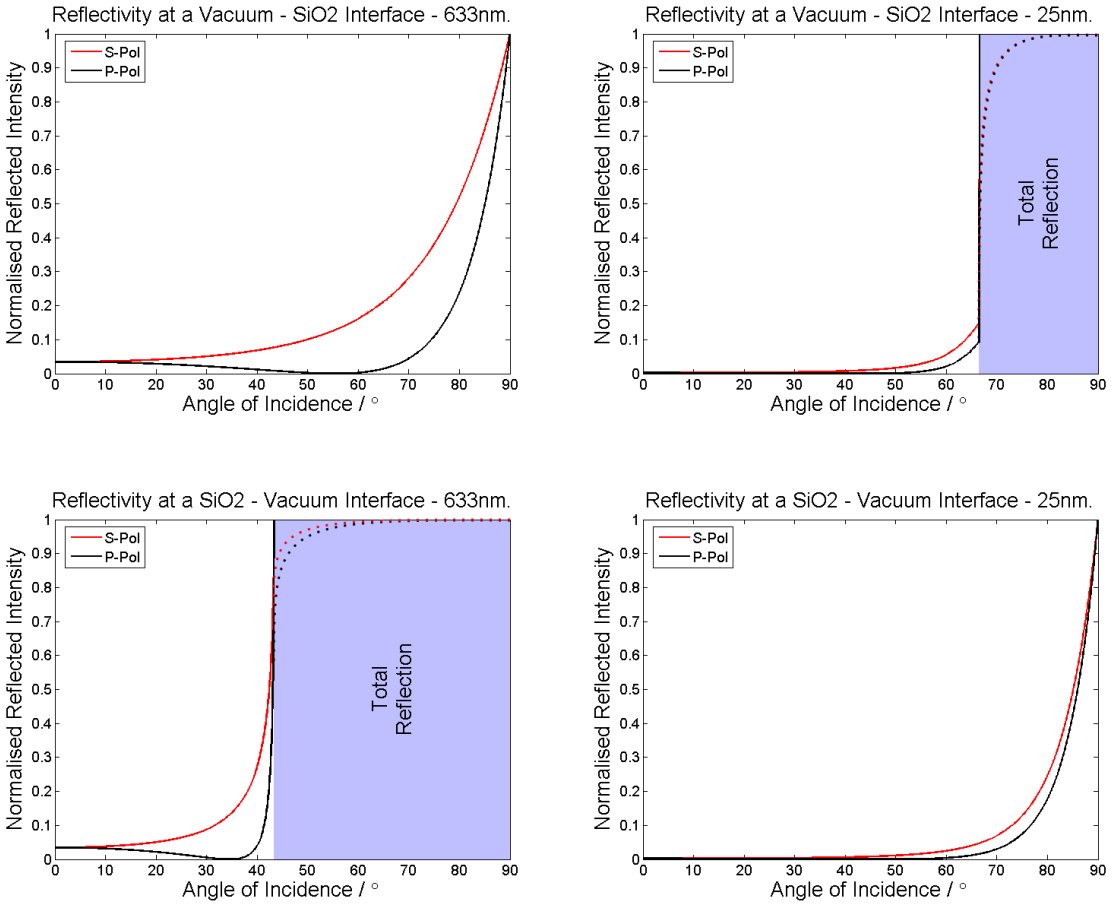

FIGURE 2.31: Angular reflectivity at... Upper Left: A vacuum-SiO_2 interface at 633 nm. Upper Right: A vacuum-SiO_2 interface at 25 nm. Lower Left: An SiO_2-vacuum interface at 633 nm. Lower Right: An SiO_2-vacuum interface at 25 nm.

As can be seen from figure 2.31 the observed behaviour for soft x-rays is approximately opposite to that of visible light. This is a simple result of the fact that the real part of the refractive index of SiO_2 is above unity for visible light and slightly below unity for x-rays.

Figure 2.31 also shows that the XUV reflectivity remains high over a large range of incidence angles. For example at 25 nm SiO_2 maintains a high reflectivity from $0 - 23°$ from grazing incidence ($90 - 67°$ from normal). The performance of gold is even better, with high reflectivity up to 40° from grazing incidence. It is because of this large angular range that grazing incidence optics are of practical use in the x-ray region (i.e. experiments do not have to be impractically large in size).

As an example application of grazing incidence reflection; parallel reflectors with an adjustable separation could be used to provide translation control of the x-ray beam. Two pairs of reflectors would give full X-Y scanning ability.

Reflection From Silicon

Reflection from silicon has been proposed as a method of simultaneously gaining translation control as mentioned above whilst also attenuating the IR. As can be seen from the two plots in figure 2.32 total reflection occurs for the x-rays at angles of incidence greater than $\sim$ 73° but does not occur for the IR. There is also a minimum in the reflectivity of P-polarised IR that occurs at $\sim$ 75° (Brewster's angle). Since the generation laser is linearly polarised it is simple to arrange for the IR to be incident on the silicon with P-polarisation - this produces the best possible contrast ratio between IR and x-ray.

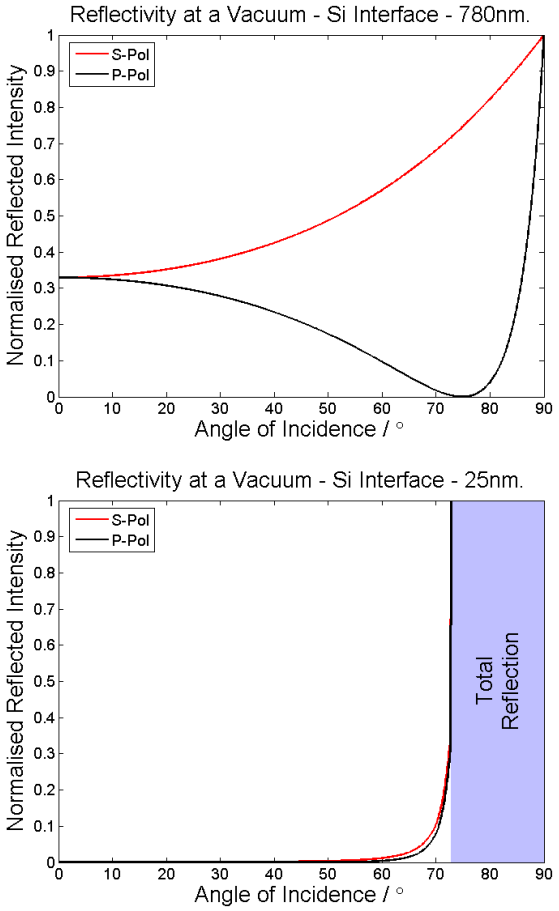

FIGURE 2.32: Angular reflectivity at a vacuum to silicon interface - S and P polarisations are both shown. The upper plot shows the reflectivity curves for 780 nm whilst the lower plot is calculated for 25 nm.

Silicon is readily available in wafer format with an atomically flat surface - this makes it an excellent choice for use as a reflector. Additionally, heat transfer from silicon and

the damage threshold of silicon have been the subject of much previous work.

Note: it is improbable that this technique would replace the foil filtering system because it does not prevent scattered light from reaching the detector - rather this would be an additional step to further increase the signal to noise ratio.

Kirkpatrick-Baez Mirrors

Kirkpatrick-Baez Mirrors can be used to focus an x-ray beam. The system consists of two sets of orthogonal mirrors - spatially separated - so that the beam is reflected from one surface after the other. (For example the first mirror might focus the beam in the X-direction then the second mirror focuses in Y).

The simplest case uses just two mirrors. In order to increase the reflecting surface area and hence the collecting aperture either the length of the optics can be increased or multiple mirrors can be nested in each set. For this project a single mirror pair is likely to be sufficient because of the small physical size of the beam. Assuming a 1 mm diameter beam and 15° grazing angle incidence (optimal for silicon) only $\sim$ 4 mm of mirror length would be required.

The great advantage of this design is that each mirror only has to curve in one plane in order to achieve a full focus. It is therefore possible to use optically flat reflective sheets (such as a silicon wafer) and then apply a physical stress to bend the sheet into a curve. The profile of the curve can be controlled to some extent by the width of wafer segment. The radius of curvature can be adjusted by the forces applied.

In order to achieve the best focus the two mirrors must represent segments of the ellipsoid whose foci are the virtual waist point within the source and the desired x-ray focus point. For a perfectly collimated beam this would be a paraboloid.

Figure 2.33 shows a 3D model of a commercially available Kirkpatrick-Baez mirror system.

Tapered Capillaries

For synchrotron beams (which are spatially very large) and for x-ray telescopes it has been possible to implement focusing mirrors that consist of many concentric curved surfaces. The input beam undergoes a grazing angle reflection at the surface and is brought to a focus. Figure 2.34 shows a schematic of the mirrors used in the European Space Agency's x-ray space telescope XMM-Newton (X-ray Multi-Mirror Mission).

The x-ray beam produced by a capillary HHG source is typically only $\sim$ 1mm in diameter. Fabrication of multiple concentric shell mirrors on this scale would be extremely

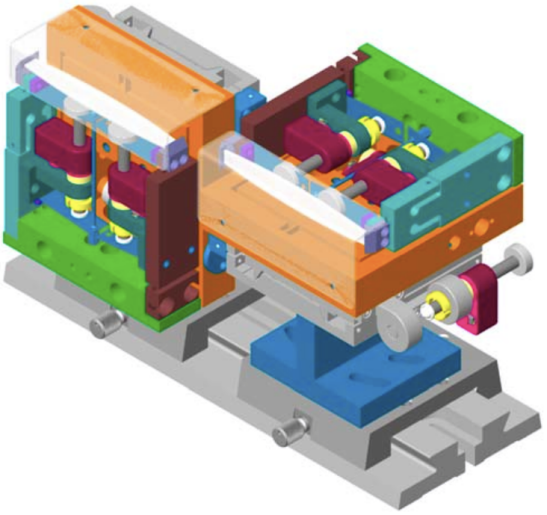

FIGURE 2.33: 3D model of a Kirkpatrick-Baez mirror system (KB-SR1) produced by Xradia - this figure was extracted from their description brochure[36].

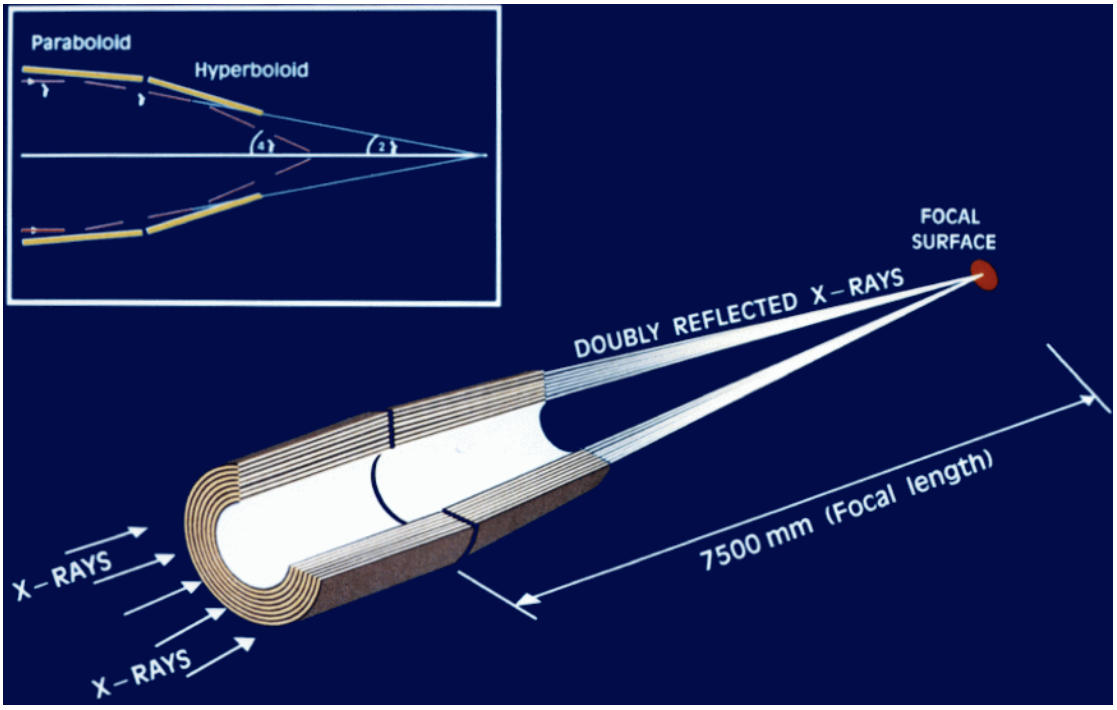

FIGURE 2.34: XMM-Newton employs a two step optical system consisting of a set of parabolic shells followed by a set of hyperbolic shells. This image is taken from ESA's XMM-Newton users handbook (v2.6)[37].

challenging and probably impossible with current technology. It is possible, however, to produce a simplified and yet analogous focusing system which uses only a single surface.

Drawing on the expertise of other groups at the University of Southampton[38] it is possible to heat and stretch glass capillaries so as to produce an approximately parabolic profile of inner radius against length. The inner surface of the tapered capillary (which is automatically very smooth) can then be used as a grazing incidence reflector in order to produce a focus. Figure 2.35 shows an SEM image of one such taper.

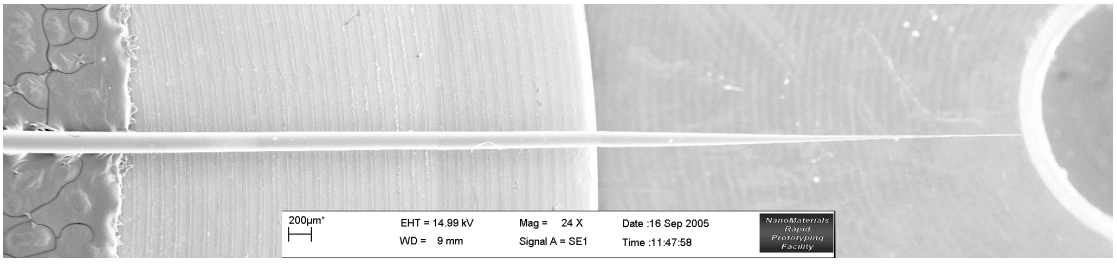

FIGURE 2.35: Composite image of a tapered capillary used for x-ray focusing - original images produced using an SEM.

Measurements of the taper width can be made at evenly spaced points along the image in order to obtain an approximation of the taper profile. This data can then be inserted into a ray-tracing model in order to simulate the output and to provide feedback to the fabrication process. The plot on the left of figure 2.36 shows a simple ray-trace of the taper profile extracted from figure 2.35. Rays shown in red escape from the taper with one or zero reflections - rays shown in green undergo multiple reflections and are ignored in this simple 2d model. The geometry of the taper in the left plot is not correct and as a result the throughput will be very low.

The central plot in figure 2.36 shows the ray tracing calculation for an improved taper design. The profile of this taper is closer to the optimal parabolic shape and so more of the incident rays are able to contribute at the focus - efficiency is greatly improved. The plot on the right of figure 2.36 simulates a taper with the optimum parabolic shape - consequently the reflected rays are brought to a point focus ensuring the maximum possible intensity. Any deviation from the optimal shape produces a fan of converging rays (and a larger, less intense focus) as illustrated by the other two plots.

Comparing the central plot of figure 2.36 to the one on the left; the total throughput is also improved by the fact that the taper has a wider opening (it is truncated further up). Some truncation is unavoidable due to the fabrication method and in order to allow access to the focus, however, the right hand plot demonstrates the effect of truncating too much of the taper. Note that only the fraction of the incident beam that is reflected

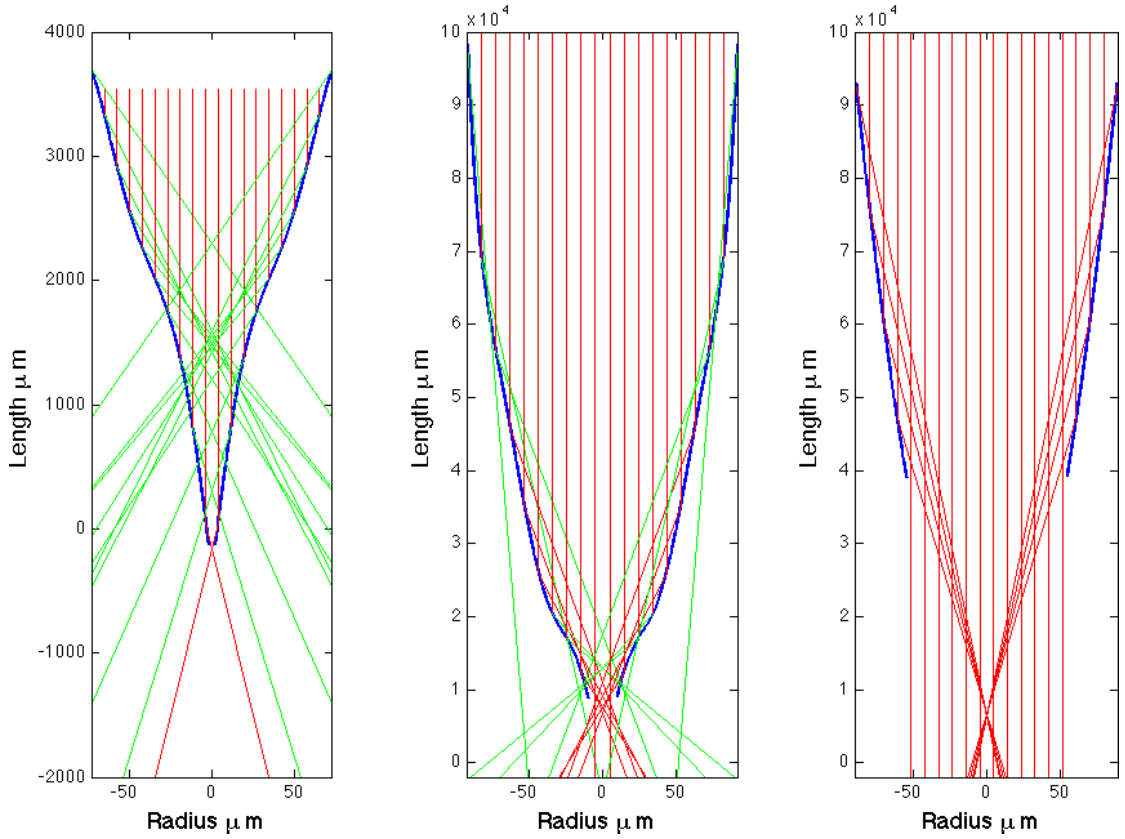

FIGURE 2.36: Left: Ray trace of the taper profile extracted from the SEM image in figure 2.35. Centre: Ray trace of an improved taper profile - extracted from optical microscope images. Right: Ray trace of a truncated parabolic taper profile. Note: the aspect ratio of these plots has been deliberately distorted to make the data easier to see.

from the taper will be brought to a focus. Depending upon where the taper is truncated there may be a considerable fraction of the beam that passes straight through the taper without reflection. This unreflected light is to be avoided where possible since it does not contribute to intensity at the focus (where the sample will be placed) but could contribute to noise at the detector.

2.3.3.2 Multilayer Mirrors

In the XUV region of the spectrum absorption of all materials is very high and under near normal incidence reflectivity is very low (see section 2.3.2). Under these conditions the use of single layer reflecting optics is usually impractical.

Multilayer mirror coatings can be used to enhance reflectivity in regions of the spectrum where absorption would typically be very high[39]. The technology used in these mirrors is very similar to that of anti-reflection coatings and dielectric mirrors. The

mirror surface is composed of multiple layers thus producing many successive interfaces. Each interface produces a partial reflection. In an anti-reflection coating these successive reflections are arranged to destructively interfere; reducing the reflected power. In a multilayer mirror the layer thicknesses are carefully calculated so that the reflections add together constructively - enhancing the reflected power (figure 2.37 shows a diagrammatic representation of this process).

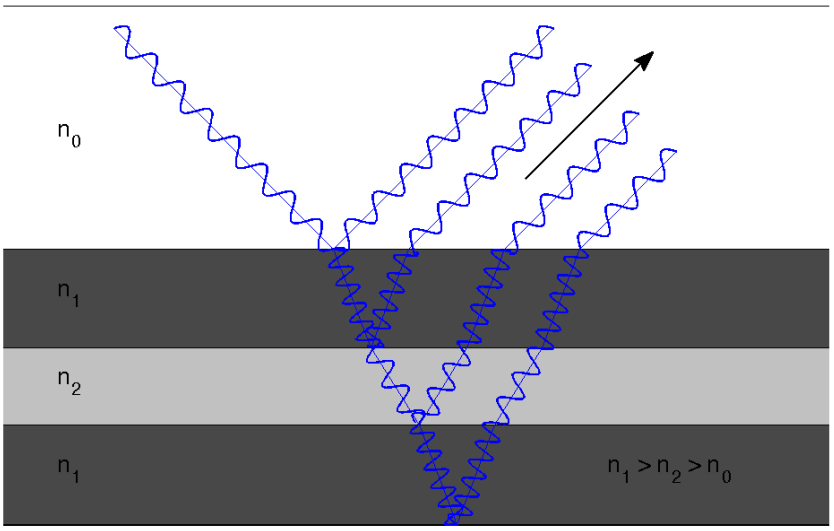

FIGURE 2.37: A partial reflection occurs at each interface within a multilayer mirror (typical optics consist of between 30 and 100 layers). By carefully designing the layer indices and thicknesses the reflections from each interface can be made to constructively interfere - enhancing reflected power. Note that a 180° phase change occurs on reflection when an increase in refractive index occurs across the interface.

For soft x-ray beams molybdenum and silicon are often used as the two layer materials - deposited with layer thicknesses of just a few nanometers. Over a narrow spectral band reflectivities of over 70% are now possible[40].

It is obvious that the operation of these multilayer mirrors will be highly dependent both on the wavelength of the radiation and the angle of incidence. Despite this the angular range of operation of the multilayers is sufficient to allow curved mirror substrates to be used in order to focus the x-ray beam. These optics are now readily available and this is an area which is being explored as part of this project.

Just as multilayer mirrors can be used for dispersion compensation in laser systems the same concept is now being applied to multilayer mirrors for the XUV[41]. This level of control over the temporal profile of the XUV beam, whilst currently superfluous to this project, might become more important in the future if single, attosecond pulses are to

be used. In any case this example serves to demonstrate the level of sophistication that is becoming possible with these multilayer mirrors.

2.3.3.3 Zone Plates

Although not yet used on this project zone plates have certainly been considered and may well be used in the future. As yet zone plates are still undergoing considerable development; fabrication is very difficult and custom-made designs would be required. These factors lead to very high purchase prices and to date these have outweighed the advantage that may be gained. Nevertheless this area has been observed with interest and deserves mention here.

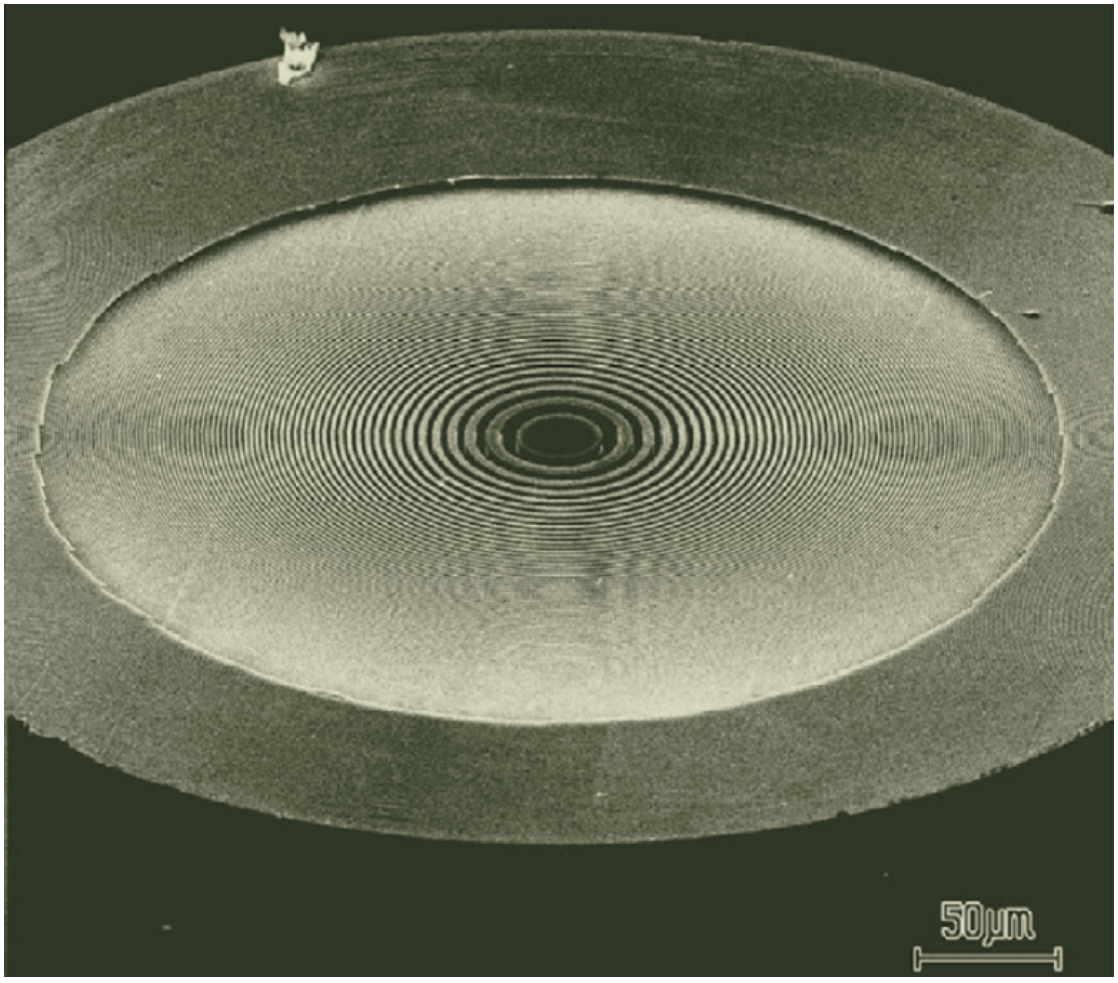

FIGURE 2.38: SEM image of an intensity mask Fresnel zone plate - image taken from[42].

Zone plates are often employed to collimate or focus x-ray beams from more energetic sources such as synchrotrons. These devices typically consist of carefully designed patterns of concentric rings (see figure 2.38). The rings apply either an intensity or phase mask to the input beam in such a way as to simulate the effect of a lens. Considering the path length between the desired focus position and the plane of the zone plate it is clear that annular zones (Fresnel zones) can be defined which alternate between having a net 'in-phase' and net 'out-of-phase' contribution at the focus. In an intensity mask zone plate any Fresnel zone that would make an out of phase contribution at the focus is attenuated. This produces a focus but leads to a reduction in the x-ray throughput. In a phase mask zone plate the thickness or refractive index profile of the zone plate is used to induce a π phase shift between each subsequent Fresnel zone. All Fresnel zones can therefore contribute (in phase) light to the focus. This is clearly more efficient than an intensity mask zone plate since more of the input energy is retained.

Towards the edge of a zone plate the alternation between in phase and out of phase contributions becomes very rapid - consequently the Fresnel zones become very narrow. In fact the imaging resolution of the zone plate is approximately defined by the outer zone width (which also impacts upon the focal length and the clear aperture that can be achieved). Outer zone widths (and resolutions) as small as 15 nm have been attained using electron beam lithography[43].

Unfortunately conventional zone plates appear to be approaching a resolution limit that is enforced by material physics rather than simply by fabrication techniques. Particularly for higher energy sources - in order to attenuate the out of phase zones sufficiently a considerable thickness of material is required. This causes the outer zones to have a high aspect ratio of thickness to width. The thickness of these high aspect ratio zone plates leads to multiple reflections and other diffraction phenomena that are not properly taken into account in simplified theoretical models. Reflection[44] or volume[45] zone plates may offer means of overcoming this limit.

Most development of zone plate technology is aimed at very high energy sources such as synchrotrons. Materials behave very differently in the soft-x-ray region and zone plate designs would have to be modified accordingly. The very high material absorption present in the soft x-ray region might mean that reflection geometries are more appropriate; or conversely it may permit the use of lower aspect ratio high resolution transmission zone plates. Since zone plate development is not one of the stated aims of this project we await commercial availability of phase mask zone plates (suitable for XUV) before investing in this technology.

2.4 X-ray detection

This section will briefly describe the x-ray detection techniques that have been used as part of this project.

2.4.1 Photodiode

The first detection method used was also the simplest. A photodiode was purchased (from UDT sensors, a 1 cm^2 XUV-100C) and mounted inside the vacuum system on a linear translation arm to allow insertion into and removal from the x-ray beam. An electrical feedthrough was used to pass the signal to the detection electronics.

The main component of the photodiode is a p-n junction which is created by depositing a layer of negatively doped semiconductor over a positively doped substrate. (The upper layer is thin enough to allow absorption to occur at the junction). Thin wires are then bonded to either side of the p-n junction, and (in our case) the device was encased in a ceramic package to aid handling.

The behaviour of the photodiode depends upon the bias voltage that is applied to it. In most of the experiments conducted for this project the photodiode was used in an unbiased mode.

When unbiased, additional electrons from the n-type material and additional positive charge carriers from the p-type material diffuse across the junction plane and combine with ions fixed in the crystal lattice near to the junction. This creates a region called the depletion zone in which very few charge carriers exist and across which there is a voltage difference imposed by the contact between dissimilarly doped materials. When photons (with sufficient energy) are absorbed in the depletion zone an electron is promoted from the valence band up into the conduction band. This automatically produces a free electron and a positive charge carrier which flow (because of the contact voltage) into the n-type and p-type materials respectively.

As more photons arrive, more electrons and positive charge carriers are produced. A voltage is therefore also produced which causes a current to flow away from the junction and through the detection circuit. This mode of operation is referred to as photovoltaic. The advantage of this mode is that very little extraneous signal is produced due to thermal effects or external voltages. The limitations are that the measured current will only remain linear with photon flux for very small signals - if too large a signal is applied then the voltage generated can overcome the p-n junction without flowing around the

detection circuit. In this project where very low signal levels were expected the improved signal to noise ratio offered by this method made it the best option.

In photoconductive mode a reverse bias is typically applied to the p-n junction, i.e. a positive voltage is applied to n-type material and a negative voltage to the p-type material. This pulls the electrons and positive charge carriers away from the junction and has the effect of increasing the volume of the depletion zone. Once equilibrium has been reached there will be a small current flowing across the photodiode because of thermal effects producing electron-hole pairs (dark current). An additional current signal will be produced by electron-hole pairs that are generated by absorbed photons.

The reverse bias voltage prevents current from leaking (in the forwards direction) across the p-n junction and so the response of the photodiode in this mode is linear over a far wider range. The response is also rather quicker when a reverse bias is applied. The increased volume of the depletion zone has the effect of increasing sensitivity slightly in this scheme. Unfortunately dark current is also increased; this can make detection of very small signals impossible and is the reason why this approach was typically not taken on this project.

The physical area of the photodiodes used in this project was typically $\sim$ 1 cm^2. Obviously larger detectors would make alignment easier (and one could also be certain that the entire beam was being detected). However, the dark current also increases with detector area - it is therefore advisable to use a detector that is just a little larger than the expected beam size.

In the XUV region of the spectrum - since absorption is typically very high - the number of electron-hole pairs produced per photon is approximately equal to the photon energy divided by the electron-hole creation energy. For silicon based detectors this energy is $\sim$ 3.63 eV and since 25 nm equates to $\sim$ 50 eV, quantum efficiencies of approximately 13 may be expected.

The current produced by the photodiode was measured using a lock-in amplifier. (The current input of the lock-in amplifier was used because it offers a lower input impedance which is more appropriate for weak signals). Using a lock-in amplifier helped to eliminate any signal that did not occur at the expected 1 kHz frequency - such as room lights and dark current. Unfortunately the major source of noise signal is the laser and this occurs at the same frequency as the x-ray signal. It may be possible to oscillate the gas pressure in the generation capillary at a different rate to the laser - this would then allow the lock-in amplifier to distinguish the x-ray signal from the laser noise. To date this method has not been applied; spectral filtering techniques have been favoured instead.

2.4.2 Andor CCD

The photodiode (as described above) reads out a continuously varying single value that is proportional to the number of photons absorbed within it's area per unit time. In contrast a CCD camera contains a rectangular array of detectors (defined lithographically) and is capable of producing an image by reporting the number of photons arriving at each detector pixel. Unlike a photodiode a CCD camera does not read out continuously. During the specified integration time each pixel acts as a capacitor and accumulates charge by collecting the photoelectrons produced when light is absorbed in the semiconductor. The quantum efficiency is a measure of how many electrons are produced for each photon that is absorbed - this figure varies with photon energy. As can be seen from figure 2.39 at 25 nm only $\sim$ 40 % of photons are converted to electrons:

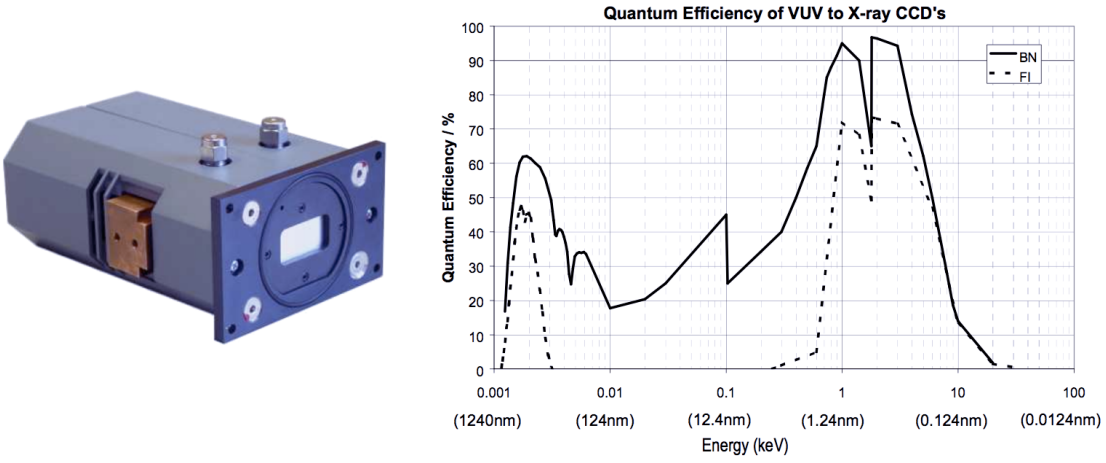

FIGURE 2.39: A plot of quantum efficiency as a function of illumination energy for the Andor camera. Note that the model used for this project was back illuminated and is represented by the solid line. This data and image are reproduced from the Andor DX434 specification sheet[46].

It is, in fact, the readout method which is described when a camera is referred to as a Charge Coupled Device. Beneath the detector array a series of semiconducting channels are implanted into the substrate. By applying alternating voltages it is possible to operate the detector array as a shift register. Using this idea, the image is read out from the detector array line by line, and within each line, one pixel at a time:

Beginning with the first line of the detector-array, charge that has been accumulated in each pixel is transferred one place over into the neighbouring pixel where it is again stored. The pixel at the end of the line transfers its charge into a current amplifier where the number of electrons is counted and recorded. This process is repeated until the whole line has been read out (let us say that it is a horizontal line and is read out

to the left). The remaining stored charges in the detector-array are now all transferred down one place so that the next line can be read out at the detector.

The Andor camera used in this project contains an array of 1024×1024 pixels. The advantage of CCD readout is that only one charge counting device is required as opposed to over a million if each pixel was read out individually. Consequently a higher proportion of the purchase price of the camera goes into the detection electronics and the quality is therefore substantially higher. Quoted readout noise for the model used in this project was just 2 electrons.

The CCD readout technique does, however, have some drawbacks. If no mechanical shutter is used the detector-array can continue to accumulate charge during readout. If exposure time is comparable to readout time this will result in smearing (particularly of bright parts of the image). Readout time for the Andor camera (full frame) is approximately 20 ms so the minimum acceptable shutter-less exposure time is approximately 200 ms and longer exposures tend to produce much sharper images.

Detection of the x-rays produced by a single laser shot is highly desirable as studying the time average of many laser shots can sometimes give a distorted view. In order to achieve single shot operation external shutter devices are needed. Initially the laser system incorporated a small aperture shutter at a beam waist between the regen and the compressor (small apertures allow short shutter times). Unfortunately, operating this shutter destabilised the temporal properties of the laser as it prevented thermal equilibrium from being reached in the external amplifier Ti:sapphire crystal. The optimum location for the shutter is in the vacuum chamber immediately after the generation capillary, however, this would require a vacuum-compatible shutter and a heat removal system. Single shot operation was demonstrated using two chopper wheels (operating at slightly different frequencies) in conjunction with a large aperture shutter - the complexity of this setup proved to be impractical for day to day use.

The data outputted by the Andor camera is digitised as a 16 bit number giving up to 65536 potential counts. In reality however the detector typically saturates in the region of $30000 - 40000$ counts. This limit is imposed by the pixel well depth, which depends primarily upon the pixel area (13×13 μm) but is also controlled to an extent by the bias voltage (gain setting). If the well depth is exceeded charge can spill over into neighbouring pixels and the image will contain regions of uniform signal that contain no detail.

In order to reduce dark current to the minimum level possible the Andor camera includes a thermo-electric cooler to lower the temperature of the detector chip and electronics. Since air cooling is not possible inside a vacuum chamber heat is exchanged from the

detector into a large copper block - this block was then attached via a flexible copper braid to a water cooling system. Additionally an uprated power supply was purchased to further increase the efficiency of the cooling. These modifications allowed the camera to operate at −50 °C despite the heating effect of translation stages and the laser. As can be seen from figure 2.40 operating at −50 °C reduces the dark noise by almost three orders of magnitude compared to that at 0 °C.

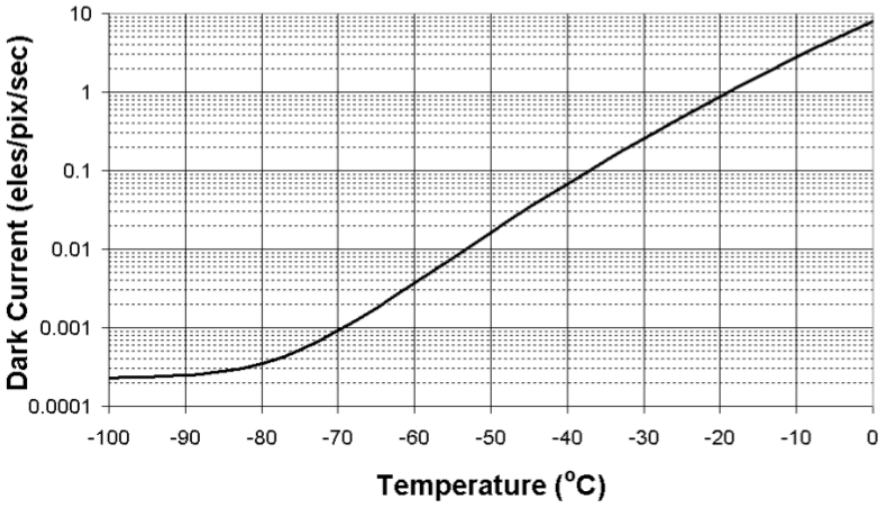

FIGURE 2.40: A plot of dark current as a function of detector temperature. This data is reproduced from the Andor DX434 specification sheet[46].

The highest intensity resolution is obtained when the pixel well is almost full for the brightest point on the image - exposure time and gain should be chosen accordingly. In some cases (particularly if the shutter could not be used) it became necessary to insert additional aluminium foil filters in order to attenuate the x-ray signal, thus allowing longer exposure times.

In some experiments (for example in diffraction measurements) the signal that is of interest (the diffracted x-rays) have a much lower intensity than the direct beam. Under these conditions it may be appropriate to mask the brightest part of the beam so that the entire dynamic range of the camera can be used to resolve the weaker signal.

2.4.3 Spectrometer

The grazing incidence x-ray spectrometer purchased for use on this project was supplied by Shultz Scientific (see figure 2.41. The housing features a vacuum flange that enables it to be attached to the existing vacuum system. At the input there is a slit holder - several different width slits were supplied and can be exchanged depending upon whether spectral resolution or signal level is preferable.

FIGURE 2.41: Grazing incidence spectrometer: 1, Entrance slit. 2, Grating. 3, MCP. 4, Lens system. 5, CCD camera. 6, Turbo pump. 7, Vacuum gauge.

50 cm beyond the entrance slit the x-ray beam is incident on a grating. The grating is again mounted in a holder to allow alternate grating periods to be used depending upon the spectral range that is to be studied. Gratings included are 300, 600 and 1200 lines/mm which provide the following spectral ranges: $60-20$, $30-10$ and $15-5$ nm respectively.

The gratings substrates are not flat; rather they represents a concave section of a sphere of 1m radius. This spectrometer uses an 'on-Rowland-circle' geometry which means that the curvature of the grating is designed to re-image the slit at the detector - consequently the slit-grating and grating-detector separations are equal to half the radius of curvature of the grating. Note that the gratings are used at grazing incidence and have a blaze angle of 2.4° to maximise the throughput.

After travelling a further 50 cm the now spatially dispersed x-ray beam arrives at a micro-channel-plate (MCP) detector. This plate (which is manufactured from a material of high electrical resistivity) is also curved to match the Rowland circle geometry and has

passing through it an array of 12 μm diameter channels. A high voltage is applied across the MCP which accelerates any photoelectrons that are ejected when x-rays collide with the channel walls. Each micro-channel then acts as a tiny photomultiplier tube and amplifies the electron signal as it cascades through the MCP.

A second high voltage accelerates the electrons that are ejected from the micro-channels towards a phosphor screen. The visible light emitted by the phosphor screen is brought to a flat viewing screen by an optical fibre plate. A lens system and scientific camera (Princeton Pixis CCD) are then used to collect the optical signal. The signal is spectrally resolved in the horizontal plane but maintains its spatial properties in the vertical direction. Consequently the spectrum can either be viewed as an image or can be 'binned' to produce a line plot.

The spectrometer was supplied with a gating pulse generator which can be used to apply a 1 kV potential across the MCP for a very short time (few hundred ns). This was intended for making time resolved experiments with continuous, high-intensity UV sources. Synchronising the gate pulse and the laser allowed some measurements to be made, however, the spectrometer was later modified to allow a continuous potential of 750 V to be applied across the MCP. The use of continuous voltage allowed us to integrate signal over many successive laser shots - this was necessary as the efficiency of the spectrometer was considerably lower than that of the direct x-ray detection methods. The voltage between the MCP and the phosphor screen was in all cases kept at a constant 5 kV.

The points made in section 2.4.2 regarding dynamic range and saturation are equally applicable to the Pixis camera, although operation is simplified as the Pixis camera does include a shutter and short exposures are unlikely to be required.

Chapter 3

Capillary High Harmonic Generation

This chapter discusses several of the key areas in which, through experimentation, we have expanded our understanding of the process of HHG in a capillary. Section 3.2 describes the steps that were taken in order to achieve an efficient launch of the laser into the generation capillary; following these steps throughputs of up to 95 % have been reached. Section 3.3 provides details of two series of experiments: In the first, gas pressure and the input laser power are varied whilst the total x-ray flux is measured on an x-ray diode. In the second, pressure and power were again varied whilst the x-ray output spectrum was observed. These experiments demonstrate a peak in total x-ray flux at approximately 120 mbar whilst phase matched generation for the highest harmonics occurs at $\sim$ 70 mbar. Section 3.4 is concerned with an experiment to determine the effect of input laser pulse spectral phase on x-ray generation. Control of x-ray flux and shifting of spectral peak positions is demonstrated. This chapter begins however with a brief explanation of some of the concepts underlying the HHG process:

3.1 Theory of HHG

HHG is a high order non-linear process in which intense laser radiation is up-converted to much shorter wavelengths. In order to attain the very high electric field strengths that are necessary to observe this effect, a femtosecond pulsed laser source is typically required. This usually implies operation in the infrared to take advantage of the large gain bandwidth that Ti:sapphire provides. The highest frequency to which the light is up-converted is dependent primarily on the peak electric field strength but with current laser systems it is possible to generate radiation down into the XUV and soft x-ray

regimes. In this work (using a $\sim$ 35 fs, 1 mJ pulse) generation of radiation down to approximately 25 nm is observed (with argon or other noble gasses as the conversion medium).

The process that leads to this frequency conversion can be explained in simple terms using the 'three step' model[47]. This model brings together the macroscopic external measurements that can be made (for example the x-ray spectrum) with an explanation of the processes that occur on a microscopic level to individual atoms within the interaction medium.

3.1.1 Three Step Model

This semi-classical model describes the response of a single atom to an intense, linearly polarised, oscillating electric field. We consider the problem from the reference frame of the atomic nucleus so that as time advances the laser pulse propagates past the atom:

Prior to the arrival of the laser pulse the atom is in its ground state with the electrons trapped by the Coulomb potential. As the laser pulse propagates past the atom the magnitude of the electric field oscillations increases rapidly until the peak of the pulse is reached. The actual electric field experienced by the electrons is the superposition of the Coulomb potential and the laser electric field. If the amplitude of the laser electric field is sufficiently high it can distort the potential well to such a degree that the weakest bound electron undergoes tunnel ionisation. This is step 1.

Although ionised, because of its close proximity, the electron remains associated with its parent ion. The electron is correctly described as occupying a continuum energy level of the atom. In this state the electron is now free to move under the influence of the electric field. The trajectory followed by the electron will depend upon the phase of the laser electric field at the instant of ionisation. In the case that contributes to HHG, the electron is accelerated away from the nucleus until the sign of the laser electric field is reversed. A dipole moment (or polarisation) is therefore induced in the atom. Classically this can be thought of as a displacement between the weakest bound electron and its parent ion. The electron is then accelerated back towards the ion core, accumulating kinetic energy as it goes. This is step 2.

In the final step of the model the electron is recaptured by its parent ion. Note that the probability of this process occurring is low - which impacts upon the efficiency of the HHG process. Upon re-collision the electron must give up its ionisation potential energy and any kinetic energy that it has accumulated. This energy is given out in the form of a high energy photon - this is the XUV output that is observed.

This simple semi-classical model for the microscopic behavior of a single atom explains a great many of the macroscopic features observed in the HHG output. For example the spectral output from the HHG process is found to consist of a series of peaks each occurring at the odd harmonics of the input laser frequency. This is explained by the fact that electrons can only be ionised during a brief time window in each half optical cycle whilst the laser amplitude is sufficiently high. Only a subset of those electrons that are ionised will follow trajectories that cause them to return to the nucleus with high kinetic energy. The recombination and consequently the x-ray emission will therefore only occur in very short bursts each half optical cycle. It is this temporal profile that (via the time-frequency Fourier relation) leads to the observed odd harmonics in the spectral domain.

The model above also explains the formula for the maximum achievable photon energy E_{max}. The 'cutoff' energy E_{max} is related to the ionisation potential of the interaction medium (I_p) and the ponderomotive potential of the laser (U_p):

$$E_{max} = I_p + 3.17U_p \tag{3.1}$$

U_p is a measure of the kinetic energy stored in the electron oscillations and is given by:

$$U_p = \frac{q_e^2 E^2}{4m_e\omega_0^2} \tag{3.2}$$

where q_e is the electron charge, m_e is the electron mass, E is the amplitude of the laser electric field and ω_0 is its angular frequency[48, 49].

The numerical factor in equation 3.1 is defined by the electron trajectories that return to the nucleus with the maximum possible kinetic energy[50].

As a final example, the model explains why drastic reduction in XUV output is observed when elliptically polarised light is used - the returning electron trajectories are deviated from the parent ion core and thus the probability of recombination is reduced.

Predictions made from this semi-classical model are in good agreement with those made from the more rigorous quantum theory[51].

3.1.2 Beyond the three step model

In the semi classical model described in section 3.1.1 the polarisation induced in an atom by an external laser field is explained as a physical displacement between the electron and

its parent ion (where both are considered to be inelastic particles). Quantum mechanically this is described in terms of the expectation value for position of the wave-packet of the outermost electron relative to that of the parent ion. For an ionised electron undergoing approximately free motion in the electric field the most probable quantum electron trajectories are the same as the classically calculated ones. This explains why the three step model works so well. Classical models, however, cannot correctly describe the ionisation and electron re-combination processes. The three step model therefore incorporates the concept of quantum mechanical tunnel ionisation (which is why it is termed a semi-classical model).

The full quantum model[52] reveals important subtleties in the HHG process. For example, the electron wavepacket is subject to wave effects such as diffusion and interference as it proceeds along its trajectory. Furthermore, the phase of the returning electron wavepacket plays an important role in the recombination which affects the generation efficiency, the coherence and the subsequent propagation of the x-ray beam[53].

3.1.3 Regimes of electric field strength

The induced polarisation can be used to describe the interaction between light and a transparent medium. I.e. the electric field of the light induces an oscillating dipole in the atom and this oscillating dipole then re-emits the radiation. The polarisation induced for a given electric field varies between materials, depending upon, for example, the amount of shielding provided by inner electrons. The propensity of a material to accept a polarisation is given by the electrical susceptibility χ which is related to the refractive index of the material by:

$$n = \sqrt{(1 + \chi)} \tag{3.3}$$

When the electric field amplitude E is very low the electrons oscillate with simple harmonic motion. This is the regime of linear optics. In this case the induced polarisation P is given by:

$$P = \epsilon_0 \chi E \tag{3.4}$$

As the electric field amplitude is increased so too does the excursion of the electrons. For large enough electric fields, neighbouring atoms may begin to influence the motion of the electrons, this will cause their response to differ from the pure linear response that is observed for low powers. This regime is referred to as that of perturbative nonlinear

optics. In order to describe the nonlinear response of the electrons the single value for χ is replaced with a Taylor series expansion. Expression 3.4 now becomes:

$$P = \epsilon_0\chi^{(1)}E + \epsilon_0\chi^{(2)}E^2 + \epsilon_0\chi^{(3)}E^3 + \epsilon_0\chi^{(4)}E^4 + \dots \tag{3.5}$$

A great range of nonlinear optical phenomena, such as, second harmonic generation (SHG), self-phase modulation (SPM) and difference frequency generation (DFG) all take place in this regime and can be explained by considering the additional frequency components created by the nonlinear response of the oscillating dipole[54, 55].

The restoring force for the anharmonic oscillations of the electron wave-packet originates from the Coulomb potential. The perturbative approach to nonlinear optics begins to lose its validity when the magnitude of the laser electric field becomes comparable to the Coulomb potential. This is because there is an increasing probability of optical field ionisation. The Keldysh parameter γ provides a convenient check of the regime of operation:

$$\gamma = \sqrt{\frac{I_P}{2U_P}} \tag{3.6}$$

For values of the Keldysh parameter $\gamma \gg 1$ it is only possible to ionise by simultaneous absorption of multiple photons - this is the multi-photon ionisation regime and perturbative nonlinear optics may be applicable. The ionisation rate in this case is proportional to the N^{th} power of the laser intensity where N is the number of photons required to meet the ionisation potential[56].

When the Keldysh parameter $\gamma \lesssim 1$ the dominant mechanism for ionisation is optical tunnel ionisation. I.e. the Coulomb potential barrier is sufficiently narrowed by the laser electric field that there is an appreciable probability of the electron escaping the barrier by quantum mechanical tunneling. This is the regime in which HHG occurs and since tunnel ionisation is an integral part of the HHG process and the HHG yield is strongly dependent upon the rate of ionisation. The rate of ionisation can be calculated using the theory of Keldysh[57] or with slightly improved accuracy using the approach of Ammosov, Delone and Krainov (ADK)[58].

For $\gamma \ll 1$ the laser E-field distorts the Coulomb potential to such a degree that the barrier is pulled below the ground state level of the weakest bound electron. The electron is then ionised directly by over-the-barrier (OTB) ionisation.

At still higher electric fields the ionised electrons are essentially free moving. The effect of the Coulomb potential has become negligible and the electrons simply follow the

electric field linearly. Since there are no nonlinear frequency components added - this regime cannot be used for HHG - it is therefore possible to saturate the HHG yield if too high an electric field is used.

The nonlinear response of the electrons ultimately returns as the electrons reach relativistic speeds. It is now no longer possible for the electrons to keep up with the oscillations in the electric field. This effect is exploited in a free electron laser (FEL) where a beam of electrons is accelerated through a set of magnets with alternating poles. The effect of the magnetic field is to 'wiggle' the relativistic beam of electrons laterally as it passes through the device. From the viewpoint of the electron the magnetic field appears to oscillate so quickly that the electrons would have to exceed the speed of light in order to respond linearly to the applied field. Since this is not possible the motion of the electrons again becomes nonlinear which leads to generation of additional spectral components. The spectral range of operation of a FEL is astonishing, ranging from the microwave to the soft x-ray regions.

3.1.4 Phase matching

The efficiency of HHG is significantly improved by proper phase matching. When properly phase matched, x-rays emitted from spatially separated atoms can add together constructively allowing total x-ray flux to accumulate with distance. This is possible because the laser (which stimulates the emission of both x-rays) provides a coherent link between the phase of emission at both locations. Phase matching therefore requires the x-ray and laser propagation constants to satisfy the following relation:

$$k_q = qk_f \qquad \text{or alternatively we can define}\ldots \qquad \Delta k = k_q - qk_f \tag{3.7}$$

where q is the harmonic order, k_q is k-vector of the q^{th} harmonic and k_f is the k-vector of the laser. Δk can therefore be used as a measure of how close to phase matching a particular harmonic is.

When $\Delta k \neq 0$, this means that the laser radiation and generated x-rays will slip in and out of phase as they propagate. The x-rays will only accumulate over the distance it takes to introduce a phase shift of π radians. This distance is referred to as the coherence length (l_c), i.e. $l_c = \pi \ / \ \Delta k$. If the length of the HHG medium is increased to twice the coherence length the x-rays generated in the first coherence length will be 180° out of phase with those produced in the second coherence length. Consequently the x-ray output oscillates as a function of the length of the HHG medium. Peaks in x-ray output occur at odd multiples of the coherence length and zeros occur at even multiples. The

x-ray output never exceeds the value reached after the first coherence length. Away from phase matching HHG efficiency is limited by the coherence length and is therefore highly inefficient.

X-ray output would be drastically improved if it could be arranged that $\Delta k = 0$ along the whole length of the HHG medium (i.e. infinite coherence length). Under these conditions the x-rays and laser radiation are perfectly phase matched such that the emission from subsequent atoms is always in phase with previously generated x-rays. The x-rays will therefore accumulate along the length of the generating medium. The number of atoms that contribute to the output increases linearly with the length of the generation medium consequently the x-ray electric field amplitude also increases linearly. Since intensity $I = |E^2|$ the x-ray intensity builds up as the square of the distance propagated through the generation medium.

In practice the rate of x-ray intensity build-up is reduced by absorption. When close to phase matching (i.e. for long coherence lengths) the limiting factor becomes the absorption length (l_a) within the generation material. This is given by the expression $l_a = 1/\sigma\rho$, where σ is the ionisation cross section and ρ is the gas density. If the length of the HHG medium is greater than a few absorption lengths then x-rays from the start of the generation region will be re-absorbed before they reach the output. This re-absorption leads to saturation of the x-ray intensity build-up whereby increasing the generation length significantly beyond the absorption length yields a diminishing increase in intensity.

HHG output is maximised when the generation medium length l is long enough to achieve saturation. Both l_c and l_a play important roles in phase matching with whichever is the shorter distance becoming the limiting factor to the x-ray generation. The ratio of l_c to l_a determines the level of x-ray intensity that is reached at saturation. According to Constant *et al.*[59] HHG efficiency is optimised when $l > 3l_a$ and $l_c > 5l_a$.

Phase matching of an HHG source is typically achieved using one of the following two methods:

In gas jet systems the focus position is translated relative to the gas jet in order to take advantage of the Gouy phase shift[60, 61] (a phase shift of π that occurs about the focus). Somewhere in the vicinity of the focus will be a region in which the phase change due to the Gouy shift is sufficient to counteract the walk-off between the x-ray and laser phases - thus facilitating phase matching. This technique can typically only phase match over a very restricted region.

In a capillary waveguide HHG system there is the potential to maintain phase matching over a far longer interaction region. In this scheme the waveguiding effect of the capillary

modifies the propagation constant of the laser sufficiently that phase matching can be achieved by fine control of the gas pressure. According to[62] the propagation constants for the laser and for the x-ray can be calculated using:

$$k \approx \frac{2\pi}{\lambda} + \frac{2\pi P(1-\eta)\delta(\lambda)}{\lambda} + (1-\eta)n_2 I - P\eta N_{atm} r_e \lambda - \frac{u_{11}^2 \lambda}{4\pi a^2} \tag{3.8}$$

where λ is the wavelength under consideration, P is the pressure in atmospheres and η is the ionisation fraction. Now, considering each term on the right individually:

The first term simply represent the normal phase evolution under a vacuum.

In the second term $\delta(\lambda)$ represents the linear dispersion due to the neutral gas. The factor of $P(1-\eta)$ weights this term depending on how much of the gas remains un-ionised.

In the third term the factors n_2 and I are the nonlinear refractive index of the neutral gas and the intensity of the laser pulse. This term is also weighted according to the ionisation fraction. This term describes the dispersion due to nonlinearity in the neutral gas.

The fourth term calculates the dispersion from the plasma. N_{atm} is the number density at one atmosphere and r_e is the classical electron radius. The classical electron radius is related to the Bohr radius (a_0) by the fine structure constant (α): $r_e = \alpha^2 a_0$.

The final term describes the dispersion added by the capillary waveguide of radius a. This term can be omitted for the x-rays because their wavelength is so small in comparison to the capillary radius - effectively the x-ray do not 'see' the capillary. u_{11} is the first root of the J_0 Bessel function (this could be substituted by u_{nm} in order to describe the laser propagating in higher order modes).

Equation 3.8 is used to calculate k for the laser centre wavelength and for each harmonic order. Over the duration of the laser pulse the first and final terms should remain approximately constant[†]. The second, third and fourth terms all contain η and will therefore vary drastically over the duration of the laser pulse as progressively more of the generation medium is ionised. The third term (nonlinear refractive index n_2) typically makes a rather small contribution to the phase matching. The primary interplay that leads to phase matching is therefore between the neutral gas dispersion (second term) and the plasma dispersion (fourth term) and since it affects only the laser light, the waveguide dispersion (fifth term).

[†]Note that this is only true if there is no spectral shifting of the x-ray or laser as it propagates along the capillary. In practice blue-shifting of the laser is frequently observed as described in section 3.4.

At low intensities (and therefore low ionisation fractions η) phase matching is achieved by balancing the linear dispersion of the neutral atoms (second term) with the effect of the capillary waveguide (fifth term). As η increases so too does the plasma dispersion (fourth term) whilst at the same time the linear dispersion from the neutral atoms (second term) is reduced. Since, initially, the neutral dispersion term is larger in magnitude than the plasma term the change in η has a larger affect on the magnitude of the neutral term. Consequently phase matching is maintained as η increases by increasing the pressure.

Higher harmonics are generated when the laser electric field is more intense - this corresponds to a time when the ionisation fraction η is also high. As a result the higher harmonics are phase matched towards higher pressures.

Once the ionisation fraction reaches a critical level (when the magnitudes of the neutral and plasma terms are approximately equal) it is no longer possible to achieve phase matching by increasing the pressure. Above this critical level of ionisation, increasing the pressure produces a larger increase in the plasma term than it does in the neutral term. Since the sign of the plasma term is the same as that of the waveguide term the net result is to move away from phase matching. The critical ionisation fraction η_c can be calculated using the following equation[63]:

$$\eta_c(\lambda_L) = \left[\frac{\lambda_L^2 N_{atm} r_e}{2\pi\Delta\delta(\lambda_L)} + 1\right]^{-1} \tag{3.9}$$

Where λ_L is the laser wavelength, and $\Delta\delta(\lambda_L)$ is the difference between the refractive indices of the laser centre wavelength and the harmonics. According to calculations based on ADK ionisation theory the critical ionisation fraction for phase matching in argon is 4.8 %[62] (assuming a laser centre wavelength of 800 nm).

A phase matching cutoff (the highest harmonic that can be phase matched) is imposed by η_c the critical ionisation fraction. Note that steepening the leading edge of the laser pulse, or using a shorter pulse duration allows a higher electric field (and therefore harmonic) to be reached for the same ionisation fraction. The critical ionisation fraction is typically achieved on the leading edge of the laser pulse. The phase matching cutoff is therefore well below the theoretical cutoff of the peak of the laser pulse[64].

An additional effect of increasing the ionisation fraction is to increase l_a the absorption length. This is because the x-ray absorption of the plasma is very much lower than that of the neutral gas. Again, since the higher harmonics are generated at higher ionisation fractions it is these harmonics that will see the greatest reduction in absorption. Clearly the reduction in absorption allows the harmonic light to accumulate over a longer distance and therefore to reach higher levels of output.

3.1.5 Quasi phase matching

In quasi phase matching (QPM) it is accepted that there is a non-zero value for Δk. As discussed in section 3.1.4 the fundamental and the generated light will therefore slip in and out of phase, leading to an x-ray intensity that does not build-up but rather oscillates with the length of the generation medium. If a periodic modulation of the generation efficiency is introduced, the x-ray intensity can build-up (in a stepwise fashion) along the length of the generation medium.

Without QPM x-rays generated in the second coherence length are out of phase with those generated in the first. Suppose that HHG could be switched off during the second coherence length (and every subsequent even multiple of the coherence length). All x-ray generation would then contribute positively to increasing the x-ray output.

QPM does not require complete cessation of HHG or exact intervals of the coherence length - these are merely the conditions for maximum rate of build-up in x-ray output with length. Even very small modulations in the HHG efficiency will be sufficient for QPM to achieve a build-up along the length of the generation medium. The period of the modulation is also not critical. An increase in x-ray intensity with generation length can always be achieved provided there is a net positive contribution to the x-ray output over each cycle of the modulation.

Since Δk need not be zero, QPM can operate at ionisation fractions $\eta > \eta_c$. This allows far more efficient use of the energy of the laser pulse; rather than just the leading edge - the whole pulse can now contribute.

A number of schemes for modulation of the HHG efficiency have been demonstrated. A variety of different parameters have been used to introduce the modulation for example: the gas pressure[65][†], the laser polarisation[66], the laser intensity with the aid of counter propagating pulse trains[67, 68] or through capillary mode beating[69].

During the course of this work an acoustic QMP scheme was attempted. A cylindrical piezo transducer was mounted on the HHG capillary and a periodic voltage was applied. The intention was to set up a stationary acoustic pressure wave along the length of the capillary - varying the drive frequency would provide tunable QPM. Unfortunately initial tests proved unsuccessful. It was not possible to couple sufficient energy into the pressure wave and no effect could be observed in the x-ray spectrum.

[†]The modulation of HHG efficiency in this case may in fact be due to modulation of the laser intensity rather than the pressure as stated by the authors.

3.1.6 Motivation for use of capillary geometry.

The capillary radius provides an additional parameter that may be used for optimisation of the phase matching process. Aside from the important role it plays in phase matching, the confinement of the capillary waveguide provides a number of other benefits - it is these factors that motivated the use of a capillary in this work:

The capillary allows high laser field strength to be maintained over a long distance. This increases the volume of the generation medium that may contribute to the HHG output in comparison to the gas jet geometry.

Since the laser propagates in a capillary mode the peak laser power is maintained at an approximately constant level throughout the generation medium. In the case of the gas jet there is a rapid change in peak laser power outside of the Rayleigh range, considerably complicating the phase matching.

The capillary helps to overcome defocusing effects. Since the intensity of the laser beam is highest at the centre, nonlinear self-focusing may occur, however, this effect is typically far outweighed by the similar but opposite effect of plasma defocusing[†]. In a gas jet this defocusing would further shorten the region where laser intensity is sufficient for HHG. In the capillary geometry, plasma defocusing shifts energy into higher order capillary modes. Although not optimal, energy from these higher order modes may still contribute towards the HHG output rather than simply being discarded.

The capillary geometry also offers the potential for very neat integration with fibre based systems. For example, a fibre coupled input laser could be used. Perhaps there is even the possibility of a vacuum-core, grazing-incidence, x-ray fibre output with a focusing taper tip.

The capillary offers a useful platform for the development of QPM schemes. The beating between capillary modes may even provide a means of modulating the laser intensity for this purpose[70].

[†]Since the refractive index of the plasma is considerably lower than that of the neutral gas ionisation at the centre of the beam causes a strong defocusing effect.

3.2 Launch

Understanding and optimising the launch conditions for an HHG capillary represented a critical step in the process of obtaining efficient x-ray generation. In the following subsections capillary modes 3.2.1, focusing conditions 3.2.2, the importance of laser mode quality 3.2.3 are discussed.

3.2.1 Capillary Modes

In order to describe the process of launching light into a waveguide it is useful to introduce the concept of modes. This concept will only be covered briefly here for a fuller discussion the reader is referred to the work of Marcatili and Schmeltzer[71]. There are, however, some unusual features of the modes of a hollow capillary which it is worth emphasising.

The purpose of waveguide modes is to provide a set of electric field distributions into which any arbitrary E-field pattern can be decomposed. Propagation calculations need then only be performed for the modes - the propagation of arbitrary E-fields being described as a linear superposition of modes with the proper intensities and phases.

The appropriate mode set to use for the decomposition will depend upon the geometry in which the light propagates. For example in free space the cylindrical transverse electro-magnetic modes (TEM) modes may be used. These modes are termed transverse because they have no component of the electric or magnetic fields in the direction of propagation. In the case of optical fibres the waveguide is often assumed to be weakly guiding (the refractive index of the core $n_1 \approx n_2$ that of the cladding). Polarised light that is launched into such a fibre is expected to maintain its linear polarisation. This leads to a scalar approximation (where the electric field is polarised only in the x direction and the magnetic field is polarised only in the y direction) and results in the LP (linear polarised) modes. In general the electric and magnetic fields are vectors and have non-zero components of electric and magnetic field radially, azimuthally and axially; giving rise to the electrical hybrid (EH) modes. The stronger the guiding - the larger the effect of the axial component of electric and magnetic fields.

Clearly in the case of a hollow dielectric waveguide (as in our experiment) there is a large difference between the refractive index of the core and the cladding (= strong guiding) it is therefore important to use the EH modes. Before the development of optical fibres the modes of hollow capillaries were mooted as a means of optical communication; these were studied in detail by Marcatili [71].

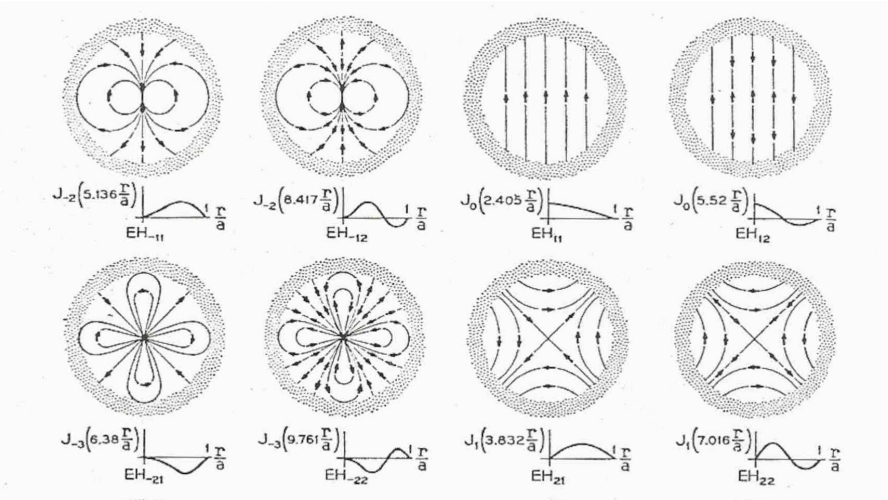

FIGURE 3.1: Electric field line diagrams and radial intensity profile for the lowest order electrical hybrid modes of a hollow dielectric waveguide. Taken from reference [71].

Figure 3.1 shows the electric field diagrams of the lowest order EH modes. Notice that the EH_{11} mode has an approximately Gaussian intensity distribution and also that it maintains a linearly polarised electric field. Because of these features it should be possible to achieve an efficient launch into this mode using only a lens (matching the beam from our CPA laser system to this mode profile).

The fact that the EH_{11} mode is linearly polarised is very important as linear polarisation is required for efficient HHG. If polarisation deviates far from being linear, electron trajectories do not necessarily return to the parent ion core. This reduces the probability of recombination and hence the x-ray flux.

For typical capillary materials the EH_{11} mode also has the advantage of offering the lowest propagation loss[†], this ensures that total throughput from the capillary, and hence pulse energy at the exit of the capillary, are kept as high as possible.

There is an important distinction to be made between the modes of an optical fibre and those of a hollow capillary. When launching into a hollow capillary it is inevitable that there will be a small amount of energy launched into higher order modes, see subsection 3.2.2. Single mode operation as utilised in optical fibres is not possible. The fundamental difference between optical fibres and capillary waveguides is that the refractive index of the capillary wall is expected to be higher than that of the hollow core. In optical fibres this situation is usually reversed. The expression below 3.10 shows how to calculate the cutoff wavelength λ_c for a step index fibre (operating in a particular

[†]EH_{11} offers the lowest loss for capillary wall refractive indices of up to 2.02. For refractive indices greater than 2.02 the TE_{01} mode exhibits the lowest attenuation. The refractive index which minimises attenuation is $\sqrt{3}$. These figures are quoted from [71] and assume a vacuum inside the capillary.

optical mode denoted by the indices l and m). For any wavelength above this cutoff value the longitudinal wavenumber (k_z) for the (lm) mode becomes imaginary and its field decays exponentially along the waveguide.

$$\lambda_c(lm) \;=\; \frac{2\pi a}{V_c(lm)}\sqrt{n_i^2 \;-\; n_e^2} \tag{3.10}$$

Here a is the radius at which the boundary between the inner and outer materials occurs, n_i and n_e are the refractive indices of the internal and external materials respectively. $V_c(lm)$ is the cutoff normalised frequency for the mode (lm). For the EH_{11} mode (and consequently for single mode operation) this should be set to 2.405 (the first root of the Bessel function J_0). For an optical fibre $n_i \;>\; n_e$ the square root therefore has a real value and a cutoff wavelength is observed at λ_c. For a hollow capillary $n_i \;<\; n_e$ the square root now has an imaginary value and so too does λ_c, consequently no cutoff is observed experimentally.

3.2.2 Focusing

Since linear polarised light and low propagation loss are desirable we aim to launch as efficiently as possible into the EH_{11} mode of the generation capillary. Optimal launch efficiency is obtained when the input laser E-field is as closely matched as possible to the profile of the EH_{11} mode. In this subsection the effect of focused laser spot size (relative to capillary bore) on launch efficiency is discussed; this is based on theoretical work by Nubling *et al.* [72]. On-axis coupling between a Gaussian input laser spot and the EH modes of the capillary is calculated as an overlap integral. Since only on-axis coupling is considered only the EH_{1m} modes can be excited. The spatial profiles of these EH_{1m} modes are approximated by the zero order Bessel function:

$$E(r) = E_0 J_0(u_m \frac{r}{a}) \tag{3.11}$$

the Gaussian spot is expressed:

$$E(r) = E_0 \exp(-r^2/w^2) \tag{3.12}$$

the overlap integral for coupling efficiency into each mode (η_m) is thus:

$$\eta_m = \frac{\lfloor \int_0^a \exp(-r^2/w^2)\; J_0(u_m \frac{r}{a}) r \; dr \rfloor^2}{\int_0^\infty \exp(-2r^2/w^2) r \; dr \;\; \int_0^a J_o^2(u_m \frac{r}{a}) r \; dr} \tag{3.13}$$

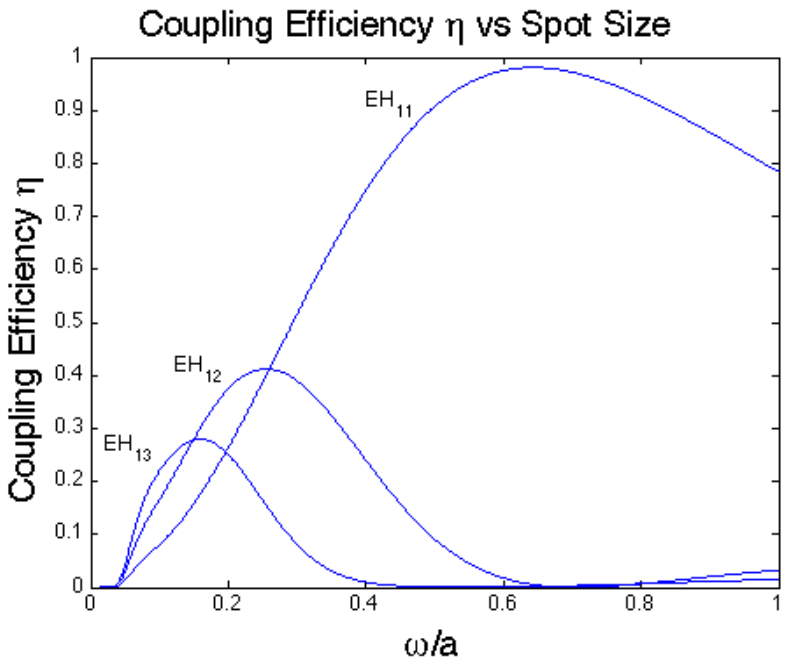

FIGURE 3.2: Calculated coupling efficiency for the first three EH_{1m} modes as a function of input spot size.

where u_m is the m^{th} root of J_0 (the zero order Bessel function), r is the radial coordinate and varies from 0 to a (the capillary radius) and w is the $1/e^2$ radius of the input laser intensity profile.

By evaluating this integral for the first three EH_{1m} modes it is possible to create a plot of coupling efficiency (per mode) against the input spot size (relative to the capillary bore). This plot is shown in figure 3.2:

Optimal coupling into the EH_{11} mode occurs for $w/a \simeq 0.64$. In this situation coupling efficiency is 98.1 % into EH_{11} but still 0.5 % is coupled into the EH_{12} mode. Assuming a perfect Gaussian beam the following relation can be used to calculate the lens focal length f required for the best launch:

$$f = \frac{\pi D w}{2\lambda} \quad \Longrightarrow \quad \frac{\pi \times 15e^{-3} \times (0.64 \times 75e^{-6})}{2 \times 780e^{-9}} = 1.45 \tag{3.14}$$

where D is the $1/e^2$ diameter of the input beam when it hits the lens and λ is the wavelength of illumination. w is the $1/e^2$ radius of the focused laser intensity profile; setting this to $0.64 \times a$ selects the optimum coupling situation. Using our measured beam diameter of 15 mm and a capillary bore radius of 75 μm we find that the optimum lens focal length is 1.45 m

3.2.3 Laser M^2

As discussed in sub section 2.2.4 the quality of a laser beam is of critical importance in nonlinear optics because of its effect on peak laser power. In the case of capillary HHG the laser mode quality is doubly important because it also has a significant impact on the launch into the generation capillary.

M^2 values greater than 1 allow laser beams that do not have a perfect Gaussian profile to be modelled using the Gaussian beam propagation equations. (The equations are calculated for a hypothetical, underlying Gaussian whose radius is everywhere 1/M times the measured beam radius). Beams with $M^2 > 1$ therefore diverge more quickly than their Gaussian counterparts. M^2 is an all encompassing parameter and there are many possible aberrations that can combine to cause this increased divergence. The effect of poor beam quality on capillary launch will vary considerably depending on the nature of the aberration and the exact alignment of the input spot relative to the capillary axis.

For example consider a beam which passes through the focusing lens slightly off-axis; this will result in comatic aberration. In the focal plane the input spot will take on a cone-shaped intensity profile. Supposing the centre of this intensity profile is aligned to the capillary axis - part of the input energy will be coupled into the EH_{1m} modes. The energy at the edges of the cone shape is away from the capillary axis -part of this energy will be launched into lossy high order modes, however, if the coma is severe enough this energy could even cause damage to the capillary facet[†].

Consider now the effect of chromatic aberration in the focusing lens - light of different wavelengths is brought to a focus at different distances from the lens. In the absence of other aberrations and in the case of an on-axis launch the EH_{1m} modes will be excited, however, since the input spot size will be a function of wavelength so too will be the mode coupling. Red and blue edges of the pulse may be launched into lossy high order modes whilst the central wavelengths are launched into the EH_{11} mode.

Each of the aberrations discussed above will launch into a different mixture of capillary modes. This can be expressed as a set of mode coefficients. These coefficients can be calculated in the case of sub-optimal launch conditions by extending the concept of equation 3.13 so that an on-axis Gaussian input profile is no longer assumed. The overlap integral between an arbitrary input beam profile and the capillary mode profiles can then be evaluated numerically.

[†]We have found that damage to the capillary facet further reduces overall coupling efficiency.

Experimental observation of the capillary output mode reveals how sensitive this is to input beam profile and to capillary alignment. Poor quality laser beams are likely to couple a large fraction of their power into higher order modes. Not only do these higher order modes experience more loss - they also lead to a highly unstable system where very small changes in alignment cause large alterations in the output mode profile. This results in HHG emission that is critically dependent on capillary alignment, making the collection of consistent data sets very difficult. Further notes on the capillary alignment process can be found in appendix A.

If a poor quality beam must be used then the focusing conditions should be amended to take into account the increase in focused spot size; i.e. a shorter focal length lens should be used. The M^2 value can be incorporated into equation 3.14 so that it becomes:

$$f = \frac{\pi D w}{2\lambda M^2} \tag{3.15}$$

It was found experimentally that (using the original laser system) an f = 0.6 m lens[†] produced the best coupling efficiency ($\sim$ 85%). Using the upgraded laser system efficiencies approaching $\sim$ 95 % have been achieved using an f = 1 m lens. Re-arranging equation 3.15 and inserting the experimental values for D, w and λ we can produce an estimate of the M^2 values for the two laser systems. For the original laser system we obtain $M^2 \approx 2.4$; for the new system we obtain $M^2 \approx 1.4$.

The M^2 values calculated above are rather higher than the experimentally measured ones which are $\sim$ 1.25 and $\sim$ 1.1 for the old and new laser systems respectively. There are several factors that contribute to the observed differences:

Firstly, the above calculation assumes collimated illumination of the lens - in reality the laser beam is diverging when it arrives at the lens. This pushes the focal plane further away from the lens - further increasing f.

Secondly, the theoretical calculation does not include any factors such as random beam walking due to air fluctuations or asymmetry of the focal spot. In practice these factors mean that it is easier to align the capillary if a slightly smaller than optimal spot size is used. This leaves a slight margin for error - reducing continuum generation and damage to the capillary facet.

If the measured values for M^2 (which may in fact be slightly optimistic) are correct then the effective focal lengths in the two cases must be 1.16 and 1.3 m for old and new lasers

[†]The focal length figures for the lenses are quoted at 633 nm whilst our laser operates at $\sim$ 780 nm. Since the dispersion of UV grade fused silica is very low the change in focal length is only $\sim$ 1% (i.e. the focus shifts $\sim$ 1 cm further from the f = 1m lens at 780 nm).

respectively. In the case of the new laser system this is not far from the measured lens to focus distance. For the old system, however, it seems that either the lens selection was not optimal or the M^2 was worse than the measured value suggested - this would explain the lower launch efficiency.

The best course of action is to directly measure the laser divergence and incorporate this into the calculation of the required focal length. Careful optimisation of laser alignment in coordination with M^2 monitoring is also essential. In addition any superfluous optics or material that may degrade the beam quality should be removed from the optical path. Ensure M^2 is as low as possible and then start with the calculated focal length. In practice it may be necessary to adjust the focusing conditions slightly in order to account for the beam walking and to achieve the highest possible HHG efficiency.

3.3 Pressure and Power dependence

After designing and constructing our capillary HHG apparatus the first experiment undertaken was to measure the x-ray output as a function of the argon gas pressure passed into the capillary. At this stage it was not known whether any x-ray generation was occurring at all. The conjecture was that the laser output from the capillary would remain approximately constant regardless of the gas pressure whilst the x-ray output would disappear if the gas interaction medium was removed. A pressure-dependent signal would therefore indicate that x-rays were being generated by the HHG process.

The setup of the experiment is as follows:

- The HHG capillary was held under vacuum in the manner described in section 2.3.1. This setup incorporates electronic control of the gas pressure passed to the capillary.
- The laser was passed through a lens and through a glass window at the front of the vacuum system to bring the beam to a focus at the capillary facet. (Fused silica was used both for the lens and the window in order to minimise the spectral phase added to the beam).
- The capillary was aligned to the laser beam in accordance with the discussion in section 3.2 and appendix A.
- An aluminum filter was used to attenuate the laser beam (as described in section 2.3.2.1. This was used in conjunction with a light baffle plate similar to the system described in 2.3.2.2.
- An x-ray photodiode (as described in section 2.4.1) produces a signal proportional to the x-ray photon flux - this was measured using a lock-in amplifier (Stanford SR830) synchronised to the trigger output from the Q-switched Evolution pump laser.

The experiment is then carried out simply by changing the input pressure and recording the current produced by the photodiode. This experiment was performed several times and so it was automated using labview. Pressure setpoint values were sent to, and instantaneous pressure values were read from, the pressure controller via a USB to RS485 adaptor. A GPIB card and connection were used to read data from the lock-in amplifier. The lock in amplifier was configured to a long time constant - its output is then effectively a rolling average. Beginning from a state with the chamber fully evacuated the labview program operates the following loop:

The pressure setpoint is changed.

There is a time delay whilst the system reaches equilibrium.

Average readings of the current and instantaneous pressure values are taken.

After the program reaches the upper limit of the scan range the values are written out to a file.

Naively one would expect an approximately linear increase in the HHG output as a function of pressure due to the increase in the number of generating atoms present. In figure 3.3 the observed increase and subsequent decrease in HHG output as a function of pressure illustrates the effect that phase matching and absorption (as described in section 3.1) have on the HHG output.

Figure 3.3 shows four example data sets, the thin solid lines are the experimental data and the thicker dashed lines are high order polynomial fits (these fits are added purely as a guide for the eye). Unfortunately this figure is rather visually confusing however it serves to illustrate the extreme variability in the HHG output. All of the data was recorded in the same session and under the same conditions with just a small re-optimisation of the capillary alignment (movement of just a few tens of microns) between pressure scans.

Not only is there a fluctuation in the signal level from one data point to the next, but the pressure dependence profile also changes considerably from one run to the next. This is especially obvious in the case of the red curve for which the peak of phase matching is shifted to higher pressures. A possible explanation for this shift is that a fraction of the laser power was launched into a higher-order capillary mode. This would increase the contribution from the waveguide term in equation 3.8 and would therefore require higher pressures in order to bring the neutral gas term up to a sufficient level to achieve phase matching.

The point-to-point variation in the photodiode signal is believed to be caused primarily by air fluctuations. Due to its large mode area the laser is particularly susceptible to deviation as a result of pressure gradients in the air. This problem is exacerbated by the long optical path from the laser to the experiment ($\sim$ 3 m) and by the high tolerances on the capillary launch conditions. (Translating the front end of the capillary by just a few tens of microns has a large effect on the output mode and launch efficiency). In later experiments beam tubes were installed to reduce the effect of air pressure fluctuations.

The high tolerance on the capillary alignment is evident from the variation between pressure scans. Before each scan the alignment of the capillary was optimised for laser

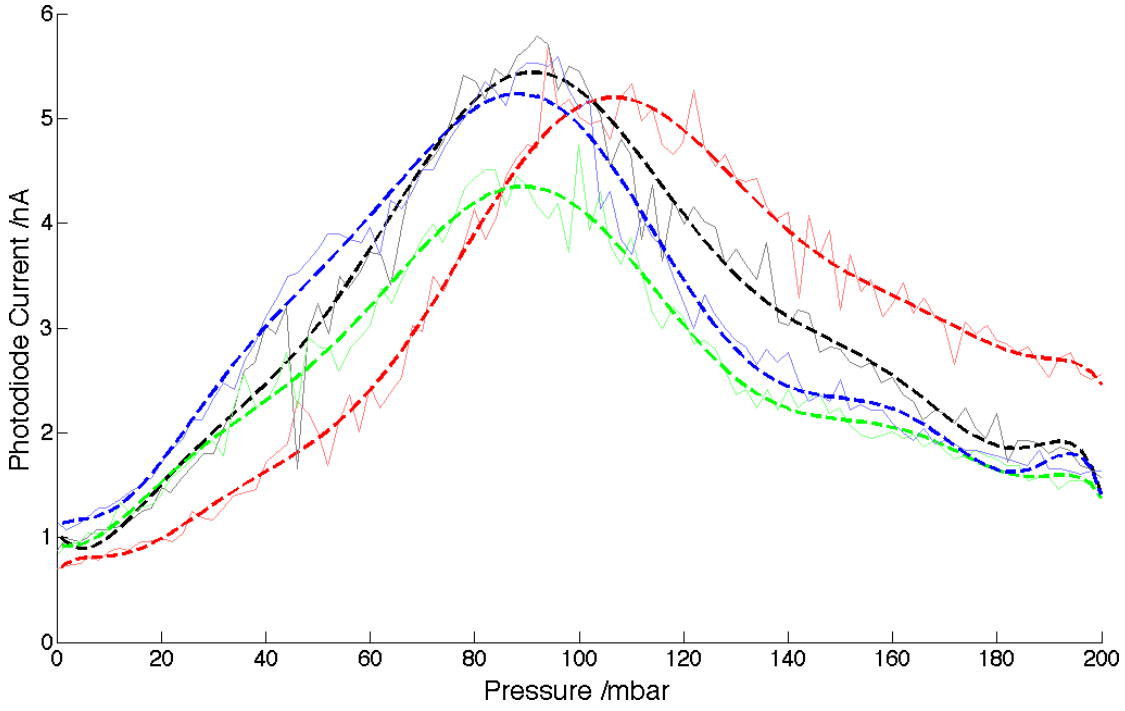

FIGURE 3.3: Photodiode signal as a function of pressure. Four different data sets are shown; each recorded at a very slightly different alignment. This highlights the variation in the output of the HHG capillary. High order polynomial fits have been added as a guide for the eye. (The apparent DC current offset is likely to be caused by laser light leaking through the aluminium filter).

output power (with no gas) and by observing the symmetry of the laser output mode. The change in alignment between each scan was in the range of a few tens of microns. These small changes are enough to significantly alter the pressure profile - this is especially apparent in the red curve which reaches a peak of generation at a higher pressure than the other data sets.

The blue lines in figure 3.4 show the results of pressure scans performed immediately after the laser had been optimised for the shortest possible pulse duration ($\sim$ 33 fs as measured by using the FROG). The red dotted line shows the average of the data sets shown in figure 3.3.

With the shorter pulse duration and for optimum alignment the photodiode signal is significantly increased for all pressure values. The peak of the photodiode signal is enhanced by a factor of 3 over that obtained in figure 3.3. It is also interesting to note that the peak of photodiode signal occurs at a significantly lower pressure.

As described in section 3.1.4 there are a number of factors at play here:

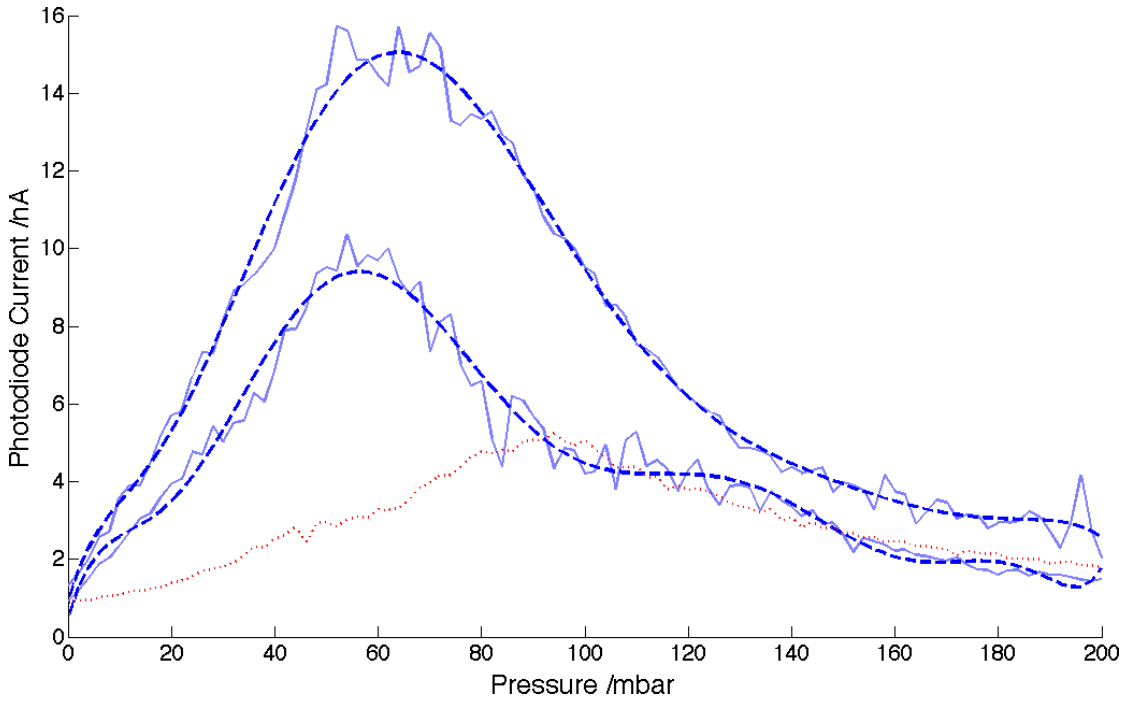

FIGURE 3.4: Photodiode signal as a function of pressure. The thin blue lines represent the experimental data from two different pressure scans. The thicker dashed lines are high order polynomial fits to this data. The red dotted line is the average of the data sets shown in figure 3.3.

The optimisation of the capillary launch conditions means that the waveguide term in the phase matching equation is reduced. A smaller contribution from the neutral gas term is now required in order to phase match. Phase matching for each harmonic therefore now occurs at a lower pressure.

The decrease in laser pulse duration means that for the critical ionisation fraction a higher laser electric field is reached. This should increase both the number of x-rays produced and the phase matching cut-off (permitting the generation of higher harmonics).

The absorption length is increased for the higher harmonics (since they propagate in a higher ionisation fraction). The higher harmonics consequently build-up over a longer distance and reach higher signal levels. Additionally the higher harmonics carry more energy and therefore produce a stronger signal on the photodiode.

All of these factors combine to produce the observed shift to lower pressure and accompanying increase in photodiode current.

The vertical scale on both figures 3.3 and 3.4 shows the signal level in nanoamps. The following relation for quantum efficiency QE is included in the photodiode documentation[73]:

$$\mathrm{QE} = E_p/3.63 \tag{3.16}$$

where E_p is the photon energy in eV. At the time these measurements were taken the spectral content of the x-ray beam was unknown. Our subsequent work using an x-ray spectrometer established that the average photon energy is $\sim$ 34 eV. (Variation of the average photon energy as a function of pressure is neglected since this is a small effect). To a first approximation it is sufficient to use this fixed average photon energy value to calculate the quantum efficiency. We therefore obtain a pressure-invariant quantum efficiency of approximately 9.4. This means that on average $\sim$ 9.4 electrons will be generated per x-ray photon or equivalently a charge of $q_e\mathrm{QE}$ will be produced (where q_e is the electron charge in Coulombs). Since the Ampere is defined as one coulomb of charge per second the number of x-ray photons per second (N_p) can be calculated simply by dividing the current (I) by the charge per photon:

$$N_p \approx \frac{I}{q_e\mathrm{QE}} \tag{3.17}$$

Using this calibration the peak of the upper blue curve in figure 3.4 is found to correspond to approximately 1×10^7 photons per pulse. This is the signal arriving on the photodiode and does not take into account the attenuation due to the aluminium foil filter. The expected transmission of the aluminum foil filter is $\sim$ 0.6 with an additional factor of 0.82 to taking into account effect of the supporting mesh[74]. The total transmission figure for the foil is therefore approximately 50 %. If we assume that the entire x-ray beam is collected on the 1 cm^2 detector then this gives a figure of 2 $\times 10^7$ photons per pulse at the source.

3.3.1 Pressure Dependence

As before, the laser is brought to a focus and the capillary is aligned to obtain the maximum laser power and mode symmetry. In this experiment the pressure setpoint is again varied however this time the x-ray spectrum is measured using the grazing incidence spectrometer 2.4.3, the Andor CCD was also used to image the spatial intensity profile of the x-ray beam.

Figure 3.5 shows the results of this pressure scan in the form of an image plot. Each horizontal line represents a spectrum recorded at a different pressure. Dark red regions in the image are the most intense whilst dark blue are the least.

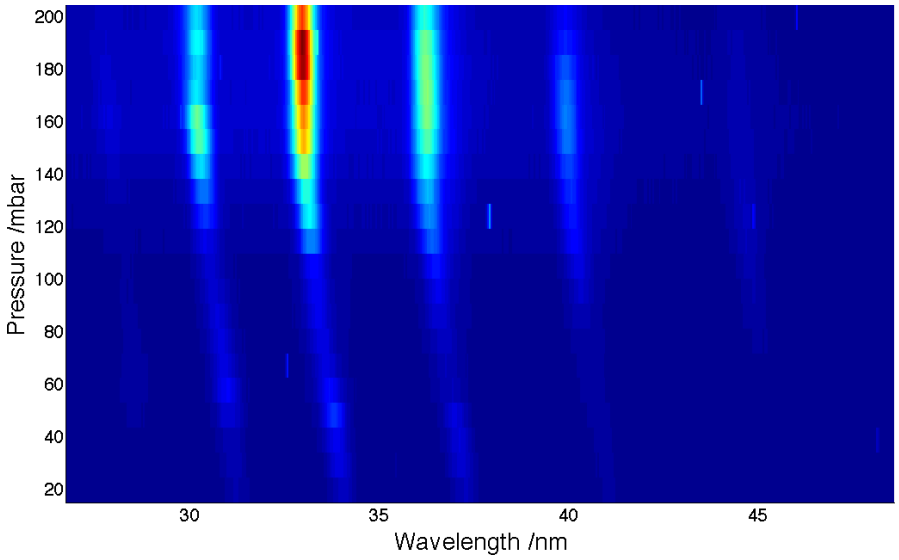

FIGURE 3.5: Image plot showing a series of spectra plotted against argon pressure.

Starting at low pressures (a few mbars) there is very little x-ray signal at all. As pressure is increased to $\sim$ 40 mbars the total x-ray flux also increases and more harmonics begin to contribute to the output each with roughly equal magnitudes. As pressure continues to increase the total x-ray flux reaches a peak. At the peak in generation there is a significant weighting in favour of of the higher harmonics. Above the peak in generation the total flux is reduced and the weighting between harmonics shifts back toward the lower orders.

As described in connection to the photodiode results, the exact pressure of the peak in x-ray generation varies considerably depending upon experimental conditions and alignment. The data set shown in figure 3.5 continues to increase in x-ray flux right up to $\sim$ 180 mbar. In many cases the peak of generation is observed to occur at pressures as low as 60 or 70 mbars in argon. It is believed that the main cause for this variation is the quality of the launch into the capillary, which in turn affects both the ionisation fraction and the waveguide term in the phase matching equation. Recording the data for figure 3.5 took a considerable amount of time; it would therefore not be surprising if during the experiment the quality of launch was degraded. This would lead to phase matching at higher pressures and may have slightly reduced the number of harmonics that were generated.

The observed shift in the envelope of the harmonics is in agreement with predictions from phase matching theory. As the pressure is increased phase matching occurs for higher and higher ionisation fractions until the critical ionisation fraction is reached. When the pressure has reached the level required to phase match at the critical ionisation fraction

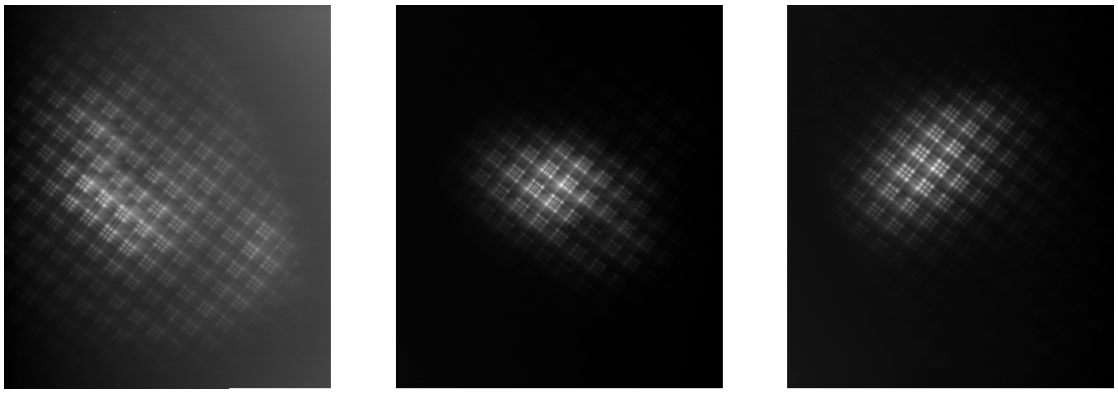

FIGURE 3.6: Series of x-ray CCD images showing the spatial intensity distribution as a function of pressure. Pressures are (from left to right) 10, 60 and 200 mbars.

generation is more intense and predominantly of the higher harmonics. Increasing the pressure further does not result in phase matching but rather increases Δk, decreasing x-ray output.

A striking feature of figure 3.5 is that as the pressure is increased each individual harmonic tends to shift to shorter wavelengths. This is frequently observed in HHG experiments and is referred to as ionisation or plasma induced blue-shifting. This effect, observed here on the harmonic emission, is in fact due to non-linear blue shifting of the fundamental[2, 75]. This is confirmed by subsequent experiments described in sections 3.3.2 and 3.4.

Observing the x-ray mode profile on the Andor CCD the beam initially appears to be rather divergent at low pressures, filling a large area of the CCD image. By the peak of total generation the mode profile is significantly reduced in size with a more spot-like appearance. As pressure is increased still further the spot increases slightly in area, but not to the same degree as was evident at the lower pressures. Figure 3.6 shows a sequence of CCD images that illustrate this point.

These results demonstrate phase matching:

At low pressures on-axis x-ray generation is phase matched only for the lowest order harmonics. For the higher harmonics there is a significant phase mismatch Δk between the x-ray and laser wavevectors. Because of the sign of the phase mismatch (phase velocity is higher for the x-rays than for the laser pulse) higher order harmonics that propagate at a particular angle to the capillary axis may build up through Cerenkov phase matching[62]. The brighter regions in the low pressure CCD image (figure 3.6, left) are probably due to co-linear phase matching of the lowest order harmonics. The less intense signal, which covers a large area of the CCD at low pressures is the Cerenkov

phase matched signal. Note that since the Cerenkov phase matched x-rays are propagating off-axis they can only accumulate over a short distance at the end of the capillary; this explains why they are less intense than the co-linear phase matched x-rays. At very low pressures (when the intensity of the phase matched x-rays is low) it may also be possible to observe the weak signal from non-phase matched generation. I.e. some of the x-rays generated in the last coherence length of the capillary will escape without being absorbed. These x-rays will be emitted over a wide range of angles, restricted only by the physical geometry of the capillary, however, they will be very weak since emission cannot build-up as it would if phase matched.

As pressure is increased the phase matching conditions will be reached for successive harmonics - beginning with the lowest orders. As described in section 3.1.4, when the pressure is just right to achieve $\Delta k = 0$ for any particular harmonic the coherence length becomes infinite. Under these conditions it is the argon absorption length that places the constraint on the volume that may contribute to the output. Since the contributing volume increases significantly when phase matched, so too does the output intensity of that particular harmonic. Furthermore, the constructive interference along the axis of the capillary gives the beam a well defined direction of propagation. X-rays that are emitted at an angle to the capillary axis have a different effective phase velocity to those that are on-axis and therefore cannot be phase matched under the same conditions. The output of each harmonic will therefore produce a tightly collimated beam as it is phase matched. This effect is most obvious when the higher harmonics are phase matched because these are also the most intense, the whole beam appears to be reduced in area (as shown in the central, 60 mbar image in figure 3.6). Phase matching of lower harmonic orders results in intensity maps such as that shown in figure 3.7 which demonstrates a phase matched central spot (probably due principally to just one harmonic) within the less intense background of the main part of the beam.

Once pressure passes the peak of total x-ray generation the weighting of harmonics begins to shift back towards the longer wavelengths, the total flux declines and the x-ray output spot becomes slightly larger. (Unfortunately we could not exceed 200 mbar for fear of over-loading the turbo pumps). Relating this behaviour back to the phase matching equation 3.8: The peak of x-ray generation occurs when the pressure is just sufficient to phase match the critical ionisation fraction. Increasing the pressure beyond this level results in a negative phase mismatch Δk between the x-ray and laser wavevectors (below this pressure the phase mismatch is positive). For pressures above this level phase matching can therefore only be maintained for lower harmonics and at lower ionisation fractions (nearer to the leading edge of the pulse). This explains the shift in the weighting of the harmonics back towards the lower orders and the decline in harmonic intensity.

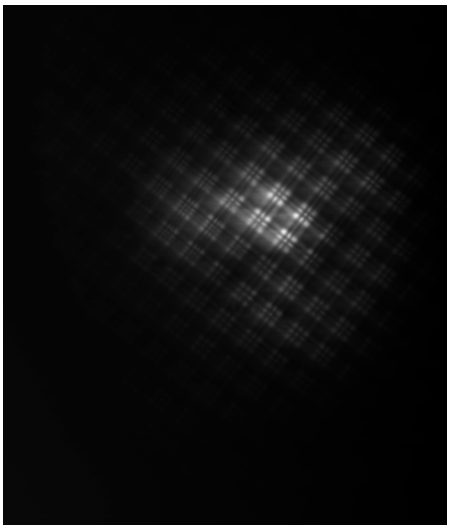

FIGURE 3.7: X-ray CCD camera image, recorded at 80 mbar. This shows an intense central core surrounded by a weaker more divergent spot. The central core is believed to be primarily due to one particular harmonic that is well phase matched.

The negative phase mismatch between the x-rays and the laser pulse makes Cerenkov phase matching impossible.

3.3.2 Power Dependence

In this experiment the setup is as before except that the variable is now the laser power. The capillary gas pressure was kept fixed at 60 mbar. The input laser power was varied from 665 mW (which was the highest achievable at the time) to 255 mW (below which no appreciable generation was present).

The results of the experiment are shown in figure 3.8. An image plot is again used to allow comparison with the results from the pressure variation 3.5:

From experimental observation, the effect of gradually increasing the average laser power coupled into the capillary is similar to that of gradually increasing the gas pressure towards the peak of phase matching. As more laser power is coupled into the capillary the total x-ray flux increases, however it increases fastest for the higher harmonics and so intensity seems to shift towards shorter wavelength generation. This profile conforms to the expected shift in cutoff with increasing laser peak power, however, this mechanism cannot be responsible for the observed profile for the following reason: In figure 3.5 when the pressure was varied the peak laser power was kept continuous - the variation in the profile must therefore be related to phase matching rather than a shift in cut-off.

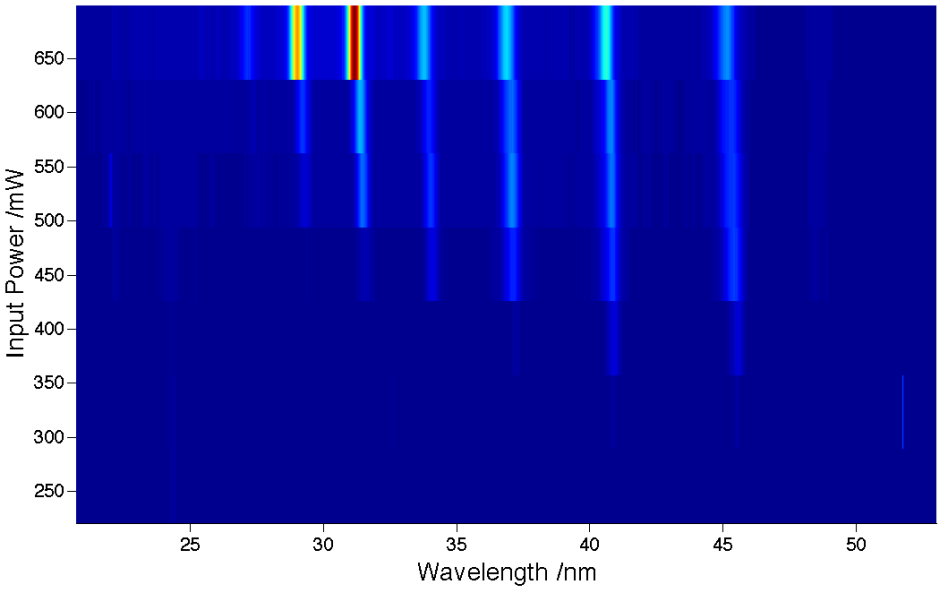

FIGURE 3.8: X-ray spectra (horizontal axis) recorded against input laser power (vertical axis) The effect is similar to that of varying pressure as shown in figure 3.5.

Because of the extreme similarity between the profiles observed in figures 3.5 and 3.8 it seems likely that phase matching (through modification of the ionisation fraction) is responsible for the power dependence profile i.e:

At very low laser powers the ionisation fraction at the peak of the pulse is also very low. With the argon pressure set to 60 mbar the contribution to the phase matching equation (equation 3.8) from the neutral gas term is so large that it prevents phase matching for all harmonics. (The low ionisation fraction means that the rate of x-ray generation is expected to be very low in any case).

As the laser power is increased so too is the ionisation fraction at the peak of the pulse. Part of the contribution from the neutral gas is now counteracted by the plasma term. Phase matching will become possible for the lower order harmonics (since these require the smallest contribution from the plasma term) and then for the higher harmonics as the ionisation fraction is increased.

Once the ionisation fraction at the peak of the pulse exceeds the level required for generation of a particular harmonic, generation of that harmonic simply takes place at an earlier time (on the leading edge of the pulse). This explains why phase matched generation can be observed across several harmonic orders.

In this experiment 60 mbars corresponds to the peak in phase matching at maximum launch power (and hence to the pressure needed to phase match at the critical ionisation

fraction η_c). If a lower (or higher) pressure had been selected then at full power the pressure would have meant too small (or too large) a contribution from the neutral gas and plasma terms. In both cases this would have prevented generation of the highest harmonics and reduced overall x-ray output.

This data shows that phase matching is the dominant (and limiting) effect in our current regime of operation. This lends weight to the argument that quasi phase matching or some other alternative scheme may be able to 'unlock' potential from our system which is currently not exploited - allowing significantly higher harmonics to be generated.

3.4 Dispersion Dependence

The Dazzler (see subsection 2.1.4) enables arbitrary control of the spectral phase of the input laser pulse. However, in this work the investigation was limited to a study of the second and third order spectral phase.

Before recording the data, the expected result was to find spectral shifting of the output harmonics as a function of input pulse spectral phase. At this stage it was our conjecture that the generation was taking place at the leading edge of the pulse where the ionisation fraction had risen to approximately 4%, allowing phase matching to occur. Crucially this model predicts that changing from negative to positive quadratic spectral phase will result in a change in the instantaneous frequency at which the 4% phase matching level is reached. We expected to see shifting of the harmonic peaks proportional to the change in instantaneous frequency at the 4% point. I.e. In a pulse that has been chirped by adding positive group delay dispersion (GDD = quadratic spectral phase) the 'red' spectral components will lead the pulse. If negative GDD is added the leading edge will contain the 'blue' components and will therefore have a higher instantaneous frequency (lower instantaneous wavelength). This effect is demonstrated in figure 3.9.

The calculation for figure 3.9 is performed as follows: The electric field for a 30 fs Gaussian pulse (with centre wavelength 780 nm) is calculated. The Fourier transform of this electric field is taken in order to obtain the spectrum. The spectral phase produced by an arbitrary dispersion can then be calculated using a Taylor expansion (see equation 3.18). (In figure 3.9 the addition of 300 fs^2 of GDD was simulated). This spectral phase is then applied to the calculated spectrum. The back Fourier transform returns the electric field of the chirped pulse. The instantaneous wavelength is simply calculated from the zero crossing points.

$$\varphi(\omega) = \varphi_0 + \varphi_1(\omega - \omega_0) + \frac{\varphi_2(\omega - \omega_0)^2}{2} + \frac{\varphi_3(\omega - \omega_0)^3}{3!} + \frac{\varphi_4(\omega - \omega_0)^4}{4!}) \qquad (3.18)$$

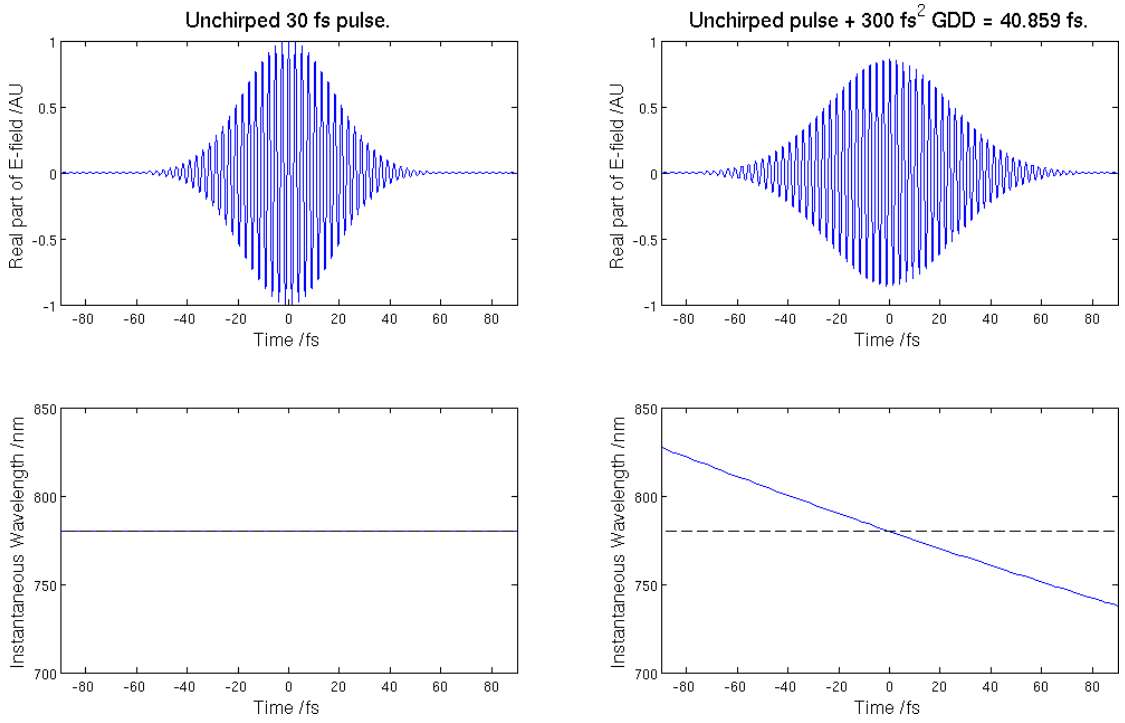

FIGURE 3.9: Top left: Real part of the electric field of a 30 fs FWHM, Gaussian pulse, centred at 780 nm. Top right: After adding 300 fs^2 of GDD the pulse is increased to $\sim$ 41 fs FWHM. Lower left and right: These plots show the variation in instantaneous wavelength with time. Note that the leading edge of the pulse (negative times) contains the longer wavelength 'red' components since normal dispersion (positive GDD) was added.

Where φ is the spectral phase, ω is the angular frequency, ω_0 is the central angular frequency. $\varphi_{1..4}$ are respectively: absolute phase, delay, quadratic phase and cubic phase. Quadratic phase and cubic phase are often referred to as second-order dispersion (GDD) or third-order dispersion (TOD).

Figure 3.10 shows an ADK based calculation[58, 76, 77] of the ionisation fraction induced by two laser pulses of different duration but the same time-averaged power. The steeper gradient at the leading edge of the unchirped pulse means that the 4% ionisation fraction required for phase matching the highest harmonics is achieved at an earlier time in the pulse. For the unchirped 30 fs pulse 4% ionisation is reached at -8 fs whilst this level is not reached until -4 fs in the chirped pulse[†]. In the unchirped pulse spectral phase is flat; the instantaneous frequency is therefore continuous throughout the pulse and corresponds to the centre wavelength. The longer pulse was chirped by the addition of quadratic spectral phase - this leads to a linear change in instantaneous frequency throughout the pulse. Since positive quadratic spectral phase (GDD) was

[†]Times are quoted relative to the centre of the pulse; therefore the minus sign denotes the leading edge of the pulse.

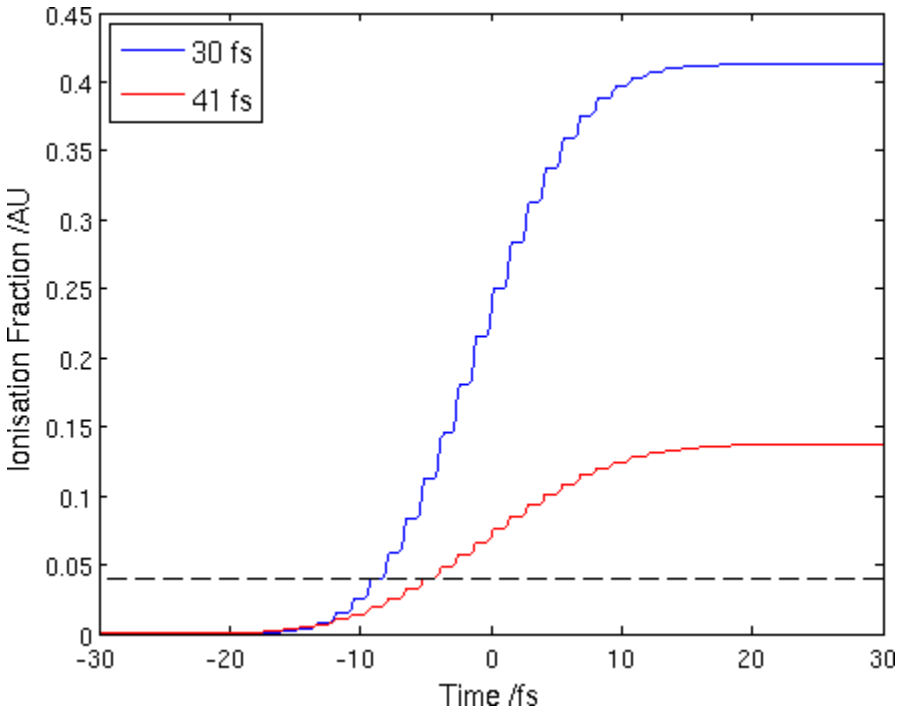

FIGURE 3.10: ADK based calculation of the ionisation fraction for two different pulse durations. Other laser parameters are: centre wavelength 780 nm, repetition rate 1 kHz, average power 650 mW, spot size 51 μm as required for launch into a 150 μm bore capillary. This leads to peak intensities of 2.8×10^{14} and 2×10^{14} W.cm^{-2} for the 30 fs and 41 fs pulses respectively. The 4 % ionisation fraction is reached at approximately -8 fs and -4 fs for the 30 fs and 41 fs pulses respectively.

added the longer wavelengths lead the pulse. From figure 3.9 we can see that at -8 fs the instantaneous wavelength of the chirped pulse is approximately 785 nm.

Whilst we observe spectral shifting of the harmonic peaks in both the quadratic and cubic spectral phase experiments, the predictions of our initial model were not verified, see figure 3.11 and 3.12. In both cases we observed harmonic peaks shifting towards the blue as the relative zero in spectral phase was approached. The peaks then shifted back to the red for both negative and positive values of spectral phase, see figure 3.11 and 3.12. From the instantaneous frequency calculation the spectral shift should be towards the red for positive GDD and towards the blue for negative GDD. Furthermore the magnitude of the observed spectral shift is greater than would be expected from the change in instantaneous frequency. A shift of 5 nm in the fundamental would only result in a shift of 0.2 nm in the 31^{st} harmonic where as we observe spectral shifts of $\sim$ 2 nm. It is also clear that the sign of the cubic spectral phase does not affect the observed shifting either.

This forced us to reconsider our theory. It was concluded that the only effect symmetric in both quadratic and cubic spectral phase would be the peak power and that this must

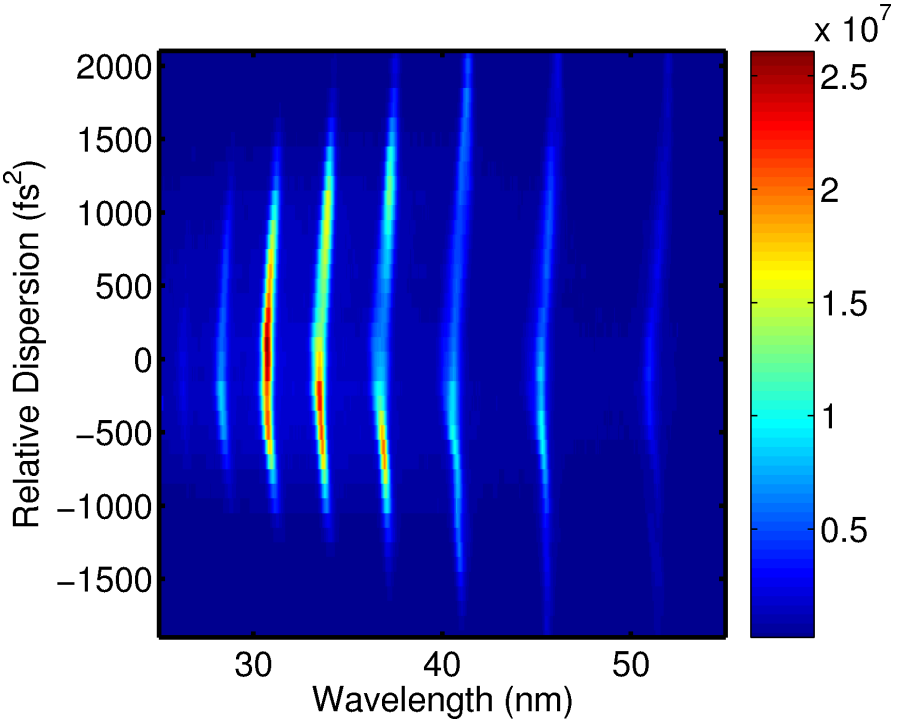

FIGURE 3.11: Image plot showing the x-ray spectra (horizontal) as a function of quadratic spectral phase (vertical).

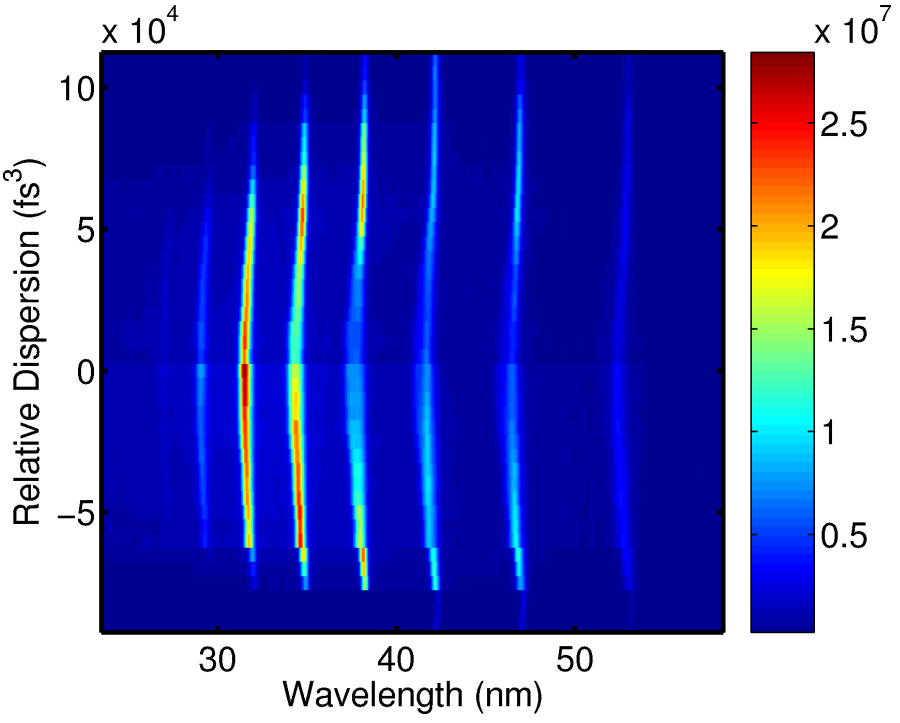

FIGURE 3.12: Image plot showing the x-ray spectra (horizontal) as a function of cubic spectral phase (vertical).

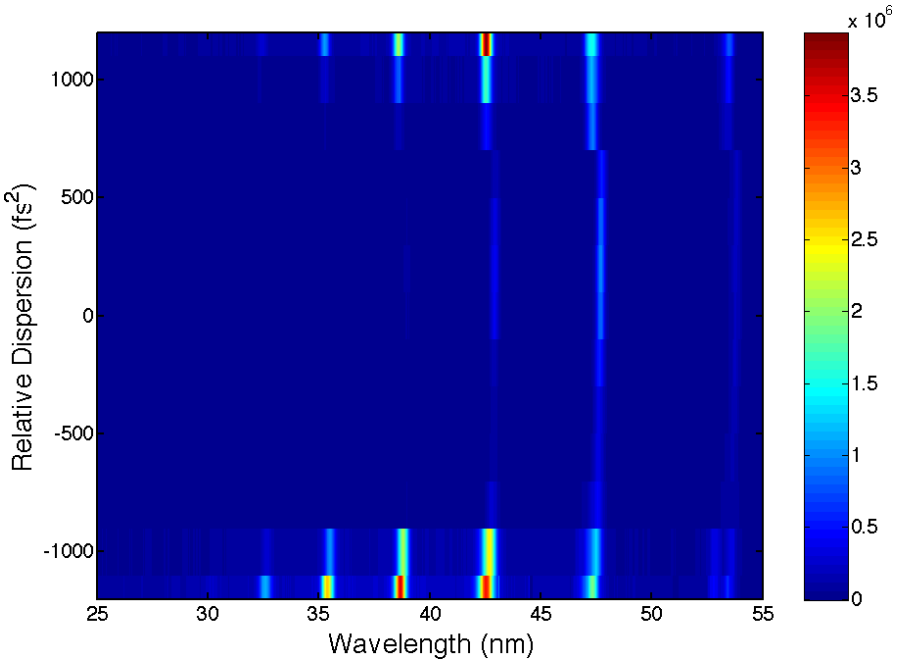

FIGURE 3.13: Image plot showing x-ray spectrum as a function of quadratic spectral phase. In this data set average laser power is adjusted to maintain a constant peak power. As a result ionisation is increased for large quadratic spectral phase offsets .

therefore be responsible for the spectral shifting in both cases. It was also concluded that generation could not be limited to the very leading edge of the pulse but rather must be from a dome-shaped region, with generation being pushed to the outer edges of the capillary as the most intense part of the pulse arrives.

It is now proposed that the observed harmonic peak shifting is simply due to a change in peak power caused by the spectral phase. In order to test this hypothesis two sets of quadratic spectral phase data will be compared. The first data set is that shown in figure 3.11. In this data set the average laser power is kept constant. This means that peak power will be decreased as quadratic spectral phase is offset from the relative zero. In the second experiment quadratic spectral phase is again scanned but this time a calculation similar to that shown in figure 3.9 is used to vary the average laser power in such a way as to maintain a constant (although lower) peak laser power.

In the first data set (figure 3.11) blueshifting is observed for harmonic peaks that are close to the relative zero of quadratic spectral phase. The harmonics produced by pulses with large (positive or negative) quadratic spectral phase are not shifted. i.e. The pulses with the highest peak power are blueshifted. See figure 3.11.

Figure 3.13 the results of the second data set (constant peak power). There is a very marked difference in the behavior. In figure 3.13

In this second data set (constant peak laser power) the pulses with large (positive or negative) quadratic spectral phase have their harmonic peaks blueshifted relative to the harmonic peaks generated by the zero spectral phase pulse. See figure 3.13.

The explanation for this results was found by studying FROG measurements of the fundamental laser beam taken at the input and output of the capillary. The frog traces showed that the pulses were leaving the capillary with their spectra broadened and shifted to shorter wavelengths. We conclude [2] that this blueshifting is caused by the time dependent variation in the argon refractive index. i.e. throughout the pulse more and more plasma is created, decreasing the average refractive index in the capillary. This therefore constitutes a nonlinear refractive index for the argon.

In the first experiment the ionisation fraction increases near to the zero of quadratic spectral phase. This is because the pulse becomes shorter in time without changing average power - it therefore attains a higher peak power.

In the second experiment in order to maintain constant peak power the zero spectral phase pulses had to be greatly attenuated. At the edges of the spectral phase range the pulse is very long in time - in order to maintain the same peak power as the other plots the total pulse energy must be greatly increased. The total energy available for ionisation and therefore the ionisation level reached is therefore much higher for the spectra with large spectral phase offsets for these pulses as compared with those of higher quadratic spectral phase. This causes the apparent blueshifting of harmonic peaks observed for pulses with large quadratic spectral phase.

Chapter 4

X-ray diffraction

The main aim of this section is to describe and to show the results of an experiment in which x-ray diffraction patterns are analysed in order to recover spatially resolved, spectral information about a capillary high harmonic x-ray source. This experiment constitutes a major block of the work undertaken during this PhD thesis. The main outcome of this work is that axial, spectral variation in the output of a capillary HHG source is directly observed for the first time. Furthermore a technique has been demonstrated that could provide critical information for the development of enhanced models of the HHG process. This technique is in fact a general one and with minor modifications could also be applied outside of the field of HHG.

The specific details of the experiment and the data analysis are provided in section 4.3. Before going into the experimental details, section 4.1 gives an overview of diffraction theory whilst section 4.2 discusses the most commonly used approximations.

4.1 Overview of Diffraction

Over the following paragraphs a very general overview of the subject of diffraction is given, then in sub-sections 4.2.1 and 4.2.2 the Fraunhofer and Fresnel approximations and their relative merits are discussed. Towards the end of subsection 4.2.2 specific details of the diffraction calculations (as used in the spectral recovery technique) are given. Throughout this section the theoretical discussion is based on that of Born & Wolf[78] however pertinent details are included here for completeness.

Any object that is introduced into the path of light causes a local disturbance to the electromagnetic waves and alters their subsequent propagation. The general term for this effect is 'scattering' and is used when the shape and composition of the scattering object are complex or unknown (as is typically the case for very small objects). The result of scattering is a very complex pattern which may be difficult to interpret.

When the scattering geometry is simplified it becomes possible to observe large scale interference effects in the scattered light. Under these circumstances more specific terms for this effect exist: By convention we use the term 'diffraction' when most of the E-field remains outside of the object (i.e. visible light is diffracted by a metal razor blade) - or 'refraction' when a significant part of the E-field passes through the object (i.e. refraction deviates visible light in a prism). 'Reflection' is a special case of diffraction where the propagation direction of the light is significantly deviated but the spatial profile of the beam is substantially unchanged.

It is important to realise that in general all of the processes (reflection, refraction, diffraction and absorption) are present to some degree - considering any of these effects individually does not tell the whole story. In all cases the mechanism for the interaction between light and matter is the same. That mechanism is the variation of optical properties that the electromagnetic wave experiences as it passes from one medium to the next.

The linear optical properties of a material are described by its refractive index n[†]. In general this property is complex and varies as a function of the illumination frequency as discussed in sub-section 2.3.2.1. This frequency dependent variation explains how the observed behaviour of visible light and x-rays can be so very different despite the fact that their propagation is governed by the same laws.

Scattering can be understood and its effects calculated by solving Maxwell's equations for the boundary conditions defined by the spatial refractive index profile of the problem.

[†]The concept of refractive index can also be extended to describe the nonlinear properties of a material by introducing n_2 the nonlinear refractive index.

In general Maxwell's equations cannot be solved analytically - even for relatively simple geometries, however, approximate numerical methods such as finite element modeling (FEM) can be used.

Powerful computing resources are now readily available and the FEM has become a good, practical option for obtaining approximate solutions to electric field problems. In this method the simulated region is separated into many smaller segments (in 2D calculations these segments are typically triangular, in 3D they are usually tetrahedral). Maxwell's partial differential and integral equations are then exchanged for ordinary differential equations that approximate their behaviour but can be solved using standard techniques - this is the key approximation taken in this approach. Figure 4.1 shows an example of a 2D FEM mesh that simulates the interference of light emitted from two parallel waveguides.

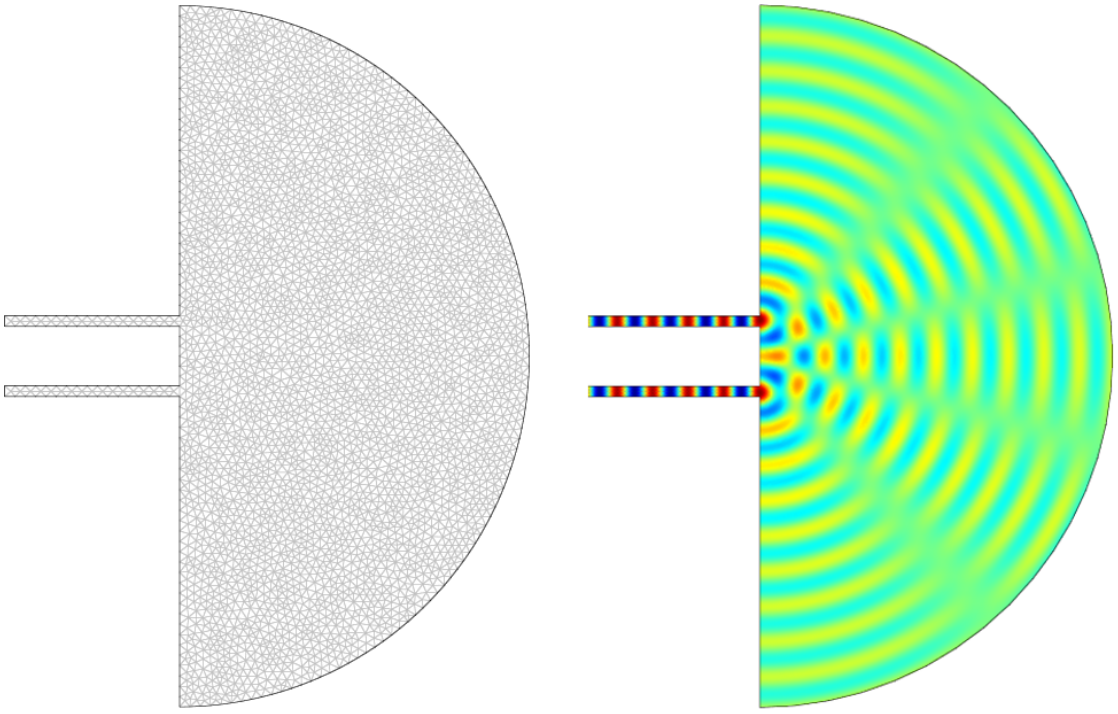

FIGURE 4.1: (Left: Finite element mesh for a 2D simulation of interference between light emitted from two waveguides. The boundary conditions are set so that the end of each waveguide is a plane wave source. Right: The intensity map obtained by solving for these boundary conditions. This is actually the output of a demonstration file included with the Comsol Multiphysics package.

FEM allows computing power to be concentrated on regions of particular interest (for example, near to interfaces where the field may change rapidly). By refining the mesh in these regions into smaller and smaller segments resolution in these key areas can be improved without the overhead of unnecessary calculation in regions of slowly varying field. Unfortunately the required computing time still expands rapidly with volume since

a minimum sampling resolution must be maintained. This means that FEM is perhaps less applicable to problems which entail propagation across large open regions.

FEM allows approximate solutions to Maxwell's equations to be found in complicated geometries where it would be difficult to apply other techniques. It is particularly appropriate for calculation of the electric field in close proximity to complex structures, for example, it is a great aid to those designing optical devices such as photonic crystals.

It is not always convenient or necessary to immediately adopt a numerical approach (such as FEM) to calculate diffraction patterns. In many cases, although a full analytic solution of Maxwell's equations is not possible, by making appropriate approximations a set of simplified diffraction equations may be formulated. These semi-analytic solutions will be faster to compute and may also yield mathematical relations that develop our understanding of the problem. Some of the approximations taken in these simplified theories are outlined below.

In a full solution to Maxwell's equations the electric field both inside and outside of the scattering object will be found. The diffraction pattern will be slightly modified as a result of the refractive index experienced within the scattering object. In most approximate diffraction calculations the scattering object is considered to be a perfect conductor - the electric field inside the object is eliminated and so refraction and absorption cannot occur. Under these circumstances it is only the shape of the scattering object that determines the diffraction pattern.

The perfect conductor approximation is used throughout the remainder of this chapter. This project is chiefly concerned with diffraction since it is possible to accurately measure (and compare with numerical simulations) the large scale interference fringes produced by the interaction of the HHG x-ray beam with simple, micron-scale, objects. The perfect conductor approximation should not produce large errors since the absorption of all materials is very high in the XUV; consequently the electric field propagating through the scattering object will be very low. The absorption will of course modify the intensity of the diffraction pattern slightly, although the shape should be largely unchanged.

In any discussion of the interaction between light and matter it is helpful to introduce the concept of the Huygens-Fresnel principle. Expressed in words this states that: the propagation of light can be described by the coherent sum of spherical wavelets re-radiated from every point on the advancing wavefront. This principle is traditionally illustrated by the diffraction of light through an aperture, see figure 4.2.

The Huygens-Fresnel principle reflects the fact that the free-space solutions to Maxwell's equations are spherical waves. (Plane waves are also a solution although these can be

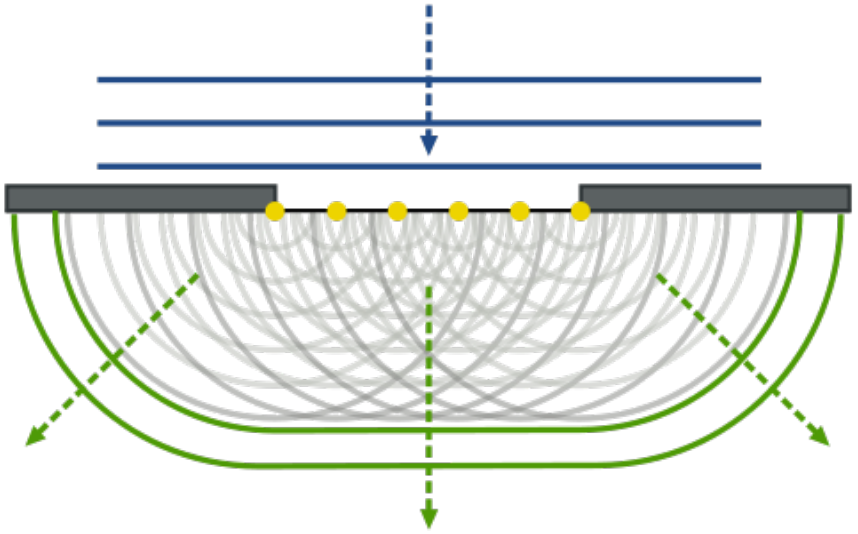

FIGURE 4.2: Illustration of the Huygens-Fresnel principle applied to diffraction through a slit. The spherical waves re-radiated from every point along the wavefront combine to approximate the diffracted light.

thought of as infinite radius spherical waves). Solutions to any set of boundary conditions can therefore be constructed by a superposition of spherical waves.

Kirchhoff provided a firm mathematical formulation of the Huygens-Fresnel principle with his scalar field theory (extending earlier work by Helmholtz carried out in the field of acoustics). A full discussion of this subject can be found in [78] however the general concept is as follows. Solutions to Maxwell's vector wave equation require calculation of the E-field components in each coordinate direction (E_x, E_y, E_z). Each component (and its differential) must satisfy the conditions of continuity at the boundary. The problem is dramatically simplified by considering the scalar field U (rather than the vector field E) and finding the solution to the scalar wave equation. The value of U is everywhere equal to the magnitude of the vector E-field in the direction of polarisation. The imaginary part of U describes the phase. Taking this scalar field approximation will in most practical examples have a minimal impact upon the accuracy of the calculated result. Only interactions where the direction of polarisation is changing rapidly will be significantly altered.

Scalar field theory can be applied to the problem of diffraction through an aperture by employing the mathematical concept - Green's theorem. Essentially, in this approach the complex amplitude U at a point P_1 is found by integrating the contributions from every point on a surface that completely encloses P_1. In order to do this the field value and its differential must be known at every point on the boundary surface.

When applied to diffraction through an aperture Kirchhoff devised the following approach: An expression for the field in the aperture plane (as produced by a point light source at P_{-1}) is written. The differential of this field is then calculated. The field on the reverse side of the aperture screen (in the geometric shadow) is assumed to be zero - as is its differential. A sphere (centred at P_1 and intersecting with the aperture screen) completes the boundary. The field on the inside of this sphere (due to the source at P_{-1}) is assumed to be as it would be if the aperture screen were not present. However, as the radius of this sphere tends to infinity, the contribution from this part of the integral vanishes. The problem is then reduced to an integration of the contributions to the field at point P_1, from each area element $\mathrm{d}S$ in the aperture. By also neglecting terms that contain a $1/r$ or $1/s$ dependence, the Kirchhoff diffraction equation is finally obtained:

$$U(P_1) = -\frac{iA}{2\lambda} \iint_{\mathscr{A}} \frac{e^{ik(r+s)}}{rs} [\cos(n,r) - \cos(n,s)] \, \mathrm{d}S \tag{4.1}$$

The geometry of this diffraction problem is illustrated in figure 4.3. $U(P_1)$ is the complex amplitude arriving at point P_1 on the screen. The double integral is performed over the area of the aperture (as denoted by the subscript $\mathscr{A}$) of which dS is an element. A is an amplitude scaling factor. r and s are the direct distances from the source to a given point in the aperture and from that point to a given point on the screen respectively. The angles (n,r) and (n,s) arise from the differentiation of U with respect to the normal to the aperture n. The sign convention of the angles is important to the calculation and is therefore illustrated in figure 4.4. (n,r) is the acute angle between n and r whilst (n,s) is the obtuse angle between n and s.

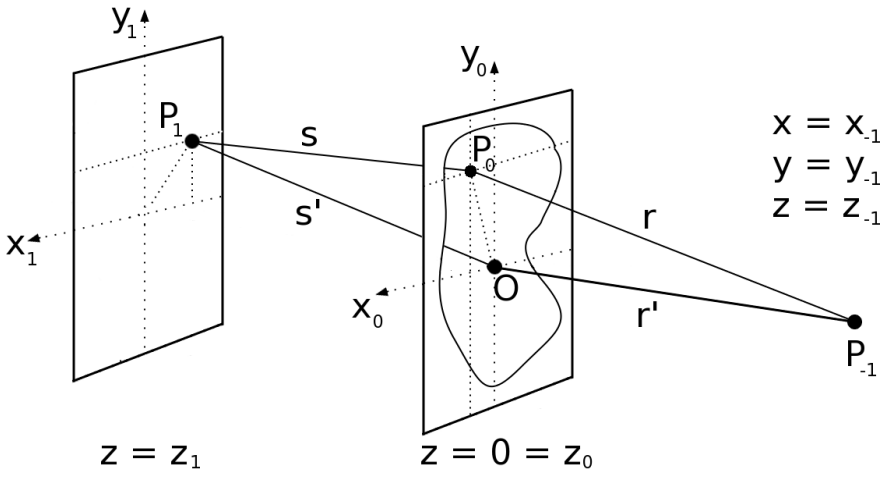

FIGURE 4.3: Diagram showing the geometry of simple aperture diffraction problem.

Although Kirchhoff's diffraction theory relies upon several assumptions, each of these has a relatively minor effect on the accuracy of the calculated diffraction patterns. The major error is likely to arise from transmission and absorption occurring at the rim of the aperture - this would in reality result in slight modifications to the local field - however in this approximation such effects are ignored. Consequently the results of Kirchhoff diffraction calculations will be acceptable provided the region of study is outside of a few wavelengths from the aperture.

4.2 Fresnel and Fraunhofer Diffraction

The two most common diffraction theories add more approximations to Kirchhoff's in order to simplify the calculation still further. Fraunhofer and Fresnel diffraction are widely known, however, for completeness and because they are applied in the subsequent chapters, they will be described briefly in sub-sections 4.2.1 and 4.2.2 respectively. Selecting which of these approximations to use is very important to the accuracy of the diffraction calculation. The Fresnel number (F) provides a means of deciding which approximate diffraction equation is appropriate to a given problem:

$$F = \frac{a^2}{L\lambda} \tag{4.2}$$

In this equation a is the characteristic size of the object (radius of the aperture for example), λ is the wavelength of illumination and L is the distance between the object and the plane in which the diffraction pattern is to be measured. For very large values of F (when the object to measurement plane separation is small relative to the size of the object) the diffraction pattern will tend towards a simple shadow mask. Under these 'Near Field' conditions accurate calculation of the exact diffraction pattern is typically not possible using the normal approximate methods - FEM techniques may be more appropriate. For values of $F \gtrsim 1$ Fresnel diffraction will produce good results and is the optimum selection. For values of $F \ll 1$ (when the object to measurement plane separation is very large in relation to the object size) Fraunhofer diffraction will normally be sufficient and is computationally very fast - although Fresnel diffraction will still be valid and may offer higher accuracy.

4.2.1 Fraunhofer Diffraction

Throughout this sub-section please refer to figure 4.3; this shows the construction of a generic diffraction problem and serves to illustrate the notation used below.

The Fraunhofer approach is based on the assertion that the distances from source (P_{-1}) to the aperture plane and from the aperture plane to the measurement plane are very large compared to the size of the aperture. (This is why Fraunhofer calculations are often referred to as 'Far Field'). The distances s' and r' are therefore approximated to infinity - effectively assuming plane wave illumination of the aperture and the measurement plane[†].

[†]This approximation can be simulated in the lab by placing lenses either side of the aperture. The first lens collimates the light from the source (producing a virtual source at infinity) - the second lens focuses the diffracted light bringing the image in from infinity, nearer to the aperture plane.

Mathematically this 'Far Field' assertion modifies the Kirchhoff diffraction equation (4.1) in the following ways:

If s' and r' are infinite - all rays arriving at (or leaving) the aperture will be approximately parallel - the cosine terms can therefore be taken out of the integral. Additionally s and r are also approximately equal, the cosines of angles (n, r) and (n, s) will therefore also be approximately equal in magnitude but opposite in sign[‡]. The cosine term in the integral $[\cos(n, r) - \cos(n, s)]$ is therefore replaced by the constant $2\cos\delta$ (where δ is the angle between $P_{-1}P_1$ and the normal to the aperture plane as shown in figure 4.4). Note that using this approximation allows the source P_{-1} and the measurement point P_1 to move away from the z-axis, however, it places the restriction that the centre of the aperture must be close to the axis between the source and the measurement point. In other words, the light need not be normal to the aperture plane, but it may not deviate through a large angle as it is diffracted by the aperture. If the light is deviated by a large angle then point Q moves a significant distance away from the origin O and the approximation loses validity.

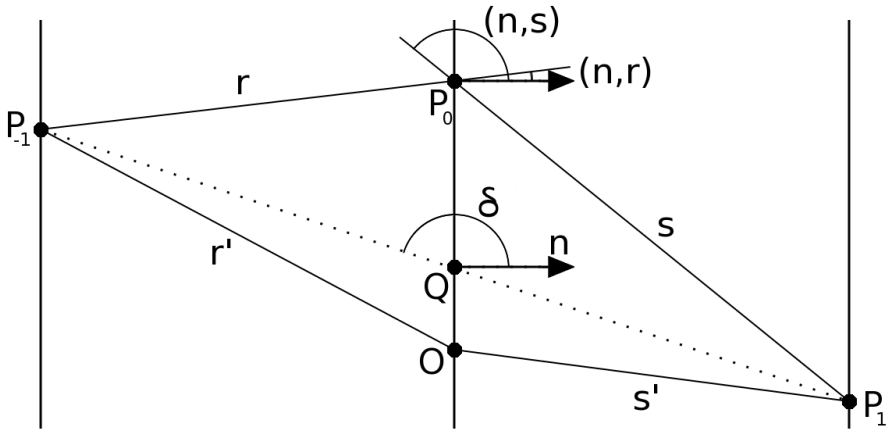

FIGURE 4.4: Diagram showing the sign convention used for the angles (n, r), (n, s) and δ. This figure also demonstrates why the centre of the aperture (at O) must be close to the axis between the source P_{-1} and the measurement point P_1. Note that this figure shows a 2D example where P_{-1} and P_1 are in the same y-plane; typically this is not the case and x_1 would for example be replaced with $\sqrt{x_1^2 + y_1^2}$.

Since s' and r' are approximately infinite - s and r must be too. Both s and r will therefore vary slowly over the aperture and the factor $1/rs$ is considered a constant and taken out of the integration. Furthermore $1/rs$ and can be replaced with $1/r's'$ which is approximately the same. Note - the factors r and s in the exponential cannot be approximated in this way since even a very small change here would have a large effect on the phase.

[‡]The angles themselves will differ by $\sim 180°$ because of the sign convention (see figure 4.4).

In order to approximate the factors r and s in the exponential they are first expressed in Cartesian coordinates and then approximated using a power series - taking just s as an example:

$$s^2 = (x_1 - x_0)^2 + (y_1 - y_0)^2 + z_1^2 \tag{4.3}$$

and

$$s'^2 = x_1^2 + y_1^2 + z_1^2 \tag{4.4}$$

Substituting 4.4 into 4.3 produces the following expression:

$$s^2 = s'^2 + x_0^2 + y_0^2 - 2(x_1x_0 + y_1y_0) \tag{4.5}$$

Equation 4.5 may then be simplified by substituting:

$$\rho^2 = \left[x_0^2 + y_0^2 - 2(x_1x_0 + y_1y_0)\right] \tag{4.6}$$

...and therefore becomes:

$$s = \sqrt{s'^2 + \rho^2} = s'\sqrt{1 + \frac{\rho^2}{s'^2}} \tag{4.7}$$

Substituting

$$x = \rho^2/s'^2 \tag{4.8}$$

and using the identity:

$$\sqrt{1+x} = \sum_{n=0}^{\infty} \frac{(-1)^n(2n)!}{(1-2n)n!^2 4^n} x^n = \left[1 + \frac{x}{2} - \frac{x^2}{8} + \frac{x^3}{16}\ldots\right] \tag{4.9}$$

...a power series expansion for s is obtained:

$$s \sim s' + \frac{s'}{2}\left[\frac{x_0^2 + y_0^2 - 2(x_1x_0 + y_1y_0)}{s'^2}\right] - \frac{s'}{8}\left[\frac{x_0^2 + y_0^2 - 2(x_1x_0 + y_1y_0)}{s'^2}\right]^2 + \ldots \tag{4.10}$$

Using the same method a similar series expansion can be written to describe r.

It is evident that these functions will rapidly become extremely complicated as more terms are added - especially if the cross terms are multiplied out. Fortunately, in the Fraunhofer method all terms in the power series expansions (of both r and s) that have quadratic or higher order dependence are neglected. This simplifies the mathematics immensely.

Combining all of the Fraunhofer approximations, substituting the truncated series expansion into the exponential, and removing constant terms from the integral - Kirchhoff's diffraction equation 4.1 can therefore be re-written as:

$$U(P_1) = -\frac{i\cos\delta}{\lambda}\frac{Ae^{ik(r'+s')}}{r's'}\iint_{\mathscr{A}} e^{-ik\left(\left[\frac{x_0x_{-1}+y_0y_{-1}}{r'}\right]+\left[\frac{x_1x_0+y_1y_0}{s'}\right]\right)}\,\mathrm{d}x_0\mathrm{d}y_o \tag{4.11}$$

Notice that the terms with r' dependence relate to the field created in the aperture plane by the point light source at P_{-1}. The terms with s' dependence relate to the transfer of field from the aperture plane to the measurement point P_1.

The most important statement to make in this subsection about Fraunhofer diffraction follows immediately from the assumption that r and s are approximately infinite:

> Since plane wave illumination of the aperture is assumed, the location and nature of the source (other than its monochromatic wavelength) have no bearing on the diffraction pattern created.

In practice the r' terms can therefore be replaced with any arbitrary expression for the field in the aperture plane. Normally this will consist of an expression for the field $E(x_0, y_0)$ (which for a plane wave is just a constant) and a separate expression $G(x_0, y_0)$ for the effect of the aperture (i.e. 1 for all positions within the aperture and 0 for everywhere on the screen). In practical use the constants in front of the integral are of limited consequence - the item of interest being the shape of the diffraction pattern as described by the terms inside the integral rather than the absolute scaling or time dependence. It is therefore possible to write equation 4.11 as:

$$U(x_1, y_1) \;\propto\; \iint_{-\infty}^{+\infty} e^{-ik\left(\frac{x_1x_0+y_1y_0}{s'}\right)} G(x_0, y_o)U(x_0, y_0)\;\mathrm{d}x_0\mathrm{d}y_0 \tag{4.12}$$

Notice that putting the aperture function explicitly inside the integral allows the limits of integration to be extended to $\pm\infty$.

Very importantly, the form of the diffraction equation shown in 4.12 can now be recognised as a Fourier integral. This means that Fraunhofer diffraction patterns can be calculated very quickly and easily using fast Fourier transform (FFT) algorithms; a technique that is very widely applied.

To be more precise there is a Fourier relationship between the spatial coordinates in the aperture plane (x_0) and the k-vector components of the diffracted light ($k_x = kx_1/s'$).

This preserves the normal relationship where the conjugate variables of a Fourier transform have reciprocal units. After this final step (and using $\mathscr{F}$ to denote the Fourier transform) the Fraunhofer diffraction equation is approximated to:

$$U(k_x, k_y) \propto \mathscr{F}\{G(x_0, y_o)U(x_0, y_0)\}\ \mathrm{d}x_0\mathrm{d}y_0 \tag{4.13}$$

To summarise this sub-section on Fraunhofer diffraction:

- In the Far Field approximation the source-to-aperture and aperture-to-screen distances are assumed to be approximately infinite. This implies plane wave illumination of the aperture and the screen.
- Small angles and some distances are approximated in accordance with the Far Field approximation.
- Only the constant and linear terms of a power series expansion of the path length $P_{-1}P_0P_1$ (source - area element in aperture -measurement point) are retained.
- Specific details about the source location are therefore not fully taken into account in the diffraction calculation.
- The approximation used for distances become less accurate as the measurement plane (or the source) approaches the aperture plane.
- The angular approximations become less accurate off-axis or as the size of the aperture increases.
- Fraunhofer diffraction will only produce acceptable results when the experimental conditions are a close match to the Far Field conditions. Within these limitations Fraunhofer diffraction remains an immensely practical tool since it can be calculated so simply using the Fourier transform.

4.2.2 Fresnel Diffraction

In this discussion of Fresnel diffraction theory I will highlight only the key points - for a full explanation one should consult [78].

The approximations taken in Fresnel diffraction are very similar to those taken in the Fraunhofer case. The main difference is that in Fraunhofer diffraction the quadratic and higher terms in the series expansions for r and s are neglected - in Fresnel diffraction all terms up to at least quadratic order are retained. Equation 4.11 therefore becomes:

$$U(P_1) = -\frac{i\cos\delta}{\lambda}\frac{Ae^{ik(r'+s')}}{r's'} \times$$
$$\iint_{\mathscr{A}} e^{ik\left(\left[\frac{x_0^2+y_0^2-2(x_0x_{-1}+y_0y_{-1})}{2r'}\right]+\left[\frac{x_0^2+y_0^2-2(x_1x_0+y_1y_0)}{2s'}\right]+\ldots\right)} \mathrm{d}x_0\mathrm{d}y_o \tag{4.14}$$

The effect of the additional terms can be demonstrated by plotting the distance r as it is calculated under each approximation. In order to illustrate this point clearly the distances are converted into phase values so that the wavefronts are plotted. A phase wrap factor of 5×10^5 is used so that the wavefronts can we displayed over a length scale that is very much longer than the wavelength of light.

The top row of figure 4.5 shows the phase, as calculated using exact geometry in the case of a point light source both on and off the optic axis. The wavefronts produced are spherical and do not change shape when the point light source is displaced. This is the behavior we expect for real light waves. In the middle row the distances are calculated using the Fresnel approximation. On axis, the phase calculated using the Fresnel approximation closely matches that of the spherical wave - even for measurements made quite close to the source. The Fresnel approximation does however lose validity when the measurement position is at a very large angle to the z-axis. Furthermore the accuracy of the approximation is reduced when the point source is displaced from the z-axis - this is because the series expansion for r has only linear dependence on x_{-1}. In all cases the Fresnel approximation considerably out performs the Fraunhofer one. Figure 4.5 clearly shows that it takes several meters before the Fraunhofer phase even begins to resemble that of the spherical waves. As in the case of the Fresnel approximation performance degrades rapidly as a function of angle and with offset from the z-axis.

The added complexity of the additional terms makes it inconvenient to write out the Fresnel diffraction equation in full. For this reason the series expansions for r and s are normally combined to produce the function $f(x_0, y_0)$ as shown in equation 4.15. Note that the very first terms (r' and s') are not included in $f(x_0, y_0)$.

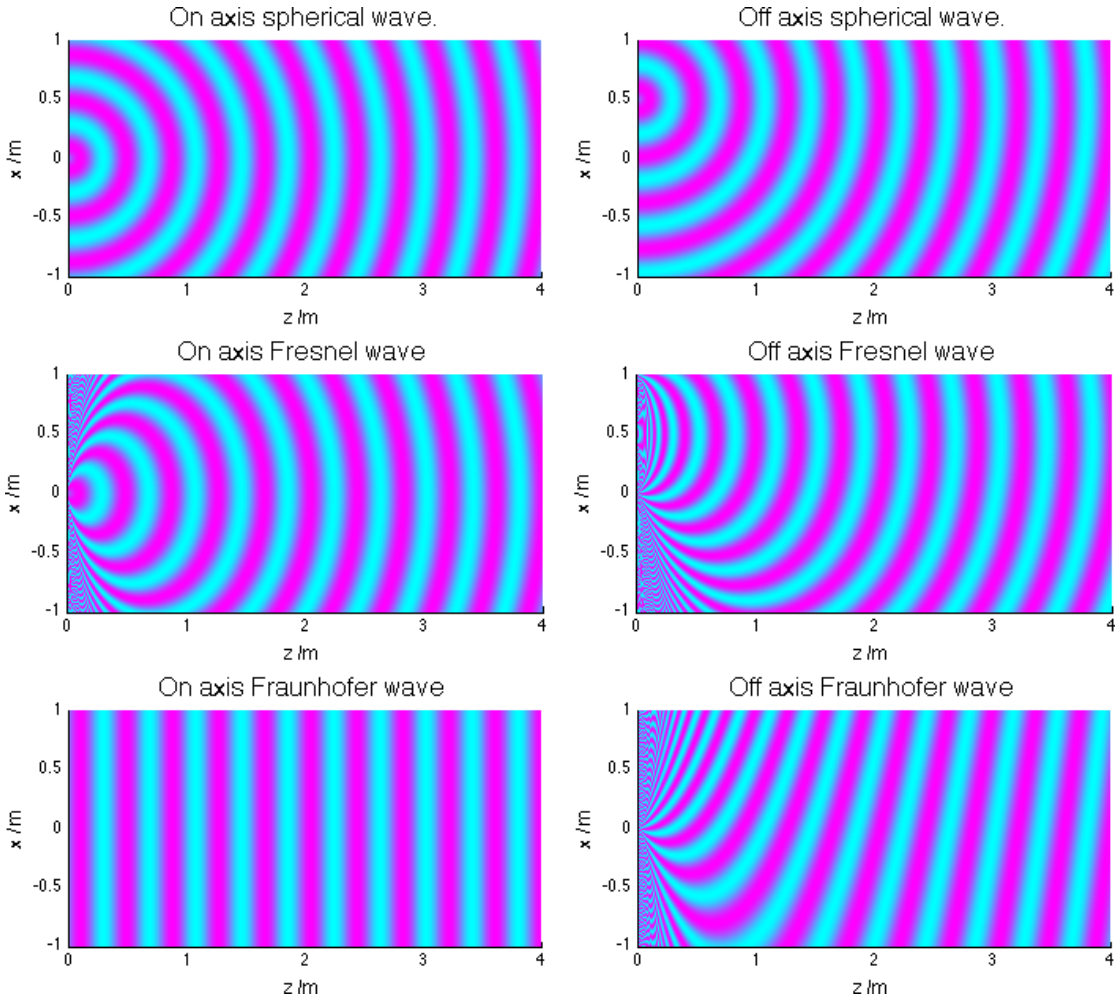

FIGURE 4.5: Comparison between the phase produced by a point light source as calculated under different approximations. Left column: Point light source is on the optical axis. Right column: Point source is off axis. Top row: No approximation - distances calculated with exact geometry. Middle row: Fresnel approximation. Bottom row: Fraunhofer approximation.

$$f(x_0, y_0) = \frac{r'}{2}\left[\frac{x_0^2 + y_0^2 - 2(x_{-1}x_0 + y_{-1}y_0)}{r'^2}\right] + \frac{s'}{2}\left[\frac{x_0^2 + y_0^2 - 2(x_1x_0 + y_1y_0)}{s'^2}\right] + \dots$$
$$- \frac{r'}{8}\left[\frac{x_0^2 + y_0^2 - 2(x_{-1}x_0 + y_{-1}y_0)}{r'^2}\right]^2 - \frac{s'}{8}\left[\frac{x_0^2 + y_0^2 - 2(x_1x_0 + y_1y_0)}{s'^2}\right]^2 + \dots \tag{4.15}$$

It is convenient to separate the real and imaginary parts of the integral in equation 4.14 so that it becomes:

$$U(P_1) = B\,(C + iS) \tag{4.16}$$

Where:

$$\begin{aligned} B &= -\frac{i\cos\delta}{\lambda}\frac{Ae^{ik(r'+s')}}{r's'} \\ C &= \iint_{\mathscr{A}} \cos\left(kf(x_0,y_0)\right)\ \mathrm{d}x_o\mathrm{d}y_0 \\ S &= \iint_{\mathscr{A}} \sin\left(kf(x_0,y_0)\right)\ \mathrm{d}x_o\mathrm{d}y_0 \end{aligned} \tag{4.17}$$

By selecting a coordinate system such that P_{-1} and P_1 are both in the same axis plane the linear terms in the series expansion can be eliminated. If only the quadratic terms are retained then by introducing new variables of integration equations 4.17 can be converted into the following form:

$$\begin{aligned} C &= b\iint_{\mathscr{A}} \cos\left(\frac{\pi}{2}\left(u^2+v^2\right)\right)\ \mathrm{d}u\mathrm{d}v \\ S &= b\iint_{\mathscr{A}} \sin\left(\frac{\pi}{2}\left(u^2+v^2\right)\right)\ \mathrm{d}u\mathrm{d}v \end{aligned} \tag{4.18}$$

Where the parameters u and v are:

$$\begin{aligned} \frac{\pi}{2}u^2 &= \frac{\pi}{\lambda}\left(\frac{1}{r'}+\frac{1}{s'}\right)x_0^2\cos^2\delta \\ \frac{\pi}{2}v^2 &= \frac{\pi}{\lambda}\left(\frac{1}{r'}+\frac{1}{s'}\right)y_0^2 \end{aligned} \tag{4.19}$$

As before, δ is the angle between $(P_{-1}P_1)$ and the normal to the aperture plane. b is given by:

$$b = \frac{\mathrm{d}x_0\mathrm{d}y_0}{\mathrm{d}u\mathrm{d}v} = \frac{\lambda}{2}\frac{1}{\left(\frac{1}{r'}+\frac{1}{s'}\right)\cos\delta} \tag{4.20}$$

In most cases it will not be necessary to calculate b as the measured diffraction pattern will be proportional to the terms inside the integral. Most Fresnel diffraction problems are therefore reduced to evaluation of equations 4.18 over an appropriate domain of integration.

In the following sub-sections some examples of Fresnel diffraction problems are provided. Most pertinent is the diffraction from a square aperture (4.2.2.3) as this calculation is applied later in section 4.3.

4.2.2.1 Fresnel diffraction at an edge

The geometry used to calculate Fresnel diffraction at the straight edge of a semi-infinite plane is illustrated in figure 4.6:

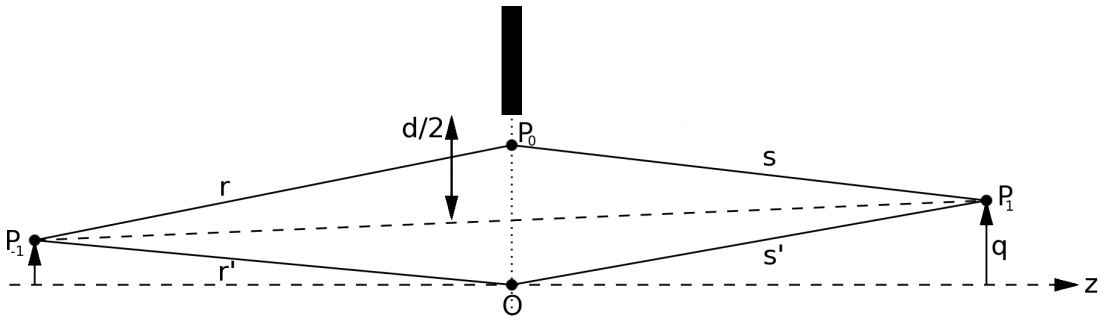

FIGURE 4.6: Diagram showing the geometry used in calculating the Fresnel diffraction pattern produced by an infinite straight edge.

The distance $(P_{-1}P_1)$ is approximated to $r' + s'$ and the distance from the edge of the plane to the z-axis is approximated to be the same as that from the edge of the plane to the line $(P_{-1}P_1)$. These approximations are adhered to perfectly when the line $(P_{-1}P_1)$ passes through the origin O. The accuracy of the approximation is degraded when the source and/or measurement point move away from the z-axis.

Only light that does not fall on the plane can contribute to the field at P_1. If the edge of the screen is displaced from the line $(P_{-1}P_1)$ by a distance $d/2$ in the x_0 direction then the limits for the integration are $-\infty < y_0 < +\infty$ and $-\infty < x_0 < d$.

By rearranging and taking the square root of equation 4.19 we produce an expression for u as a function of x_0.

$$u = \sqrt{\frac{2(r' + s')}{r's'\lambda}} x_0 \cos\delta \tag{4.21}$$

Under normal incidence the factor $\cos\delta$ may be ignored. Putting $d/2$ in as the value for x_0 gives the upper limit of the integral over the u parameter - let this be called w. A similar expression for v could be produced, however, integration over the parameter v is, in fact, not necessary due to the symmetry of the problem. The integration of equations 4.18 can be separated into two distinct regions of u - from $-\infty$ to 0 and from 0 to w. The second region is described by the Fresnel sine and cosine integrals:

$$\mathscr{C}(u) = \int_0^w \cos\left(\frac{\pi}{2}u^2\right)\ \mathrm{d}u$$

$$\mathscr{S}(u) = \int_0^w \sin\left(\frac{\pi}{2}u^2\right)\ \mathrm{d}u \tag{4.22}$$

These integrals frequently occur in diffraction problems. There are a number of ways in which they may be calculated numerically - including a rapidly converging recurrence relation. The plots on the left of figure 4.7 demonstrate the behaviour of these functions.

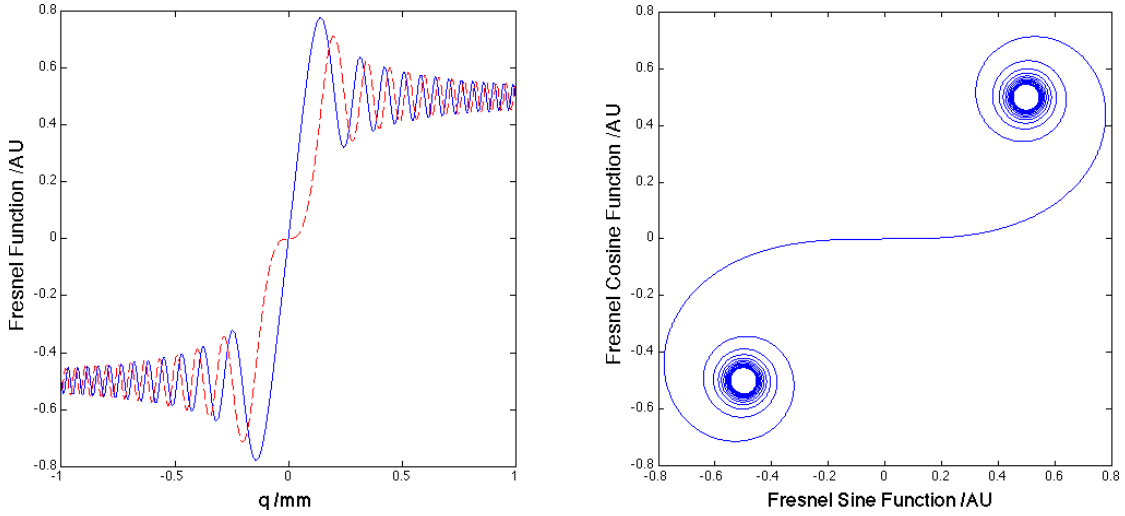

FIGURE 4.7: Left: Plot showing the behaviour of the Fresnel $\mathscr{C}(u)$ function (blue-solid) and the Fresnel $\mathscr{S}(u)$ function (red-dashed). These are both plotted against q here since u is a function of q. Right: Plot of the spiral of Cornu created by plotting $(\mathscr{C}(u) + i\mathscr{S}(u))$ in the form of an Argand diagram.

As is evident from the plot on the left of figure 4.7 the values of the Fresnel integrals perform a damped oscillation on either side of the edge. The plot on the right of figure 4.7 shows the spiral of Cornu. This is obtained simply by plotting $\mathscr{C}(v)$ along the vertical (real) axis and $\mathscr{S}(v)$ along the horizontal (imaginary) axis. Each point on the spiral therefore corresponds to a particular value of v and also to a particular value of q (the position on the screen). If the range of values plotted could extend to $\pm\infty$ then the spiral would reach converge at $[-0.5, -0.5]$ and $[0.5, 0.5]$. It is clear from these figures that the following identities hold:

$$\int_0^{\infty} \cos\left(\frac{\pi}{2}u^2\right) \mathrm{d}u = \frac{1}{2} \qquad \int_{-\infty}^{0} \cos\left(\frac{\pi}{2}u^2\right) \mathrm{d}u = \frac{1}{2}$$

$$\int_0^{\infty} \sin\left(\frac{\pi}{2}u^2\right) \mathrm{d}u = \frac{1}{2} \qquad \int_{-\infty}^{0} \sin\left(\frac{\pi}{2}u^2\right) \mathrm{d}u = \frac{1}{2} \tag{4.23}$$

Equations 4.18 can now be evaluated in terms of these identities (4.23) and the Fresnel integrals (4.22):

$$\int_{-\infty}^{w} \cos\left(\frac{\pi}{2}v^2\right) \mathrm{d}v = \int_{-\infty}^{0} \cos\left(\frac{\pi}{2}v^2\right) \mathrm{d}v + \int_0^{w} \cos\left(\frac{\pi}{2}v^2\right) \mathrm{d}v = \frac{1}{2} + \mathscr{C}(w)$$

$$\int_{-\infty}^{w} \sin\left(\frac{\pi}{2}v^2\right) \mathrm{d}v = \int_{-\infty}^{0} \sin\left(\frac{\pi}{2}v^2\right) \mathrm{d}v + \int_0^{w} \sin\left(\frac{\pi}{2}v^2\right) \mathrm{d}v = \frac{1}{2} + \mathscr{S}(w) \tag{4.24}$$

These relations can in turn be substituted into equation 4.16 to obtain the field amplitude $U(P)$. The intensity distribution produced on the screen is then calculated by taking the square modulus of $U(P_1)$. This produces the following expression for the intensity distribution:

$$I(P_1) = \frac{|A|^2}{2(r'+s')^2}\left(\left[\frac{1}{2} + \mathscr{C}(q)\right]^2 + \left[\frac{1}{2} + \mathscr{S}(q)\right]^2\right) \tag{4.25}$$

In many cases the absolute field amplitude of the source is not known and the absolute intensity on the screen is of little interest - calculation of the scaling factor at the start of equation 4.25 may not therefore be necessary. The shape of the intensity distribution is simply proportional to the bracketed terms as shown in figure 4.8.

The spiral of Cornu can be used to evaluate $I(P)$ graphically. A line drawn from the eye of the spiral at $[-0.5, -0.5]$ to any other position on the spiral has a length proportional to the square root of the intensity at the corresponding position q. The oscillation in intensity occurs as measurement point on the curve v cycles around the positive eye.

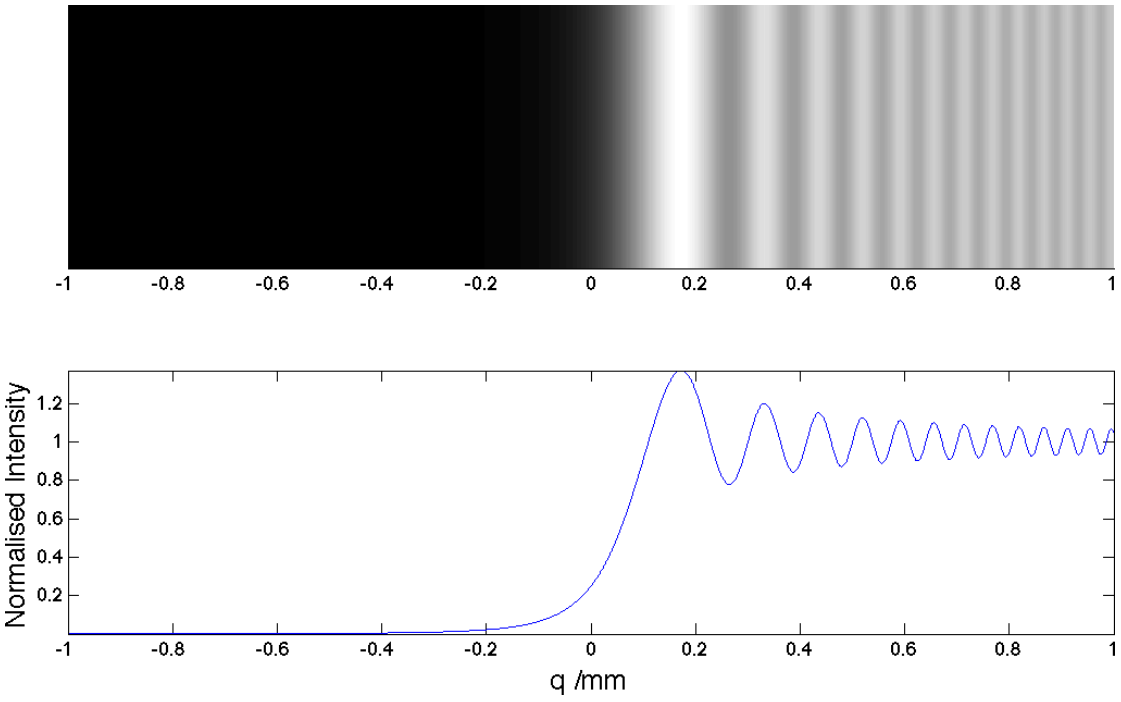

FIGURE 4.8: Image plot showing the calculated Fresnel diffraction pattern produced by the edge of a semi-infinite plane. The conditions used here are: $\lambda = 25$ nm, $r\prime$ and $s\prime$ are both approximately 0.5 m and q takes the range $\pm$ 1 mm.

4.2.2.2 Fresnel diffraction through a slit

Only light passing through the slit contributes to that arriving at each point P_1 on the screen. The domain of integration is therefore $-d/2 < x_0 < d/2$. In order to evaluate this integral the diffraction from each edge is calculated separately using the same method as above. (i.e. the diffracted field of an edge at $-d/2$ and the diffracted field of an edge at $d/2$ are both calculated). The parameter w can be calculated for each edge using:

$$\begin{aligned} w_1 &= \sqrt{\frac{2(r'+s')}{r's'\lambda}}\frac{d}{2}\cos\delta \\ w_2 &= \sqrt{\frac{2(r'+s')}{r's'\lambda}}\frac{-d}{2}\cos\delta \end{aligned} \tag{4.26}$$

Calculating the diffracted field from the first edge we find that the region $-d/2$ to ∞ is illuminated - for the second edge the illuminated region is $d/2$ to ∞. Babinet's principle is now applied and the field from the second region is subtracted from that of the first.

This allows only the region $-d/2$ to $d/2$ to contribute. Figure 4.9 illustrates this process graphically. The expression for the total field is therefore:

$$U(q) = \frac{U_0}{2}\left(\left[\frac{1}{2} + \mathscr{C}(w_1) + i\left(\frac{1}{2} + \mathscr{S}(w_1)\right)\right] - \left[\frac{1}{2} + \mathscr{C}(w_2) + i\left(\frac{1}{2} + \mathscr{S}(w_2)\right)\right]\right) \tag{4.27}$$

The factor $U_0/2$ is simply a normalisation constant (U_0 is the field amplitude arriving at P in the absence of any obstructions). The $1/2$ terms inside the brackets cancel each other out so extending the integration to $-\infty$ is actually not necessary - although it does make the physics clearer. Simplifying 4.27 it becomes:

$$U(q) = \frac{U_0}{2}\left(\left[\mathscr{C}(w_1) + i\left(\mathscr{S}(w_1)\right)\right] - \left[\mathscr{C}(w_2) + i\left(\mathscr{S}(w_2)\right)\right]\right) \tag{4.28}$$

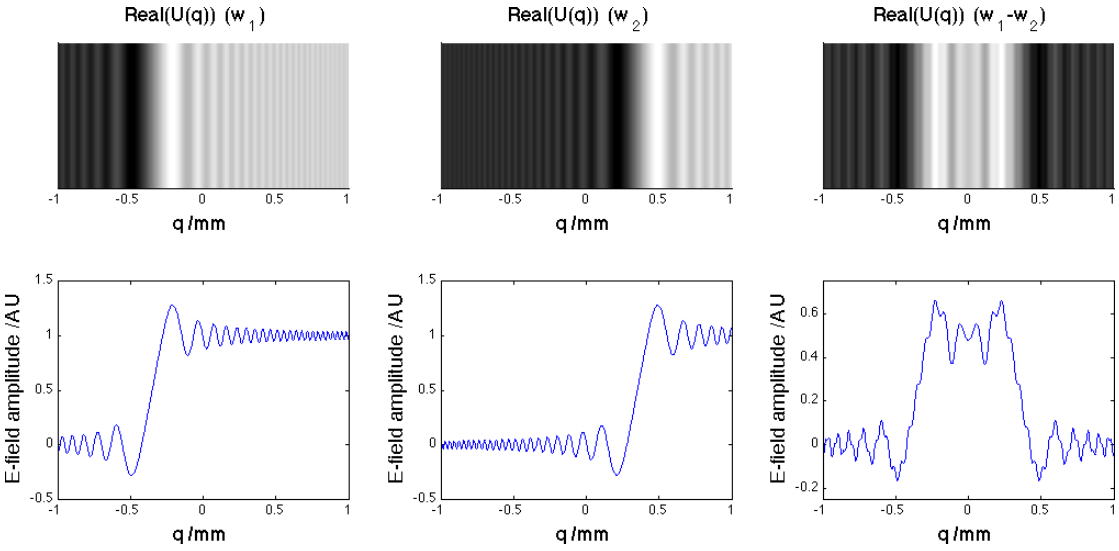

FIGURE 4.9: The diffracted fields of two spatially separated infinite edges may be subtracted (according to Babinet's principle) in order to simulate in infinite slit. Image and line plots both show the real part of the E-field amplitude as a function of position q. Left and middle show the contributions from each edge. Right shows the amplitude resulting from the subtraction.

In order to be compared with experimental measurements this calculated E-field must first be converted into an intensity distribution - this is done simply by multiplying $U(q)$ by its complex conjugate. The expression for $I(q)$ is shown in equation 4.29 and the resulting intensity distribution is shown in figure 4.10.

$$I(q) = \frac{I_0}{4}\left(\left[\mathscr{C}(w_1) - \mathscr{C}(w_2)\right]^2 + \left[\mathscr{S}(w_1) - \mathscr{S}(w_2)\right]^2\right) \tag{4.29}$$

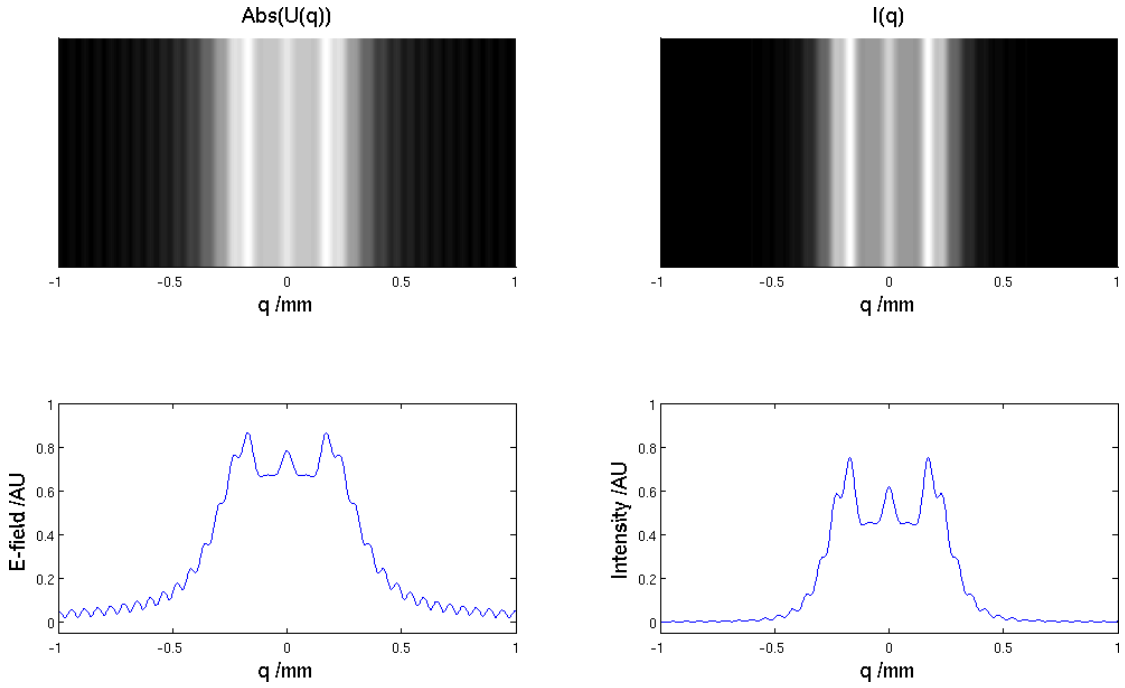

FIGURE 4.10: The calculated E-field (left) is multiplied by its complex conjugate to recover the intensity distribution (right). This simulation is calculated for a 300 μm wide slit diffracting 25 nm radiation. The source-slit and slit-screen distances are both $\sim$ 0.5 m. Note that the plots on the left differ from those in the right hand column of figure 4.9. This is because here the absolute value of the E-field is plotted whereas previously the real part was plotted.

4.2.2.3 Fresnel diffraction by a rectangular aperture

In order to calculate the diffraction pattern produced by a square aperture, two orthogonal, transverse spatial dimensions are required. On the screen let these be denoted as q and q'. The parameter v (the dummy variable for integration in the y direction) can no longer be ignored. Equations 4.18 must now be integrated over both spatial dimensions (both parameters) and therefore become:

$$C = \int_{w_2}^{w_1} \int_{w_2'}^{w_1'} \cos\left(\frac{\pi}{2}\left(u^2 + v^2\right)\right) \, \mathrm{d}u\mathrm{d}v$$

$$S = \int_{w_2}^{w_1} \int_{w_2'}^{w_1'} \sin\left(\frac{\pi}{2}\left(u^2 + v^2\right)\right) \, \mathrm{d}u\mathrm{d}v \tag{4.30}$$

As before the limits of integration (w) are found by inserting the edge position into equation 4.26. (i.e. w_1' is the limit for the first edge in the y_0 dimension). In the case

of a square aperture $w_1 = w_1'$ and $w_1 = w_1'$ - for a rectangular aperture this is not the case. By making the following substitutions:

$$\int_{w_2}^{w_1} \cos\left(\frac{\pi}{2}u^2\right)\,\mathrm{d}u = \mathscr{C}(w_1) - \mathscr{C}(w_2) \qquad \int_{w_2'}^{w_1'} \cos\left(\frac{\pi}{2}v^2\right)\,\mathrm{d}v = \mathscr{C}(w_1') - \mathscr{C}(w_2')$$

$$\int_{w_2}^{w_1} \sin\left(\frac{\pi}{2}u^2\right)\,\mathrm{d}u = \mathscr{S}(w_1) - \mathscr{S}(w_2) \qquad \int_{w_2'}^{w_1'} \sin\left(\frac{\pi}{2}v^2\right)\,\mathrm{d}v = \mathscr{S}(w_1') - \mathscr{S}(w_2') \tag{4.31}$$

...and by using the trigonometric identities for the sum of two angles equations 4.30 for C and S become:

$$C = \left([\mathscr{C}(w_1) - \mathscr{C}(w_2)]\left[\mathscr{C}(w_1') - \mathscr{C}(w_2')\right]\right) - \left([\mathscr{S}(w_1) - \mathscr{S}(w_2)]\left[\mathscr{S}(w_1') - \mathscr{S}(w_2')\right]\right)$$

$$S = \left([\mathscr{S}(w_1) - \mathscr{S}(w_2)]\left[\mathscr{C}(w_1') - \mathscr{C}(w_2')\right]\right) + \left([\mathscr{C}(w_1) - \mathscr{C}(w_2)]\left[\mathscr{S}(w_1') - \mathscr{S}(w_2')\right]\right) \tag{4.32}$$

If equations 4.32 are then inserted into the expression

$$I = |B|^2\left(C^2 + S^2\right)$$

after simplification we find that:

$$\begin{aligned} I(P) = \frac{I_0}{4}&\left([\mathscr{C}(w_1) - \mathscr{C}(w_2)]^2 + [\mathscr{S}(w_1) - \mathscr{S}(w_2)]^2\right) \\ &\times \left(\left[\mathscr{C}(w_1') - \mathscr{C}(w_2')\right]^2 + \left[\mathscr{S}(w_1') - \mathscr{S}(w_2')\right]^2\right) \end{aligned} \tag{4.33}$$

Notice that the two bracketed terms are in fact the intensity distributions produced by two orthogonal slits of width $w_1 - w_2$ and $w_1' - w_2'$. In the case of a square aperture the intensity values calculated for each of the two slits are identical. Re-using the same values rather than re-calculating will save time in a numerical calculation.

In order to numerically calculate the full diffraction pattern produced on the screen equation 4.33 must be evaluated for each position on the screen $P_1(q, q')$. As asserted earlier and as illustrated in figure 4.6 the line $(P_{-1}P_1)$ has been approximated to be the z-axis. The w parameters for each edge were calculated by setting x_0 as $d/2$ and $-d/2$ in equation 4.26. This centered the aperture on the z-axis. In order to calculate the field that would be produced on the screen at a given offset from the z-axis the w parameters can simply be re-calculated. Effectively this translates the aperture relative to the source and measurement position. Equation 4.26 can be modified with an expression for the

edge position in terms of the q (the position on the screen). It therefore becomes:

$$\begin{aligned} w_1 &= \left[\frac{d}{2} + \frac{qr'}{s'+r'}\right]\sqrt{\frac{2(r'+s')}{r's'\lambda}} \\ w_2 &= \left[\frac{-d}{2} + \frac{qr'}{s'+r'}\right]\sqrt{\frac{2(r'+s')}{r's'\lambda}} \end{aligned} \tag{4.34}$$

Notice that the factor $\cos\delta$ has been approximated to 1 which implies that the illumination of the aperture is in all cases close to normal incidence.

A Matlab computer program was written in order to calculate the diffraction patterns produced by square apertures (this is included in appendix B.2). A 1D array q is created with all of the screen positions for which intensity values are required - for example these could represent the transverse position of pixels of a CCD. By inserting the array q and the other parameters of the diffraction problem into equation 4.34 arrays for w_1 and w_2 are produced. Functions to numerically evaluate the Fresnel sine and cosine integrals (using the adaptive Simpson quadrature) were written. These functions were used to evaluate $\mathscr{C}(w_1)$, $\mathscr{C}(w_2)$, $\mathscr{S}(w_1)$ and $\mathscr{S}(w_2)$ for each position q. These values can then be combined to evaluate the bracketed terms in equation 4.33. Note that to perform this calculation for a rectangular aperture, w parameter arrays would have to be calculated for all four edges separately. In the case of the square aperture the second bracketed term is simply the matrix transpose of the first. The multiplication of these two 1D arrays must reflect the fact that the arrays are orthogonal in order to produce a 2D array output. In order to do this the built-in Kronecker tensor function was used.

Figure 4.11 shows the results of a sequence of diffraction calculations each with a slightly different aperture size. It is possible to see the transition from Fraunhofer (for very small apertures) to the Fresnel region (for larger apertures).

The calculations in the upper row of figure 4.11 represent conditions where the Fraunhofer approximations could have been used. On the bottom row the conditions place the calculation firmly in the Fresnel regime and the characteristic checker-board pattern is produced. In the central row the conditions transition from Fraunhofer to Fresnel regimes. All of these images are scaled so that the brightest point in the calculation takes the highest value in the available dynamic range. Unfortunately this scaling prevents us from seeing fine structure that is present in the diffraction pattern. Figure 4.12 shows the central image from figure 4.11 plotted on a logarithmic intensity scale and using false colour shading.

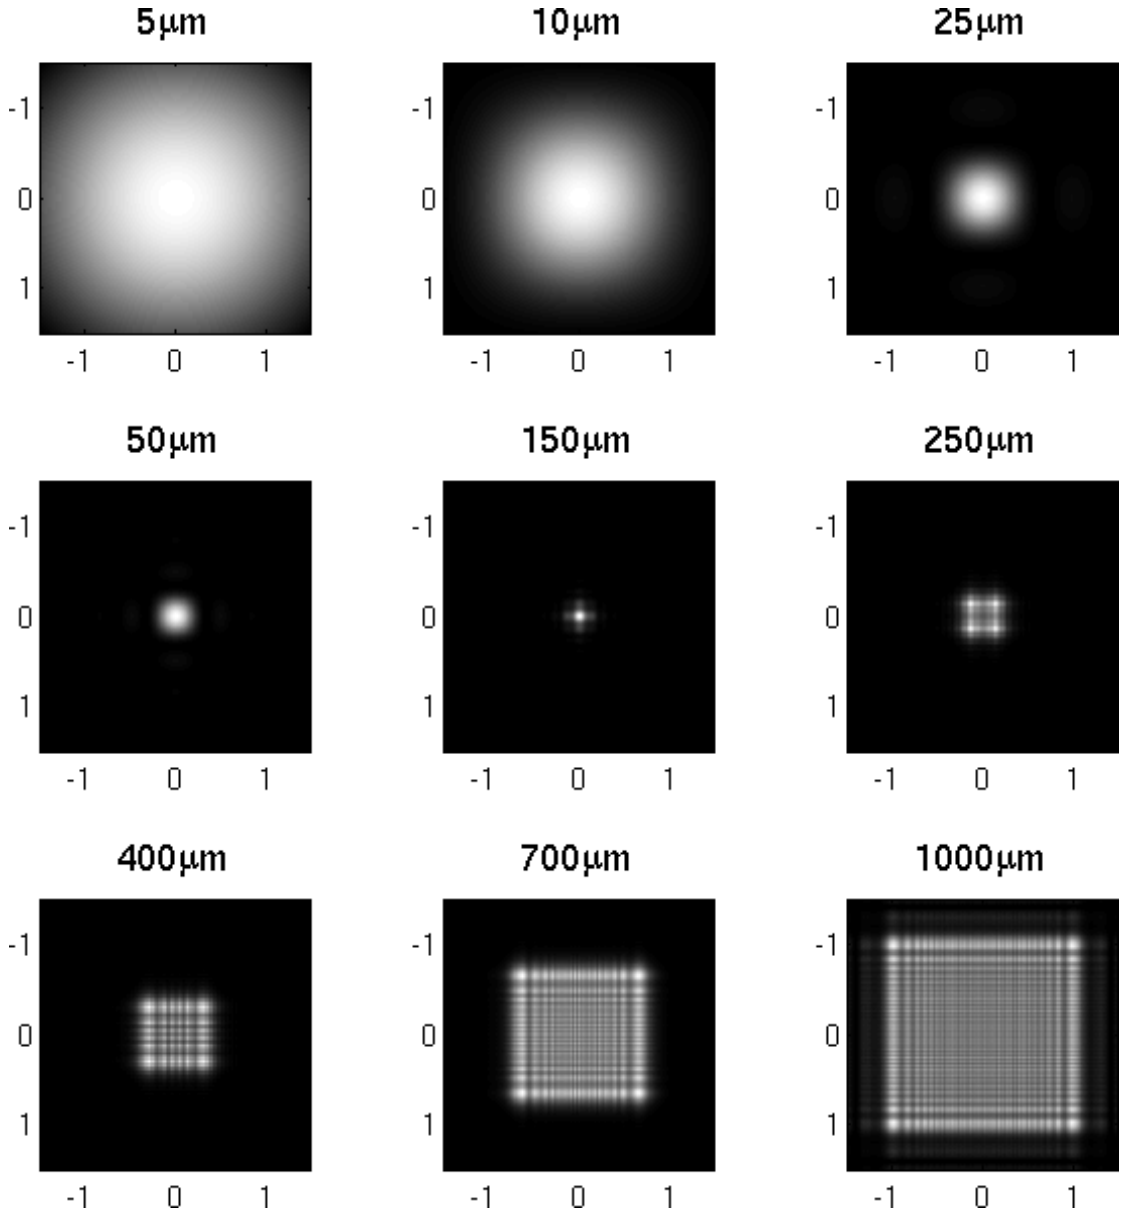

FIGURE 4.11: Calculated intensity profiles for a selection of square apertures (aperture width is written in the title above each plot). The simulated wavelength was 25 nm, the source-slit and slit-screen distances are both $\sim$ 0.5 m. The axis scales represent translation on the detector screen in mm.

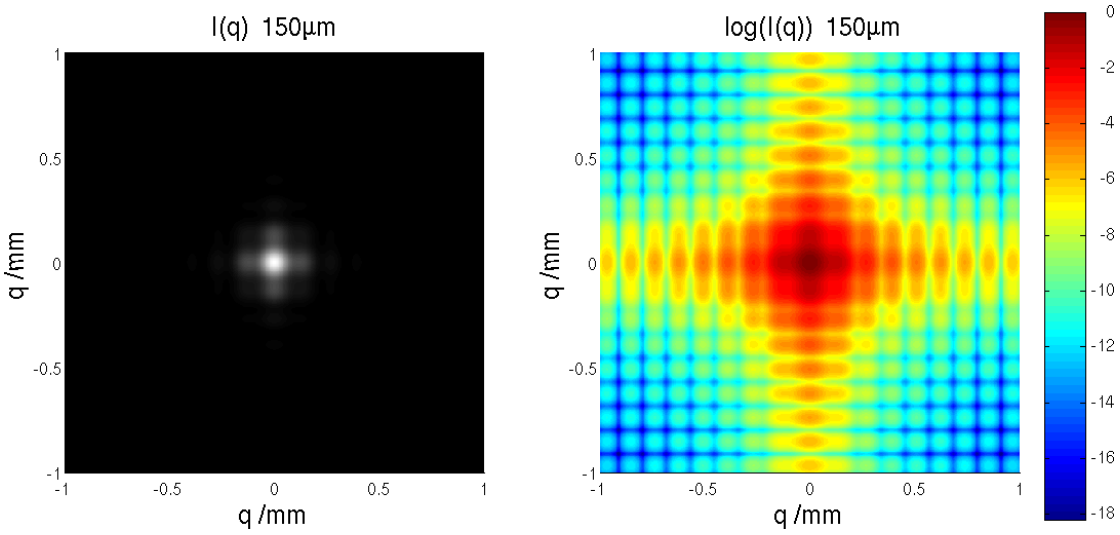

FIGURE 4.12: Diffraction pattern produced by 25 nm light passing through a 150 μm square aperture. Left: plotted on a linear scale and displayed in monochrome. Right: plotted on a logarithmic intensity scale using a false colour scale (dark red is most intense, dark blue is least intense, the markers on the colorbar give the scale in dB).

The results of these calculations are in agreement with the expected behaviour. For example, rearranging the equation for Fresnel number (4.2) with an assumed Fresnel number of 1, we find a simple expression for the aperture radius at which the transition from Fresnel to Fraunhofer regime occurs:

$$a = \sqrt{L\lambda} \tag{4.35}$$

Inserting the values $L = 0.5$ m and $\lambda = 25$ nm we find that the transition occurs for an aperture radius of approximately 120 μm (or an aperture width of 240 μm) - just as was found in figure 4.11.

Results from this diffraction pattern calculation will be further compared with experimental observations and utilised in the data analysis described in section 4.3.

4.3 Spatio-Spectral Reconstruction

This section describes a technique for analysing Fresnel diffraction patterns in order to extract spectral information about the light source. Furthermore, by analysing an array of diffraction patterns, spatially-resolved, spectral information is recovered. This work was carried out in the XUV using an HHG source although there is no reason why this technique could not be applied in other regions of the spectrum. This work was published in Nature Physics[1].

Subsections 4.3.3 and 4.3.5 deal with two different versions of the reconstruction algorithm, each designed to work best for sources of different types. The least squares fitting routine described in section 4.3.3 is most suited to sources consisting of discrete spectral components; this works well with the harmonic spectrum produced by HHG and consequently is the more fully developed. Subsection 4.3.5 discusses reconstruction using the maximum entropy approach. This approach is more suited to smoothly varying, continuum-like, spectra such as might be found on a synchrotron source; despite this, attempts at retrieving HHG spectra were met with some success.

Before going into specific details of the two reconstruction algorithms some general details of the experimental setup and the approximations used are provided below in sub-sections 4.3.1and 4.3.2.

4.3.1 Experimental Setup

The setup of this experiment is shown diagrammatically in figure 4.13. A 5 cm long 150 μm bore capillary with a 3 cm long active region was used to generate a beam of soft x-rays. As described in chapter 3 this gives rise to a spectrum consisting of discrete peaks corresponding to the odd harmonics of the input laser pulse. Using argon as the interaction medium we typically observe the 17^{th} to the 31^{st} harmonics of the Ti:sapphire laser pulse; this produces peaks in the x-ray spectrum between 46 and 25 nm.

To prevent absorption of the XUV beam the optical path must be evacuated as discussed in section 2.3.1. The x-ray diffraction sample is inserted into the beam approximately 0.5 m after the exit of the HHG capillary. In this experiment the diffraction sample is a Ni wire grid consisting of $\sim$ 38 μm thick bars which cross at 90° in order to define an array of 340 μm square apertures. The grid is also used to support a 200 nm thick Al foil which reflects most of the laser radiation but partially transmits the x-rays (see subsection 2.3.2).

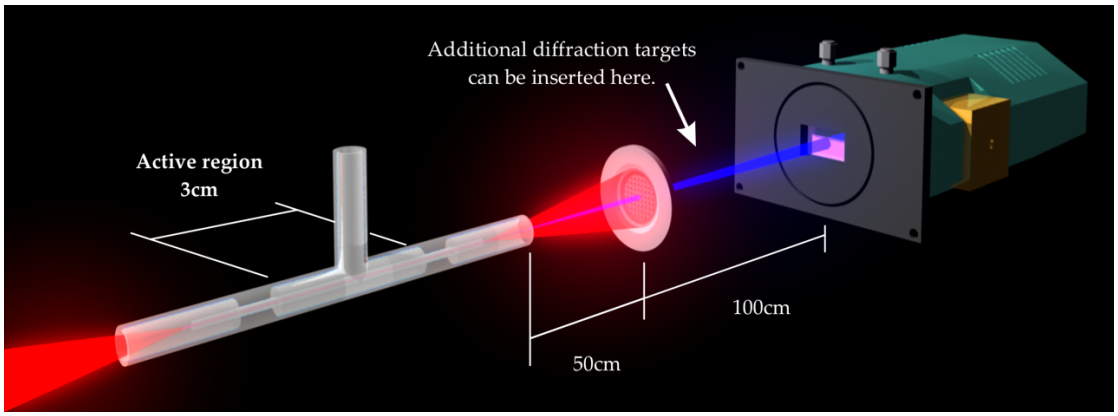

FIGURE 4.13: Diagram showing the setup used to create Fresnel diffraction patterns using a wire mesh as a diffraction target.

Not shown in figure 4.13 but of great importance is the filter holding device. As described in section 2.3.2.2 this acts as a baffle plate, preventing light from passing around the edge of the foil. This reduces the scattered light that arrives at the detector and consequently improves the signal-to-noise ratio. This holder was carefully designed to be as "light-tight" as possible whilst still providing a bypass path for vacuum pumping - without this a pressure differential would build up and eventually burst the foil.

The diffracted x-ray beam is allowed to propagate approximately 1 m beyond the diffraction sample before detection. This allows the diffracted beam to expand to a larger size on the detector (improving resolution); but more importantly it ensures that the plane of the detector is inside the Fresnel region (see subsection 4.2.2). As discussed in section 4.2 this can be confirmed with a simple calculation of the Fresnel number. Recalling equation 4.2 and inserting typical experimental values for each parameter:

$$F = \frac{a^2}{\lambda R}$$

The aperture radius a is set to 170 μm (340 μm width), λ to 25 nm and the aperture-detector distance R to 0.5 m). A Fresnel number of 2.31 is obtained - placing the detector firmly in the Fresnel regime ($\ll 1$ = Fraunhofer, ≥ 1 = Fresnel, $\gg 1 \approx$ shadow mask). In fact for this aperture size and aperture-detector separation all wavelengths below $\sim$ 58 nm (the 13^{th} harmonic) are in the Fresnel region.

It is vital to this experiment that the diffraction pattern is recorded in the Fresnel regime. If the diffraction pattern is measured in the Fraunhofer regime (very small Fresnel number) no correlation can be made between different spatial regions of the beam and different parts of the diffraction pattern. Spatial resolution would therefore

be lost. This can be readily understood by considering the fact that propagation into the far field is equivalent to a 2D Fourier transform. Taking the 2D Fourier transform of an image converts it into a spatial frequency spectrum. Whilst the x-y coordinates describe transverse position in the aperture plane they would describe spatial frequency in the plane of a screen in the 'far field'. Another way to explain this is that in the 'far field' the beam has diffracted enough that light from any given position in the aperture is distributed evenly over the entire screen. Conversely, in the near field (large Fresnel numbers) the diffraction pattern tends towards a direct shadow mask. Not only is this near field diffraction pattern difficult and time consuming to calculate (typically requiring finite element modeling) as the pattern approaches the shadow mask limit the interference fringes get closer together and become harder to detect. Eventually when no fringe pattern can be detected all spectral resolution of this technique is lost.

The spatial intensity distribution of the diffracted x-ray beam is measured using an Andor CCD (described in subsection 2.4.2). The specific camera used in this experiment is a DX-434-BN, it has 1024 $\times$ 1024, 13 μm square pixels. Although higher spatial resolution could be achieved by using a camera with smaller pixels their relatively large size ensures high sensitivity to low level signals. The larger pixel size also provides a deep quantum well. A deep quantum well means that saturation occurs at a higher level, improving the dynamic range of the camera. Good dynamic range is especially important when measuring diffraction patterns as it allows details to be resolved simultaneously in the dark and light regions of the image. The camera also features 16 bit digitisation to take advantage of the large dynamic range of the CCD. The CCD was thermo-electrically cooled to -30°, which simultaneously reduces the dark current and increases the signal gain; when combined, these effects significantly improve the signal-to-noise ratio achieved. It is very important to use a detector that features both high sensitivity and large dynamic range, as without these features it is impossible to record the subtle, low-intensity features contained within the diffraction pattern.

4.3.2 Initial Findings

This section begins by describing the initial findings which prompted this experiment. It then goes on to discuss checks that were made on the validity of approximations used in the spatio-spectral analysis technique.

An example of the measured data is displayed in figure 4.14. Upon close inspection, the measured diffraction pattern was found to vary from one aperture to the next (as highlighted by the insets). The diffraction patterns produced by each aperture were found to change as the gas pressure of the HHG source was altered. Furthermore the

variation in the observed diffraction pattern was not fixed to any particular location, rather it moved in reaction to experimental conditions. It was therefore proposed that this spatial variation in the observed diffraction pattern was a result of changes in the spectral content of the x-ray beam.

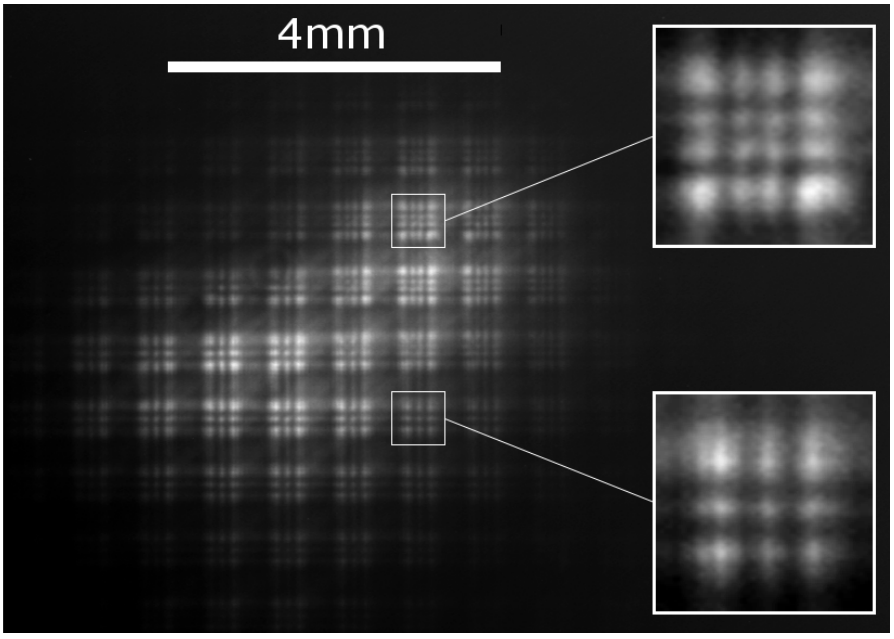

FIGURE 4.14: Example of a measured diffraction pattern (340 μm square apertures separated by 38 μm bars). The insets show the variation in the diffraction pattern produced by different apertures. The scale bar represents 4 mm in the CCD plane.

Since the location of the modified diffraction pattern was seen to move relative to the grid it seemed unlikely that the grid (which was stationary) could be the cause. Nonetheless, in order to completely eliminate the possibility that this effect was caused by variation in the grid geometry, electron microscope images of the several grids were studied. A computer program was written to assist in the analysis of these images. The code for this software and some explanation of its operation is included in appendix B.1. An example of the output from this analysis software is shown in figure 4.15.

The SEM image analysis software detects the edges of the mesh wires using the Sobel method. The red lines are then fitted to the edges in sets of four so that the width of the aperture, the width of the bars and the angular offset are used as fitting parameters. Variation in aperture size was found to be approximately 1% (this measurement was mainly limited by the accuracy of the SEM geometry calibration).

A 1 % level of aperture size variation is not sufficient to cause diffraction pattern changes as large as were observed. Using Fresnel diffraction simulations (as described in subsection 4.2.2.3) an average pixel error of less than 1.5 % (of full scale) is found to be

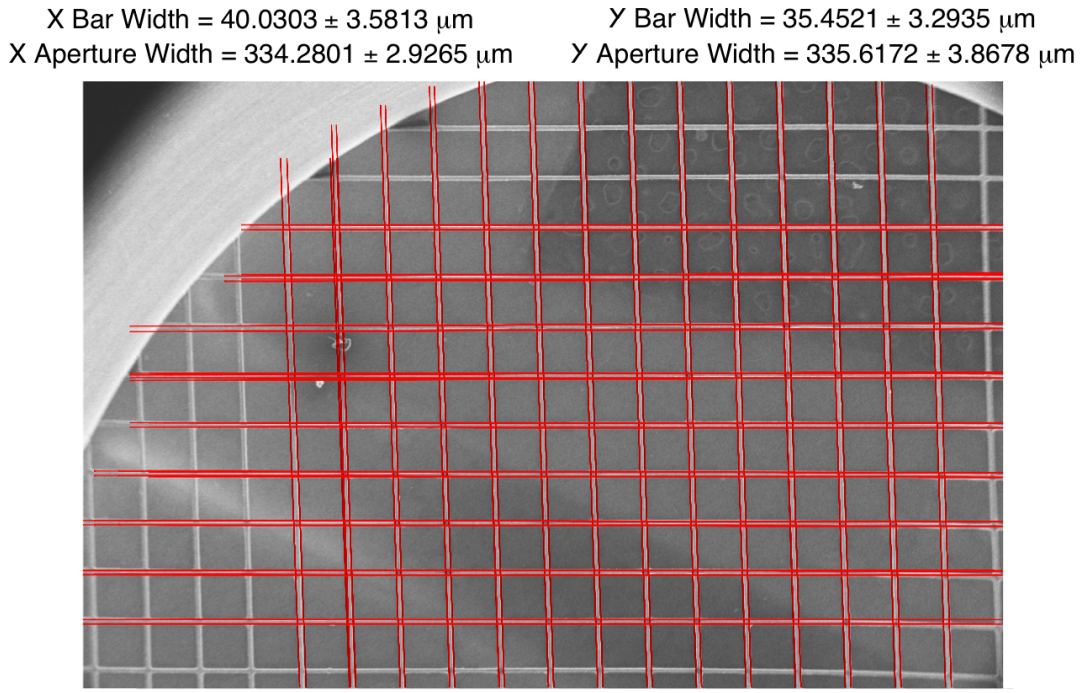

FIGURE 4.15: Typical output from the SEM image analysis software described in appendix B.1. This image shows analysis of the grid geometry of a Ni wire mesh diffraction target. Errors in aperture size are approximately 1%.

introduced by a 1 % change in aperture size. This is demonstrated in figure 4.16 which shows calculated diffraction patterns for square apertures of 340 and 345 μm diameter (A and B respectively). For this level of pixel error there is very little visible alteration to the diffraction pattern.

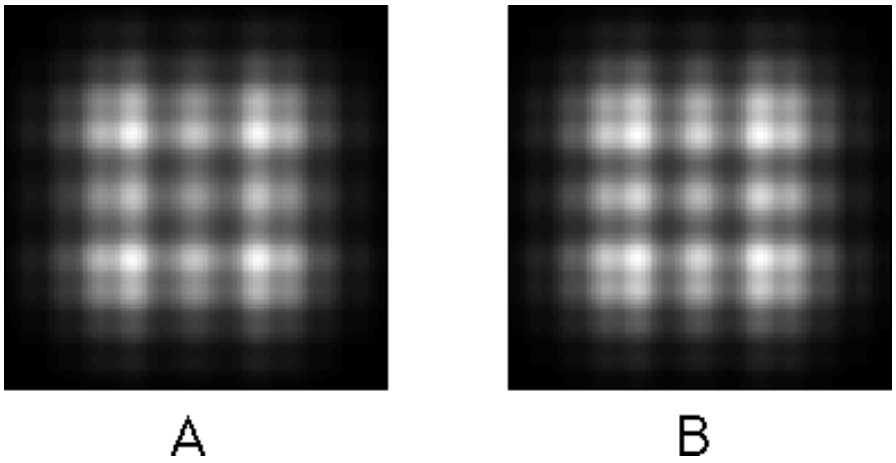

FIGURE 4.16: Calculated diffraction patterns for $\lambda = 41$ nm. Pattern A has a 340 μm square aperture whilst pattern B is 345 μm square.

A similar calculation can be performed in order to establish that spectral changes are a potential source of diffraction pattern variation. Furthermore, such a calculation will provide information about the ability to distinguish between diffraction patterns produced by different wavelengths. Figure 4.17 shows results from such a calculation. Diffraction patterns are calculated at evenly spaced wavelength intervals. The diffraction pattern for the harmonic under test is then subtracted from each diffraction pattern in turn and the average pixel error is plotted against the wavelength.

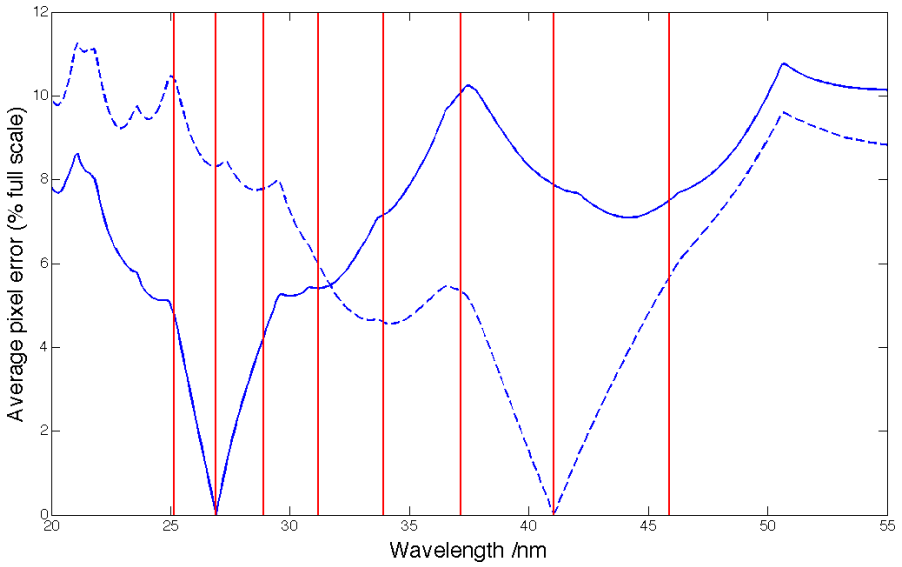

FIGURE 4.17: Plot showing the wavelength selectivity for the 19^{th} (dashed) and 29^{th} (solid) harmonics, in terms of pixel error (% full scale). The red lines indicate the harmonics that are typically generated by the HHG system.

Although it becomes more difficult to distinguish between harmonics as the wavelength is reduced, even in the worst case the pixel error between harmonics is greater than 4 %. In most cases, the difference between adjacent harmonics is greater than 5 %. For the spectral range that is typically generated by the HHG system it seems that the 340 μm square aperture will be appropriate for the spectral retrieval technique. By calculating the wavelength selectivity in this manner it should be possible to optimise the aperture size for different spectral regions.

In this work the diffraction patterns produced by each aperture are considered separately[†]. Each aperture forms a different diffraction pattern (depending upon the spectral content incident on it) and this falls onto a region of the detector uniquely associated with that aperture. Light diffracted through angles large enough that it arrives on a

[†]It may be possible to develop a genetic algorithm that recovers the spectrum of the whole beam in one go. An estimate of the beam's spatial-spectral profile would be made and then evolved until a good fit to the measured data is achieved. Whilst this is in principle possible, each loop iteration would require a complex and lengthy calculation and the idea is not discussed further here.

region of the detector associated with a neighbouring aperture is neglected. The validity of this approach can be tested by measuring the difference between calculated diffraction patterns for an isolated aperture and for an array of apertures. Figure 4.18 shows just such a comparison with simulated illumination by the 27^{th} harmonic. In both calculations the source is on axis with the centre of the image; the fact that apertures at the edge of the array are not on axis with the source is incorporated into the calculation. A coherent sum is used to combine the contributions of all apertures so that inter-aperture interference effects will be visible if present.

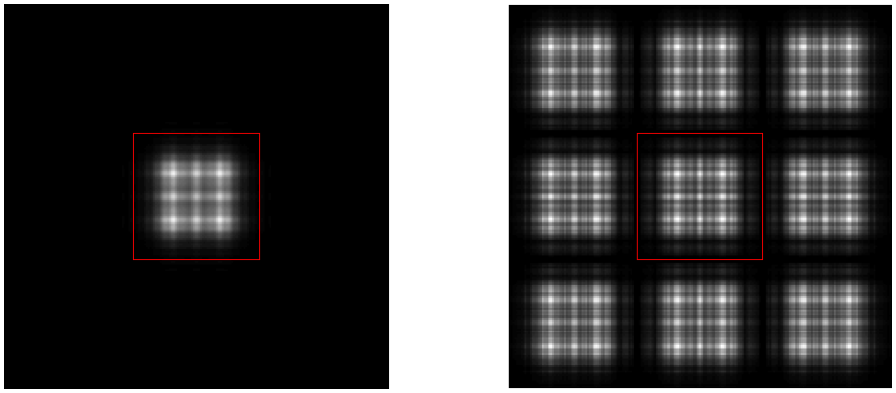

FIGURE 4.18: Image on left shows the calculated diffraction pattern for an isolated aperture illuminated by the 27^{th} harmonic. The right hand image shows calculated diffraction pattern for an array of identical apertures under the same illumination conditions. The average pixel error between the two boxed regions is $\sim$ 7.5 % of the full scale.

The average pixel error as a result of light from surrounding apertures is found to be $\sim$ 7.5 % of the full scale (see figure 4.18). This error is slightly reduced for higher harmonics since the energy of the beam is diffracted through a smaller angle for such a beam. As shown in figure 4.18 there are now some visible differences between the two boxed regions. The interference pattern appears to have been sharpened causing the addition of fine dark lines. Despite this the underlying character of the pattern remains much the same. When measuring the diffraction pattern in the experiment very fine features such as these dark lines will tend to blur out because of random noise - consequently the measured diffraction pattern is likely to be closer in appearance to that calculated for the isolated aperture.

Note that in the simulated data the dark regions between the diffraction patterns appear narrower than in the experimental data (see figure 4.14). It is thought that this effect is due to the thickness of the Ni wire grid (which is not taken into account in the simulated diffraction pattern). Increasing the dark space between diffraction patterns reduces the errors due to light from neighbouring apertures. If the bar width in the simulation is

increased to give the same mark-to-space ratio as is observed experimentally then the average pixel error is reduced to approximately 5 % (of full scale). Given the considerable simplification that is obtained by neglecting this inter-aperture light, this level of error is tolerable. In future experiments, however, it will clearly be beneficial to obtain custom-made grids either with slightly wider, or higher aspect-ratio bars.

The overall level of the error caused by inter-aperture light could also be reduced by measuring the diffraction patterns at higher Fresnel numbers. This can be achieved either by positioning the camera closer to the aperture or by increasing the aperture size. As the camera is moved towards the aperture there is a trade-off between improving spatial fidelity due to less inter-aperture light and reducing spectral resolution because fringes become less distinct and fewer CCD pixels are illuminated per aperture. If the aperture size or bar width is increased the spatial fidelity is improved at the expense of spatial resolution and spectral resolution.

4.3.3 Least squares fitting

In this spectral reconstruction algorithm the aim is to compare the measured Fresnel diffraction pattern with simulated ones calculated for a sequence of different wavelengths. A least squares fitting routine is then used to determine the proportions with which the simulated diffraction patterns must be added together in order to best approximate the measured pattern.

The assumptions made in this calculation are explained in more detail either in section 4.3.2, or below; however for quick reference they are listed here:

- The diffraction grid is assumed to be uniform and consist of 340 μm square apertures separated by opaque 38 μm bars.
- The x-ray spectrum is assumed to consist of odd harmonics of the laser centre wavelength.
- The spectral width of each harmonic is ignored, the calculation uses only the central wavelength.
- The diffraction pattern is calculated using the Fresnel approximation.
- The diffraction pattern of each harmonic is calculated separately - these are then combined using an incoherent sum. I.e. phase is ignored.
- Each aperture is treated as though it is centred on the x-ray beam - the radius of curvature is included but spatial offset from the axis is not.

- The diffraction pattern for each aperture is treated separately. I.e. light diffracted from one aperture into its neighbour is ignored.
- The intensity of the harmonics is assumed to be uniform across each aperture.
- Some intensity that falls in the geometric shadow of the bars is ignored - the calculation uses mainly the more intense diffraction that falls directly behind each aperture.

Using the prior knowledge that the HHG source produces discrete spectral peaks, the problem can be further simplified. Only the diffraction patterns produced by the centre wavelength of each harmonic peak need to be calculated - the measured pattern is then decomposed into contributions from each harmonic. Effectively the shape and spectral width of each harmonic peak has been neglected. In this way the problem is reduced to that of finding the appropriate intensity coefficients for each harmonic peak in order to achieve a best fit to the measured data.

Approximating a spectrally broad beam to its centre wavelength is an approach which is often used in optics. The effect of a source's spectral width is typically to average out some of the sharper features of the diffraction pattern. This can be demonstrated very simply; figure 4.19 shows a comparison between the diffraction pattern produced by a single spectral component at 33 nm (left) and the diffraction pattern produced by a 1 nm[†] wide (FWHM) intensity distribution also centred at 33 nm (right). Visually the diffraction patterns remain similar, although much of the fine detail is averaged out in the broad spectrum case and the intensity profile has a flatter top. The average pixel error between the two patterns is ~ 3 % of the full scale.

A second minor simplification arises from the fact that the CCD detector measures only the time averaged intensity of the E-field incident on it. This means that it is possible to calculate the diffraction patterns produced by each of the spectral components separately and then use an incoherent sum (i.e. adding intensities rather than E-fields) in order to calculate the combined diffraction pattern. This approach means that interference fringes in the diffraction pattern that are caused by phase differences between spectral components will be neglected. This is perfectly valid since the beat period between adjacent spectral peaks is very short (1.3 fs for the 29^{th} and 31^{st} harmonics). Therefore, any details in the diffraction pattern that result from interference between different spectral components will tend to average out over the integration time of the CCD.

The remaining part of this sub-section will give a detailed step-by-step explanation of the spectral analysis technique; for quick reference the key points are listed below:

[†]Typically the FWHM peak width for an HHG source would be slightly less than 1 nm.

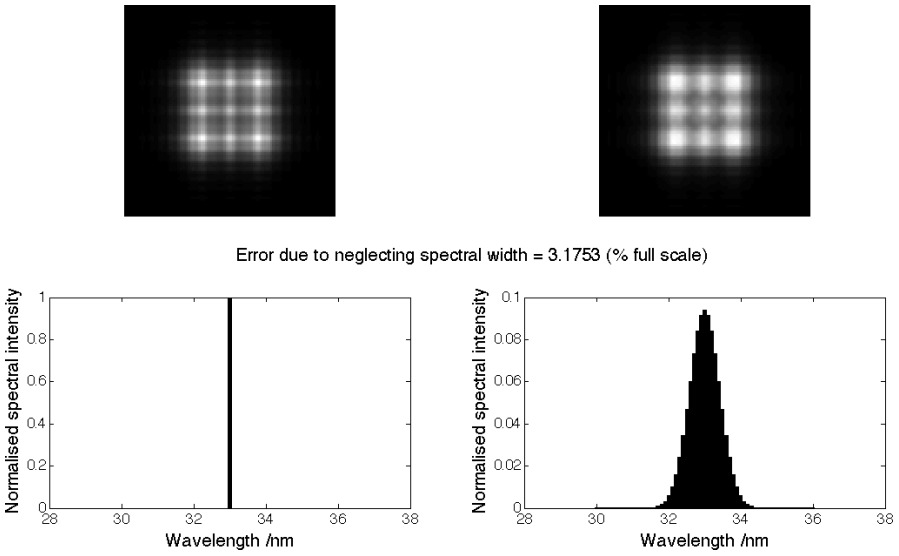

FIGURE 4.19: Diffraction patterns produced by an aperture illuminated by a discrete spectral line at 33 nm (left); and by a 1 nm FWHM spectral intensity distribution centred at 33 nm (right). The average pixel error between the two patterns is ~ 3 % of the full scale.

- Process measured data into usable form.
- Generate simulated diffraction patterns.

...then considering each aperture separately:

- Execute least squares fit to obtain harmonic coefficients.

...then finally:

- Combine data from all apertures into a single data array.
- Present data in a suitable format.

Before the spectral reconstruction can begin the CCD data must be processed into a suitable form. The Ni wire mesh has an arbitrary rotation relative to the CCD camera pixels. Linear interpolation is used to re-sample the data so that the grid wires are aligned to the CCD pixels, which simplifies the next step. Once aligned the diffraction patterns produced by each aperture are extracted from the main image and stored in the measured data matrix. These steps were performed by the matlab code described in appendix B.3.1.

Once the diffraction patterns for each aperture have been isolated they are then sorted in order of contrast ratio. i.e. Patterns that have a higher signal-to-noise ratio are

ranked above apertures that contain little information. Some parts of the CCD image are not illuminated by the x-ray beam - consequently they cannot contain any useful information. A user adjustable limit on the number of apertures to be extracted from the CCD image eliminates these lower ranking apertures.

The illumination of each aperture is typically not uniform. Intensity gradients at the aperture lead to asymmetry in the diffraction pattern. This asymmetry cannot be replicated in the diffraction pattern simulation, and would therefore hinder the fitting routine in finding matches between the measured and simulated data. The intensity gradient across the aperture is of no use in the spectral recovery technique - the aim is to produce a single set of spectral coefficients for each aperture, not to resolve details inside the aperture. Fortunately, the measured diffraction data contains a considerable amount of duplicated data because of its four fold symmetry. To gain advantage from this redundant data the average is taken of four copies of the measured data in different orientations. This produces a symmetrised version of the original data (see figure 4.20).

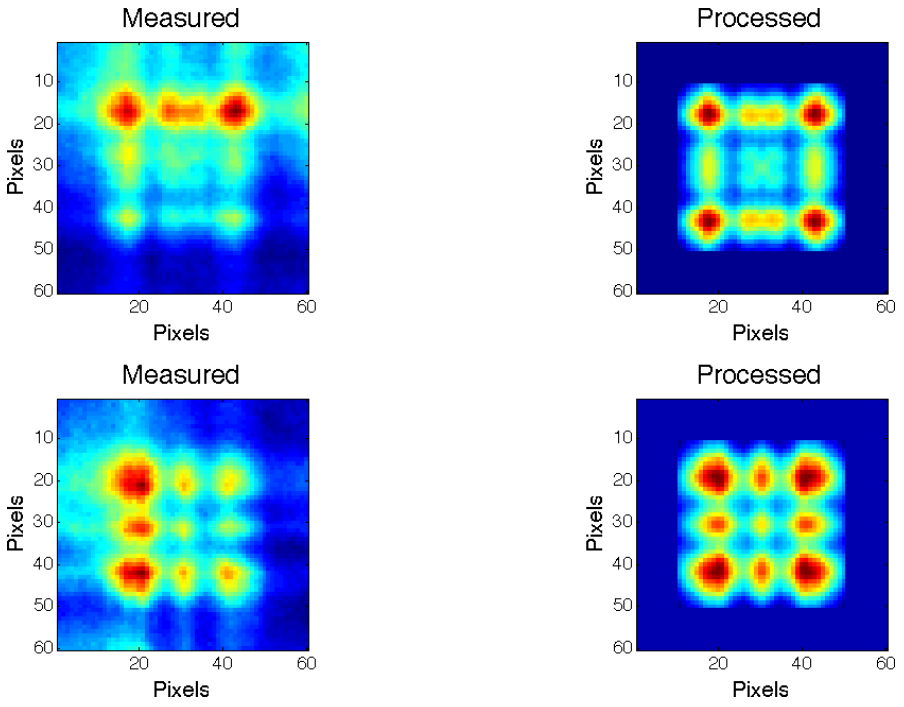

FIGURE 4.20: Comparison between raw input data and processed data for two different apertures in the same CCD image. Note that the processed data has been symmetrised and the background signal has been removed from the dark area surrounding the diffraction pattern. Despite the processing it is clear that the underlying shape of the diffraction pattern has been maintained.

As shown in figure 4.20 the signal in the 'dark area' surrounding the main part of the diffraction pattern has also been removed. This removed data is affected by the surrounding apertures to a greater degree than the central part - its inclusion would

have a negative effect on the spatial fidelity of the technique. If left in place the fitting algorithm interprets this data primarily as very long wavelength light resulting in a large value for the lowest simulated harmonic.

Several simulated diffraction patterns are generated for comparison with the measured data. In accordance with the HHG source the diffraction patterns produced by the odd harmonics of the input laser centre wavelength are calculated. The number of simulated diffraction patterns generated can affect the spectral range of the reconstruction technique. In this work odd harmonics 1–51 were used - this range considerably exceeds that which is generated by the HHG source. In this case the limits of the spectral range will be determined by other factors such as CCD resolution.

Each diffraction pattern is calculated using the technique described in section 4.2.2.3 and the code included in appendix B.2. The aperture size and source-detector distances are set to match the experimental conditions. The output resolution of this simulated data must also be set to match that of the CCD detector (each diffraction pattern and its surrounding 'dark area' was typically 60 × 60 pixels). As a programming convenience the 2D image data matrix for each simulated diffraction pattern is reshaped into a 1D vector (i.e. 60 × 60 becomes 1 × 3600). The data vectors for all harmonics can then be concatenated to form a 2D simulated data matrix (this is shown graphically in figure 4.21). This step does not affect the data values in any way but it allows the algorithm to execute faster and makes handling the data easier.

The spectral reconstruction is performed by finding the set of coefficients with which the diffraction patterns in the simulated data matrix must be added in order to match the diffraction pattern for a single aperture as measured on the CCD. (This process is then repeated for each aperture contained in the CCD image). This routine can be expressed mathematically as shown in equation 4.36:

$$P_m(x, y, i, j) = \sum_{h=1}^{\infty} P_s(x, y, i, j, h) I(i, j, h) \tag{4.36}$$

Here P_m is the measured diffraction pattern, P_s is the simulated diffraction data and I represents the intensity coefficients with which the simulated data must be added. x and y are the spatial coordinates within the $(i, j)^{th}$ aperture and h is an odd integer since only odd harmonics are generated.

This problem therefore defines a linear system of equations which must be solved in order to retrieve the spectral information. The system defined here is highly overdetermined since the measured data contains 3600 data points (60 × 60 pixels) and there are only 26 unknown coefficients (the spectral intensity value for each harmonic). This means

that there is no unique solution to the problem and consequently the output must be considered a "best fit" to the measured data rather than a direct inversion. Figure 4.21 shows a pictorial representation of this linear system problem.

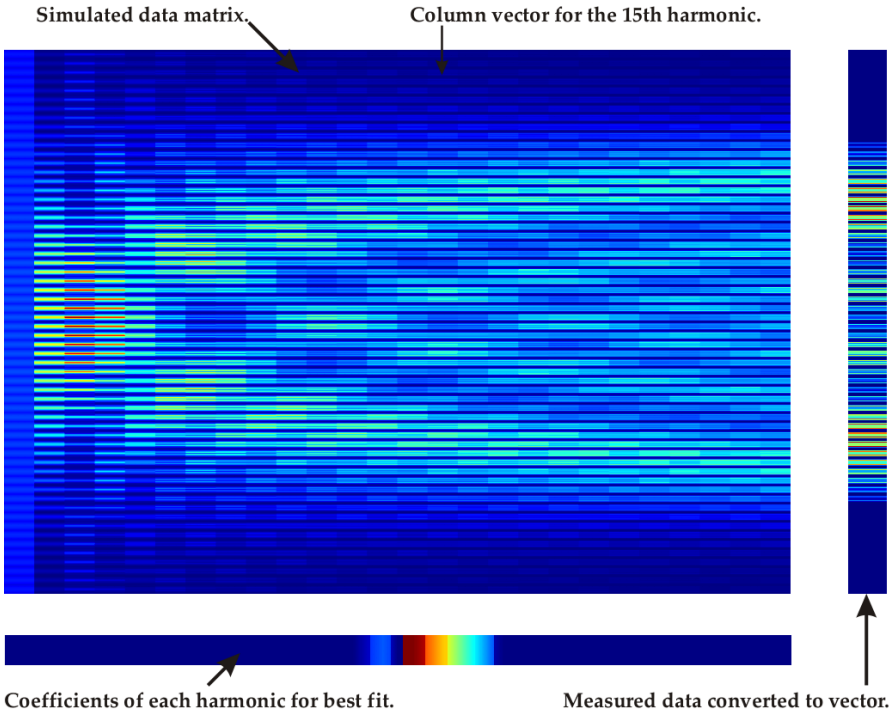

FIGURE 4.21: Graphical representation of the least squares fitting routine. Each vertical column in the main block contains the simulated diffraction data calculated for a different harmonic. (The 2D 60 × 60 grid of simulated pixels has simply been wrapped to produce a 1 × 3600 column). The measured diffraction pattern (60 × 60 pixels) is also re-sampled into a column vector format and is shown on the right. The row vector at the bottom of the figure shows the coefficients for each harmonic that are required to obtain the best fit to the measured data.

The large block in figure 4.21 represents the simulated data matrix; each column of which contains the calculated diffraction pattern of the aperture under illumination by a different harmonic. The column vector on the right represents the processed diffraction pattern data from one aperture - as extracted from part of the CCD image. The row vector at the bottom of figure 4.21 shows the set of coefficients with which the columns of the simulated data matrix must be added in order to provide the best fit to the measured data. This row vector is therefore the spectral data that we wish to obtain.

The matlab programming environment contains a number of linear system solvers which would be capable of solving this least squares fitting problem. The function 'lsqnonneg.m' was used here as this constrains the coefficient values it outputs to be positive. The extra constraint will improve the robustness of the spectral analysis algorithm and is justified since negative values for spectral intensity cannot represent a real physical

solution to this problem. By making use of this 'built-in' function the fitting process is essentially reduced to a single line of code:

```
spectrum = lsqnonneg(simulated_diffraction,data_in);
```

Where `simulated_diffraction` is the simulated data matrix, `data_in` is the processed CCD data and `spectrum` is the resulting set of harmonic coefficients. The solution provided by 'lsqnonneg.m' minimises the quantity $norm(P_s(x,y,h)I(h) - P_m(x,y))$ where the indices i and j have been dropped since the diffraction patterns of each apertures are dealt with individually. The function $norm(x)$ accepts a vector or matrix input and returns a scalar value that reflects the average magnitude of x. 'lsqnonneg.m' therefore minimises the error between the measured and simulated data and effectively provides a least squares fit.

The fitting routine operates in a loop and outputs a set of spectral coefficients for each of the apertures contained in the CCD image. These coefficient sets can be viewed individually but it is more interesting to study the variation from one aperture to the next. Consequently, for ease of handling the resulting spectral coefficients are concatenated into a single 4D matrix which then represents the spatio-spectral profile of the x-ray beam. (The four dimensions are: i and j the aperture coordinates, h the harmonic number and I the spectral intensity coefficient value).

4.3.4 Results

The most obvious way of checking the performance of the spectral recovery algorithm is to test the results against measurements made with the grazing incidence spectrometer. In order to do this the spectral coefficients for all apertures in the CCD image are summed. The resulting totals for the harmonic intensity coefficients are analogous to a spectrum of the whole beam (sampled only at the harmonic wavelengths).

Some examples of the comparison between spectrometer and diffraction analysis data are shown in figure 4.22. This data was recorded with the HHG gas pressure set to a variety of different values in order to generate spectral profiles with several different shapes. In each sub-plot the blue curve displays the spectrum as recorded by the grazing incidence spectrometer and the red lines indicate the retrieved harmonic intensity coefficients.

The data shown in figure 4.22 is representative of the typical output from the diffraction pattern analysis technique. The agreement between the spectrometer data and the harmonic coefficients is far from perfect however there are several reasons why this should be so.

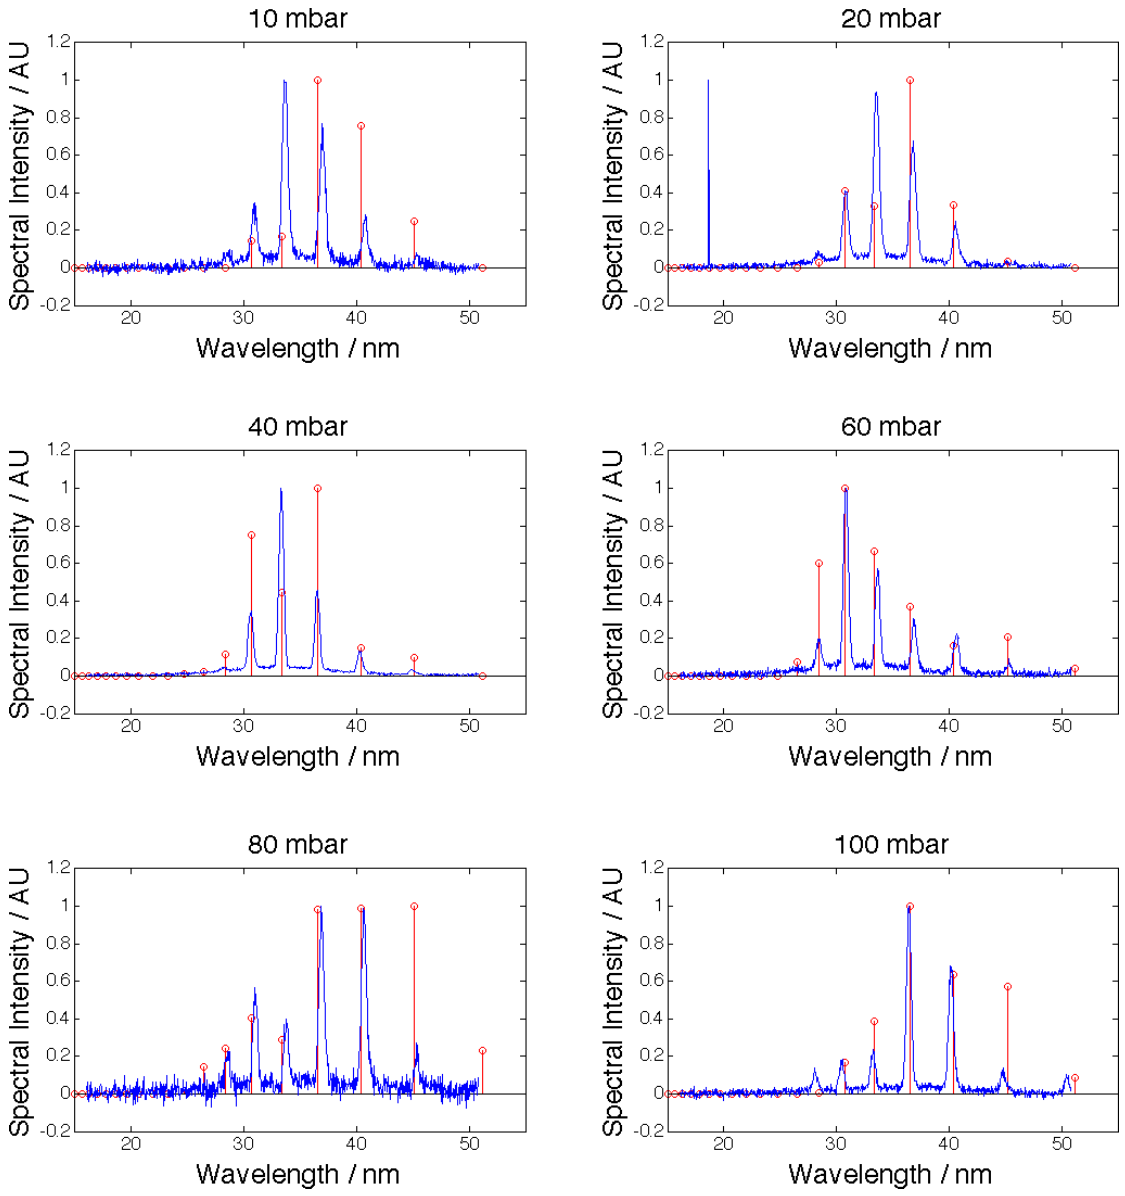

FIGURE 4.22: Comparison between spectrum measured with a grazing incidence spectrometer (blue line plot) and spectral coefficients recovered from the diffraction pattern (red stem plot). Argon pressures in the HHG system ranges from 10 to 100 mbar as labeled.

Firstly, it is not possible to record the diffraction pattern and the spectrum simultaneously. In order to pass the x-ray beam to the spectrometer the CCD camera must be moved out of the beam using an electronically actuated translation stage. Operation of this stage is slow and also introduces heat into the vacuum chamber, affecting the performance of the CCD camera. Because of this several minutes may have passed between the collection of the spectrum and the diffraction pattern. The output of the HHG capillary is extremely sensitive to a number of factors - it is therefore possible for the output to change slightly between measurements.

Secondly, the sampling regions are not necessarily the same. By summing the harmonic coefficients obtained for all apertures in the CCD image the sampling region essentially covers the entire beam. In the case of the spectrometer a slit is used and this selects only a portion of the beam (which portion is unknown).

Considering these points, and the errors expected because of the approximations made in the reconstruction technique, the level of agreement shown in the data is in fact highly encouraging. In all cases the analysis algorithm distributes the harmonic intensity over the same wavelength range as the spectrometer. This is despite the fact that the algorithm has a completely unweighted choice between all odd harmonics from the 1^{st} to the 51^{st}.

One effect that is visible on several of the plots concerns the intensity of the 17^{th} harmonic at $\sim$ 45 nm. The 10, 80 and 100 mbar data sets all feature considerably higher intensity from the diffraction analysis than is measured with the spectrometer. Unfortunately, the linearity of response from the spectrometer has never been accurately tested. It is possible, for example, that the grating efficiency rolls off towards the edge of the spectral range.

The errors in all of the plots seem to be of the order of $\pm$ one harmonic. Looking at the 40 mbar data for example it seems as though the spectral intensity from the central peak in the spectrum may have been decomposed into the two neighbouring harmonic coefficients. A similar effect is perhaps also visible in the 20 mbar data. It is not surprising that this type of error should occur, particularly in the light of the wavelength selectivity calculation that was shown earlier in figure 4.17. In future experiments it may well be possible to substantially improve on the wavelength selectivity and eliminate this type of error by careful optimisation of the aperture size and spacing.

Note that in this work the CCD image was integrated over 100 laser shots, however, single shot operation is potentially possible. Because of grating and MCP efficiencies and the use of a slit - the traces recorded with a grazing incidence spectrometer in figure 4.22 had to be integrated over 60000 laser shots.

In fact, the spectral reconstruction yields far more information than can be shown in a simple spectrum; the spectral data is spatially resolved with a resolution equal to the aperture spacing of the diffraction grid. Some examples of the retrieved data are shown in figures 4.23 to 4.28. In these examples the 4D spatio-spectral data is visualised using the matlab command `contourslice` (shown on the left). The CCD data from which the spatio-spectral data was recovered is shown in the column on the right.

Note that the analysis routine was run on the same data set that was plotted in figure 4.22. This is to permit the reader to cross reference between the two visualisations.

The `contourslice` plots on the left of figures 4.23 to 4.28 can be understood as follows: The spatial variation (in the i, j plane) of the spectral intensity of each harmonic is displayed in the form of a contour plot. (The contours are color coded - red represents the highest spectral intensity). The contour plots are then layered on top of one another so that the vertical axis represents harmonic order.

The CCD image is displayed both at the bottom layer of the `contourslice` plot and also to the right of the diffraction analysis output in the form of a false colour image plot. This is to help the reader to discern how features in the spectral analysis correspond to regions of the CCD image. The red and white triangles indicate the relative orientation of the two plots since the `contourslice` plots have been rotated in order to emphasise particular features of each spectral reconstruction. This visualisation of the data is particularly insightful when viewed live on the computer screen. The ability to rotate and view the data from any angle really helps to identify interesting features of the data.

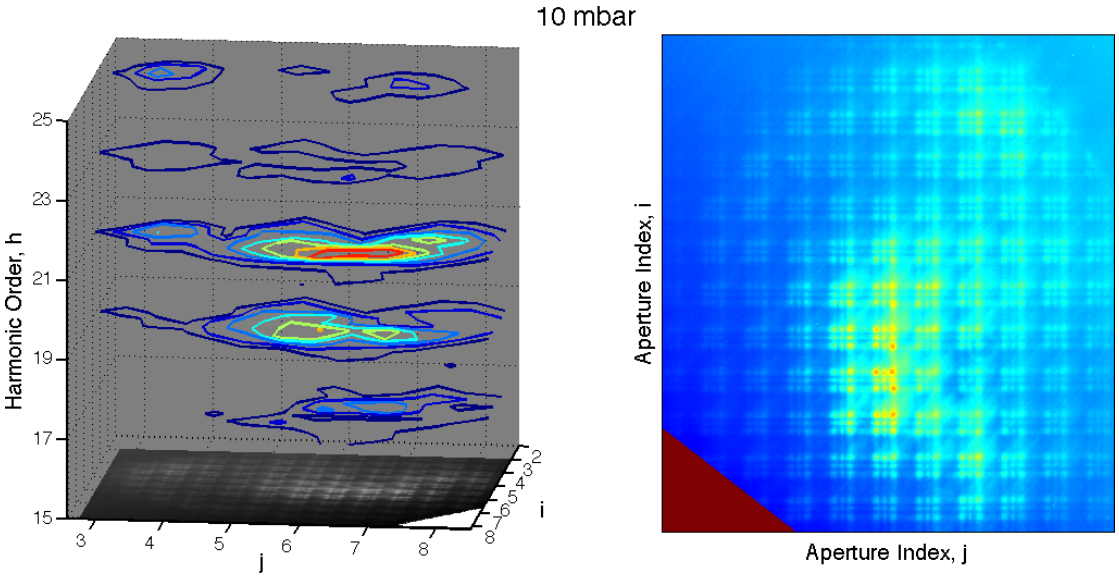

FIGURE 4.23: Left: Examples of the retrieved spatio-spectral data. Right: Rotated CCD image from which the data is recovered. HHG gas pressure set to 10 mbar

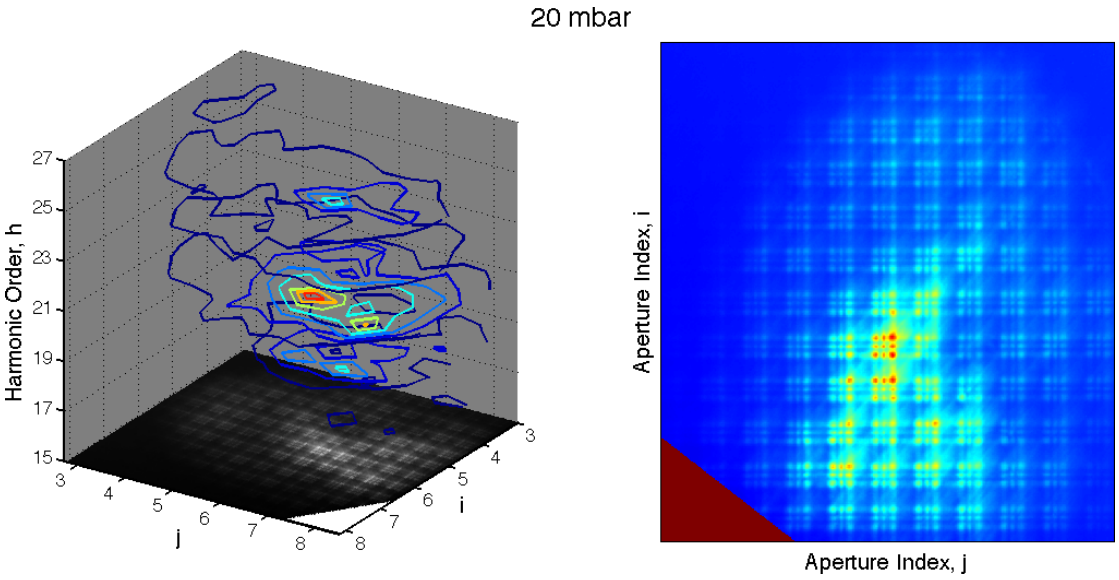

FIGURE 4.24: Left: Examples of the retrieved spatio-spectral data. Right: Rotated CCD image from which the data is recovered. HHG gas pressure set to 20 mbar

Taking the 10 mbar data as an example (figure 4.23) when viewed at an oblique angle from the side the `contourslice` plot shows that the x-ray beam profile has separated into three lobes. An isolated lobe on the left (top right of the CCD image) and two lobes that are merging together on the right (centre and lower right of the CCD image). The contour lines indicate that the isolated lobe consists primarily of the higher harmonics (25^{th} to 21^{st} in this case) whilst the two other lobes contain mostly the lower harmonics (21^{st} to 17^{th}). The 40 mbar data shows a similar separation of the x-ray beam into two distinct regions; in this case, however, the spectral content of the two regions remains much the same.

Notice that as the gas pressure increases so too does the intensity of the x-ray beam. In the 10 and 20 mbar data the intensity is distributed evenly over a rather large area with a slightly more intense core. The lower intensity x-rays are thought to be the result of non-phase matched generation at the exit of the capillary whilst the more intense core is the result of phase matching. As the pressure is increased - the phase matching conditions are met and the central part of the beam becomes so intense that the non-phase matched x-rays no longer appear so prominent.

Comparing the data for 60 and 80 mbar it clearly shows that at 60 mbar the HHG source is phase matching for the higher harmonics whilst at 80 mbar generation of the lower order harmonics is optimised. This is in agreement with observations made using the grazing incidence spectrometer.

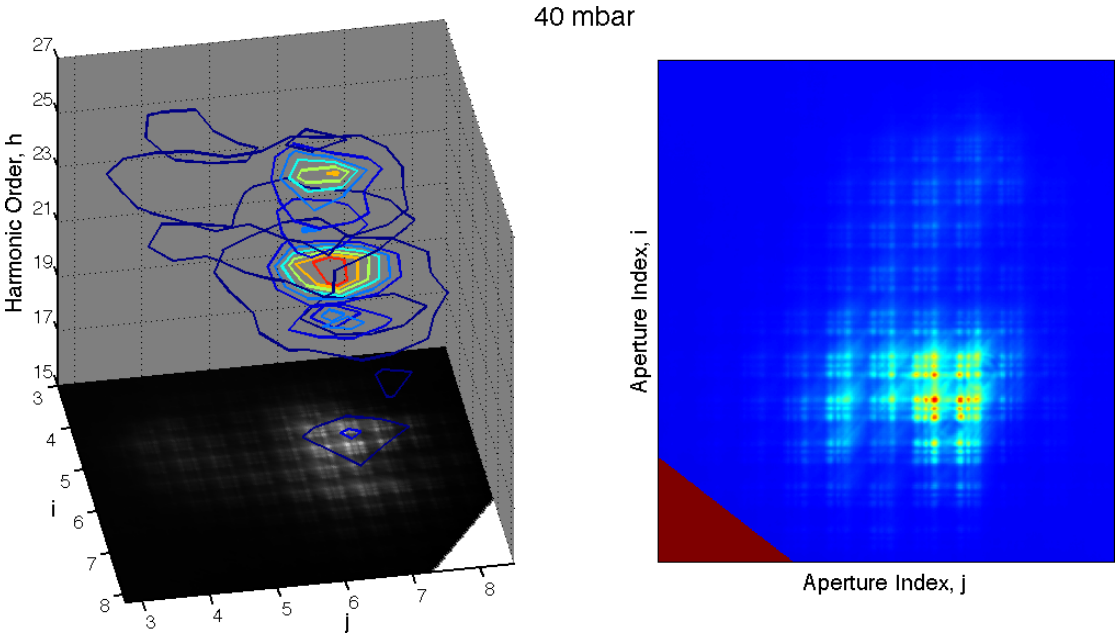

FIGURE 4.25: Left: Examples of the retrieved spatio-spectral data. Right: Rotated CCD image from which the data is recovered. HHG gas pressure set to 40 mbar

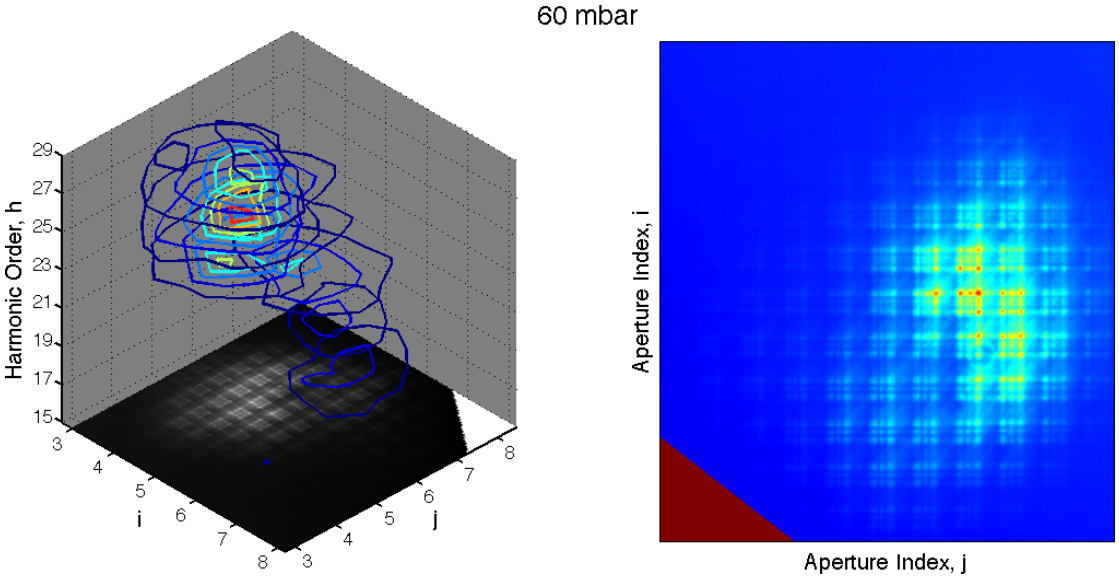

FIGURE 4.26: Left: Examples of the retrieved spatio-spectral data. Right: Rotated CCD image from which the data is recovered. HHG gas pressure set to 60 mbar

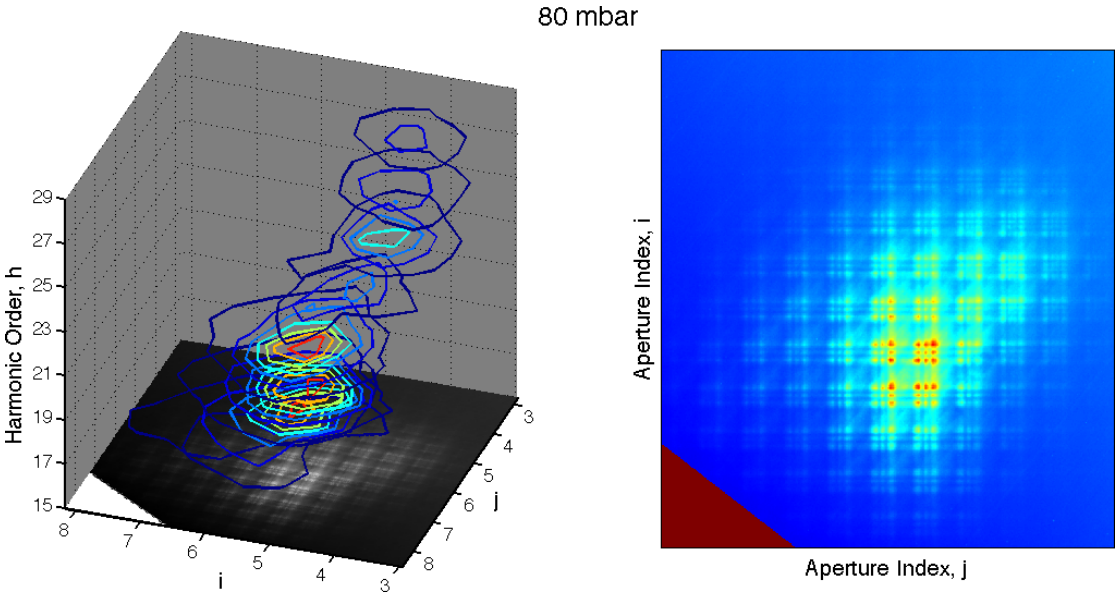

FIGURE 4.27: Left: Examples of the retrieved spatio-spectral data. Right: Rotated CCD image from which the data is recovered. HHG gas pressure set to 80 mbar

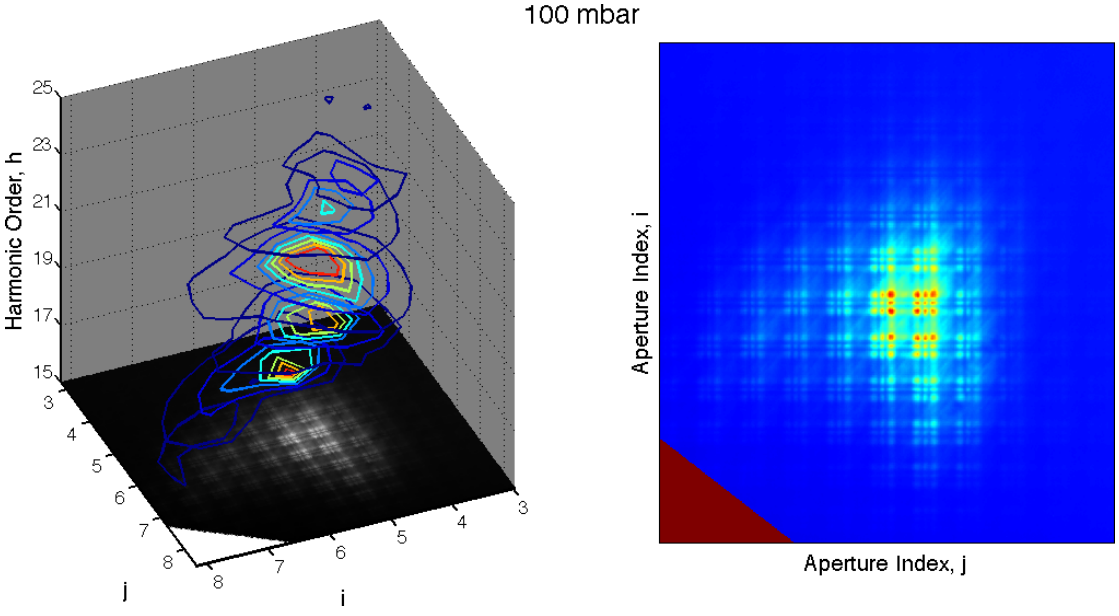

FIGURE 4.28: Left: Examples of the retrieved spatio-spectral data. Right: Rotated CCD image from which the data is recovered. HHG gas pressure set to 100 mbar

A common feature of many of the diffraction analysis data sets is that they display a spatial chirp on the x-ray beam. This can be seen clearly in figures 4.26 to 4.28 as the contour lines for each harmonic appear to have a transverse displacement from one another.

The spatial chirp evident in the spectral analysis of the x-ray beam is confirmed as a genuine physical effect by measurements using the grazing incidence spectrometer in imaging mode. In this mode the light arriving at different parts of the spectrometer input slit are imaged to different parts of the MCP (as usual). The camera is then simply read out in imaging mode rather than binning the data as it would when collecting a normal spectrum.

Figure 4.29 shows several spectrometer images (recorded at different pressures). On each image two separate areas are defined that correspond to the spectrum arriving at different heights on the slit. The spectra from the two regions can clearly be seen to be different - indicating spatial chirp.

A puzzling feature of this data is that the spatial chirp appears to be at least partially pressure dependent. It is possible that this spatial chirp is resulting from asymmetry in the generation process, for example a pressure gradient at the holes where gas is passed into the capillary.

A more likely cause of spatial chirp in the x-ray beam would be spatial chirp at the laser focus. (However this cannot explain the pressure dependence). These data sets were recorded before the laser system upgrade and spatial chirp was known to be present on that beam (this was confirmed by translating an optical fibre coupled to a spectrometer across the laser beam). The spatial chirp on the laser is thought to have been caused primarily by sub optimal compressor alignment. The laser focus was in fact found to be slightly elliptical in shape which was causing problems in attaining good capillary launch efficiencies. The possibility of improving the laser mode profile was one of the main motivations for the laser upgrade.

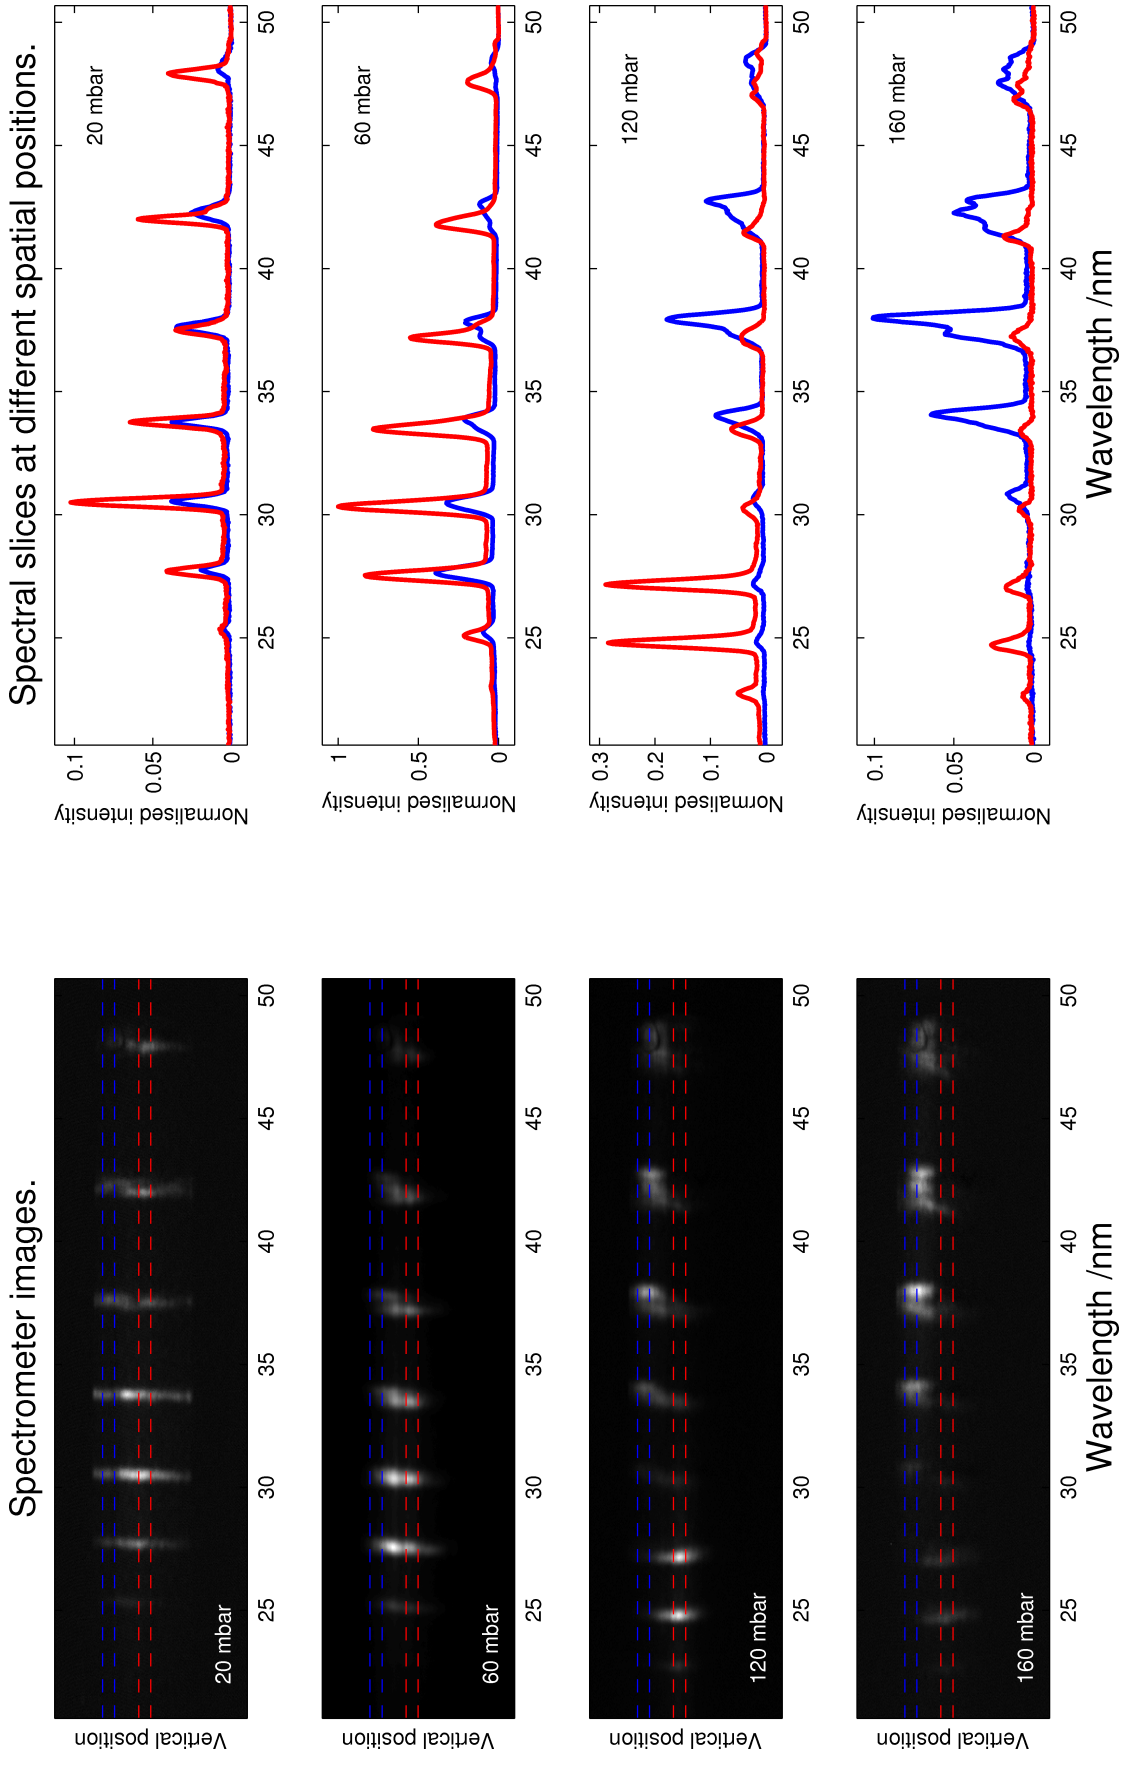

FIGURE 4.29: Measurements made using an imaging x-ray spectrometer. Different spectra are found to arrive at different parts of the slit.

An alternative way of viewing the diffraction analysis data is by converting it into a sequence of intensity maps each representing the spatial intensity profile of a different harmonic. Viewing the data in this way is more compatible with paper hard copy and also allows more precise analysis of subtle features.

Figure 4.30 shows the example data set reproduced in this sequence representation. The harmonic order is printed below each pane and the pressure is stated in the title above. Notice that the resolution is rather low because each pixel in the intensity map corresponds to an aperture in the CCD image.

The spectral intensity within each sequence uses the same normalisation - this allow the viewer to discern which harmonics are more intense. The same normalisation could not be applied over all pressures as this would make it impossible to see the lower intensity signal from the 10 and 20 mbar sequences.

The spatial chirp described above is now visible as a shifting of the centre of the intensity map from one pane to the next within each sequence.

Of particular interest are the spectral intensity maps of the 23^{rd} and 25^{th} harmonics at 10 mbar. These appear to show an approximately annular intensity profile which is not evident on the CCD images shown on the right of figures 4.23 and 4.24 and cannot be easily discerned (but is present) in the contour line plots. Further examples of this annular shape are shown in figure 4.31.

In figure 4.31 the two diffraction patterns at each pressure were recorded at different dazzler settings - giving the input laser a slightly different pulse duration in each case. This appears to have had very little discernible effect on the spectrum of the x-rays except that in the last plot the generation extends a little more intensity into the higher harmonics. The similarity between the recovered data for the two pressure pairs highlights the reproducibility and stability of the spectral analysis technique.

It is very interesting to note that at 40 mbar (which is below the optimum phase matching pressure) the annular patterns occur on the higher energy harmonics. At 180 mbar, which is well above the optimum, the annular patterns occur on the low energy harmonics (the optimum pressure for high energy x-ray generation is $\sim$ 70 mbar).

This represents very important information and to the best of our knowledge is the first experimental observation of an annular spectral intensity profile at the output of an HHG capillary. Information about the spatio-spectral profile of the x-ray beam is required to further the modeling, and our understanding of the HHG process that occurs inside the generation capillary. It could for example provide supporting experimental evidence for the theory that radial phase matching leads to generation of different harmonics in

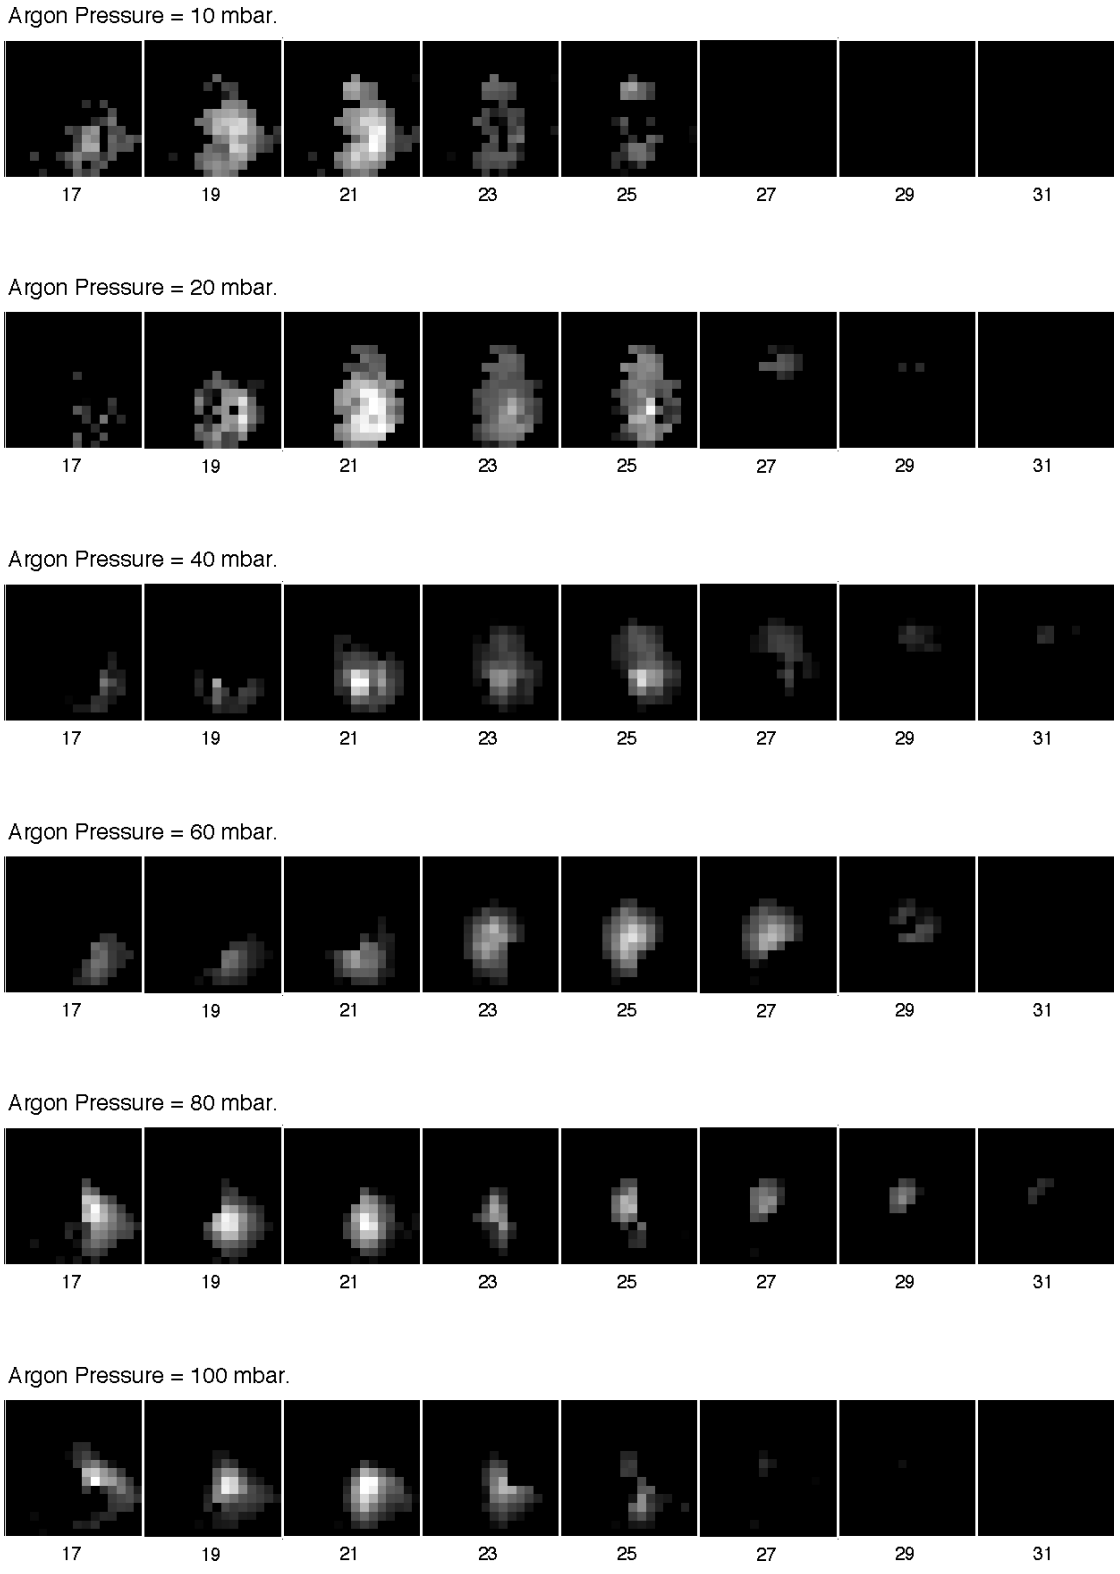

FIGURE 4.30: Spectral intensity map sequences recorded at a several different pressures. Upper row was extracted from a CCD image taken at 10 mbar whilst the image used for the lower row was recorded at 100 mbar.

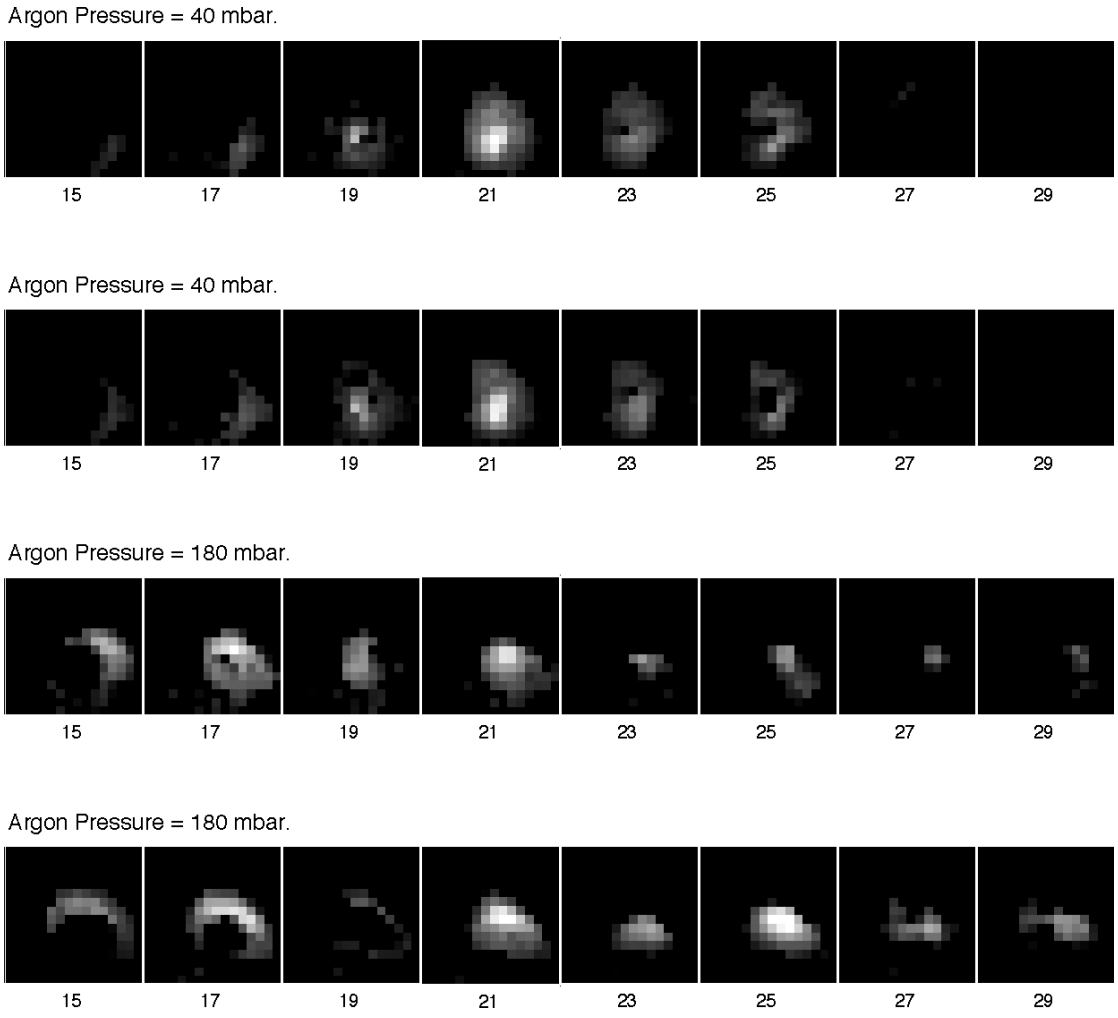

FIGURE 4.31: Spectral intensity map sequences extracted from four different CCD images. Each sequence contains panes that appear to have an annular intensity profile.

annular regions of different radii (at any particular time in the pulse). These annular spectral intensity profiles and the spatial chirp described above are precisely the type of spatially resolved spectral information which it was hoped this technique could resolve.

4.3.5 Maximum entropy

The least squares spectral reconstruction technique described in subsection 4.3.3 takes advantage of prior knowledge that the spectra to be recovered consists of a series of peaks at the odd harmonics of the input laser wavelength. This places a very strong constraint on the spectral reconstruction and dramatically improves the accuracy of the technique. An alternative to the least squares fitting algorithm was sought in order that the spectral reconstruction technique could be extended for use with sources other than HHG. Synchrotron and plasma sources for example exhibit smoothly varying, continuous

spectra so the constraints and approximations employed in the least squares fitting algorithm are not appropriate.

The maximum entropy method[79] is a data analysis technique that is most often used for image enhancement in astronomy. In this application the aim is to extract the most information possible from noisy images with a limited resolution without asserting anything that is not justified by the original data.

All measured data sets have a limited test resolution and contain errors. Consequently the data set simply defines a range in which the real values must reside. The maximum entropy method provides a way of selecting one of these possible solutions. Uniquely, the maximum entropy solution fits all of the known constraints of a given problem whilst remaining as uniform as possible. Any other possible solution could be argued to contain a bias which cannot be justified by the measured data.

In the statistical sense entropy is defined as follows: When a data point is known with absolute certainty its entropy is said to be zero. Conversely, a data point that provides little additional information has high entropy. For example, consider the problem of interpolating additional data values for a line. If there are just two existing data points then these points define a straight line. If the new data point is added anywhere other than on that line then it is making an unfounded assertion and changing the shape of the data. If the new point is placed on the line, the shape of the data is unchanged and the new data point adds no new information - entropy is said to have been maximised.

The maximum entropy method is applied to the problem of spectral reconstruction in a manner very similar to that taken with the least squares fitting routine in subsection 4.3.3. Just as before each aperture in the grid is considered separately. The approach is still to generate simulated diffraction patterns and find the coefficients with which they must be added to reproduce the measured data. The simulated diffraction patterns used in this method are not calculated for the harmonics of the laser but rather they are calculated at regular intervals in wavelength. As before (and for the same reasons as discussed in subsection 4.3.3) an incoherent sum is used to combine the simulated diffraction patterns, equation 4.36 therefore becomes:

$$P_m(x,y,i,j) = \sum_{\lambda=\lambda_{min}}^{\lambda_{max}} P_s(x,y,i,j,\lambda) I(i,j,\lambda) \tag{4.37}$$

where λ is the wavelength of the simulated diffraction data and λ_{min} and λ_{max} are the lower and upper limits of the simulated wavelength range respectively. Just as before P_m is the measured diffraction data, P_s is the simulated diffraction data and I is the

coefficient of spectral intensity with which the simulated diffraction patterns are to be added.

Now considering the diffraction pattern produced by just one aperture.

An estimate of the error in the measured data must be made. (An error value is assigned to each pixel in the measured data). This error defines the range of possible solutions to the fitting problem. If the error is set unrealistically low then it may not be possible to find a solution that fits the measured data at all - the algorithm will therefore fail to converge. If the errors are set too high then the problem is not sufficiently constrained and the algorithm will converge to a uniform value that does not contain any detailed information. Approximate shot noise (N_s) values for a CCD detector can be found using the following simple relation (thermal and readout noise are neglected):

$$N_s = \sqrt{S} \qquad \text{where } S \text{ is the measured signal in electron counts.}$$

The error values for the central part of the diffraction pattern were estimated using this relation and a small additional factor was added to account for systematic errors. In order to assist the maximum entropy routine the error values for the 'dark area' surrounding the diffraction pattern were set very high. This effectively tells the algorithm to ignore these data points. This was necessary and valid since the simulated data contains no intensity in this region - it would therefore be impossible for it to simulate the noise that was present in the measured data.

Processing of the measured data is kept to the absolute minimum possible however it is necessary to subtract the background offset (i.e. the minimum value in the measured data). This ensures that the data values begin at zero (and eliminates negative values if present) which helps the algorithm to converge. In addition an averaging filter (with a large window size) was applied to the data. This removes the intensity gradients that are caused by non-uniform illumination of the aperture but keeps intact the rest of the diffraction pattern.

Before attempting any fits to experimentally measured data the maximum entropy algorithm was first tested on simulated data. Test patterns were created simply by adding together the simulated diffraction patterns with selected coefficients. For simple, smoothly varying patterns the recovery algorithm worked well. Example tests are shown in figures 4.32 and 4.33:

As the complexity of the input data increases the maximum entropy algorithm starts to have difficulty accurately recovering the input spectral profile (see figures 4.34 and 4.35).

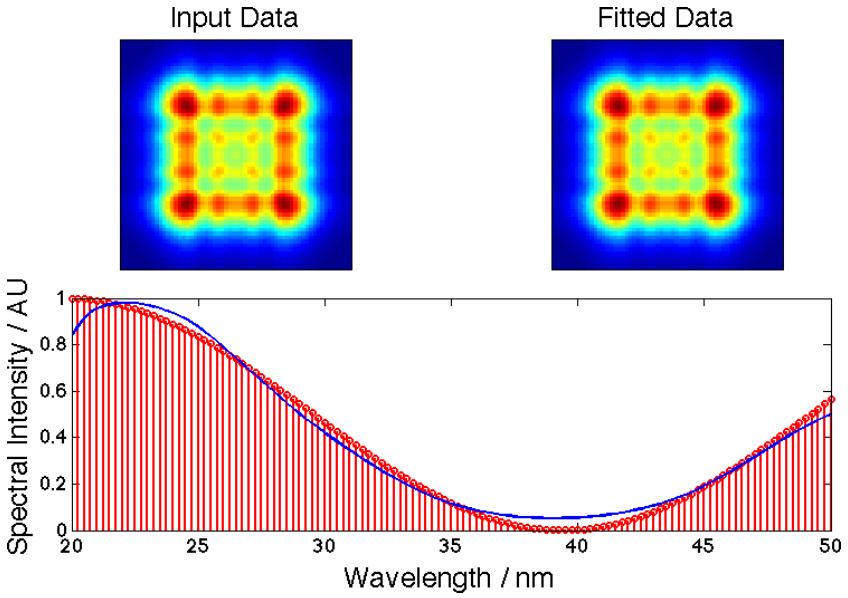

FIGURE 4.32: Test of the maximum entropy analysis algorithm. The red lines on the lower plot show the coefficient values that were used to create the input data diffraction pattern. The blue line on the lower plot shows the spectral coefficients that were recovered by the maximum entropy algorithm. Multiplying the simulated data by these coefficients and then summing produced the fitted data diffraction pattern. For this simple test pattern the fit of the recovered spectrum to the input spectrum is good.

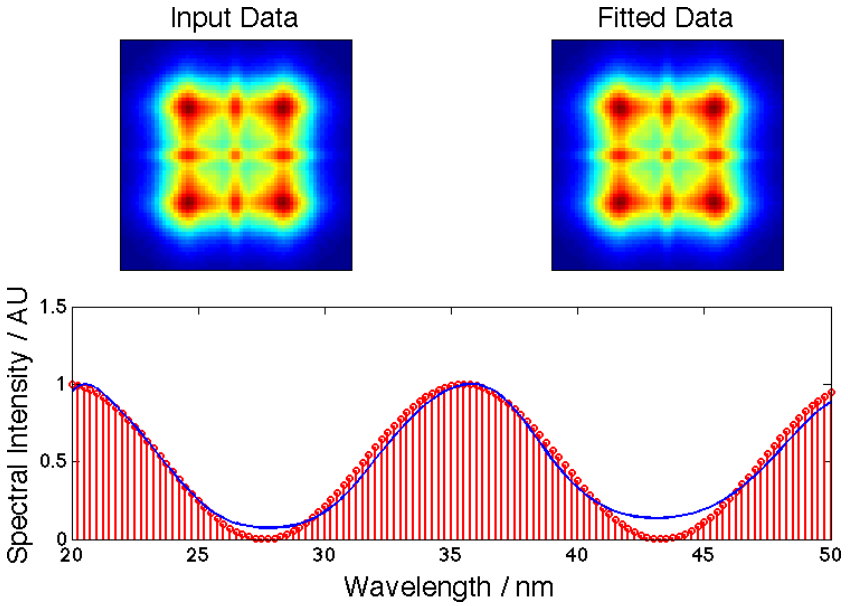

FIGURE 4.33: Layout of this figure is the same as that of 4.32. The agreement with the input data is also good for this simple test pattern.

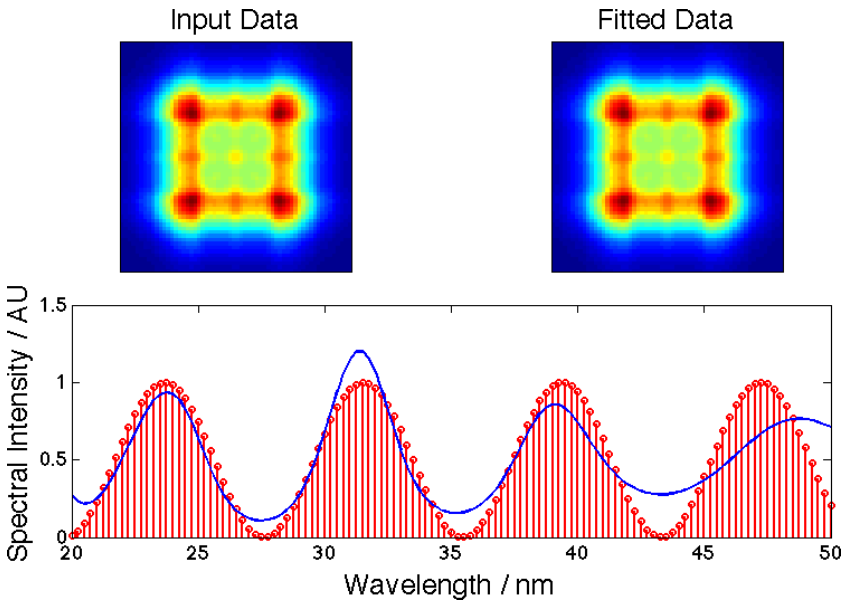

FIGURE 4.34: Fitting accuracy for a more complicated pattern is slightly reduced.

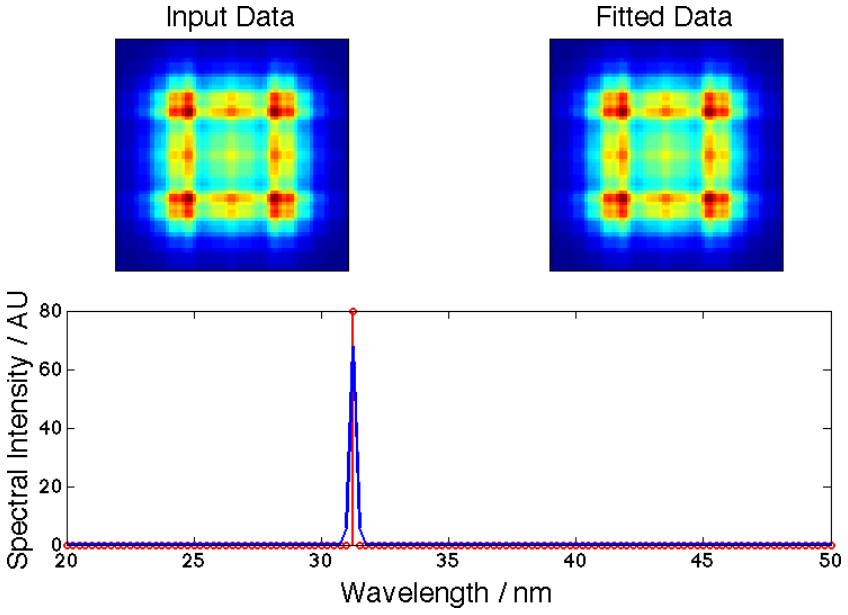

FIGURE 4.35: This figure shows the recovered spectrum that is obtained when a single spectral coefficient is input into the algorithm.

When the diffraction pattern produced by a set of isolated spectral coefficients (positioned at the harmonic wavelengths) is entered into the algorithm the individual peaks cannot be resolved. Instead a smooth distribution in the approximate region of the coefficients is obtained (see figure 4.36):

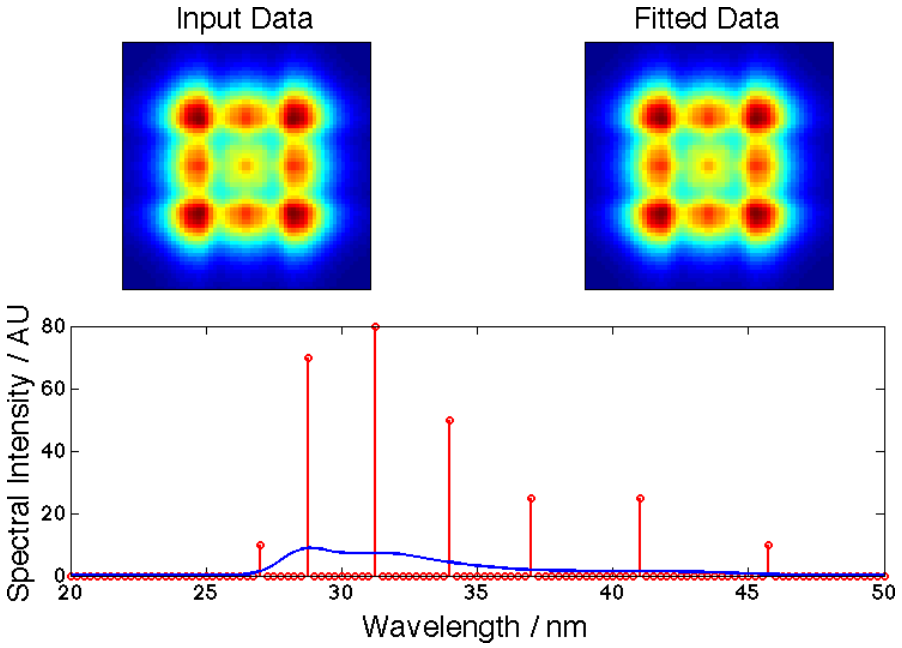

FIGURE 4.36: When attempting to fit to a simulated harmonic spectrum composed of narrow spectral peaks the maximum entropy algorithm interprets the peaks as noise. Rather than reconstructing the harmonic peak structure the algorithm returns the spectral envelope of the x-ray emission.

Despite this result we attempted the fitting routine with some of our experimentally measured diffraction patterns (still just looking at a single aperture). The fitting results are shown below in figure 4.37. The input and fitted data are as described above - the error data is simply calculated by subtracting the fitted data from the input data:

Although the maximum entropy method was able to achieve very close matches between the input and fitted data - the spectrum did not appear to be accurately reconstructed. Typically the spectrum had a smooth intensity profile and therefore contained no information about the harmonics. After some careful optimisation of the error estimate, figure 4.38 was produced. This shows that the maximum entropy method was very close to being able to resolve the individual harmonic peaks.

In conclusion the maximum entropy method proved not to be suitable for use with the HHG source because of its inability to retrieve sharp features in the fitted spectrum. This outcome was anticipated from the outset, however, the intention was to test the maximum entropy method to see if it could extend the spatio-spectral retrieval technique

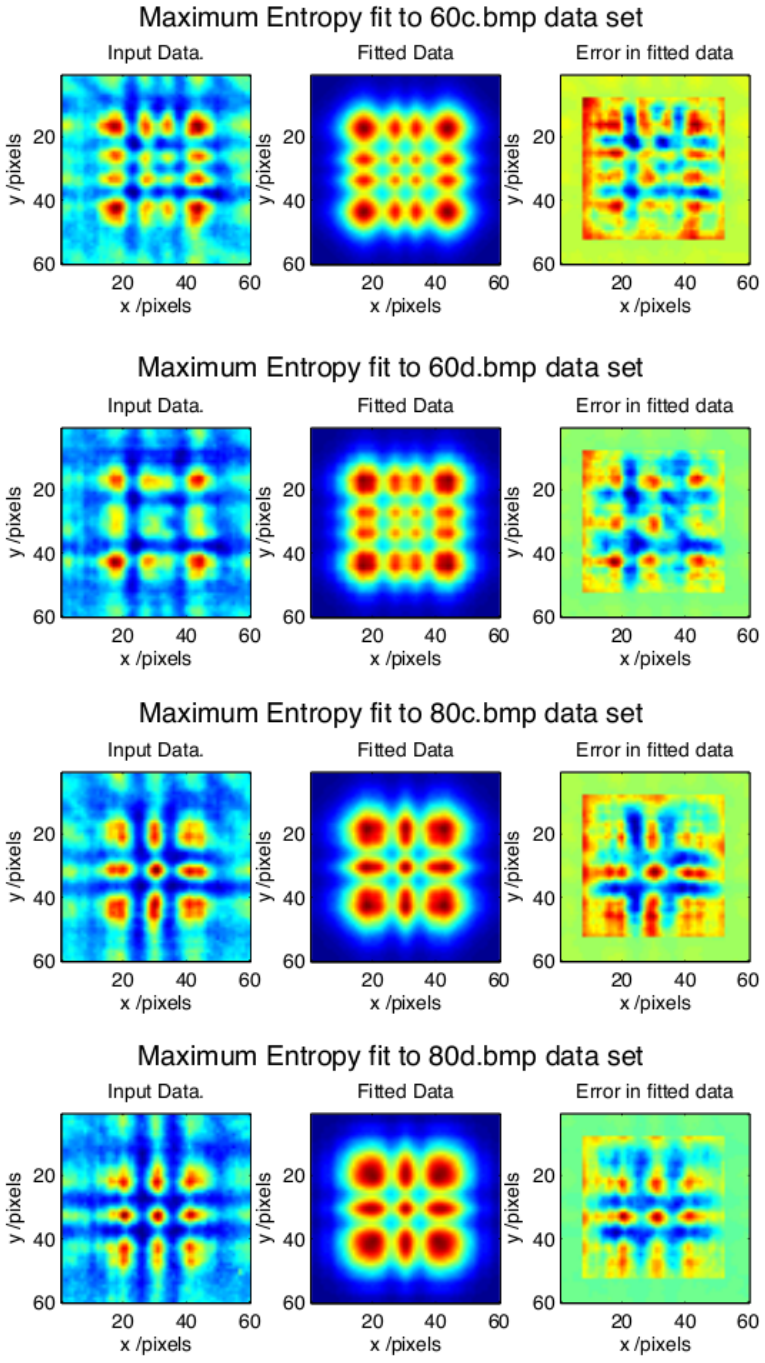

FIGURE 4.37: Here the maximum entropy method is applied to some measured diffraction patterns. The agreement between measured and reconstructed diffraction patterns is qualitatively rather good. There is some residual error although this is slightly exaggerated by the colormap.

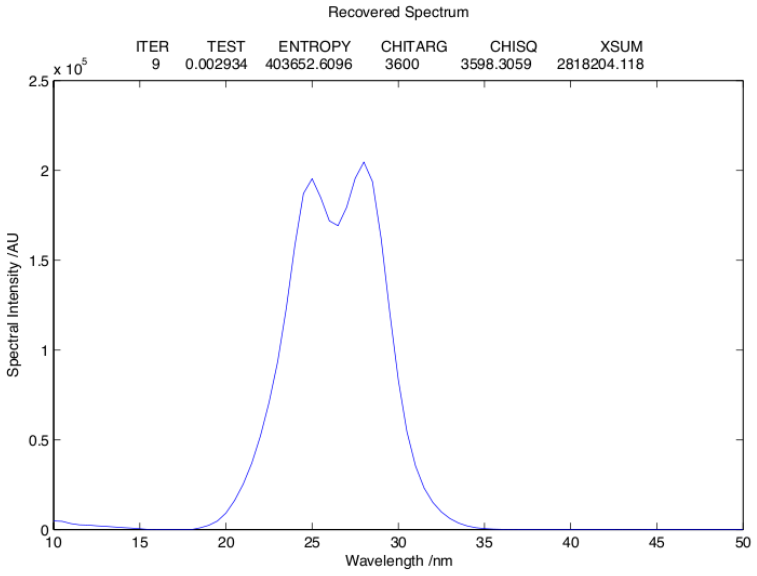

FIGURE 4.38: The dip in the peak of this fitted curve suggests that the maximum entropy method is on the verge of being able to resolve individual harmonic peaks from the measured diffraction data.

for use with other sources. The results from tests using smooth spectral profiles (4.32, 4.33 and 4.34) were very promising. Further testing would need to be performed but it seems feasible that this method could be used to analyse the diffraction patterns produced by a synchrotron source.

Chapter 5

Pulse Compression

Pulsed laser systems are continually being developed in order to produce ever shorter pulse durations. The birth of attosecond science (a sibling to HHG) now takes this development into a new regime. There are several aspirations that provide the drive for this development however two in particular stand out as key:

temporal resolution *and* peak power.

There is considerable interest in attaining the highest possible temporal resolution as this allows processes to be observed and analysed at a level that has never before been possible. As a consequence there is much new physics to be uncovered each time an advance in temporal resolution unlocks a new region of study. For example the advent of attosecond pulse systems in combination with carrier-envelope phase stabilisation techniques enables fields such as attosecond dynamics to arise.

The second motive for short pulse laser systems is that they make incredibly high electric field strengths available for use in the laboratory. The interaction between these intense fields and matter (nonlinear optics) often produces unanticipated but fascinating results. New physics - and new applications of that physics are constantly being developed in this field.

Whilst attosecond and HHG have provided routes to shorter pulses and higher photon energies there is a very considerable penalty to pay in average power. The obstacle that prevents development of directly generated pulses shorter than $\sim$ 5 fs is fundamentally the gain bandwidth of Ti:sapphire[80]. Techniques will no doubt be developed to overcome this barrier (meta-materials may for example offer increased gain bandwidths) however, pulse compression techniques remain an attractive route to extremely short pulses, particularly when considered as an upgrade to existing laser systems.

5.1 Capillary Compressor and Filamentation

Temporal compression of an already transform-limited pulse requires two steps; this concept is common to all pulse compression schemes. Firstly, additional spectral content must be added to the pulse. This is typically achieved through self phase modulation. The second step is to de-chirp the pulse in order to obtain a temporally short pulse. Standard approaches used for this step are gratings, chirped mirrors, or prisms. Clearly the spectral content must be added to the pulse with a well defined spectral phase in order that it can be effectively removed on recompression.

Self phase modulation is a $\chi^{(3)}$ process so a gas phase medium is normally used for the spectral broadening step (typically a noble gas such as argon). Several different geometries are used for this step: photonic crystal fibres[81], guided capillary modes[82, 83] or self guiding filamentation[84].

In the filamentation scheme there are two competing effects - plasma de-focusing and self focusing due to the Kerr effect. If the launch, pressure and power parameters are set to approximately the correct values a filament is automatically formed - the laser effectively burns its own waveguide by ionising the gas.

Currently filamentation appears to offer advantages in terms of throughput and ease of alignment and also offers the possibility of self compression simply by truncating the filament after the right length - eliminating the requirement for external recompression optics[84].

A possible drawback of filamentation in comparison to guided capillary geometries is that pressure fluctuations within the gas volume can easily lead to beam pointing instability. This could cause difficulties in integration with optical devices further along the beam path.

5.2 Cascaded $\chi^{(2)}$ compression

This section describes the application of cascaded $\chi^{(2)}$ pulse compression to a laser centred at 790 nm. Previous theoretical work and more recent models do not agree about the viability of such a proposition. Chiefly this experimental was work carried out to provide feedback for further modeling. This work follows that of F. W. Wise *et al.* and was guided by the theory calculations of N. G. R. Broderick.

As mentioned in section 5.1 the spectral broadening step in pulse compression schemes is usually performed by self phase modulation in a $\chi^{(3)}$ nonlinear material. In 1999

Wise *et al.*[85] showed that cascaded $\chi^{(2)}$ nonlinearity[86] could also be used to produce an effective $\chi^{(3)}$ nonlinearity and hence could be used for pulse compression. In later work, a negatively chirped input pulse was used to achieve modest compression at 800 nm[87]. (The normal, linear dispersion of the material was used to de-chirp the pulse from its initially negatively chirped state). All of these examples utilized some form of external dispersive optics to correct the frequency shifts produced by SPM and the material dispersion.

A slightly different approach is taken in ref[88] (carried out at 1260 nm) and in the work presented here at 790 nm. Cascaded $\chi^{(2)}$ nonlinearity is used to produce an effective $\chi^{(3)}$, thus facilitating SPM. However, the magnitude and sign of the SPM-induced frequency shift are controlled by means of n_2 (the effective nonlinear refractive index) which can be directly manipulated by setting the crystal angle. Essentially this provides a direct method of controlling the spectral phase with which the additional spectral components are added to the pulse.

With a slight negative detune from phase matching enough negative frequency shift is produced to cancel out the effect of material dispersion; achieving compression of the fundamental field (FF) in a single, bulk-material, step. This method therefore has tremendous potential as it would enable the construction of highly compact compression systems, eliminating the need for long gas filled tubes.

5.3 Experiment

A telescope is used to bring a beam of 25 μJ pulses, with approximately 90 fs duration, to a weak focus (spot size $\sim$ 250 μm) at the centre of a 2.5 cm long beta-barium-borate crystal (BBO). (Peak intensity is estimated to be approximately 250 GW/cm^2). The focusing conditions used here are of critical importance to the experiment. Theoretical calculations suggest that a level of $\sim$ 500 GW/cm^2 must be reached in order to produce the desired nonlinear effect, however, great care must be taken to avoid exceeding the damage threshold of the BBO crystal ($\sim$ 800 GW/cm^2). Figure 5.1 shows some example output of a Gaussian beam simulation that was developed in order to determine the optimum focusing conditions.

The focusing conditions must also be optimised to provide a long Rayleigh range. This ensures that, when the BBO crystal is positioned at focus, the illumination of the crystal is approximately plane wave. The aim being for the experimental conditions to match those of the theoretical calculation as closely as possible. The necessity for a long Rayleigh range means that rather long focal lengths had to be used - making the optical

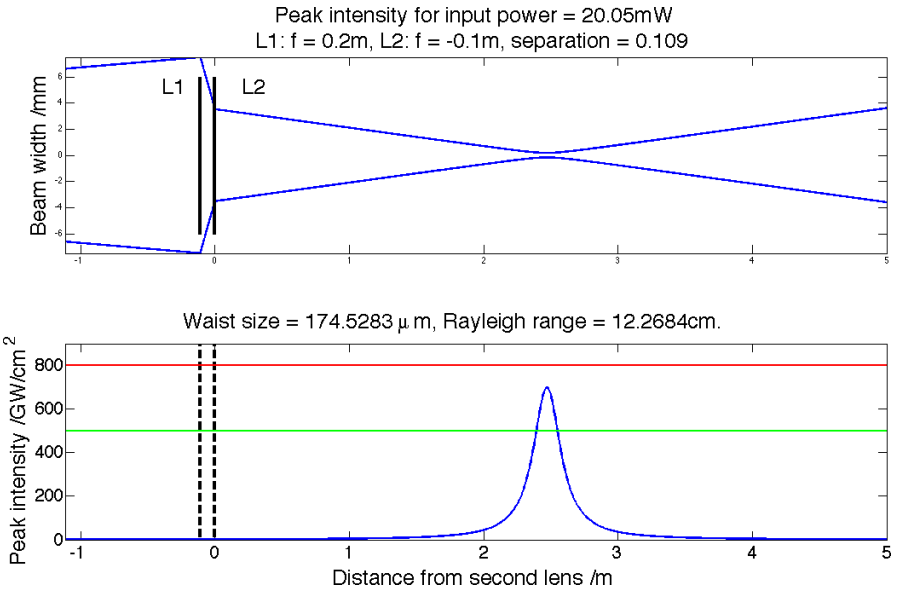

FIGURE 5.1: Results of a Gaussian beam focusing simulation. This was used to find suitable focusing conditions.

path for this experiment several metres long. It is anticipated that in the final application this constraint will be relaxed.

The type I BBO crystal (cut at 22°) is mounted on a rotation stage allowing offset from the phase matched peak to be measured and set accurately. At the beginning of each experiment the angle tuning curve was recorded in order to find the offset for the angular calibration (The crystal was kept in a desiccator between experiments to prevent it absorbing moisture from the air). A typical example of this phase matching curve is shown in figure 5.2.

The input laser spectrum is also show in figure 5.2. Since the laser spectrum is broad it has a significant affect on the measured phase matching curve. In order to verify that this 'broadband phase matching' was properly taken into account in the theory, the spectrum of second harmonic light was recorded as a function of crystal angle. Figure 5.3 confirmed that the phase matching calculation used in the modelling was in fact correct.

FROG measurements of the pulse duration are to be taken before and after the BBO crystal. However, the second harmonic light cannot be allowed to enter the FROG because this would disturb the measurement. At the output of the BBO crystal the FF and SH beams are spatially offset (the amount of offset was observed to vary significantly with input laser power) this means that a carefully positioned iris can be used to block most of the SH beam. A long pass filter could also be used here if required.

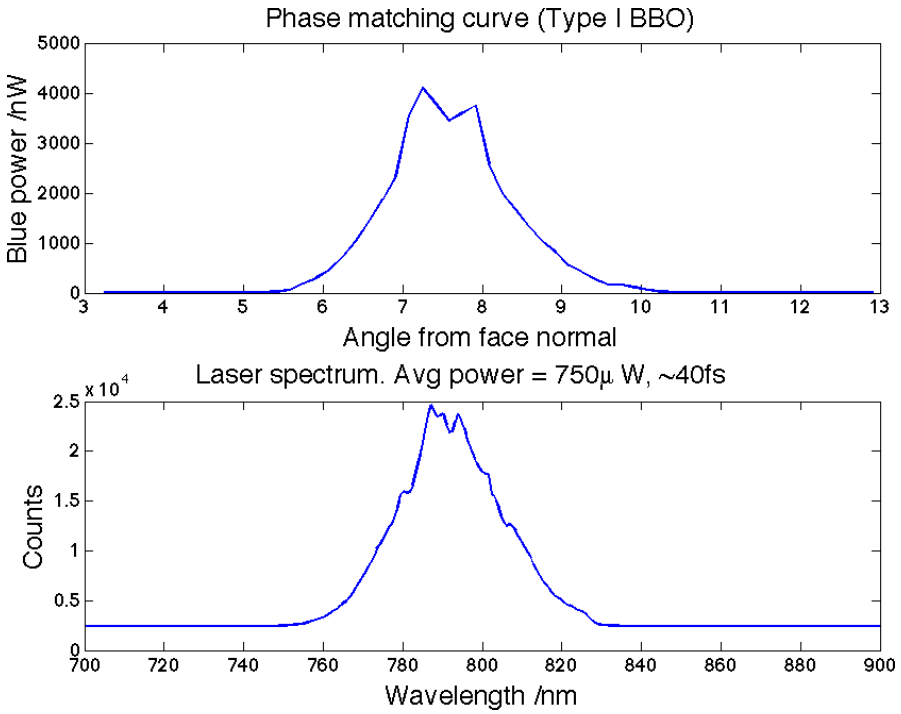

FIGURE 5.2: Upper plot: Second harmonic signal as a function of BBO crystal angle. Lower plot: Laser spectrum.

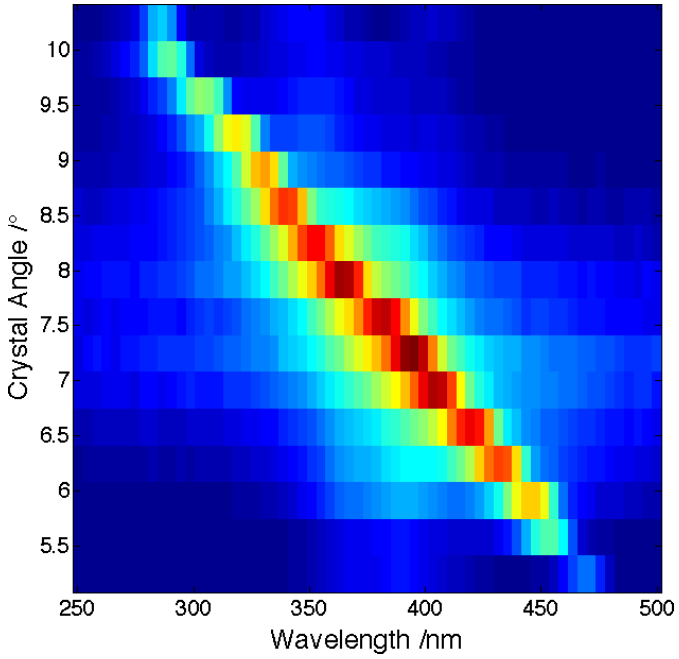

FIGURE 5.3: Image plot showing the spectrum of the second harmonic at the crystal exit as a function of crystal tuning angle.

A second telescope is then used to collimate the output FF before it is sent into the FROG for temporal and spectral analysis. The experiment was configured to allow FROG traces of FF to be recorded either with or without the BBO crystal in place; we also recorded FROG traces of the direct laser beam (bypassing the telescopes). A spectrometer was used to verify the accuracy of the FROG spectral reconstruction and optionally to provide higher spectral resolution measurements when needed.

5.4 Results and analysis

Figure 5.4 shows the power dependence of the output pulse duration (intensity FWHM) for phase matching detune of ± 1.5°. The process observed is clearly nonlinear showing a minimum pulse duration of approximately 50 fs for negative detuning with ∼ 12 mW of input laser power.

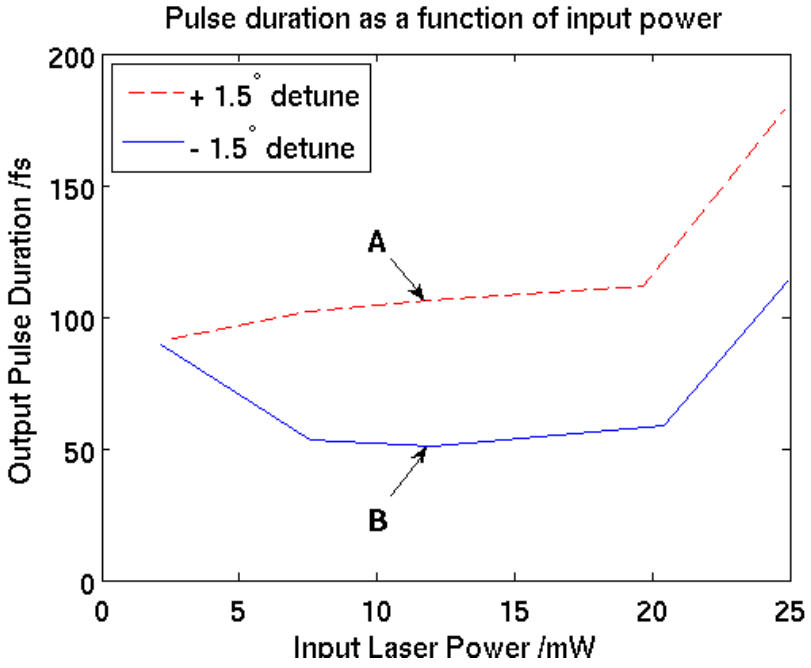

FIGURE 5.4: Output pulse duration (intensity FWHM) is plotted as a function of input laser power both for positive and negative detuning of 1.5° from phase matching.

The FROG traces that correspond to the conditions indicated by the arrows in figure 5.4 are shown in figure 5.5. Image A (on the left of figure 5.5) contains the FROG trace recorded for positive detuning at ∼ 12 mW; this is slightly broader in time than the trace recorded when the BBO crystal was removed and exhibits very little spectral broadening. Image B (on the right of figure 5.5) shows considerable spectral broadening (relative to the measurements taken with the BBO crystal removed) and evidence of soliton shedding.

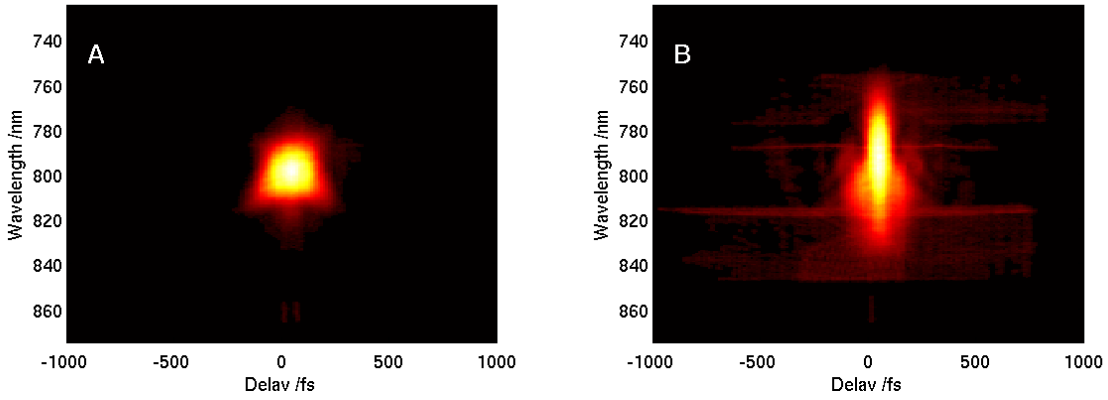

FIGURE 5.5: Images A and B show the measured FROG trace recorded at $\sim$ 12 mW for $\pm$ 1.5° detuning from the phase matched angle.

Alternatively the input power can be kept fixed and the crystal tuning angle scanned (recording the FROG trace at each point). The results of just such a scan are shown in the sequence of images in figure 5.6:

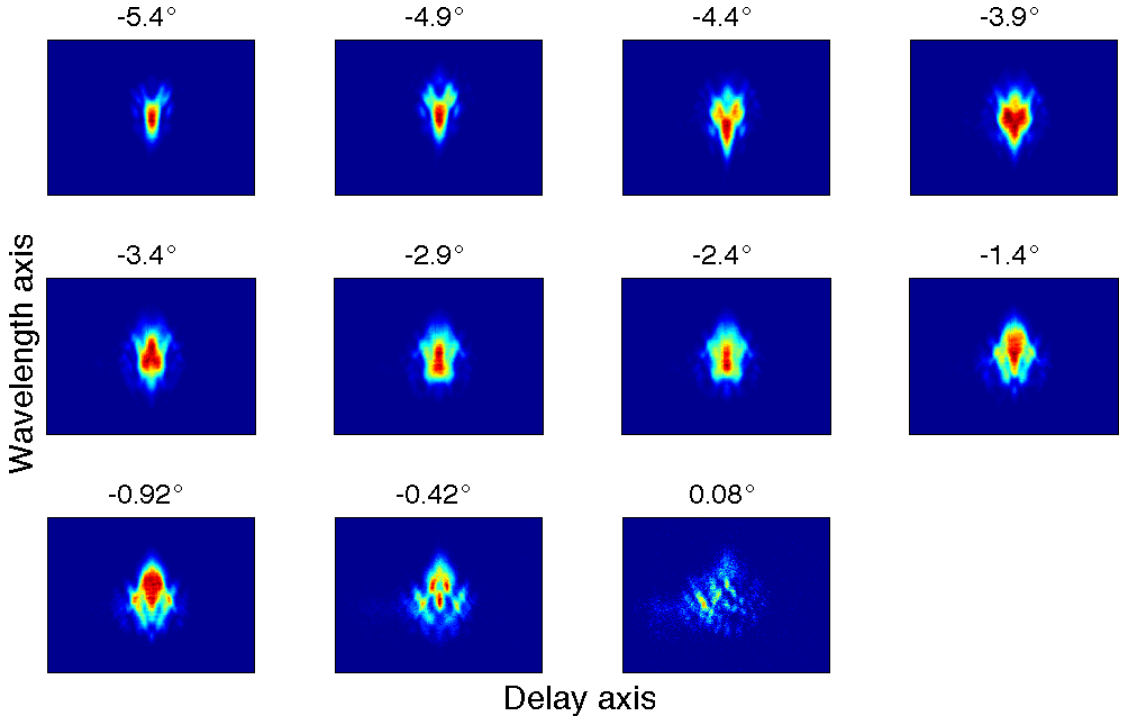

FIGURE 5.6: Sequence of FROG traces showing the evolution of the pulse shape with crystal tuning angle. Crystal angle (relative to face normal) is marked above each trace. The horizontal axis for each trace is the delay axis. The vertical axis of each trace represents wavelength.

As can be seen in figure 5.6 the pulse duration progressively decreases with negative offset from the phase matching angle. Data could not be recorded for the positive

detuning angles because the beam became highly divergent and was accompanied by a lot of second harmonic generation.

5.5 Conclusions

The modeling of these experiments is still at an early stage. Although there is good qualitative agreement between the theory and the experiments; there is more work to be done before quantitative agreement is obtained.

The model used to describe this effect takes the slowly varying envelope approximation developed by Wise *et al.*[89] but then adds new terms to represent $\chi^{(2)}$ self-steepening as well as higher order dispersion and spatio-temporal coupling. This approach is valid for pulses down to a few optical cycles and so can be used for the pulse widths here.

As the laser pulse frequency increases, the self-steepening terms become an increasingly important part of the model. In the simulations used here the self-steepening terms have become equally important as the nonlinear terms. Models that exclude these terms cannot therefore give a true description of the effect.

Soliton effect compression can be observed when the sign of the nonlinearity is opposite to that of the material dispersion. This can always be achieved since the sign of the nonlinearity depends on the detuning from phase-matching which can be positive or negative.

Wise *et al.* showed soliton compression in BBO for pulse wavelengths of approximately 1 micron. In this set of experiments we extend that earlier work to shorter wavelengths using pulses centred at 780 nm. This set of experiments, combined with those of Wise, clearly demonstrate the versatility of this single-step technique and extends the wavelength range over which soliton effect compression can be seen in BBO.

Chapter 6

Conclusions

During the course of this study an apparatus for capillary HHG has been designed, constructed and tested. In fact there has been a continual program of improvement in the HHG rig throughout the project. Improvements such as increased stability and better control over capillary alignment have had knock-on effects in terms of generation efficiency and beam stability.

Although optimisation of the source has represented a considerable part of the day to day work, the focus here should be on the science that has been undertaken. In the following paragraphs the key scientific findings of this work will be re-iterated and conclusions will be drawn.

The novel M^2 meter:

Optimising the laser beam quality is of critical importance to the HHG process. This device was conceived and constructed completely from scratch and a user friendly software tool was written to accompany the device. Despite their vital importance, beam quality measurements are often neglected because they can be so time consuming. This tool saves a lot of time and makes taking an accurate beam measurement less of an ordeal. Consequently this device is of real merit.

Since writing section 2.2.4 the M^2 meter has remained in frequent use. The software has been upgraded to permit an arbitrary number of beam samples to be collected. Typically it is now possible to fit 7–9 samples on the current sensor which improves the accuracy of the M^2 measurement.

A possible area for improvement in this design is the reduction of unwanted scattered light. The M^2 meter is currently operated almost exclusively in the 'Gaussian fit mode'

because the 'second moment mode' is very sensitive to any such stray light. Replacing the two separate reflective optics with an etalon would eliminate two interfaces that are not directly involved in the measurement - reducing unwanted multiple reflections. This would have the drawback of reduced flexibility (it would not be possible to change the reflector separation) but could offer subsidiary benefits in stability, ease of use and compactness.

If the scattered light could be reduced sufficiently to allow reliable operation in the 'second moment mode' this would improve the accuracy of the M^2 measurement, the generality to different beam shapes and could reduce computation time allowing faster frame rates.

HHG output, variation with experimental parameters:

This study began by measuring the total x-ray output as a function of gas pressure (see section 3.3. Initially this served simply as verification that x-rays were indeed being produced and that they could be distinguished from the laser output with the aid of aluminium foil filters. Subsequently the observed peak in the pressure profile proved that we could achieve pressure tuned phase matching in our HHG capillary. Finally, variation of the pressure profile with capillary launch mode shows that direct modification of the waveguiding term in the phase matching equation is possible.

Once the x-ray spectrometer had been installed our study of pressure tuned phase matching could be extended with the inclusion of spectral resolution. Again, it was found that the pressure at which the peak of phase matching occurred could be varied by careful adjustment of the launch conditions. In fact this, combined with adjustment of the pressure, provided a degree of control over the relative spectral intensities of the harmonics. Our measurements of the pressure tuned phase matching profile of argon and of other gasses are in qualitative agreement with the trends predicted by 1D phase matching models[62] as discussed in section 3.1.4.

A further degree of control over the x-ray spectrum was obtained when the Dazzler (see section 2.1.4 was utilised. It was initially observed that applying a change in quadratic spectral phase resulted in a shift in the centre wavelength of each harmonic. A systematic study of the x-ray spectrum as a function of quadratic and cubic spectral phase was then undertaken. Prior to these results it had been postulated that the observed spectral shifting resulted from a change in instantaneous frequency of the laser at the time of generation. It was proposed that by changing the spectral phase the instantaneous frequency as a function of time (relative to the centre of the pulse) could be directly modified. This was shown not to be the mechanism responsible for the spectral shifting

since the sign of the spectral shift was reversed about the zero of quadratic spectral phase. It was concluded that the spectral shifting of the harmonics was due to ionisation-induced blue shifting of the laser[2]. This is a nonlinear phenomenon and is therefore strongly dependent upon peak laser power. The peak laser power was of course changed by altering the spectral phase. This was verified by an experiment in which the quadratic spectral phase was adjusted whilst the peak laser intensity was kept constant.

The spectral phase experiments made us question the validity of the 1D theoretical models. In particular we were intrigued as to the effect of phase matching at different radii within the capillary. This led to the development of a more advanced model and to closer quantitative agreement between experiment and theory[90]. This in turn led to closer analysis of the spatial profile of the x-ray output and to the development of the diffraction analysis technique.

The Fresnel diffraction analysis technique:

This section represents a large body of work. The outcome of this work was very satisfactory with a much better than expected accuracy of fitting. I am certain that this technique could be applied in other areas.

The spatial resolution of this technique depends critically upon the fact that the diffraction pattern is measured in the Fresnel regime.

The technique also relies on the knowledge that the expected spectrum is composed of odd harmonics of a known laser frequency, and that the diffraction patterns produced by each harmonic are sufficiently different that they can be distinguished.

These conditions are well met for the aperture size, source-to-aperture separation, aperture-to-camera separation and camera resolution used in this work. If time allowed it would be an interesting challenge to find the theoretical optimisation for these parameters.

The results of the diffraction analysis demonstrate the need for more complex, spatially resolved models for the HHG process. A number of interesting features were observed such as spatial chirp in the x-ray beam and generation of different harmonics in spatially separated spots. These observations were confirmed with the aid of the x-ray spectrometer (used in imaging mode). Perhaps the most interesting observation was that of generation of particular harmonics in ring-like zones. To our knowledge this was the first observation of its kind and is just the type of effect that would be predicted by full 3D HHG models.

An attempt was made to generalise the diffraction analysis code to spectra not consisting of harmonics. To this end, the maximum entropy method was employed. This code

proved successful in reconstructing smoothly-varying spectra but could not accurately recognise the harmonic peaks produced by HHG.

The problem of developing a spectral reconstruction algorithm that works accurately for any type of spectrum represents an interesting challenge. If I were to revise this code I would perhaps begin by fitting to the entire diffraction pattern simultaneously rather than studying each aperture in isolation. A genetic algorithm could be used to find the spectral and spatial intensity profiles illuminating the diffraction grid. This would make the calculation far more time-consuming but could provide improved resolution since the x-ray spectrum would be free to vary across each aperture. This might also increase the accuracy of the reconstruction since the maximum amount of data is being extracted from the measured diffraction pattern.

Pulse compression:

This work also represents a considerable body of work. The outcome has been the development of an accurate model for this nonlinear process. Despite ultimately not being the optimum choice for use on our laser system (very large crystals would be required to avoid material damage) the work carried out has had direct implications on the design of the compression system used on the fibre laser system.

Considering the future of this project the most important decisions to be faced will be those relating to the following questions: What are the best paths to pursue in order to increase total flux and pulse energy?

6.1 Paths to higher energy x-rays

One of the ultimate goals of this project is to achieve water window x-ray generation. This would enable imaging of biological samples in an aqueous environment - allowing full advantage to be taken of the high-contrast between water and carbon in this region of the spectrum. Despite having this aspiration, to date, work in this group has centred on optimising and utilising the x-rays obtained from capillary phase matched generation in argon ($\sim$ 25–50 eV). A number of routes that could facilitate generation of higher energy x-rays are discussed below:

Gas species:

HHG has not been confined entirely to argon during the course of this project [90]. In fact very weak generation was also observed with neon (extending to slightly higher

energies than could be achieved in argon ($\sim$ 56 eV). However, argon proved to offer the greatest x-ray output and was therefore typically chosen.

As illustrated by equation 3.1 the ionisation potential of the gas species strongly affects the HHG cutoff. A higher ionisation potential should in theory mean a higher cutoff energy[62]. This effect may be partially responsible for the observation of higher harmonics with neon ($I_p = 21.6$ eV) than for argon ($I_p = 15.7$ eV). Certainly the observed increase in harmonic energy by approximately 6 eV seems to be in agreement with the theory.

Unfortunately, using equation 3.1 to calculate the cutoff we find that this should occur at $\sim$ 56 eV for argon and at $\sim$ 62 eV for neon. This calculation assumes a pulse duration of 35 fs, centre wavelength of 780 nm and an average laser power of $\sim$ 300 mW at the exit of a 150 μm diameter capillary. The laser is assumed to propagate in the EH_{11} mode ($1/e^2$ radius of 51 μm) giving a peak electric field amplitude (E) of $\sim$ 4.2×10^{10} V m^{-1} and an intensity (I) of $\sim$ 2.3×10^{14} W cm^{-2}.

The discrepancy between the theoretical and observed cutoffs indicates that another factor is limiting x-ray generation under the conditions used here. As described in section 3.1.4 that limiting factor is phase matching. Changing gas species also modifies the phase matching conditions considerably. Different gasses have different refractive indices, dispersion and a different change in refractive index upon ionisation. Since generation is below cutoff in both argon and neon and since there is only a modest increase in x-ray energy when generating with neon, it is likely that phase matching remains the limiting factor.

The x-ray output intensity is strongly dependent on the amount of ionisation - tunnel ionisation is after all one of the steps in three-step-model. Choosing a gas whose ionisation potential is too high results in negligible ionisation and no measurable x-ray output. This was confirmed when we attempted HHG with helium ($I_p = 24.6$ eV) and explains the far lower x-ray intensity obtained with neon. In order to take advantage of generation in higher ionisation potential gasses, the intensity of the laser must be increased to the point where they can be readily ionised.

Furthermore it must be accepted that capillary phase matching (as employed in this work) is not possible for very high order harmonics. Even if these high order harmonics are well below the theoretical cutoff imposed by the laser power they can only be generated with very low efficiency. It is quite possible that in this work harmonics were generated up to the theoretical cutoff but that their spectral intensity was so low they could not be detected.

A similar concept to changing the gas species is x-ray generation from the ionised gas. In this scheme the laser is sufficiently intense to fully ionise the neutral gas at an early time in the pulse. HHG is then performed by tunnel ionisation of the next weakest bound electron - in argon this has an ionisation potential of 27.6 eV. Generation from argon ions has been demonstrated[91] using a 3 mJ pulse energy with pulse durations only a little shorter than were employed in this work (22 fs). This produced x-ray energies as high as 180 eV but with very low efficiency.

In conclusion changing gas species (or generating from ionised gas) will only provide an effective means of achieving higher harmonics if combined with increased peak electric field. Phase matching (as employed in this work) will not be possible for these high harmonics and so generation efficiencies are likely to be extremely low.

Reduced capillary bore:

Reducing the capillary bore constrains the electric field more tightly. An increase in field strength proportional to the reduction in area would mean higher peak powers could be achieved using the same laser system. (Focusing conditions would obviously need to be changed). The increase in peak power would give a higher theoretical cutoff frequency therefore permitting generation of higher harmonics. Reducing the capillary bore by an order of magnitude (from 150 μm to 15 μm diameter) would in theory increase the peak intensity from $\sim 2 \times 10^{14}$ to $\sim 2 \times 10^{16}$ W cm^{-2}. Equation 3.2 can be used to calculate the increase in ponderomotive potential (from 13 eV to 1.3 keV). This value can then be inserted into equation 3.1 to find the new theoretical cutoff; the cutoff wavelength is reduced from 22 nm to 0.3 nm.

Unfortunately, as described in section 3.1.4, it is phase matching rather than the theoretical cutoff that sets the limit for the highest harmonic generated. Furthermore, reducing the capillary bore increases the magnitude of the waveguide term in the phase matching equation (3.8). This means that phase matching of each harmonic occurs at a slightly lower ionisation fraction. As a consequence generation will be forced to earlier times in the pulse or to narrow regions at larger radii within the capillary. Since x-ray generation cannot be phase matched at the peak of the pulse it will be impossible to take full advantage of the increased peak power.

In practice there may also be severe problems associated with focusing the intense laser beam into such a small waveguide. Focusing would certainly have to be performed under vacuum to avoid nonlinearity and ionisation in the air. The rate of capillary degradation would be greatly increased - which would lead in turn to lower launch efficiency into the EH_{11} mode and higher loss. Gains in harmonic order due to reducing the capillary

bore (or simply to increasing the laser power) are likely to be far more modest than the predictions of equation 3.1 suggest.

An interesting possibility that could be explored in the future is that of generation with higher ionisation potential gasses (or ionised gas) in a waveguide of reduced capillary bore.

If it is accepted that capillary phase matching cannot be achieved for very high harmonics this brings into question the purpose of the capillary - why not simply use the gas jet configuration? As discussed in section 3.1.6 the capillary does allow confinement of the laser to be extended over a greater length than could be achieved by free space focusing of the beam. However, the merit of this increased interaction length is diminished by the fact that the coherence length of these higher harmonics is very short. The greatest motivation for the use of a capillary waveguide (when generating wavelengths that cannot be phase matched normally) is, therefore, as a means of implementing quasi-phase matching schemes.

Higher laser output power:

Increased laser power (assuming the same pulse duration is maintained) should offer benefits similar to those of reducing the capillary bore. However, when laser power is increased directly, the waveguiding term in the phase matching equation is not modified. Suppose the laser output power could be increased to a level sufficient to reach the same peak intensities achieved by reducing the capillary bore by an order of magnitude (power at capillary output would increase from 300 mW to 30 W)!

Since the waveguiding term in the phase matching equation is unchanged so too is the critical ionisation fraction (equation 3.9) imposed by phase matching. This critical ionisation fraction will simply be attained at an earlier time in the pulse. For the vast majority of the laser pulse duration the ionisation fraction will be too high to permit phase matched generation. Consequently emission of higher energy x-rays will have very low efficiency. Any such emission may well result in x-ray spectral intensities that fall below the detection limit of our current apparatus.

Just as was discussed in relation to the capillary bore, the benefits of increased average power (in the capillary phase matched scheme) will be far more modest than predicted by the theoretical cutoff because the higher harmonics cannot be phase matched. It is only in non-phase matched or quasi-phase matched generation that increasing laser intensity will provide significant gains in harmonic energy.

If a further laser upgrade is to be considered then a pulse compressor and carrier envelope stabilisation would both be preferable to another mJ of energy.

Pulse compression:

Pulse compression offers an alternative route to higher peak powers. Moreover, shortening the pulse also steepens the gradient of the electric field amplitude on the leading edge. This means that far higher laser electric fields are present when the ionisation fraction required for phase matching is reached. This both increases harmonic yield and permits phase matching of higher harmonics.

Shortening the pulse duration in order to achieve a higher peak electric field amplitude is therefore far more beneficial (especially in the capillary phase matched scheme) than directly increasing the laser power in order to achieve the same peak field strength. Equivalently, shortening the pulse duration but maintaining the same peak electric field amplitude will also provide improvements in the maximum energy of harmonic generation and in generation efficiency. This means that one can afford to lose a considerable amount of pulse energy when implementing a compression scheme.

There are several options available for pulse compression systems (see section 5) however, the most promising to date appears to be spectral broadening through self-phase-modulation in a gas filled capillary. A particularly elegant example of this scheme is the filamentation system whereby the laser propagates as a spatial soliton - balancing Kerr self-focusing against diffraction[92]. A careful optimisation of the length of the gas tube may yield a self compressing, self aligning system that supports very high throughput.

Spielmann *et al.*[76] used a capillary compressor to reduce their 20 fs, 1.2 mJ pulse to 5 fs and 0.5 mJ. By focusing this beam tightly through a helium filled gas cell they were able to reach peak intensities of 10^{17} W cm^{-2}. This allowed x-ray generation to reach $\sim$ 500 eV although with very low intensity as a result of the very limited coherence length.

Shortening the pulse duration whilst maintaining the same pulse energy obviously causes a dramatic increase in peak intensity and therefore facilitates generation of extremely high harmonics. Such high peak intensities can be best exploited by using high ionisation potential gas. Through generation (in a helium gas jet) with a 5 fs, 1 mJ pulse Seres *et al.* have achieved x-ray energies as high as 1.3 keV[93]. This work employed a rather elaborate laser system involving two CPA stages, two gas filled capillaries (before and after the second CPA stage) a Dazzler pulse shaper, and broadband chirped mirrors to provide the final pulse compression. The intensities reached in this experiment are

sufficiently high that generation is, in part, due to multiple ionisation of the helium. Phase matching (as employed in our capillary work) would not have been possible under these conditions, however, Seres *et al.* state that they operate in a new regime of self phase matching which can only be reached for pulse durations $<$ 10 fs and intensities $>$ 1 $\times 10^{15}$ W cm^{-2}[94]. This allows higher x-ray intensities to be attained than would be possible in non phase matched operation.

In their work Seres *et al.* have simultaneously taken advantage of reduced pulse duration, increased peak electric field, change of gas species, generation from ionised gas and a new phase matching mechanism.

Quasi phase matching:

The work presented in this thesis has focused on the capillary phase matched regime of operation. For the relatively modest intensities reached by our laser system the addition of the capillary waveguide allows phase matched generation of harmonics almost up to the theoretical cutoff. This type of phase matching provides the highest possible rate of build-up in x-ray signal and consequently a relatively intense x-ray output. Unfortunately there is an upper limit on the harmonics that can be phase matched by this scheme (imposed by the critical ionisation fraction, see section 3.1.4). For harmonics above this limit there is a rapid drop in coherence length, capillary phase matching becomes ineffective and x-ray intensity is dramatically reduced. This mode of non-phase matched x-ray generation has been employed to achieve high energy x-ray generation (see above) however, this type of work typically requires extremely sensitive detectors such as Energy Dispersive X-ray spectrometers (EDX). Signal levels may be of the order of just a few photon counts per second in comparison to 10^7 photons per second in capillary phase matched schemes. This diminishes the applicability of this type of beam to scattering experiments where the scattered intensity per solid angle is already expected to be very low.

It is in the higher intensity regime where quasi phase matching becomes a very attractive proposition. As discussed in section 3.1.5 QPM provides a means for significant build-up in x-ray intensity to be achieved even at x-ray energies where the coherence length is very short.

QPM alone is likely to offer relatively limited gains in x-ray energy if employed in the current system. This is because the capillary already facilitates phase matching almost to the theoretical cutoff. In conjunction with other modifications such as reduced capillary bore, decreased pulse duration or increased laser power QPM could offer significant

gains. The true benefit of QPM is that it allows usable levels of x-rays to be generated at high photon energies.

QPM has been demonstrated using a capillary whose inner diameter is modulated along its length; this allowed x-ray generation to extend to 250 eV in argon[91] and 280 eV in neon[95]. Although effective, this scheme is rather inflexible since the period of the capillary diameter modulation is fixed. More dynamic solutions which enable real-time tuning of the modulation length have since been shown for example multi mode beating[96] (360 eV in argon) and counter propagating laser pulses[67] (showing control of electron trajectories).

Dynamic QPM schemes are extremely attractive as they could permit tuning of x-ray emission energy and are suited to feedback optimisation routines. The concept of using an acoustic wave for QPM is very elegant; however, in the configuration use for our initial trials it was not possible to couple enough power into the acoustic wave. Although the required depth of pressure modulation is very low I am doubtful that this will be achieved in the co-axial piezo configuration. Perhaps direct modulation of the gas pressure at the inlet would provide a more efficient means of launching the acoustic wave. Quasi phase matching remains a highly promising tool and is very likely to be pursued by this group in the future.

6.2 Paths to higher flux

In order to be successful in carrying out scattering experiments from very small (potentially single molecule diffraction targets) considerably greater flux will be required, or very long integration times will have to be endured.

Laser repetition rate:

Increasing repetition rate is a very interesting route to higher flux. The fibre laser system that has been designed in parallel with this project will produce lower peak power but could operate at a much higher repetition rate. It will be very interesting to compare the output of the two systems when they are both online. Provided the fibre laser system can achieve sufficient intensities for capillary phase matched generation in argon it should offer very usable levels of x-ray output. This is achieved in a potentially more compact and lower cost device than more traditional CPA systems. At a later time it may well be possible to add secondary amplification and/or pulse compression stages to increase the peak intensities that are produced.

X-ray focusing:

Focusing of the x-rays with multilayer mirrors has already been demonstrated to be an efficient means of intensifying signal. This technique will no doubt continue to be used. Development of the hollow fibre tapers has now stopped in favour of multilayer mirrors. Whilst this makes very sound sense it is a shame to have lost the possibility of constructing a totally integrated system.

The main difficulty encountered with the taper concept was in the fabrication of tapers with sufficiently accurate profiles. For optimum performance sub x-ray wavelength ($<$ 25 nm) flatness and conformity is required over an object that is physically quite long (several centimeters). This proved to be very difficult to achieve and tapers had to be severely truncated in order to observe any throughput. Shortening the taper increases the spot size thereby reducing the intensifying effect.

Increased capillary bore:

Increasing bore size improves x-ray flux in two ways: Firstly a larger volume of gas is able to contribute towards generation. Secondly the waveguiding term in the phase matching equation is reduced (see section 3.1.4). This means that phase matching is shifted slightly towards higher ionisation fractions - allowing more intense parts of the pulse to contribute towards the phase matched x-ray output.

Unfortunately, for the current laser system, increasing capillary bore by any significant amount is likely to reduce peak power by so much that the highest harmonics can no longer be generated. Increased bore could only really be considered in conjunction with a higher power laser system.

Quasi phase matching:

Quasi phase matching allows x-ray intensity to build-up over the length of the interaction region even when there is a large phase mismatch between the laser and the generated x-rays. Not only does this extend practical generation to higher photon energies it also enhances x-ray intensity at lower energies. This is because (in traditional phase matching) a large fraction of the pulse energy cannot be utilised; at the peak of the pulse, for example, the ionisation fraction is too high and plasma dispersion cannot be compensated for. When quasi phase matching is used the whole of the pulse may contribute to the x-ray output which is consequently significantly increased.

Appendix A

Capillary Alignment

In this section the steps that should be taken in order to achieve an efficient launch into an HHG capillary using a lens based focusing system are outlined. When starting alignment from scratch:

1. First ensure that spatially and temporally the laser is producing the best possible pulse: Measurements of the laser spatial profile can be made with an M^2 meter such as that described in section 2.2.4; if this is not possible then imaging the focused laser spot with a CCD camera will give a first order indication of beam quality. Measurements of the laser pulse temporal properties should be performed as close as possible to the capillary position; or failing this any additional dispersion that is in the path to the focus must be compensated for. Carrying out both of these checks ensures that the highest possible peak laser power will be produced at the input face of the capillary.

2. The unfocused laser beam should be aligned to be parallel with the optical table and at the appropriate height for the vacuum apparatus. Laser power should be attenuated and safety goggles should of course be worn; an I.R. viewer is therefore required. Items that need to be aligned to the x-ray beam should be aligned to the centre of the unfocused laser beam; for example the slit of the grazing incidence spectrometer.

3. Depending upon the apparatus being used it is helpful if the capillary holding device can now be brought into alignment with the laser. For example the Swagelok capillary to flange adaptors (described in section 2.3.1) can be aligned to the beam by observing the light passing through the adaptor on a white card (align the one nearest the laser first).

4. Position a marker (perhaps a cross drawn on a white piece of card) at the centre of the beam in the plane where the front of the capillary will be positioned. Now position the focusing lens in the beam so that the focused spot hits the centre of the marker. Note that it is important to attenuate sufficiently or the marker will be burned. Also take care of that the back reflection from the lens is dumped in a safe way. Once the x and y positions of the lens are set the marker may be removed.

5. Accurately setting the lens z position is very important and easily overlooked part of the alignment process. If attenuation is set correctly it is possible to use the volume of the sound made by the laser as it hits a white card in order to find the plane of the focus. Adjust the lens z position so that the sound is loudest when the card is positioned in the plane of the front face of the capillary. It may be necessary to repeat step 4 if the z translation axis of the lens is not exactly parallel to the laser beam.

6. Position the capillary holders so that the beam passes exactly through their centres. Now that the lens is installed there should be no clipping of the beam at all on the input side. It is worth spending some time to do this as accurately as possible as doing so will ensure the capillary is very close to correct alignment when inserted.

7. Insert the capillary and pump down the vacuum chamber. The attenuation should be high enough to prevent any continuum generation from occurring[†]. Use and I.R. viewer to look at the output light from the capillary projected onto a white card. Note the capillary translation stage positions so that it is easy to return to the start position.

8. Move the front of the capillary by small amounts horizontally and vertically. Observe the output light profile carefully, watching for mode like patterns. Light that has propagated as a capillary mode has a distinctive appearance, it forms patterns consisting of one or more bright spots. This can be distinguished from the more diffuse light and streaked patterns which are caused by refraction through the capillary walls.

Once a capillary mode output has been obtained, careful optimisation of the launch into the EH_{11} mode can begin. At all times throughout this process it is important to check that the capillary is not placed under undue shear stress by the holders at either end. Set up the I.R. viewer or I.R. camera so that you can observe to output mode profile whilst

[†]Continuum generation is evidenced by visible light emitted from the capillary and is caused by self phase modulation in the glass walls of the capillary. This broadens the spectrum of the input laser to such an extent that it becomes easily visible by eye.

having both hands free to adjust the capillary actuators. Note, these optimisations can also be performed using a power meter although the slight delay in response makes this more difficult. If this is attempted be aware that the entire mode must be collected by the power meter and be careful to check the mode profile frequently to avoid optimising a local maxima other than the EH_{11} mode.

I. Working either in the horizontal or vertical plane, translate the back of capillary (furthest from the laser) a little way from side to side whilst watching the output mode. Then step the front of the capillary slightly and repeat the oscillation of the back of the capillary. Compare the best mode obtained for the two front positions and decide weather or not the mode was improved. If the mode was improved keep stepping the front further in that direction.

II. Frequently swap from horizontal to vertical planes of adjustment until the output mode is a single spot and is optimised for symmetry and intensity. In time you will recognise mode patterns and will learn to anticipate which direction the capillary needs to be translated.

III. A very delicate last adjustment can be made when setting the final rear position of the capillary. When a capillary has been aligned previously this is very often all that is required.

IV. The focus can also be fine tuned by observing the capillary output mode. If the focused spot is much too small a series of faint rings will be visible around the central spot, in this situation it may be best to choose a lens with a slightly longer focal length. However, if no such lens is available then translating the lens (away from the capillary) so that the capillary is slightly out of the focal plane can reduce the appearance of these rings and thus improve coupling into the EH_{11} mode. The aim should be to achieve a single output spot, with as large a size as possible and with no sign of a sharp or clipped edge (this can be evidence of a ring about to form).

The laser attenuation should only be reduced once a clean output mode with high throughput is obtained. Care should be taken because removing the attenuation may move the beam slightly and at full power this could easily damage the capillary. Remove the attenuation and watch the capillary, if there is a slight flickering of continuum generation it may be possible to alleviate this by tweaking the capillary alignment very slightly (start with the front). If however there is a lot of continuum generation and this cannot be prevented by altering capillary alignment the focused spot must be too large. Check that there is nothing in the optical path, disturbing the beam quality. If

the M^2 cannot be improved the only remedy is to return to step 4 and focus slightly tighter using a shorter focal length lens.

Appendix B

Matlab Code

B.1 SEM Image Analysis

In order to measure the uniformity of the diffraction grids used in the Spatio-Sectral analysis experiment 4.3 scanning electron microscope images of the grids were recorded. The following Matlab code was used to analyse the uniformity of the grid structure:

```
%%%%%%%% Vertical Part %%%%%%%%
clear all
close all

I = imread('17x_X_54.503_Y_68.124_noch.tif');          % Load in SEM image

% 1mm = [20,21] to [168,169] = 1000um
scale = 1000/(168.5 - 20.5);                  % Calculate pixel size in um

I = I(1:675,:);                          % Crop image to get rid of scale bar
                                         % see figure B.1a
I = imrotate(I,0,'crop');                          % Rotate image if required

BW = edge(I,'Sobel',0.07,'vertical');    % Use Sobel method to detect
                                           % vertical edges, see figure B.1b

imshow(I')                                                  % Display image
axis image                                        % Set correct aspect ratio

[XA,YA] = find(BW == 1);                         % Get coordinates of points
                                                 % on edges (see figure B.2)

ind = 1;
n = 1;
```

```
25 d_Y = find(diff(YA)>5);                % Look for large steps in Y
26                                         % coordinate value (see figure B.2)
27 temp = d_Y(1);
28 for loop = 2:1:size(d_Y)                  % Loop aver all points found
29                                          % to be on an edge
30
31     if ( (d_Y(loop) - d_Y(loop-1)) > 5 )  %
32         G(ind) = temp/n;                 % Assign each Y coordinate a
33         ind = ind + 1; n = 1;            % group number, each group is
34         temp = d_Y(loop);                % separated by a large step
35     else                                 % in Y coordinate.
36         n = n +1;                        %
37         temp = temp + d_Y(loop);         % I.E. generate a group of
38     end                                  % points for the edges of
39 end                                      % each wire.
40 G(ind) = temp/n;                         %
41
42 m0 = 0.04;                      % Initial gues value for wire gradient...
43 W0 = 45;                        % aperture width...
44 B0 = 8;                         % and bar width.
45
46 global X Y
47 startloop = 4;                             % Throw away first few groups
48 for loop = startloop:1:size(G,2)-2          % Loop over remaining groups
49     y_pos0 = loop*(size(BW,2)/size(G,2));  % Set initial guess for y
50                                            % intercept depending on group
51
52 X = XA(G(loop):G(loop+2));          % Select only data from current group
53 Y = YA(G(loop):G(loop+2));          %
54
55 %scatter(X,Y)         % Optional scatter plot of fit data, see figure B.2.
56
57 mL = 1; WL = 1; BL = 1; y_posL = 1;
58                                    % Allocate high previous values for fit
59                                    % parameters so that first loop runs
60
61 while ( (m0≠mL) && (W0≠WL) && (B0≠BL) && (y_pos0≠y_posL) )
62                   % Itterate this loop until fit parameters stop changing
63
64     mL = m0; WL = W0; BL = B0; y_posL = y_pos0;
65                % Set initial parameter values to those of last itteration
66
67     m = fminsearch(@(m) edges_fn(m,W0,B0,y_pos0),m0);
68     W = fminsearch(@(W) edges_fn(m,W,B0,y_pos0),W0);
69     B = fminsearch(@(B) edges_fn(m,W,B,y_pos0),B0);
70     y_pos = fminsearch(@(y_pos) edges_fn(m,W,B,y_pos),y_pos0);
71                  % These parts call edges_fn.m which uses the current fit
72                  % parameters to generate 4 fit lines to the fit data of
```

```
73                  % the current group.  edges_fn.m then returns the error
74                  % of this fit.  The error is minimsed by the fminsearch
75                  % for each parameter in turn.
76
77      m0 = m; W0 = W; B0 = B; y_pos0 = y_pos;
78                                          % Store the current values as the
79                                          % last itteration parameter values
80
81
82      Y1 = polyval([m,(y_pos + (W/2) +B)],X);        % Calculate the four
83      Y2 = polyval([m,(y_pos + (W/2))],X);           % lines defined by the
84      Y3 = polyval([m,(y_pos - (W/2))],X);           % current parameter
85      Y4 = polyval([m,(y_pos - (W/2) -B)],X);        % values.
86
87      error = min(([(Y1-Y).^2,(Y2-Y).^2,(Y3-Y).^2,(Y4-Y).^2])')';
88              % Calculate error for each group once the best fit is reached
89      RMSE = sqrt(mean(error));          % Calculate the RMSE between groups
90
91  end    % End of the itterative loop to find best fit lines for each group
92
93
94  Y_Position(loop-(startloop-1)) = y_pos;  % Move on to next fit data group
95
96  Aperture_Width(loop-(startloop-1)) = W;              % Record fit data
97  Bar_Width(loop-(startloop-1)) = B;                   % parameters for the
98  Angle(loop-(startloop-1)) = ((atan(m)*180)/pi);      % previous loop
99
100 hold on
101
102 Y1 = polyval([m,(y_pos + (W/2) +B)],X);           %
103 Y2 = polyval([m,(y_pos + (W/2))],X);              % Generate fitted lines
104 Y3 = polyval([m,(y_pos - (W/2))],X);              %
105 Y4 = polyval([m,(y_pos - (W/2) -B)],X);           %
106
107 plot(X,Y1,'r')                                       %
108 plot(X,Y2,'r')                                       % Plot fitted lines
109 plot(X,Y3,'r')                                       % see figure B.3
110 plot(X,Y4,'r')                                       %
111
112 drawnow                                            % Update plot each loop
113
114 end                                    % End of loop over vertical groups
115
116
117 title({['X Bar Width = ',num2str(scale*(mean(Bar_Width))),...
118 ' \pm ',num2str(scale*(std(Bar_Width))),' \mum'];...
119       ['X Aperture Width = ',num2str(scale*(mean(Aperture_Width))),...
120 ' \pm ',num2str(scale*(std(Aperture_Width))),' \mum']} ,...
```

```
121         'fontsize',16)
122                         % Display calculated parameter values in figure title
123                         % See figure B.4
124
125
126
127
128 %%%%%%%% Horizontal Part %%%%%%%%
129
130 figure
131 clear G RMSE X Y XA YA Y1 Y2 Y3 Y4 error ind loop temp
132 I = I';
133
134                          % This part simply rotates the image and repeats the
135                          % same code as was used for the vertical lines...
```

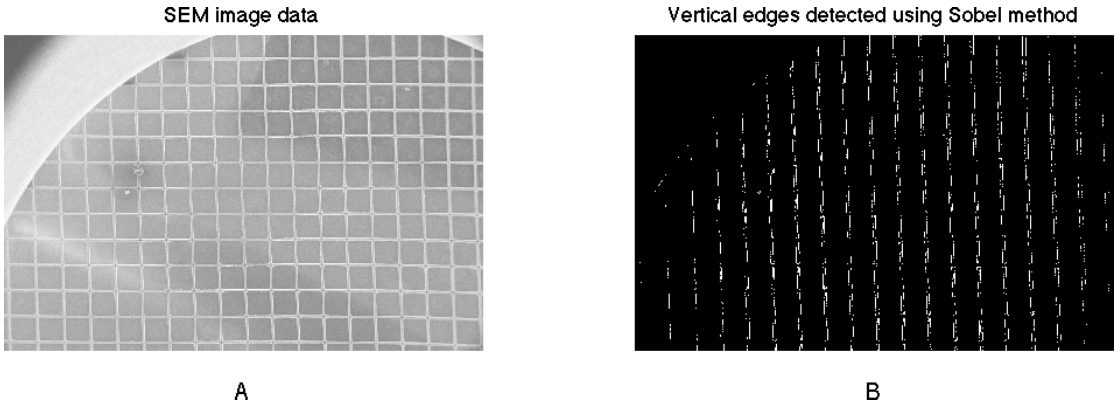

FIGURE B.1: A: Plot of data measured using a scanning electron microscope. B: The Sobel method is used to detect the vertical edges in image A.

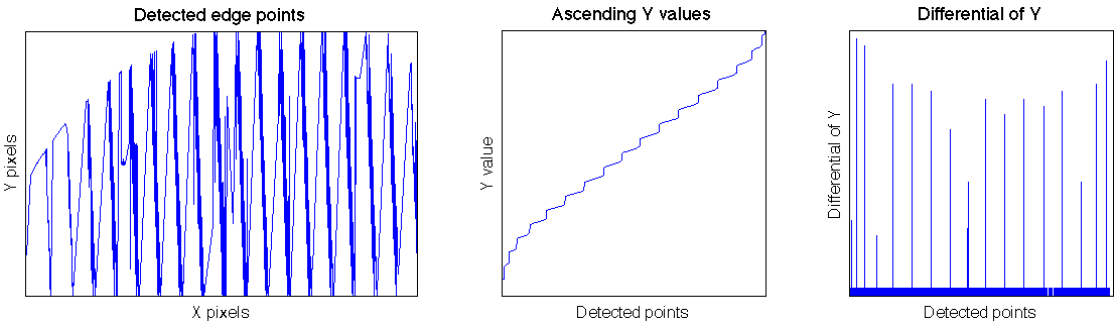

FIGURE B.2: Left: X,Y plot of all data points determined to be on edges. Middle: Plot of Y data values in ascending order. Right: Plot the differential of the ascending Y data values, the spikes clearly separate the data into groups corresponding to each edge.

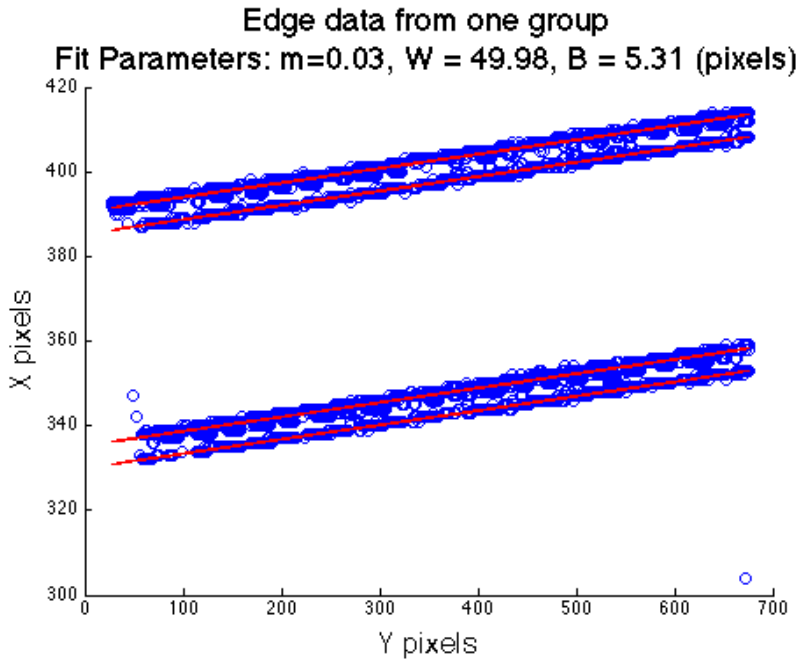

FIGURE B.3: X,Y scatter plot of the data contained in one group. The red lines show the best fit to the data after the itterative loop, the parameter values are shown in the title. Notice that data is transposed whilst the fit is performed.

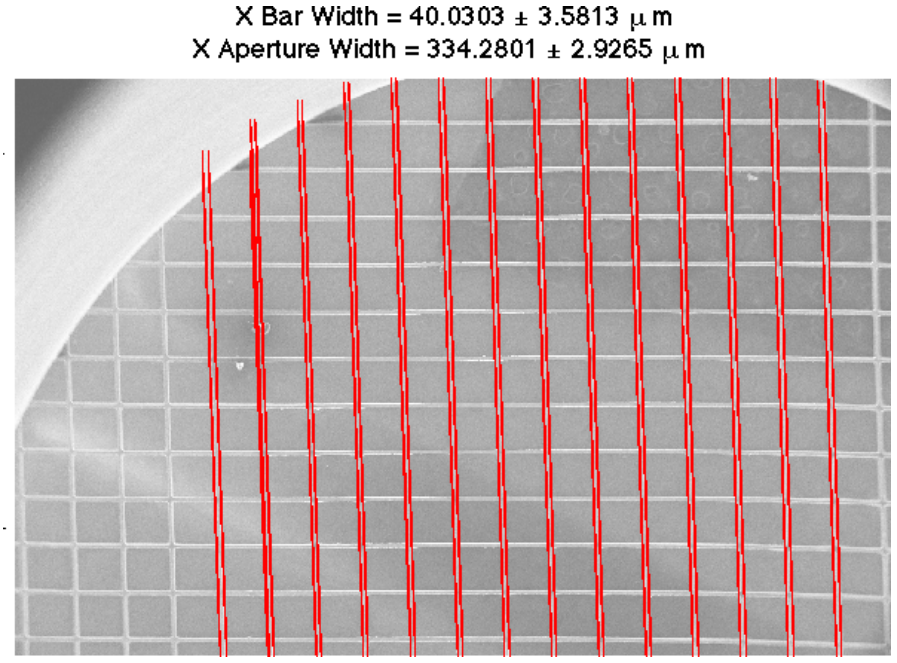

FIGURE B.4: Example output with vertical lines fitted. Average values and errors are quoted in the title these are now scaled to units of microns.

B.2 Diffraction Simulation

This function was written in order to simulate the diffraction pattern produced by a square aperture:

```
function Int2D = CornuVariable(wl,output_size)
% Plots the diffraction pattern for a given wavelength(nm)
% wl is the wavelength in nm
% output_size is the number of pixels across the output

% Set hard wired parameters here
d=340e-6;               % Aperture diameter /m
r=0.512;                % Source-Aperture distance /m
s=1.2-0.515;            % Aperture-Screen distance /m


% Convert wavelenth from nm to m
wl=wl*1e-9;

% Produce an array of screen positions q
qmax=1e-3/2;                        % Set detector size /m
qstep=(2*qmax)/(output_size-1);     % Set pixel size /m
q=-qmax:qstep:qmax;                 % Calculate array of screen positions


% Calculate parameter values for each screen position q
k=(q*r)./(r+s);                     % Edge displacement changes with q
w1=((d/2)+k)*( ( (2*(r+s))./(r*s*wl) )^0.5);    % Parameters for each edge
w2=((-d/2)+k)*( ( (2*(r+s))./(r*s*wl) )^0.5);   % of the aperture.


% Evaluate the Fresnel integrals for the parameters v1 and v2
P1=fresnelc(w1)+(i*fresnels(w1));       % This calls the Fresnel sin and
P2=fresnelc(w2)+(i*fresnels(w2));       % cos functions


% Subtract the field diffracted by the second edge from that of the first
% this applies Babinet's principle to recover the E-field for a slit.
Pt=(P2-P1)/2;


% Convert from E-Field to intensity
Int=real(Pt).^2+imag(Pt).^2;

% Multiply all values of slit intensity by all values of its transpose.
% This produces a 2D array of values from 2 1D arrays.
```

```
42 Int2D = kron(Int,Int');
43
44
45 % Plot the diffracion pattern
46 imagesc(Int2D);
47 colormap gray;
```

The Fresnel sine and cosine integrals were calculated numerically using the following functions:

Fresnel sine integral:

```
1 function sx = fresnels(x)
2 eqn =inline('sin(pi*0.5.*(v.^2))');
3 sx(1) = quad(eqn,0,x(1));
4 for a=2:1:size(x,2)
5     delt(a) = quad(eqn,x(a-1),x(a));
6     sx(a) = sx(a-1) + delt(a);
7 end
```

Fresnel cosine integral:

```
1 function cx = fresnelc(x)
2 eqn =inline('cos(pi*0.5.*(v.^2))');
3 cx(1) = quad(eqn,0,x(1));
4 for a=2:1:size(x,2)
5     delt(a) = quad(eqn,x(a-1),x(a));
6     cx(a) = cx(a-1) + delt(a);
7 end
```

An example output of `CornuVariable(25,1000);` is shown in figure B.5:

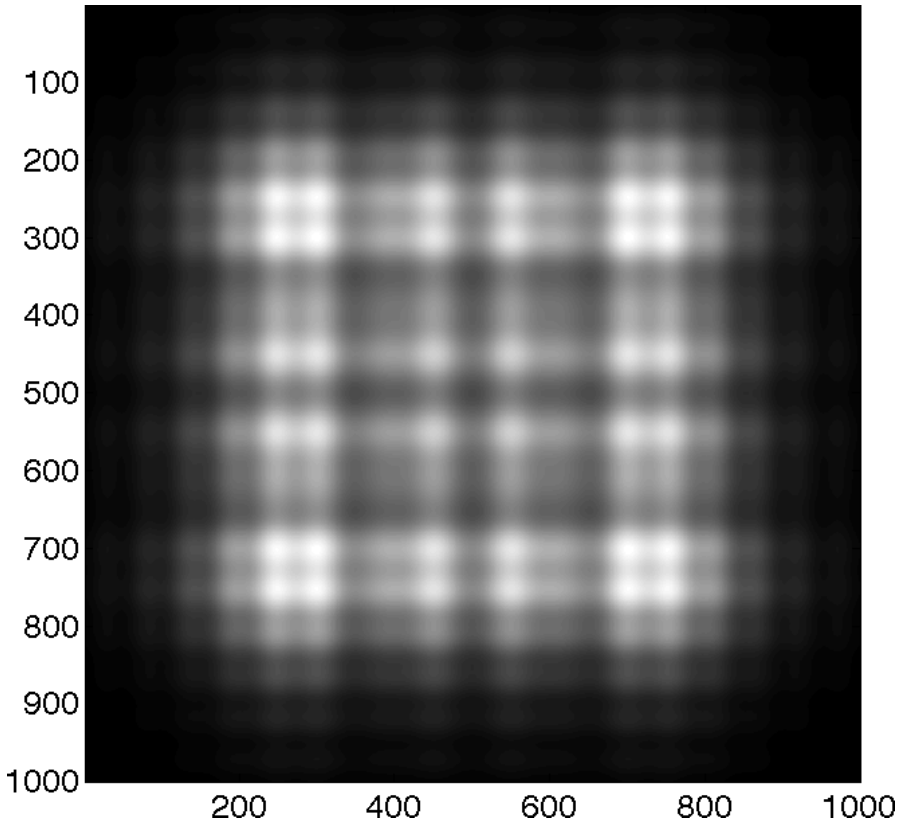

FIGURE B.5: Example output of CornuVariable.m. The axis units are simply array indices (the output size was set to 1000).

B.3 Least Squares Fitting Routine

This appendix contains the code used in the spatio-spectral reconstruction technique.

B.3.1 Input data processing

The code below was used to convert the raw CCD data into a form suitable for analysis. Values for the image rotation and the aperture spacing were entered by hand.

```
1  function [img, x, y, output_symmetric, offset] = mkgrid2_BG_v2(fname)
2
3  num_grids = 70;        % set the number of grids to be extracted per image
4
5  img = imread(fname,'bmp');          % load the CCD data
6  img = imrotate_M(img,-37.8);        % rotate image using nearest neighbour
7
8  % imagesc(img);
9  % axis image
10 % hold on
11
12 % Manually enter these values to extract the data for each aperture
13 subimage_size = 60;     %60 % Set the number of pixels across each pattern
14 gap = 1;                %1
15 offset_x = 54;          %54
16 offset_y = 15;          %16
17 subimage_size = subimage_size - 1;
18
19 offset.x = offset_x;                        % Store these values in a
20 offset.y = offset_y;                        % globally defined structure
21 offset.gap = gap;                           % array.  This allows them to
22 offset.subimage_size = subimage_size+1;     % be passed to other functions
23
24 % Optional plotting commands to aid in the setting the values above
25 % the output of these commands is shown in figure B.6
26 % imagesc(img)
27 % for a = offset_x+1:subimage_size+gap:size(img,1)
28 %     line([a,a],[1,size(img,2)],'color','k');
29 % end
30 % for b = offset_x+gap+1:subimage_size+gap:size(img,1)
31 %     line([b,b],[1,size(img,2)],'color','k');
32 % end
33 % for c = offset_y+1:subimage_size+gap:size(img,2)
34 %     line([1,size(img,1)],[c,c],'color','k');
35 % end
36 % for d = offset_y+gap+1:subimage_size+gap:size(img,2)
```

```
%     line([1,size(img,1)],[d,d],'color','k');
% end
% axis image

a = offset_x+1:subimage_size+gap:size(img,1);       % Convert values above
b = offset_x+gap+1:subimage_size+gap:size(img,1);   % into coordinate
c = offset_y+1:subimage_size+gap:size(img,2);       % ranges for extraction
d = offset_y+gap+1:subimage_size+gap:size(img,2);   % from the CCD data.



for m = 1:1:size(c,2)-1          % loop over number of apertures in x-dir
    for n = 1:1:size(a,2)-1      % loop over number of apertures in x-dir

        data(:,:,n,m) = img(d(m):c(m+1),b(n):a(n+1));   % Extract data

        min_pixel = min(min(data(:,:,n,m)));            % The data from
        max_pixel = max(max(data(:,:,n,m)));            % each apreture
        if ((max_pixel == 0) || (min_pixel == 0))       % is given a score
            contrast_ratio(n,m) = 0;                    % contrast-ratio.
        else                                            % This is a measure
            contrast_ratio(n,m) = max_pixel/min_pixel;  % of how usefull
        end                                             % the data is.
    end
end
temp2 = reshape(contrast_ratio,[],1);   % Convert from 2D to 1D


temp = sum(sum(data,2),1);       % The sum of intensities for
temp = squeeze(temp);            % each aperture is found.

s = size(temp,1);                % Finds the number of aperures

[i_of_square, ind_of_square] = sort(temp2); % This sorts the apertures in
                                            % order of their contrast_ratio
                                            % score - outputs in indices.


% This section creates the x and y aperture index arrays.
% The apertures are ordered according to their contrast_ratio.
counter = 1;
for loop = size(ind_of_square,1):-1:size(ind_of_square,1)-num_grids;
    y(counter) = floor((ind_of_square(loop)-1)/s) + 1;
    x(counter) = rem((ind_of_square(loop) - 1),s) + 1;
    counter = counter + 1;
end

% Pre-assign 3D output array
```

```
85 output_symmetric=...
86     zeros(offset.subimage_size,offset.subimage_size,num_grids);
87
88 % This loop selects only the apertures with the best contrast_ratio
89 for loop = 1:1:num_grids
90     A = data(:,:,x(loop),y(loop));  % This just gets the aperture
91     output(:,:,loop) = A;           % data in order of contrast_ratio.
92
93     % This line symmetrises the measured data
94     output_symmetric_temp(:,:)=(A+fliplr(A)+flipud(A)+fliplr(flipud(A)));
95
96     % This line finds the average background level of each image
97     % The background level is found by taking the mean of the data
98     % surrounding the diffraction pattern (i.e. the dark area).
99     background = (sum(sum(output_symmetric_temp(:,:))) -...
100        sum(sum(output_symmetric_temp(11:50,11:50))))/(60^2-40^2);
101
102    % The background signal level is subtracted from each pattern
103    output_symmetric_temp(:,:) = output_symmetric_temp(:,:) - background;
104
105    % The data from the dark area (which now averages to 0 is stripped
106    output_symmetric(11:50,11:50,loop)=output_symmetric_temp(11:50,11:50);
107    % A comparison between output and output_symmetric is shown in fig B.7
108 end
```

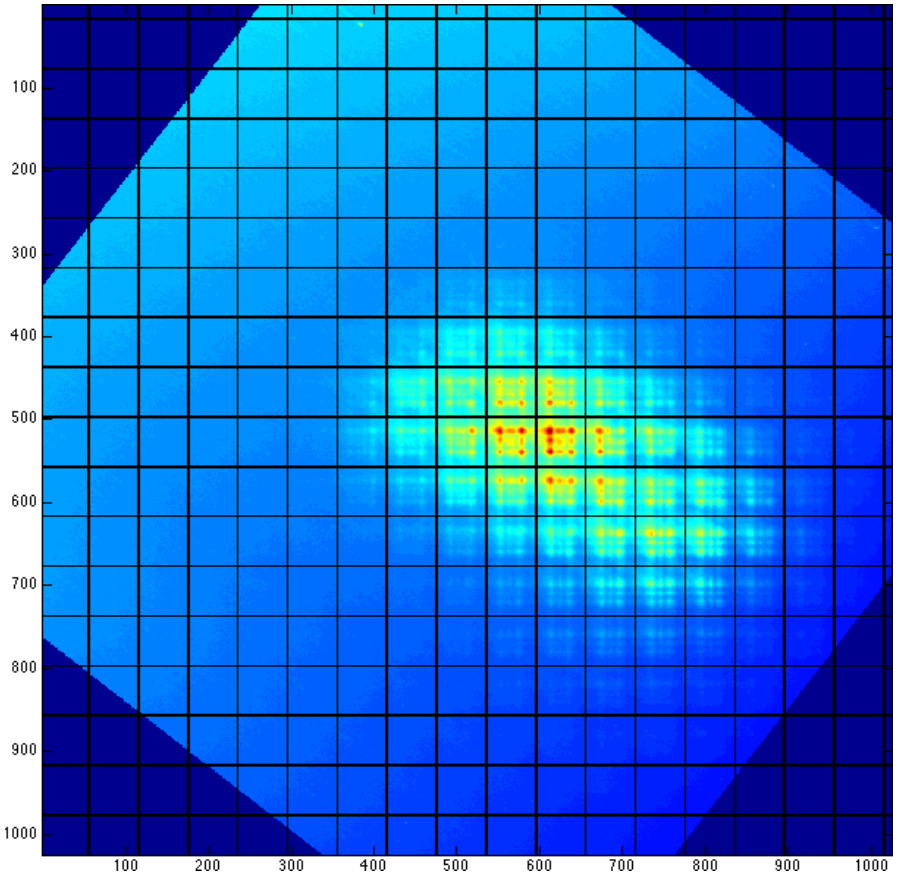

FIGURE B.6: Example output when the optional plotting commands are executed - this is for visual checking that the aperture spacing values have been set correctly.

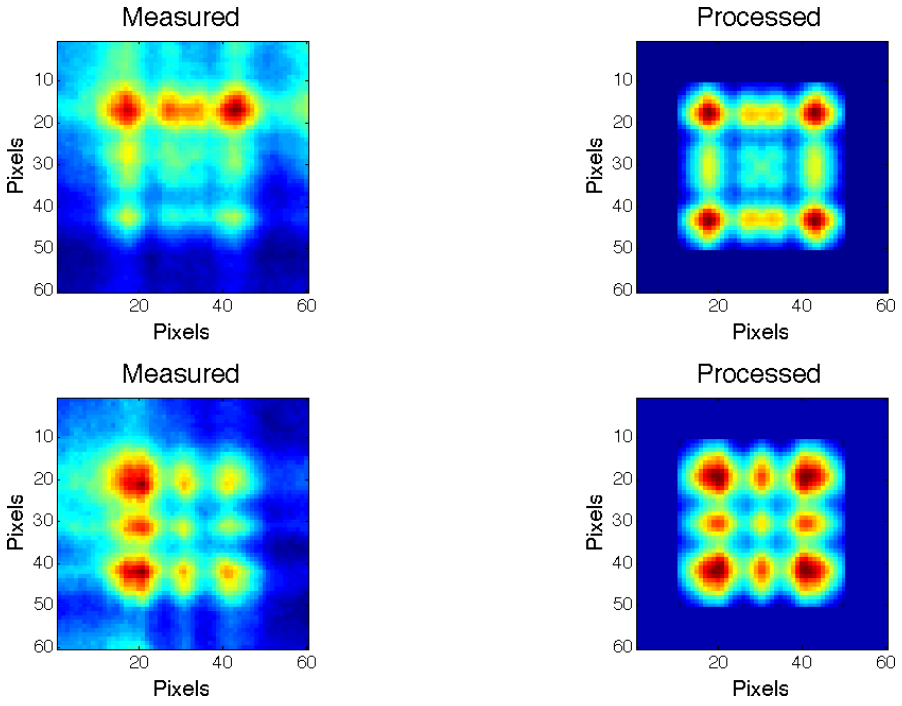

FIGURE B.7: Comparison between raw input data and processed data for two different apertures. Note that the processed data has been symmetrised and the background signal has been removed from the dark area surrounding the diffraction pattern.

Bibliography

[1] M. Praeger, A. M. De Paula, C. A. Froud, E. T. F. Rogers, S. L. Stebbings, W. S. Brocklesby, J. J. Baumberg, D. C. Hanna, and J. G. Frey. Spatially resolved soft x-ray spectrometry from single-image diffraction. *Nature Physics*, 3:176–179, 2007.

[2] C. A. Froud, E. T. F. Rogers, D. C. Hanna, W. S. Brocklesby, M. Praeger, A. M. de Paula, J. J. Baumberg, and J. G. Frey. Soft-x-ray wavelength shift induced by ionization effects in a capillary. *Optics Letters*, 31:374–376, 2006.

[3] G. Sines and Y. A. Sakellarakis. Lenses in antiquity. *American Journal of Archaeology*, 91(2):191–196, Apr 1987.

[4] H. E. Burton. The optics of Euclid. *Journal of the Optical Society of America*, 35 (5):357–357, 1945.

[5] Max Herzberger. Optics from Euclid to Huygens. *Applied Optics*, 5(9):1383–1393, 1966.

[6] I. Newton, Translated by, I. B. Cohen, and A. Whitman. *The Principia : Mathematical Principles of Natural Philosophy*. University of California Press, October 1999. ISBN 0520088174.

[7] T. Melvill. Observations on light and colours. *Journal of the Royal Astronomical Society of Canada*, 8, Aug 1914.

[8] J. C. Maxwell. A dynamical theory of the electromagnetic field. *Philosophical Transactions of the Royal Society of London*, 155:459–513, 1865.

[9] W. C. Roentgen. On a new kind of rays. *Science*, 3(59):227–231, 1896.

[10] L. Gerward and A. Rassat. Paul villard's discovery of gamma rays - a centenary. *Comptes Rendus de l'Acadmie des Sciences - Series IV - Physics*, 1(7):965 – 973, 2000.

[11] T. Lyman. Victor Schumann. *Astrophysical Journal*, 39:1–+, Jan 1914.

[12] D. Attwood. *Soft X-Rays and Extreme Ultraviolet Radiation: Principles and Applications*. Cambridge University Press, Mar 2007. ISBN 052102997X.

[13] A. L. Schawlow and C. H. Townes. Infrared and optical masers. *Physical Review*, 112(6):1940–1949, Dec 1958.

[14] V. Yanovsky, V. Chvykov, G. Kalinchenko, P. Rousseau, T. Planchon, T. Matsuoka, A. Maksimchuk, J. Nees, G. Cheriaux, G. Mourou, and K. Krushelnick. Ultra-high intensity- 300-tw laser at 0.1 hz repetition rate. *Optics Express*, 16(3):2109–2114, 2008.

[15] M. Nisoli and G. Sansone. New frontiers in attosecond science. *Progress in Quantum Electronics*, 33(1):17–59, Jan 2009.

[16] S. Khan. Free-electron lasers. *Journal of Modern Optics*, 55(21):3469–3512, 2008.

[17] R. A. Bartels, A. Paul, H. Green, H. C. Kapteyn, M. M. Murnane, S. Backus, I. P. Christov, Y. Liu, D. Attwood, and C. Jacobsen. Generation of spatially coherent light at extreme ultraviolet wavelengths. *Science*, 297(5580):376–378, 2002.

[18] S. Eisebitt, J. Luning, W. F. Schlotter, M. Lorgen, O. Hellwig, W. Eberhardt, and J. Stohr. Lensless imaging of magnetic nanostructures by x-ray spectro-holography. *Nature*, 432(7019):885–888, Dec 2004.

[19] H. N. Chapman, A. Barty, M. J. Bogan, S. Boutet, M. Frank, S. P. Hau-Riege, S. Marchesini, B. W. Woods, S. Bajt, H. Benner, R. A. London, E. Ploenjes, M. Kuhlmann, R. Treusch, S. Duesterer, T. Tschentscher, J. R. Schneider, E. Spiller, T. Moeller, C. Bostedt, M. Hoener, D. A. Shapiro, K. O. Hodgson, D. Van der Spoel, F. Burmeister, M. Bergh, C. Caleman, G. Huldt, M. M. Seibert, F. R. N. C. Maia, R. W. Lee, A. Szoeke, N. Timneanu, and J. Hajdu. Femtosecond diffractive imaging with a soft-x-ray free-electron laser. *Nature Physics*, 2(12): 839–843, Dec 2006.

[20] R. Neutze, R. Wouts, D. van der Spoel, E. Weckert, and J. Hajdu. Potential for biomolecular imaging with femtosecond x-ray pulses. *Nature*, 406:752–757, Aug 2000.

[21] O. Martinez. Design of high-power ultrashort pulse amplifiers by expansion and recompression. *Quantum Electronics, IEEE Journal of*, 23(8):1385–1387, Aug 1987.

[22] F. Verluise, V. Laude, Z. Cheng, C. Spielmann, and P. Tournois. Amplitude and phase control of ultrashort pulses by use of an acousto-optic programmable dispersive filter: pulse compression and shaping. *Optics Letters*, 25:575–577, 2000.

[23] D. Kaplan and P. Tournois. Theory and performance of the acousto optic programmable dispersive filter used for femtosecond laser pulse shaping. *Journal De Physique IV*, 12:69–75, 2002.

[24] Newport Spectra-Physics. Tsunami brochure. "http://www.newport.com/file_store/Data_Sheet/Tsunami__Brochure.pdf", 2002.

[25] Newport Spectra-Physics. Spitfire pro datasheet. "http://www.newport.com/file_store/Data_Sheet/Spitfire_Pro_Data_Sheet_(print_version_-_1_2_MB).pdf", 2007.

[26] R. Trebino. *Frequency-resolved optical gating: the measurement of ultrashort laser pulses*. Kluwer Academic, Boston, Aug 2002. ISBN 9781402070662.

[27] Light Speed Technologies. Timewarp. "http://www.light-speed-tech.com/TimeWarp.htm", 2004.

[28] APE. Autocorrelator pulsecheck datasheet. "http://www.ape-berlin.de/pdf/Autocorrelator-PulseCheck.pdf", 2008.

[29] R. Trebino. GRENOUILLE Tutorial. "http://www.swampoptics.com/PDFs/tutorials_GRENOUILLE.pdf", 2003.

[30] A. E. Siegman. How to (maybe) measure laser beam quality. "http://www.stanford.edu/~siegman/beam_quality_seminar.pdf", Apr 2004.

[31] Lawrence Berkeley National Laboratory. Center for x-ray optics. "http://www-cxro.lbl.gov/", Since 1995.

[32] A. van Brakel, C. Grivas, M. N. Petrovich, and D. J. Richardson. Micro-channels machined in microstructured optical fibers by femtosecond laser. *Optics Express*, 15:8731–8736, 2007.

[33] R. P. Feynman, R. B. Leighton, and M. L. Sands. *The Feynman Lectures on Physics, Volume II, 32-6*. Pearson/Addison-Wesley, San Francisco, definitive ed edition, Jul 2005. ISBN 9780805390476.

[34] M. J. Weber. *Handbook of optical materials*. CRC Press, Boca Raton, Sep 2002. ISBN 9780849335129.

[35] A. van der Lee. Grazing incidence specular reflectivity: theory, experiment, and applications. *Solid State Sciences*, 2:257–278, 2000.

[36] Xradia Inc. Kirkpatrick-baez focusing system. "http://www.xradia.com/Products/kbsr1.html", Jul 2002.

[37] M. Ehle, I. de la Calle, M. Diaz Trigo, R. Gonzalez Riestra, L. Loiseau, and P. Rodriguez. Xmm-newton users handbook v2.6. "http://xmm.esac.esa.int/external/xmm_user_support/documentation/index.shtml", Jul 2008.

[38] G. Brambilla, V. Finazzi, and D. J. Richardson. Ultra-low-loss optical fiber nanotapers. *Optics Express*, 12:2258–2263, 2004.

[39] J. H. Underwood and T. W. Barbee. Layered synthetic microstructures as bragg diffractors for x-rays and extreme ultraviolet - theory and predicted performance. *Applied Optics*, 20:3027–3034, 1981.

[40] S. Bajt, J. B. Alameda, T. W. Barbee, W. M. Clift, J. A. Folta, B. Kaufmann, and E. A. Spiller. Improved reflectance and stability of mo-si multilayers. *Optical Engineering*, 41(8):1797–1804, Aug 2002.

[41] M. Suman, F. Frassetto, P. Nicolosi, and M. G. Pelizzo. Design of aperiodic multilayer structures for attosecond pulses in the extreme ultraviolet. *Applied Optics*, 46:8159–8169, 2007.

[42] V. Aristov, A. Isoyan, V. Kohn, A. Kuyumchyan, E. Shulakov, A. Snigirev, and I. Snigireva. Study of optical properties of x-ray system based on two zone plates. *Nuclear Instruments & Methods In Physics Research Section A-Accelerators Spectrometers Detectors and Associated Equipment*, 575:238–241, 2007.

[43] W. L. Chao, B. D. Harteneck, J. A. Liddle, E. H. Anderson, and D. T. Attwood. Soft x-ray microscopy at a spatial resolution better than 15nm. *Nature*, 435:1210–1213, 2005.

[44] A. G. Michette, S. J. Pfauntsch, A. Erko, A. Firsov, and A. Svintsov. Nanometer focusing of x-rays with modified reflection zone plates. *Optics Communications*, 245:249–253, 2005.

[45] S. Rehbein and G. Schneider. Volume zone plate development at bessy. In *IPAP Conf. Series 7*, pages 103–106, 2005.

[46] Andor Technology. Andor dx434 specification sheet. "http://www.andor.com/scientific_cameras/x-ray/models/?iProductCodeID=76", Sep 2005.

[47] P. B. Corkum. Plasma perspective on strong-field multiphoton ionization. *Physical Review Letters*, 71:1994–1997, 1993.

[48] T Brabec and F Krausz. Intense few-cycle laser fields: Frontiers of nonlinear optics. *Reviews of Modern Physics*, 72(2):545–591, Apr 2000.

[49] R. R. Freeman, P. H. Bucksbaum, and T. J. McIlrath. The ponderomotive potential of high intensity light and its role in the multiphoton ionization of atoms. *Quantum Electronics, IEEE Journal of*, 24(7):1461–1469, Jul 1988.

[50] J. L. Krause, K. J. Schafer, and K. C. Kulander. High-order harmonic-generation from atoms and ions in the high-intensity regime. *Physical Review Letters*, 68: 3535–3538, 1992.

[51] M. Lewenstein, P. Balcou, M. Y. Ivanov, A. L'Huillier, and P. B. Corkum. Theory of high-harmonic generation by low-frequency laser fields. *Physical Review A*, 49 (3):2117–2132, Mar 1994.

[52] L. Gao, X. Li, P. Fu, R. R. Freeman, and D. Guo. Nonperturbative quantum electrodynamics theory of high-order harmonic generation. *Physical Review A*, 61 (6):063407, May 2000.

[53] M. Lewenstein, P. Salières, and A. L'Huillier. Phase of the atomic polarization in high-order harmonic generation. *Physical Review A*, 52(6):4747–4754, Dec 1995.

[54] R. W. Boyd. *Nonlinear Optics*. Academic Press, London, 2003. ISBN 9780121216825.

[55] Y. R. Shen. *The Principles of Nonlinear Optics*. Wiley-Interscience, New York, 1984. ISBN 9780471889984.

[56] J. L. Krause, K. J. Schafer, and K. C. Kulander. Calculation of photoemission from atoms subject to intense laser fields. *Physical Review A*, 45(7, Part B):4998–5010, Apr 1992.

[57] V. S. Popov. Tunnel and multiphoton ionization of atoms and ions in a strong laser field (keldysh theory). *Physics-Uspekhi*, 47(9):855–885, Sep 2004.

[58] M. V. Ammosov, N. B. Delone, and V. P. Kraĭnov. Tunnel ionization of complex atoms and of atomic ions in an alternating electromagnetic field. *Soviet Physics - JETP*, 64(6):1191–1194, 1986.

[59] E. Constant, D. Garzella, P. Breger, E. Mevel, C. Dorrer, C. Le Blanc, F. Salin, and P. Agostini. Optimizing high harmonic generation in absorbing gases: Model and experiment. *Physical Review Letters*, 82(8):1668–1671, Feb 1999.

[60] C. R. Gouy. Sur une propriete nouvelle des ondes lumineuses. *C. R. Acad. Sci. Paris*, 110:1251, 1890.

[61] F. Lindner, G. G. Paulus, H. Walther, A. Baltuška, E. Goulielmakis, M. Lezius, and F. Krausz. Gouy Phase Shift for Few-Cycle Laser Pulses. *Physical Review Letters*, 92(11):113001, Mar 2004.

[62] C. G. Durfee, A. R. Rundquist, S. Backus, C. Herne, M. M. Murnane, and H. C. Kapteyn. Phase matching of high-order harmonics in hollow waveguides. *Physical Review Letters*, 83(11):2187–2190, Sep 1999.

[63] T. Popmintchev, M.-C. Chen, O. Cohen, M. E. Grisham, J. J. Rocca, M. M. Murnane, and H. C. Kapteyn. Extended phase matching of high harmonics driven by mid-infrared light. *Optics Letters*, 33(18):2128–2130, 2008.

[64] A. Paul, E. A. Gibson, X. S. Zhang, A. Lytle, T. Popmintchev, X. B. Zhou, M. M. Murnane, I. P. Christov, and H. C. Kapteyn. Phase-matching techniques for coherent soft x-ray generation. *IEEE Journal of quantum electronics*, 42(1-2):14–26, Jan 2006.

[65] A Paul, R. A. Bartels, R. Tobey, H. Green, S. Weiman, I. P. Christov, M. M. Murnane, H. C. Kapteyn, and S. Backus. Quasi-phase-matched generation of coherent extreme-ultraviolet light. *Nature*, 421(6918):51–54, Jan 2003.

[66] M. Landreman, K. O'Keeffe, T. Robinson, M. Zepf, B. Dromey, and S. M. Hooker. Comparison of parallel and perpendicular polarized counterpropagating light for suppressing high harmonic generation. *Journal of the Optical Society of America B*, 24(9):2421–2427, 2007.

[67] X. Zhang, A. L. Lytle, T. Popmintchev, X. Zhou, H. C. Kapteyn, M. M. Murnane, and O. Cohen. Quasi-phase-matching and quantum-path control of high-harmonic generation using counterpropagating light. *Nature Physics*, 3(4):270–275, Apr 2007.

[68] O. Cohen, A. L. Lytle, X. Zhang, M. M. Murnane, and H. C. Kapteyn. Optimizing quasi-phase matching of high harmonic generation using counterpropagating pulse trains. *Optics letters*, 32(20):2975–2977, Oct 2007.

[69] B. Dromey, M. Zepf, M. Landreman, and S. M. Hooker. Quasi-phasematching of harmonic generation via multimode beating in waveguides. *Optics Express*, 15(13): 7894–7900, 2007.

[70] C. A. Froud, R. T. Chapman, E. T. F. Rogers, M. Praeger, B. Mills, J. Grant-Jacob, T. J. Butcher, S. L. Stebbings, A. M. de Paula, J. G. Frey, and W. S. Brocklesby. Spatially resolved Ar* and Ar+* imaging as a diagnostic for capillary-based high harmonic generation. *Journal of Optics A: Pure and Applied Optics*, 11(5):054011 (7pp), 2009.

[71] E. A. J. Marcatili and R. A. Schmeltzer. Hollow metallic and dielectric waveguides for long distance optical transmission and lasers. *Bell Systems Technical Journal*, 43:1783–1809, 1964.

[72] R. K. Nubling and J. A. Harrington. Launch conditions and mode coupling in hollow-glass waveguides. *Optical Engineering*, 37:2454–2458, 1998.

[73] UDT Sensors. X-ray photodiode datasheet. "http://www.osioptoelectronics.com/productimages/X-ray-Photodiodes.pdf", 1990.

[74] F. R. Powell. Care and feeding of soft x-ray and extreme ultraviolet filters. *24th Annual Boulder Damage Symposium Proceedings – Laser-Induced Damage in Optical Materials: 1992*, 1848(1):503–515, 1993.

[75] N. L. Wagner, E. A. Gibson, T. Popmintchev, I. P. Christov, M. M. Murnane, and H. C. Kapteyn. Self-compression of ultrashort pulses through ionization-induced spatiotemporal reshaping. *Physical Review Letters*, 93:173902 1–4, 2004.

[76] C. Spielmann, C. Kan, N. H. Burnett, T. Brabec, M. Geissler, A. Scrinzi, M. Schnurer, and F. Krausz. Near-keV coherent X-ray generation with sub-10-fs lasers. *IEEE Journal of Selected Topics in Quantum Electronics*, 4(2):249–265, Mar 1998.

[77] Gibson E. A. *Quasi-Phase Matching of Soft X-ray Light from High-Order Harmonic Generation using Waveguide Structures*. PhD thesis, Department of Physics, University of Colorado, Boulder, 2004. URL http://jilawww.colorado.edu/pubs/thesis/gibson/gibson_thesis.pdf.

[78] M. Born and E. Wolf. *Principles of optics : electromagnetic theory of propagation, interference and diffraction of light*. Cambridge University Press, Cambridge, 7th expanded ed edition, 1999. ISBN 9780521642224.

[79] S. F. Gull and T. J. Newton. Maximum-entropy tomography. *Applied Optics*, 25: 156–160, 1986.

[80] D. H. Sutter, G. Steinmeyer, L. Gallmann, N. Matuschek, F. Morier-Genoud, U. Keller, V. Scheuer, G. Angelow, and T. Tschudi. Semiconductor saturable-absorber mirror-assisted kerr-lens mode-locked ti:sapphire laser producing pulses in the two-cycle regime. *Optics Letters*, 24:631–633, 1999.

[81] B. Schenkel, R. Paschotta, and U. Keller. Pulse compression with supercontinuum generation in microstructure fibers. *Journal of the Optical Society of America B-Optical Physics*, 22(3):687–693, Mar 2005.

[82] M. Nisoli, S. DeSilvestri, and O. Svelto. Generation of high energy 10 fs pulses by a new pulse compression technique. *Applied Physics Letters*, 68(20):2793–2795, May 1996.

[83] L. Gallmann, T. Pfeifer, P. M. Nagel, M. J. Abel, D. M. Neumark, and S. R. Leone. Comparison of the filamentation and the hollow-core fiber characteristics for pulse compression into the few-cycle regime. *Applied Physics B-Lasers and Optics*, 86: 561–566, 2007.

[84] A. Couairon, J. Biegert, C. P. Hauri, W. Kornelis, F. W. Helbing, U. Keller, and A. Mysyrowicz. Self-compression of ultra-short laser pulses down to one optical cycle by filamentation. *Journal of Modern Optics*, 53(1-2):75–85, Jan 2006.

[85] X. Liu, L. J. Qian, and F. Wise. High-energy pulse compression by use of negative phase shifts produced by the cascade $\chi^{(2)} : \chi^{(2)}$ nonlinearity . *Optics Letters*, 24 (23):1777–1779, Dec 1999.

[86] M. Kim and C. S. Yoon. A new third-order nonlinear optical process. In *Nonlinear Optics: Materials, Fundamentals, and Applications, 2000. Technical Digest*, pages 68–70. Optical Society of America, 2000.

[87] J. A. Moses, F. W. Wise, L. Vuong, T. Grow, A. Gaeta, J. Nees, K. H. Hong, B. Hou, and G. Mourou. Chirped-pulse cascaded quadratic compression of millijoule-energy, 35-fs pulses. In *Lasers and Electro-Optics, 2005. (CLEO). Conference on*, volume 2, pages 962–964. Optical Society of America, May 2005.

[88] J. Moses and F. W. Wise. Controllable self-steepening of ultrashort pulses in quadratic nonlinear media. *Physical Review Letters*, 97(7):073903 1–4, Aug 2006.

[89] J. Moses and F. W. Wise. Soliton compression in quadratic media: high-energy few-cycle pulses with a frequency-doubling crystal. *Optics Letters*, 31(12):1881–1883, Jun 2006.

[90] S. L. Stebbings, E. T. F. Rogers, A. M. de Paula, M. Praeger, C. A. Froud, B. Mills, D. C. Hanna, J. J. Baumberg, W. S. Brocklesby, and J. G. Frey. Molecular variation of capillary-produced soft x-ray high harmonics. *Journal of Physics B-Atomic Molecular and Optical Physics*, 41(14), Jul 2008.

[91] E. A. Gibson, A. Paul, N. Wagner, R. Tobey, S. Backus, I. P. Christov, M. M. Murnane, and H. C. Kapteyn. High-order harmonic generation up to 250 ev from highly ionized argon. *Physical Review Letters*, 92(3):033001+, Jan 2004.

[92] A. Mysyrowicz, A. Couairon, and U. Keller. Self-compression of optical laser pulses by filamentation. *New Journal of Physics*, 10(2):025023 (14pp), 2008.

[93] J. Seres, P. Wobrauschek, C. Streli, V. S. Yakovlev, E. Seres, F. Krausz, and C. Spielmann. Generation of coherent kev x-rays with intense femtosecond laser pulses. *New Journal of Physics*, 8(10), Oct 2006.

[94] G Tempea, M Geissler, M Schnurer, and T Brabec. Self-phase-matched high harmonic generation. *Physical Review Leters*, 84(19):4329–4332, May 2000.

[95] E. A. Gibson, A. Paul, N. Wagner, R. Tobey, D. Gaudiosi, S. Backus, I. P. Christov, A. Aquila, E. M. Gullikson, D. T. Attwood, M. M. Murnane, and H. C. Kapteyn. Coherent soft x-ray generation in the water window with quasi-phase matching. *Science*, 302(5642):95–98, 2003.

[96] M. Zepf, B. Dromey, M. Landreman, P. Foster, and S. M. Hooker. Bright quasi-phase-matched soft-x-ray harmonic radiation from argon ions. *Physical Review Letters*, 99(14):143901, 2007.