# Computed Tomography of Chemiluminescence: A 3D Time Resolved Sensor for Turbulent Combustion

Jeremy Floyd

Thesis submitted to
Imperial College London
for the degree of
Doctor of Philosophy

October 28, 2009

Department of Mechanical Engineering

## Abstract

Time resolved 3D measurements of turbulent flames are required to further understanding of combustion and support advanced simulation techniques (LES). Computed Tomography of Chemiluminescence (CTC) allows a flame's 3D chemiluminescence profile to be obtained by *inverting* a series of integral measurements. CTC provides the *instantaneous* 3D flame structure, and can also measure: excited species concentrations, equivalence ratio, heat release rate, and possibly strain rate. High resolutions require simultaneous measurements from many view points, and the cost of multiple sensors has traditionally limited spatial resolutions. However, recent improvements in *commodity* cameras makes a high resolution CTC sensor possible and is investigated in this work.

Using realistic LES *Phantoms* (known fields), the CT algorithm (ART) is shown to produce low error reconstructions even from limited noisy datasets. Error from *self-absorption* is also tested using LES Phantoms and a modification to ART that successfully corrects this error is presented. A proof-of-concept experiment using 48 non-simultaneous views is performed and successfully resolves a Matrix Burner flame to 0.01% of the domain width ($D$). ART is also extended to 3D (without *stacking*) to allow 3D camera locations and optical effects to be considered. An *optical* integral geometry (weighted double-cone) is presented that corrects for limited depth-of-field, and (even with poorly estimated camera parameters) reconstructs the Matrix Burner as well as the standard geometry.

CTC is implemented using five PicSight P32M cameras and mirrors to provide 10 simultaneous views. Measurements of the Matrix Burner and a Turbulent Opposed Jet achieve exposure times as low as 62 $\mu$s, with even shorter exposures possible. With only 10 views the spatial resolution of the reconstructions is low. However, a cosine Phantom study shows that 20–40 viewing angles are necessary to achieve high resolutions (0.01–0.04$D$). With 40 P32M cameras costing £40000, future CTC implementations can achieve high spatial and temporal resolutions.

# Acknowledgements

I would like to begin by thanking my supervisor, Dr Andreas Kempf, for all the time and effort he has afforded me and also for his boundless enthusiasm for this project. I would also like to thank Prof Peter Lindstedt for much good advice, and for his help with the excited state chemistry along with Drs Kostas Gkagkas and Roger Robinson. I particularly would like to thank Andreas and Peter again for their understanding during some particularly hard months for my family and I at the end of 2008.

Undertaking an experimental project with such limited funds for experiments would have been much more difficult were it not for the help of many generous people in the Department of Mechanical Engineering at Imperial College. I am extremely grateful to Dr Andy Heyes, Dr Joerg Feist and Prof Peter Lindstedt for the loan of much equipment and lab space, and I would also like to thank Drs Nikos Soulopoulos and 'George' Charalampous for much helpful advice. I am very grateful to the Department of Mechanical Engineering for the additional funding to buy the cameras, and would also like thank Guljar Singh and Vim Patel for their help setting up the experiments.

I am particularly indebted to my 'belayer' Philipp Geipel for his help with all the experiments in this project, and without whom none of the measurements performed in this work would have been possible. I am also extremely grateful to my department for funding my bursary through an EPSRC DTA award.

I would like thank all my friends from rooms 503 and 600, it has been a true pleasure to spend these last few years with so many good fellows. I am particularly grateful to 'Lord' Stein for his help with LES, and also for helping to proof read this document along with Nick Vaughan, Sandeep Saha, Michele Bonanni and my neighbour Pankaj Vaishnavi.

I would like to thank my beloved wife Wendy, for putting up with an absent husband, and for proof reading the longest chapter. Finally, I would also like to thank my parents for their love and support over the last 28 years, and I wish to thank my father in particular, for telling me to *go for it*.

*. . . this one's for you Dad.*

# Contents

**List of Main Symbols** **22**

**1 Introduction** **26**
1.1 The Need for 3D Flame Data . . . . . 26
1.2 Computed Tomography of Chemiluminescence . . . . . 30
1.3 Aims . . . . . 33
1.4 Document Outline . . . . . 34

**2 Background Theory** **36**
2.1 Turbulent Combustion . . . . . 36
2.1.1 Turbulence . . . . . 36
2.1.2 Combustion Modes . . . . . 38
2.1.3 Reaction Kinetics and Mechanisms . . . . . 40
2.1.4 Chemiluminescence . . . . . 41
2.2 Computed Tomography . . . . . 43
2.2.1 Reconstructing from Projections . . . . . 44
2.2.2 An Overview of Algorithms . . . . . 46
2.2.3 ART in Greater Detail . . . . . 51
2.3 Large Eddy Simulation of Combustion . . . . . 55
2.3.1 The Filtered Mass and Momentum Equations . . . . . 55
2.3.2 Modelling Combustion in LES . . . . . 56

**3 Computed Tomography of Chemiluminescence** **60**
3.1 The Sensor Concept . . . . . 60
3.2 Prior Art . . . . . 62
3.2.1 Focus on Emission Computed Tomography . . . . . 63
3.2.2 Focus on the Algorithms . . . . . 64
3.2.3 Resolutions achieved . . . . . 65
3.2.4 Conclusions . . . . . 68
3.3 The Basic ART Implementation . . . . . 69

**4 Testing Sensor Capability** **72**
4.1 LES Phantoms . . . 72
4.1.1 Simulating Excited State Species . . . 74
4.1.2 LES Configurations and Results . . . 75
4.1.3 The Excited Species Phantoms . . . 78
4.2 Reconstruction from limited angles . . . 79
4.2.1 LES Phantom Tests . . . 81
4.2.2 The Effect of Noise . . . 87
4.3 Resolution Achievable with Limited Angles . . . 89
4.3.1 Existing Theory . . . 90
4.3.2 Cosine Phantom Tests . . . 92
4.4 Proof of Concept Experiment . . . 98
4.4.1 Experimental Set-Up . . . 98
4.4.2 Matrix Burner Results . . . 100
4.4.3 Conclusion . . . 102

**5 Improving ART for CTC** **104**
5.1 Self absorption Study . . . 104
5.1.1 Investigating Error from Self absorption . . . 105
5.1.2 Modifying ART for Self Absorption . . . 111
5.1.3 Conclusions . . . 113
5.2 Enhancements to ART . . . 115
5.2.1 Nested Grids . . . 115
5.2.2 Linear Interpolation on the Forward Projection . . . 116
5.3 Extending ART to 3D . . . 118
5.3.1 3D projections . . . 119
5.3.2 Notes on the Computer Implementation . . . 120
5.3.3 Phantom Test . . . 121
5.4 Accounting for Optical Effects . . . 122
5.4.1 Projections for Insufficient Depth-of-Field . . . 123
5.4.2 Non-Parallel Projections . . . 132
5.4.3 Testing Optical Projections using the Matrix Burner . . . 133
5.4.4 Conclusion . . . 140

**6 Experimental Results and Discussion** **141**
6.1 Camera Selection . . . 141
6.1.1 Selection Criteria . . . 142
6.1.2 Chosen Camera: PicSight P32M . . . 143

6.2 Time Resolved Matrix Burner . . . 147
6.2.1 Specific CTC Set-Up . . . 148
6.2.2 Camera Registration . . . 149
6.2.3 Results . . . 153
6.3 Improving Camera Registration . . . 160
6.3.1 View Registration using Camera Calibration . . . 161
6.3.2 Registration Approach comparisons . . . 165
6.3.3 Lens Distortion . . . 168
6.4 Impinging Matrix Burner . . . 171
6.5 Premixed Turbulent Opposed Jet . . . 173
6.5.1 Experimental Set-Up . . . 174
6.5.2 Phantom Studies . . . 181
6.5.3 Experimental Results . . . 190
6.5.4 Conclusions . . . 196

**7 A Novel Application for CTC** **199**
7.1 The Measurement of Flame Local Strain Rate . . . 199
7.1.1 Laminar Opposed Jet Simulation . . . 201
7.1.2 Laminar Opposed Jet Experiment . . . 203
7.1.3 Discussion and Conclusions . . . 205

**8 Conclusion** **206**
8.1 Project Summary . . . 206
8.2 Review of Main Conclusions . . . 209
8.3 Recommendations . . . 213
8.3.1 Specific Recommendations . . . 214
8.3.2 Areas for Further Research . . . 215

**Bibliography** **217**

**A Technical Drawings** **232**

# List of Figures


2.1 Example projections at two different angles for a two cylinder object. . . 44

2.2 Example projections based on line integrals. Note how multiple projections at a common angle comprise a view. . . . . 45

2.3 A pixels contribution $w_{qpv}$ to projection $qp$ is the intersection of the projection geometry and the pixel. The object domain has been discretised with $N_v$ pixels. . . . . 48

3.1 The generic CTC set-up. $N_q$ cameras are located around a burner at radius $r$. The angle of each camera is $\theta$. . . . . 61

3.2 Resolutions achieved in other selected studies. Crosses represent laser absorption measurements, filled shapes are emission measurements, and unfilled shapes refractive index measurements. The resolved wavelength $\lambda_{res}$ has been normalised by the object domain diameter $D$ to allow comparisons. Also shown is a theoretical relationship obtained from sampling theory. . . . . 66

3.3 The computational method used to determine the 2D fractional area, $w_{qpv}$, of pixel $v$. The fractional area is equal to the area to the right of vector $AB$ ($a_{AB}$) minus (because of negative gradient) the area to the right of vector $CD$ ($a_{CD}$). The area to the right of a vector is determined from the local pixel intersection points $S_{in} = [x_{in}, y_{in}]$ and $S_{out} = [x_{in}, y_{in}]$ using eq. 3.1. . 70

3.4 Example reconstructions of the Flame D Phantom using ART and MART. The reconstructions used 128 evenly spaced angles with 128 projections per angle. . . . . 71

4.1 An axial section of the Flame D instantaneous mixture fraction field. Also shown are two downstream locations. D is the nozzle diameter of 7.2 mm. 76

4.2 Radial profiles of mixture fraction $\xi$ and OH mass fraction $Y_{OH}$ taken from the LES alongside experimental data for an axial location 15D downstream of the nozzle. . . . . 77

4.3 An axial section of the Turbulent Opposed Jet instantaneous temperature field (left) and mixture fraction (right). The top is the fuel side and these sections span 30 mm axially. . . . . 78

4.4 Axial profiles of mixture fraction $\xi$ and OH mass fraction $Y_{OH}$ taken from the LES alongside experimental data where available. . . . . 79

4.5 The 128 pixel, LES excited species Phantoms. For Flame D the section was taken 30 nozzle diameters downstream of the jet exit. For the Turbulent Opposed Jet (TOJ) the section was taken at the stagnation plane of the two jets. . . . . 80

4.6 The correlation coefficient,$r_{xy}$ of the Phantom and reconstructions for different numbers of views, $N_q$. Results for both Phantoms using each algorithm are shown. . . . . 82

4.7 The normalised error RMS, $E_{RMS}$, of the reconstruction for different numbers of views, $N_q$. Results for both Phantoms using each algorithm are shown. . . . . 83

4.8 Example ART and MART reconstructions of the Flame D Phantom, which is shown in fig. 4.5. . . . . 84

4.9 Example ART and MART reconstructions of the TOJ Phantom, which is shown in fig. 4.5. . . . . 85

4.10 The normalised RMS error of reconstructions of the TOJ Phantom for two different values of $\Delta_c$. . . . . 86

4.11 Projection noise amplification. . . . . 88

4.12 The theoretically achievable resolution for a given number of equally spaced angles, $N_q$. The derivation of the Fourier Slice Theorem relationship, $\lambda_{res} = D\pi/N_q$, can be found in references [112, 43]. The resolution, $\lambda_{res}$, has been normalised by the domain diameter/width $D$. Also shown is the resolution achieved by a fully determined ART reconstruction ($2\sqrt{2}/N_q$), and that predicted by Frieder and Herman [48]. . . . . 91

4.13 Examples Phantoms showing the different locations used for the wave centre. These examples have $256 \times 256$ pixels and wavelengths of 30 pixels. 92

4.14 An example of the resolution judging process, for $N_q = 20$ and Phantom family 256TL. At $\lambda = 20$ aliasing artefacts can be clearly observed, and at $\lambda = 35$ small artefacts are still seen between the outer rings. By $\lambda = 40$ the rings are distinct from each other and reconstruction is judged to be resolved at that wavelength. . . . . 93

4.15 An example of the *levelling off* in the cross-correlation, $r_{xy}$, observed as $\lambda$ increases. The $N_q = 20$ case is shown for the Phantoms described in table 4.2. . . . . 94
4.16 The normalised resolved wavelength achieved, $\lambda_{res}/D$, for different numbers of equally spaced views, $N_q$. The resolved wavelength is determined from ART reconstructions of radial cosine Phantoms. Shown are results for the Phantoms detailed in table 4.2. The continuous lines depict the theoretical result based on the Fourier Slice Theorem (FST) and the relation given by Frieder and Herman [69]. . . . . 95
4.17 Example reconstructions of the three 256 pixel Phantom types, at $\lambda = 20$ and $N_q = 20$. . . . . 95
4.18 Example of the 256SU Phantom. $\lambda = 5$. . . . . 96
4.19 The normalised resolved wavelength, $\lambda_{res}/D$, for different numbers of equally spaced views, $N_q$. The resolved wavelength is determined from ART reconstructions of radial cosine Phantoms. Shown are results for the Phantoms detailed in table 4.2 and also for Phantom type 256SU. The continuous lines depict the theoretical result based on the Fourier Slice Theorem (FST) and the relation given by Frieder and Herman [69] . . . . 97
4.20 The specific CTC experimental set-up used in this study. The camera is repositioned to different angles to provide multiple views. . . . . 99
4.21 *Left:* Photograph of the Matrix Burner showing the angles of views 12 (A), 19 (B) and 24 (C) seen in fig. 4.22. *Right:* Drawing of the Matrix burner. . . . . 99
4.22 Example measured views. The angles of view 12 (top), 19 (middle), and 24 (bottom) can be seen in fig. 4.21. . . . . 101
4.23 Sections from the reconstruction of the Matrix burner using 48 angles. . . 101
4.24 A volume rendering of the 3D reconstruction (left) of the $CH^*$ chemiluminescent light intensity alongside a picture (right) of the flame taken at a similar angle. The reconstruction used 48 angles with 225 projections per angle. . . . . 102
4.25 (Left) A *raw* 460 by 460 pixel section showing the $CH^*$ chemiluminescent light intensity. This section was reconstructed from 48 views each 450 projections wide. (Right) The same section reconstructed from views filtered to 225 projections wide. . . . . 103

5.1 The locations of sections used as Phantom data for Flame D. . . . . 106
5.2 The normalised Standard Deviation error $E_\sigma$ in reconstructions using ART for increasing values of nominal absorption coefficient. . . . . 108

5.3 The normalised Standard Deviation error $E_\sigma$ in reconstructions using ART for increasing values of nominal absorption coefficient and self absorption. Results are shown for the $OH^*$ Phantoms taken from different axial locations in Flame D, as shown in fig. 5.1. . . . . 109

5.4 The mass fractions of the excited and ground state species from LES. Quoted below are the mean mass fractions $\overline{Y}$ for that species. . . . . 110

5.5 The Flame D $OH^*$ Phantom (a) and its reconstructions using SMART (b), and ART (c) when the projections are subjected to strong self-absorption ($\alpha = 10^5$). These example reconstructions were performed on $128^2$ pixel grids using 128 angles and 128 rays. . . . . 112

5.6 The TOJ $OH^*$ Phantom (a) and its reconstructions using SMART (b), and ART (c) when the projections are subjected to strong self-absorption ($\alpha = 10^5$). These example reconstructions were performed on $128^2$ pixel grids using 128 angles and 128 rays. . . . . 112

5.7 The normalised standard deviation error $E_\sigma$ in reconstructions using ART and SMART for increasing values of nominal absorption coefficient. Results are shown for the Flame D $OH^*$ and $CH^*$ Phantoms. . . . . 113

5.8 The normalised standard deviation error $E_\sigma$ in reconstructions using ART and SMART for increasing values of nominal absorption coefficient and self absorption. Results are shown for the TOJ $OH^*$ and $CH^*$ Phantoms. . 114

5.9 Convergence time comparison for ART and ART using nested grids, for different reconstruction grid sizes. . . . . 115

5.10 An interpolation cell. Diamonds represent object/Phantom sample points. $\mathbf{S}_1$ and $\mathbf{S}_2$ are vectors defined by local co-ordinates. . . . . 116

5.11 The error in the RMS of reconstructions using ART (squares) and intART algorithms (diamonds). All algorithms used the same convergence and relaxation parameters. . . . . 118

5.12 the 3D coordinate system. . . . . 119

5.13 Example sections from the co-planar and inclined reconstructions as well as the Phantom. Shown is a section in the $xy$ plane at $z = 16$ (middle) along with central sections in the $yz$ (right) and $xz$ (below) planes. . . . . 122

5.14 Volume renderings of the co-planar and inclined reconstructions, and the Phantom, taken from a view 30°above the $xy$ plane and at an angle of 30°to the $y$ axis. . . . . 123

5.15 A diagram illustrating in 2D how ray-tracing describes the propagation of a spherical light wave diverging from point $v$ in the object domain, through a circular lens and aperture, and converging back to a point $p$ on the image plane. The double cone defined by the *marginal* rays (thick continuous lines) is formed about the central ray (dashed line). Note the diagram is not to scale, $f$ is the focal length of the lens. . . . . 125

5.16 Rays traced back from the pixel limits (thick and dashed lines) expand the double cone of fig. 5.15 to the shaded area as shown. . . . . 127

5.17 An illustration of the depth-of-field determined by ray-tracing. Points of light originating away from the object plane at $S_o$ are spread over a circle of diameter $D_{blur}$ on the image plane at $S_i$. The depth-of-field is determined from the near and far limits when $D_{blur}$ reaches the maximum acceptable size termed the circle-of-confusion. . . . . 128

5.18 Diagram illustrating different intersection areas formed between a circular pixel and a blur circle whose centres are separated by a distance $l$. . . . . 130

5.19 Comparison of the exact equation for the intersection area and the linear approximation used. The intersection area $A_i$ has been normalised by the pixel area and the separation $l$ of the centres of the two circles is shown as a multiple of $r_p$: the pixel radius. . . . . 131

5.20 Illustration of the non-parallel light rays. . . . . 132

5.21 Example views taken with insufficient depth-of-field. Shown (from top left) are views: 1, 6, 11, 16, 21, 27, 33, 39 and 45. . . . . 134

5.22 Shown are examples of each of the four projection types used. Also shown is the logarithm of the optical projection (e). For each example three projections, $m = 32$, $m = 64$ and $m = 96$, from view 1 ($\theta = 0$) have been composited together. . . . . 136

5.23 Sections from the reconstructions of the Matrix Burner using the four different projections (a,b,d,e). Also shown in (c) is the difference of the sections in (a) and (b). The sections are taken from the $xy$ plane at intercept $k = 29$ (upper) and in the $xz$ plane at $j = 61$. The intensities in (d) have been scaled by 4.2 for clarity. . . . . 137

5.24 Volume renderings of the Matrix Burner reconstructions for each projection type. The viewer is at an angle of 53° to the $xy$ plane. The intensities in (c) have been scaled by 7.3 for clarity. . . . . 138

5.25 Volume renderings of the Matrix Burner reconstructions for each projection type. Also shown in (c) is the difference of the renderings (a) and (b). The viewer is perpendicular to the *xy* plane. The intensities in (d) have been scaled by 3.9 for clarity. . . . . 139

6.1 An example SNR test shot using an exposure of 500 $\mu$s and a gain of 30 dB. 145
6.2 The variation of SNR with the normalised signal RMS for an exposure of 1 ms and different gain values. . . . . 145
6.3 The variation of noise with the normalised signal RMS for four different exposure lengths. The noise is taken as the variance of the sample area. . 146
6.4 The variation of SNR with the normalised signal RMS for four different exposure lengths. . . . . 147
6.5 A scale diagram of the mirror arrangement, the lens is located 600 mm from the burner centre. The dashed lines show the marginal rays of points on the edge of the domain and define the required size of the mirrors. Also shown are the locations of the two *virtual* burners seen by the camera. 149
6.6 The CTC experimental set-up. The five cameras are positioned 36 degrees apart at a radius of 600 mm from the burner (measured to the lens). The mirrors are used to provide two views for each camera, with an example ray path shown by the dashed line. Each camera is connected to the controlling PC and an electronic external trigger source. All mirrors and cameras are mounted to a single 10 mm thick aluminium plate. . . . . 150
6.7 Example registration image for view 10. The registration object sits atop the burner, so that their central axes are in line. . . . . 151
6.8 The synthetic Phantom alongside a reconstruction with ten views and no error in the viewing angles. . . . . 152
6.9 The correlation coefficient $r_{xy}$ of the Phantom and reconstructions subjected to increasing random error in the viewing angle. . . . . 152
6.10 Reconstructions with 1° (*left*) and 2.3° (*right*) peak error in the view angle. 153
6.11 The measured view 7 for each exposure used. The aperture is f/1.6. The images have been normalised by the maximum value of (a). . . . . 154
6.12 Volume renderings of the reconstruction of each exposure. The viewing angle is 53 degrees from horizontal. . . . . 155
6.13 Shown is a section of each of the reconstructions taken through the centre (plane normal [1 0.2 0]). . . . . 156
6.14 Shown is a section of each of the reconstructions at $z = 25$. . . . . 157
6.15 A volume rendering and section of the reconstruction of the 1000 $\mu$s exposure views using the $\gamma$ values in table 6.4. . . . . 158

6.16 Pinhole camera model coordinates. . . . . 162
6.17 example calibration image. . . . . 164
6.18 Diagram of the triangulation approach in the *xy* plane. . . . . 165
6.19 A vector plot of the distortion of camera 1. The vectors have been scaled by 2 for clarity. Also shown is a calibration image taken with camera 1 that has been appropriately undistorted. Please note that the vector plot is shown inverted. . . . . 169
6.20 The difference in the viewing angle of each pixel, with and without distortion for camera 1. . . . . 169
6.21 The 5 measured views of the impinging Matrix Burner. View 1 is shown at the top and view 5, at the bottom. . . . . 171
6.22 Volume renderings of the Impinging flow reconstructions. The viewing angles of each rendering are shown below the images. . . . . 172
6.23 Sections of the impinging flame reconstruction. The section indices are shown below the images. . . . . 173
6.24 Diagram of the Turbulent Opposed Jet. . . . . 174
6.25 A scale diagram of the TOJ mirror arrangement, the lens is located 1674 mm from the burner centre. The dashed lines show the marginal rays of points on the edge of the domain and define the required size of the mirrors. The diagram has been truncated to fit, as indicated by the parallel curved lines. 176
6.26 Scale drawing of the ideal view arrangement around the TOJ. The dashed lines represent the limiting marginal rays and the continuous line the central ray of the view. Only the rays from the outer mirrors are shown. This ideal CTC set-up is prevented by the upright supports of the burner, which obscure views 3, 4, 7, and 8. . . . . 177
6.27 Scale drawing of the *intended* view arrangement around the TOJ (note, the arrangement is not symmetrical). The dashed lines represent the limiting marginal rays and the continuous line the central ray of the view. Only the rays from the outer mirrors are shown. . . . . 178
6.28 Example registration images of the TOJ set-up. . . . . 179
6.29 Views of the 3D FSD Phantom from the registered angles in table 6.6. View geometries have been scaled by two. . . . . 182
6.30 Volume renderings of the FSD Phantom (top) and its reconstruction using 10 unequally spaced views (middle), and 10 equally spaced views (bottom).183
6.31 Selected sections of the FSD Phantom (top and left) and its reconstruction using 10 unequally spaced views (middle), and 10 equally spaced views (bottom and right). . . . . 184

6.32 Reconstructions of the 2D cosine Phantoms using the experimental unequal viewing angles. Three different Phantom wavelengths are shown which illustrate how the resolved wavelength is determined. . . . . 185

6.33 The correlation between filtered reconstructions and the similarly filter TOJ 2D Phantom. The cut-off wavelength, $\lambda_{co}$, has been normalised by the domain diameter $D$. . . . . 186

6.34 The 2D TOJ Phantom ($k = 5$) and its reconstructions using equal angles (middle) and experimental angles (right). All have had wavelengths less than $0.13/D$ filtered out. . . . . 186

6.35 Sections of the 3D TOJ Phantom and its reconstructions using the experimental angles (middle) and with 1°systematic error in views 9 an 10 (right). The sections are at $k = 5$. . . . . 187

6.36 The correlation of the Phantom and its reconstruction using different filtering approaches at a range of cut-off wavenumbers. . . . . 189

6.37 View 6 of the TOJ at different exposures and equivalence ratios. Note each view has a different gain: (a) gain = 30, 24, 18 dB (exposure = 1000 $\mu$s); (b) gain = 18, 30, 33 dB ($\phi = 0.8$). . . . . 190

6.38 The column sum of the measured intensity ($\sum_n I_{mn}$) for the 3 truncated views: 3, 4, and 7. Also shown is an untruncated view 6 for comparison. . 191

6.39 Volume renderings of the TOJ at different equivalence ratios. From top: $\phi = 0.7$, $\phi = 0.8$; $\phi = 0.9$. . . . . 193

6.40 Selected mid-sections (upper, $z = 16$; lower, $y = 64$) of the TOJ reconstruction for different equivalence ratios. . . . . 194

6.41 Volume renderings of the TOJ for different exposure times. From top the exposures are: 2000, 500, 250 $\mu$s. . . . . 195

6.42 Selected mid-sections (upper, $z = 16$; lower, $y = 64$) of the TOJ reconstruction for different view exposures. . . . . 196

7.1 Illustration of how the local flame normal integral of chemiluminescence can be obtained once the spatial distribution is known. The spatial distribution of chemiluminescence can be provided for arbitrary flames by CTC. The line **AB** illustrates how line-of-sight measurements alone cannot always provide the local flame normal integral of chemiluminescence. The flame normal direction is $\omega$, $\xi_{st}$ is the stoichiometric mixture fraction, and $N_B$ is the number of emitting molecules of species B. . . . . 200

7.2 The integral of the mole fraction $X_B$ across the flame normal $\omega$ plotted against the bulk strain rate $a_b$ for the excited species $B$: $OH^*$, $CH^*$, $CHO^*$, $CH_2O^*$, $C_2^*$. The data results from the 1D solution off a Laminar Opposed Jet by Gkagkas [57] presented previously in section 4.1. . . 202
7.3 A schematic drawing of the experimental set-up used to measure the integral chemiluminescence of the Laminar Opposed Jet. . . . . 203
7.4 The measured intensity of the 430 nm wavelength light, emitted by decaying $CH^*$ radicals, against the bulk strain rate. Intensity values were taken 2 mm from the centre line of the Laminar Opposed Jet Burner. The mean of 30 samples is shown, with the (error) bars representing the maximum and minimum of the 30 samples. . . . . 204

# List of Tables

2.1 Examples of commonly studied chemiluminescent transitions in combustion. The ground state is prefixed with X and the first excited state with A. The symbols $\Sigma, \Delta$ etc. represent the orbital angular momentum and the preceding superscript the spin. . . . . 42

4.1 The reactions added to the mechanism of Lindstedt and coworkers [98, 130]. Units are J/mol and further details are given in Benvenutti et al. [17] 74
4.2 Cosine Phantom Parameters. . . . . 93

6.1 Test image settings . . . . . 145
6.2 The calibrated viewing angle for each view shown alongside the ideally spaced (designed) angle. View 1 is used as the reference value. . . . . 151
6.3 Experimental Parameters . . . . . 153
6.4 $\gamma$ estimates for each view . . . . . 158
6.5 Camera Calibration results. See the main text for discussion of the uncertainty. . . . . 166
6.6 The registered relative viewing angle for each view alongside the intended angles shown in fig. 6.27. View 1 is used as the reference value. . . . . 180
6.7 Experimental Parameters . . . . . 192

# List of Main Symbols

| | |
|---|---|
| $\alpha$ | A nominal absorption coefficient |
| $\chi$ | The scalar dissipation rate |
| $\varepsilon$ | A noise signal |
| $\kappa$ | Absorption coefficient |
| $\lambda$ | Wavelength |
| $\omega$ | The flame normal spatial coordinate |
| $\sigma$ | The standard deviation |
| $\upsilon$ | An index of the voxels that lie between the voxel under consideration and the domain edge |
| $\xi$ | The mixture fraction |
| $A$ | An area |
| $a_b$ | The bulk strain rate |
| $c$ | The concentration of an absorbing species |
| $D$ | Diffusion Coefficient |
| $E_\sigma$ | The error of the standard deviation |
| $E_{RMS}$ | The normalised RMS of the error field |
| $h$ | Planck's constant |
| $I_0$, $I_1$ | Light intensity before and after passing through an absorbing medium |
| $I_\omega$ | The flame normal integral of chemiluminescence |
| $I_{ev}$ | The light intensity emitted from the pixel $v$ before passing through an absorbing medium |

$I_{pv}$ The light intensity emitted from the pixel $v$ after passing through an absorbing medium as measured by projection $p$

$k$ Wavenumber

$l_A$ length travelled through an absorbing medium

$MW_B$ The molecular weight of species $B$

$N_B$ the number of emitting molecules of species $B$.

$P_B$ The partial pressure of species $B$

$r$ A radius

$r_{xy}$ Correlation coefficient

$v$ Frequency of the emitted light

$X_B$ The mole fraction of species $B$

$Y_B$ The mass fraction of species $B$

SNR The Signal-to-noise ratio

## Optics Specific Symbols

$\delta$ The angle between a projection and the middle projection ($X = 0$, $Y = 0$) of that view

$\eta$ The view resolution at the object plane in metres

$A_i$ The area of intersection between a pixel and the blur circle

$A_p$ The area of a pixel $p$

$A_{blur}$ The area of the blur circle on the image plane

$D_{blur}$ The diameter of the blur circle formed by and out of plane point on the image plane

$f$ The focal length of the lens

$G_{qpv}$ The fractional intersection volume of a voxel $v$ and the double cone projection of pixel $p$ of the view $n$

$l$ The separation distance of the centres of a circular pixel and a blur circle

$M_t$ The transverse magnification of a lens

$N$ The focal ratio or $f$-number

$S_1$ The distance on the object side from a point to the lens/aperture

$S_2$ The distance on the image side from the lens/aperture to the focal plane corresponding to a point a distance $S_1$ from from the lens/aperture on the object side

$S_i$ The distance from the image plane to the lens/aperture

$S_o$ The distance from the object plane to the lens/aperture

## Computed Tomography Specific Symbols

$\beta$ ART relaxation coefficient

$\Delta_c$ The minimum change in some criteria for convergence

$\gamma$ The view angle made with the object $xy$ plane

$\lambda_{res}$ The resolved wavelength of a reconstruction

$\mathbf{C}_{qp}$ A point on the central axis of a 3D projection in the object domain coordinates

$\mathbf{f}$ The discrete object vector. Vectorised from $f(i,j,k)$

$\mathbf{I}$ Vector of projection values

$\mathbf{U}_{qp}$ A unit vector in the direction of the 3D view

$\mathbf{V}$ The position vector of a voxel $v$ relative to $\mathbf{C}_{qp}$

$\mathbf{w}_{qp}$ The vector containing the contribution of all voxels to the projection $I_{qp}$

$\theta$ The view angle about the object $z$ axis measured from the $y$ axis

$D$ View/object width/diameter

$f(i,j,k)$ The discrete approximation of $f(x,y,z)$

$f(x,y,z)$ The continuous *object* function

$f_a(x,y)$ An approximation to the true object function $f(x,y)$

$f_p(i,j,k)$ The Phantom field

$f_s$ View sampling frequency

$f_v$ The value of vector $\mathbf{f}$ at voxel (or pixel) $v$

$h$ The ART iteration index

$i,j,k$ Object voxel indices

$I_{qp}$ Integral value of projection $qp$

$m,n$ View pixel indices

$N_p$ Number of projections per view

$N_q$ Number of views

$N_v$ The total number of voxels (or pixels) in the object domain

$N_{qp}$ The total number of views ($N_q \times N_p$)

$p$ A view pixel of indices $m,n$. And by extension the projection index

$q$ The view index

$r_{qp}$ The radius of a cylindrical projection

$v$ An object voxel of indices $i,j,k$

$V_r$ The radial component of $\mathbf{V}$ in the cylindrical projection reference frame

$W$ *Projection* matrix

$w_{qpv}$ The contribution of the voxel $v$ to projection $I_{qp}$

$X,Y,Z$ Coordinates in the view or camera reference frame

$x,y,z$ Coordinates in the object reference frame

# Chapter 1

# Introduction

Whether to allow planes to fly further, or to allow power stations to run more profitably, improving the efficiency of combustion has long been the concern of scientists and engineers. Recently this concern has been given fresh impetus by diminishing fossil fuel supplies and the need to dramatically reduce carbon dioxide emissions [72]. While alternative low $CO_2$ energy generating technologies exist, they will take time to deploy. In the short term, the combustion of fossil fuels will remain the primary source of energy and to minimise emissions and maximise fuel reserves ever more efficient combustion processes are required.

In the longer term, rising demand will mean combustion remains an important power source, particularly in areas where sufficient alternatives do not yet exist, as for example in aviation. However, to meet this demand without further contribution to climate change new combustion technologies will have to be developed, such as bio-fuels for aviation and *clean coal* for electricity generation. These new technologies present new challenges to combustion scientists, which can only be quickly and effectively met if we continue to improve our understanding of combustion processes, and develop ever better tools for their study and design.

## 1.1 The Need for 3D Flame Data

Our understanding of combustion processes has been greatly advanced by point, line and planar experimental measurements, and in particular by recent high precision laser based techniques, such as Particle Image Velocimetry (PIV) [47, 97] and Raman/Rayleigh spectroscopy [51, 161]. However, practical combustion processes are nearly always turbulent and as such are characterised by unsteady 3D fluctuations occurring over a wide range of scales. The inherent unsteady and 3D nature of turbulent combustion means that time resolved 3D data must offer deeper insight into turbulent combustion. For example,

turbulence-chemistry interactions are 3D and fundamental to flame instability; in their investigations of these interactions Barlow and coworkers [12, 78] improved their—very advanced—multi-scalar measurements through extra measurements that determined the time resolved 3D flame orientation. The instantaneous 3D structure of a flame is also of great importance to the fundamental understanding and later modelling of combustion, as described by Bilger et al. [22]. Furthermore, in striving for ever greater efficiency we are pushed to more sophisticated burners and complex flows; for example, that feature swirl or acoustic oscillations. In such flows their 3D nature becomes even more important and 3D data from diagnostics and simulations is needed for their design and assessment.

Although 3D data is needed, it is not easy to physically measure. This is partly due to equipment limitations, but is also due to fundamental issues of just how such information can be gained. For example, filling a flow volume with intrusive thermocouples for instantaneous temperature measurements is hardly desirable, nor even practical for many combustors. Planar, line, and point measurements can provide 3D information by performing measurements sequentially at different points in the flow volume, but this approach can be very time consuming and is usually at the expense of temporal resolution. While we are somewhat forced, by the wide range of scales, to consider turbulent combustion statistically, the instantaneous spatial distribution of a flame *is* of fundamental importance and so time resolved data is needed.

Comprehensive time resolved 3D data, of all scales of a reacting flow, can be obtained from whole volume simulations based on the direct solution of fundamental governing equations. However, this Direct Numerical Simulation (DNS) approach requires such vast computational resources that in practice only very simple flows with low Reynolds number can be considered. One common approach to overcome this limitation is by time averaging the governing equations and solving for mean quantities on much coarser grids. Aside from destroying the time resolved information, the averaging introduces terms that represent the effect of turbulence at all scales and these require modelling. These RANS (Reynolds Averaged Navier-Stokes[1]) simulations depend greatly on the models employed and numerous types have been suggested. While these models can perform very well for the specific flows they have been calibrated for, they are less suitable for generic and complex flows.

Large Eddy Simulation (LES) is a promising technique that can provide time resolved 3D data of reacting flows of practical interest. By filtering out the smallest scales and instead resolving only the larger and most relevant scales, LES can be performed on a much coarser and so less computationally expensive grid than DNS. Furthermore, because a

[1]This name stems from the type of time averaging employed—Reynolds—and the common set of governing equations—the Navier-Stokes equations

significant portion of the turbulence is already resolved, LES allows very simple models to be used. As with RANS approaches, the equations governing chemical reaction cannot practically be solved in LES and models for combustion are employed. Although, as noted by Pitsch [124], one of the principle advantages of LES over RANS in this respect is the better prediction of the scalar mixing—which is of great importance in combustion modelling. The application of LES to turbulent combustion is relatively new and much attention is given to the modelling of combustion, but despite its relative immaturity LES has been applied to several reacting flows of practical interest. Nevertheless, the full predictive potential of LES has yet to be reached [124], and one of the reasons for this is the need for more experimental data to validate the large amount of 3D information predicted by LES.

Currently, the validation of LES generally relies on point, line and planar measurements of flames that can be well represented by such measurements. For example, opposed jets or single jet flames. Indeed much progress has been made from comprehensive datasets of such flames as seen, for example, in the International Workshop on Measurement and Computation of Turbulent Non-Premixed Flames (TNF) [9]. However, with more complex flames the 3D, time dependent nature of the flow becomes more important, and testing the ability of LES to predict such flows is difficult using 1D or 2D measurements alone. The need for 3D data is even more evident when transient 3D phenomena such as flame blow-off, or vortex break down are present. Even for simple flows time dependent 3D information provides additional avenues for validation, such as the qualitative checks of the flame structure [80] and possibly quantitative comparisons of 3D spatial correlations. To summarise: LES contains much additional time dependant 3D information, which can be (or should be) used for additional validation, and thus complimentary 3D experimental data is needed to aid the development of LES.

Despite the difficulties, several methods have been used to obtain 3D data of flames. Abel transforms are commonly used to obtain 3D data from single line-of-sight measurements of axisymmetric flows. However, such axisymmetry is only realistically encountered with time averaging of simple flows. Crossed Planar Laser Induced Fluorescence (PLIF) [87, 78, 12] can locate the instantaneous 3D flame normal at a single point, and stereoscopic imaging, as performed by Ng and Zhang [114], can be used to create 3D surfaces of flames. These approaches however, provide information only over a limited volume of the flame. To obtain time resolved *whole volume* information of a flame two main methods are available: fast laser sheet scanning, or Computed Tomography of line-of-sight measurements. In the first method, 3D information is achieved by rapidly scanning a laser sheet through the volume of interest. With an appropriately placed and sufficiently rapid camera, a spatial sequence of sections can be measured either by Mie

scattering, as performed by Böhm et al. [23], or by PLIF as used by Nygren et al. [118]. If the time between frames is sufficiently short then by stacking these sections together quasi-instantaneous 3D information can be recovered. As well as the temporal resolution, the resolution achievable in the third dimension is dependent on the camera/laser repetition rate, and a trade-off exists: a longer time is required to measure more sections. To achieve both high temporal and spatial resolutions very fast cameras are required. For example, Böhm et al. [23] use a state-of-the-art 30 kHz CMOS camera and still achieve only 15 sections over 15 mm in the scanning direction, compared to 256 over 12 mm in the imaging plane, and the repetition rate drops to 2 kHz. To obtain comparable resolutions in all 3 directions would require even faster cameras or an even further reduction in the repetition rate.

Computed Tomography[2] (CT) is essentially a mathematical inversion technique that can return 3D data from multiple integral (line-of-sight) measurements taken from different angles around the subject. Computed Tomography is employed successfully in many varied disciplines; for example, Medicine, Seismology, and Electron Microscopy. However, despite early interest in the 1980s [43, 133, 16, 150, 70], in combustion applications CT hasn't yet achieved the success seen in other areas of science. This is primarily due to equipment limitations, particularly for earlier work. In CT the resolution achievable is dependent upon the number of angles at which measurements can be taken. Where time resolution is desired these measurements need to be performed simultaneously, meaning many detectors (and sources) are required. Some studies used complex arrangements involving mirrors and/or optic fibres to reduce the number of detectors and sources required [43, 15] but as far as this author is aware these were never fully implemented. In combustion, CT has most commonly been applied to optical measurements and so historically the limit has been the cost of suitable cameras; i.e. low noise and high sensitivity. However, with the advent of consumer digital cameras, the capabilities of CCDs have improved while their cost has reduced dramatically and we can now reconsider this powerful technique.

In summary, time resolved 3D experimental flame data is required to both further the fundamental understanding of combustion and to support the development of advanced simulation techniques such as LES. While traditionally time resolved 3D data is difficult to experimentally obtain, the declining cost of CCDs means we can now exploit Computed Tomography to provide such 3D information. The current work is concerned with the development of a new combustion sensor that uses images of chemiluminescence, coupled to CT, to provide time resolved, 3D flame information. In the following section we expand on this sensor and discuss further its merits.

[2]Tomography: from the Greek *tomos* meaning *to cut or section*.

## 1.2 Computed Tomography of Chemiluminescence

Computed Tomography of Chemiluminescence (CTC) is the name given to the sensor developed in this work. This sensor uses Computed Tomography to reconstruct the time resolved 3D chemiluminescence profile of an object flame from multiple measurements of the naturally emitted light. The measurements are performed simultaneously at different angles using multiple cameras, with greater numbers of measurement angles allowing higher spatial resolutions in the reconstructed profiles.

Computed Tomography (CT) uses a computer algorithm to provide a discrete estimate of a field from multiple integral measurements, and is distinct from Tomography, which is a general term for any imaging of an internal section of a 3D object. More formally, if we define an *integral* measurement to be one which represents any spatial summation of a scalar field, CT is then a general technique that can be applied to any set of integral measurements, taken at different angles and with *known* geometries, to recover an estimate of the scalar field. Originally used to reconstruct sections from 1D measurements, CT can also return 3D fields by the stacking of sections *or* from 2D measurements directly[3] by recasting the problem in 3D. (Please note that we still use the term *Computed Tomography* in this later direct 3D case because the algorithms essentially remain the same, even though we are no longer extracting only a section.)

Aside from allowing 3D data to be obtained, with the correct selection of algorithm, CT does offer some distinct advantages to the practitioner. Firstly, there is considerable flexibility in the camera (detector) locations as, provided the locations are known, almost any arrangement of detectors is possible—even out of plane locations. (Though any departure from the ideal has a penalty in terms of resolution.) Furthermore, the CT algorithms can be quite robust and allow noisier measurements to be used and so place less stringent requirements on equipment. Finally, and perhaps most importantly in this case, CT has already been successfully (and extensively in some cases) employed in many other fields including Geophysics, Materials Science, Astro-Physics, and not least in Medicine.

Computed Tomography can be coupled to many different integral measurements and this work we consider Chemiluminescence measurements. Chemiluminescence describes when light is emitted as a result of chemical reaction and occurs in a narrow region close to the flame *reaction zone* [171], and thereby provides information on the flame structure. Furthermore, as the wavelength of emitted light can be related to a particular chemical species, the chemiluminescence intensity is also a measure of the concentration of the particular excited state molecule[4]. With modern equipment chemiluminescence measurements can be obtained to very high temporal resolutions, allowing turbulent flames to be

[3] This direct approach has a number of advantages as will be shown later.

[4] Further detail on chemiluminescence can be found in section 2.1.4.

studied. For example, Böhm et al. [23] achieve a 10 kHz repetition rate with a state-of-the-art intensified CMOS camera. By coupling CT with chemiluminescence measurements we can therefore obtain 3D time resolved structural flame information and, with calibration, the concentration profile of excited state species such as $OH^*$ and $CH^*$. Such 3D scalar information will compliment existing laser based approaches that provide multi-scalar and velocity measurements, and together offer greater insight into turbulent combustion.

Chemiluminescence measurements are non-intrusive and are commonly employed in the study of combustion—in particular for premixed flames where the chemiluminescence of certain species has been used to study local equivalence ratios [38, 63, 2, 120, 117] and heat release rate [63, 117]. Chemiluminescence has also been used to study transient phenomena such as flash back [23] and to detect auto-ignition [101]. Recently, much interest has been shown in the modelling of chemiluminescence and several studies [157, 100, 17, 89, 120, 117] have suggested and/or tested several different chemical reaction mechanisms. These recent studies show areas of good agreement between modelled values and calibrated intensities/concentrations, and so open up more opportunities for quantitative validation of LES. Furthermore, as the chemical mechanisms and their parameters become better understood, further relationships may be found between chemiluminescence and relevant combustion parameters, such as local strain rate.

Other line-of-sight measurements have been coupled with CT; for example, Schlieren photography as used by Schwarz [135] or laser absorption as used by Wright [170] and many others. The advantage of using chemiluminescence over other optical measurement techniques is that it occurs naturally in a flame, and so there is no need for external sources or seeding. Given that CT employs multiple sensors, then potentially, either multiple sources would be required, or some complex optical arrangement as used by Beiting [15]. By avoiding the need for sources, the complexity and the equipment requirement of CTC is greatly reduced, making the experiment much more viable. Other natural flame emissions, also do not require an external source; for example, Infra Red (IR) soot emissions, as investigated by Correia et al. [33]. Indeed, we note that the CTC sensor developed in this work could also be used to study such emissions. We however, focus on chemiluminescence, as many benchmark flames used for fundamental modelling and simulations [9] are non-sooting. To reiterate: the CTC sensor is intended for the study of lean premixed or non-sooting partially-premixed flames, although much of this work will be applicable to emission measurements of sooting flames.

In combustion applications Computed Tomography has been used many times previously to localise integral measurements and recent examples include the work of Feng

et al. [44], Gillet et al. [56], and Ito et al. [74][5]. The flexibility that CT affords in terms of the number and location of sensors has meant that many previous studies have focused on lower resolution, 2D reconstructions for the purposes of industrial control [31, 40, 54, 74, 93, 123, 170, 167], where pressure vessels typically limit the number of angles from which views can be taken. Higher resolution 3D reconstructions are more commonly considered in flow visualisation applications using interferograms and Schlieren images [1, 44, 135, 141] though the requirement for sources makes these systems more complex. High resolution time resolved 3D reconstructions of flame chemiluminescence have been performed by Ishino and Ohiwa [73] using a custom built film based 40 lens camera. However, they consider a hardly turbulent flickering flame, and the film based camera can take only a single exposure. In this work we aim to develop the CTC sensor to measure turbulent flames using CCD cameras that can obtain a sequence of images to track transient events.

The spatial resolution of a CTC sensor is related to the number of cameras that can be employed, with more camera angles making higher resolutions possible. In principle by using large numbers of state-of-the-art intensified high speed cameras the CTC sensor could achieve high enough spatial resolutions (the order of $\mu$m) *and* high enough repetition rates (the order of kHz) to track even small scale instabilities and turbulent events such as local extinction and re-ignition. The current cost of very high speed cameras makes their multiple deployment less viable at present[6]—although the declining cost of CCD/CMOS sensors suggests that it may be possible to use such cameras in the future. However, the specifications of more affordable *commodity* cameras have also improved greatly over recent years and the potential for a high spatial resolution CTC sensor exists. (Note that the term *commodity cameras* is used to refer to high quality cameras made using commodity hardware, which are produced in sufficient volumes to be significantly less expensive than low-volume scientific cameras.)

While a CTC sensor using commodity cameras may not be able to track the smallest scales of very turbulent flames, many 3D events of interest occur at larger scales; for example, acoustic instabilities and larger scale turbulent motion. Lower time resolution information would also still be useful for validation purposes as simulation approaches such as LES already filter out the smallest scales. Nevertheless, many affordable cameras are capable of short exposures, albeit at a low throughput (frames-per-second). If sufficient Signal-to-Noise Ratios (SNR) can be achieved, then single *shot*, high spatial and temporal resolution, 3D reconstructions of moderately turbulent flames are possible. Furthermore, a series of uncorrelated *shots* would allow 3D scalar statistics of the whole

[5] In section 3.2 a more detailed review of CT in combustion is given

[6] Though the monies involved would be dwarfed by some current research budgets.

flame volume to be obtained faster than with point, line and 2D measurement techniques.

In this work we are concerned with the measurement of unconfined turbulent flames in a laboratory setting for which we can expect good optical access. Such favourable flames are routinely used as benchmarks for measurement and simulation using LES, as seen for example in the TNF workshop [9]. It is not therefore, the aim of this work to minimise the number of camera angles, as has been the purpose of much previous work, but rather to achieve high spatial resolutions.

The spatial and temporal resolutions that are currently achievable by a CTC sensor are, however, yet to be demonstrated. It is part of the purpose of this work to determine whether cameras that are inexpensive enough for multiple deployment, can achieve sufficient spatial and temporal resolutions for the study of turbulent combustion. The resolutions that are required of a CTC experiment are dependent on the application, and in this work we consider LES as a target application to aid the development process. However we stress that a CTC sensor is not only intended as a tool for LES validation, and we expect CTC to be useful in many areas, both foreseen and unforeseen.

To summarise: in this work we develop a new 3D sensor for the study of turbulent combustion. Called CTC, the sensor uses computed tomography to recover time resolved 3D chemiluminescence fields from multiple simultaneous integral measurements of a turbulent flame. Such 3D data is needed for the validation of LES and for further insight into turbulent combustion. Despite the success of CT in other disciplines, in combustion its use has been limited; due in part to the cost of the multiple detectors (and sources) required. However, recently the decreasing cost and increasing capability of CCD/CMOS sensors suggests that CTC can now be deployed with a sufficient number of cameras to provide high resolution reconstructions. Whether the currently achievable resolutions are sufficient for the study of turbulent combustion is one focus of this work. To enable such a judgement we consider LES validation as a target application though we do not discount other uses. The previous studies that have used CT in a combustion context have generally considered only low numbers of camera angles for control purposes, or hardly turbulent flames. In this work we develop CTC for turbulent flames and using multiple CCDs for greater numbers of angles and so higher resolutions.

## 1.3 Aims

The principle purpose of this work is to determine whether a CTC sensor can be implemented using affordable (commodity) cameras, and whether such an implementation can achieve sufficient spatial and temporal resolutions for the study of turbulent flames. To clarify, we define *affordable* cameras as those for which multiple deployment is achiev-

able with a moderate equipment budget—less than approximately £50000 at the time of writing.

We can expand the principle purpose of this work into a number of more specific aims. Firstly, we aim to actually implement the CTC sensor and use it to measure turbulent flames. While this may seem obvious, in the literature many CT based sensors have not been implemented past the proof-of-concept stage—usually, in this authors opinion, as the result of quite ambitious optical set-ups designed to reduce the numbers of cameras and sources required. Secondly, to judge whether the sensor achieves a high enough resolution, for the measurement of turbulent flames, we first need to determine the resolution that the CTC sensor does achieve. Again, while this appears obvious, the link between the numbers of angles and the resolution achieved is not well defined and we aim to develop an appropriate method of assessment. Finally, we also aim to demonstrate that commodity cameras can provide chemiluminescence measurements at high enough temporal resolution (at high enough SNR) for the study of turbulent flames—specifically for our target application of LES validation.

As previously stated, Computed Tomography is a general mathematical tool that can be applied to all kinds of integral measurements. While the generic algorithm may be sufficient in many cases, there is some scope for improvement in the CTC sensor by tailoring the algorithms more specifically to the case of 3D reconstructions of optical integral measurements. Such improvements have been found in other work; for example, by Correia et al. [33] who implement improvements to the algorithm to correct for self-absorption of IR soot emissions. We also aim therefore, to develop and test new improvements to the CT algorithm in this work, and in particular better represent the measurement optics.

The development of the CTC sensor stems from a need for 3D data in the study of combustion, and to give context to the development we consider more specifically the need for more extensive validation data for LES. However, we also aim to develop other novel uses for the data from CTC that better exploit its 3D nature. For example, as the recent improvements in chemiluminescence modelling have shown, there is some scope to relate chemiluminescence to more combustion parameters of interest, such as the local strain rate.

## 1.4 Document Outline

This document describes the development of the CTC sensor and presents the first results of its application to a turbulent flame. The document is arranged as follows.

In the next chapter some background theory is presented that is relevant for the understanding of this work. Most attention is given to Computed Tomography as much of

this work is focused on the development of the CT algorithms, and this is the subject area that the reader is least likely to be familiar with. We also provide an overview of LES so that the reader may better understand the requirements for validation data. Large Eddy Simulations are also performed to provide data to aid the development of CTC and details are given of the specific solver used to perform the simulations.

In chapter 3 we define in greater detail the CTC sensor concept and review relevant previous studies before describing the bespoke computer implementation of the CT algorithm used in this work (the Algebraic Reconstruction Technique). Following this, in chapter 4, we perform a number of Phantom studies (studies involving known fields) using realistic excited state species fields obtained from the LES. These studies are conducted first to verify the correctness of the CT computer implementation and then to assess the quality and resolution of reconstructions performed using a realistic range of camera angles. Finally in this chapter we report the results of a proof-of-concept experiment performed on a non-axisymmetric Matrix Burner and using a consumer standard digital camera.

Chapter 5 features the results of several studies that have aimed to improve the CT algorithm, in particular to better tailor it to CTC; for example, by accounting for optical effects and allowing 3D camera locations. After first selecting and testing appropriate cameras, in chapter 6 we present time-resolved experimental results obtained using a CTC sensor implemented with 5 machine vision cameras. The selected cameras are Leutron PicSight 32M models, which are used to provide 10 instantaneous views of first the non-axisymmetric Matrix Burner, and then a premixed Turbulent Opposed Jet. In this section we also discuss limitations associated with the camera registration and test a method for improved registration. Finally, we consider a novel application of the CTC sensor, the measurement of local strain rate, and present results from modelling and simulation that support a relationship between strain rate and the local (flame front) integral of chemiluminescence intensity.

# Chapter 2

# Background Theory

The work presented in this thesis draws on aspects of several different, large disciplines. It is the aim of this chapter to provide the reader with a sufficient comprehension of these areas to allow them to understand the remainder of this work while recognising that many concepts are expanded on in later sections.

We begin by considering the wider context of this work, turbulent combustion, and also provide further details of chemiluminescence, the physical phenomena measured by CTC. Following this the theory of Computed Tomography is presented, with particular attention paid to ART, the CT algorithm used in the current project. Finally, as the chosen target application for this development, an overview of LES is given, which also includes specific details of the LES solver used later to provide Phantoms for the development of CTC.

## 2.1 Turbulent Combustion

In most engineering applications of combustion the flow field is turbulent. This is often by design; for example, to realise the benefits of increased mixing that turbulent flows provide. It is assumed that the reader is already familiar with the fundamentals of turbulence and combustion and so in this section only those parts of the theory relevant to the understanding of this work are highlighted. For further depth the reader is directed to the many good books [122, 163, 169] available on the subject of turbulent combustion.

### 2.1.1 Turbulence

Turbulence describes apparently chaotic fluctuations in the velocity of a fluid flow. These fluctuations are characteristically 3D, unsteady and occur over a range of scales, making the prediction of turbulent flows much harder than laminar flows. The fluctuations are

instigated by random perturbations that exist in all flows. If the viscosity of the fluid is sufficient to damp out these perturbations then the flow remains laminar, otherwise it becomes turbulent.

The tendency for flow to become turbulent is described by its Reynolds number, *Re*, which relates the inertial forces to the viscous forces. The Reynolds number is defined in eq. 2.1, where $U$ and $L$ are the characteristic velocity and length scales, and $\nu$ is the fluid kinematic viscosity. Higher Reynolds numbers indicate a greater tendency to be turbulent and in pipes transition from laminar to turbulent flow begins at numbers greater than approximately 2300.

$$Re = \frac{UL}{\nu} \tag{2.1}$$

Turbulent flows contain rotating structures referred to as eddies which manifest over a broad range of scales. The largest scales are of the order of the flow geometry and are referred to as the integral scales. These large scale structures are successively broken down until the eddies become small enough that the viscosity of the fluid causes them to dissipate. Through this process energy is passed down from the largest to smallest scales a process termed the *energy cascade*. The smallest scales are dependent on the viscosity and the break up rate of the eddies and are referred to as the Kolmogorov scales. This concept of energy and scales is integral to Large Eddy Simulation, as shall be seen.

#### 2.1.1.1 Mathematical Description of Turbulent Flows

Turbulent flows can be described by a set of governing partial differential equations. For incompressible flow of constant temperature only two equations are required, comprising the conservation of mass, eq. 2.2, and the conservation of momentum, eq. 2.3. In these equations $\rho$ is the fluid density, $u_i$ the velocity in $x_i$ direction ($i \in \{1,2,3\}$) and $t$ denotes time. In eq. 2.3 $P$ denotes the pressure, $\nu$ is the kinematic viscosity, $\delta$ is the Kronecker Delta, and $g$ is the gravitational constant.

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_i}(\rho u_i) = 0 \tag{2.2}$$

$$\frac{\partial}{\partial t}(\rho u_i) + \frac{\partial}{\partial x_j}(\rho u_i u_j) = \frac{\partial}{\partial x_j}\left[\rho\nu\left(\frac{\partial u_j}{\partial x_i} + \frac{\partial u_i}{\partial x_j}\right) - \frac{2}{3}\rho\nu\frac{\partial u_k}{\partial x_k}\delta_{ij}\right] - \frac{\partial P}{\partial x_i} + \rho g_i \tag{2.3}$$

For more complicated flows, such as the mixing of two fluids or flows of variable temperature, additional equations are required for a proper description. Such equations can be obtained by an appropriate substitution for $\Phi$ in the general transport equation of eq. 2.4, where $D_j$ represents the diffusive flux and $S_\Phi$ is a source term.

$$\frac{\partial \Phi}{\partial t} + \frac{\partial}{\partial x_j}(\Phi u_j) = \frac{\partial D_j(\Phi)}{\partial x_j} + S_\Phi \tag{2.4}$$

Although analytical solution of these equations in 3D is not (normally) possible, it is possible to numerically solve their discrete versions, and this approach is termed Direct Numerical Simulation (DNS). However, due to limits on computational processing power and data storage, DNS is only feasible in small temporal and spatial domains, and at low Reynolds numbers, making it unsuitable for many turbulent flows. The difficulty in solving these equations using DNS is due to the range of scales seen in turbulent reacting flows. One way that this is overcome is by averaging in time: typically Reynolds averaging. However, this approach loses the time dependent information and places great importance on modelling the effect of turbulence, which is necessary to close the system. Large Eddy Simulation (LES) directly solves only the larger scales, and so is computationally less expensive than DNS while still providing time dependent information on the most relevant scales. Further details on LES is given is later in section 2.3.

### 2.1.2 Combustion Modes

The coupling of turbulent flows with chemical reactions (in the form of combustion) presents a complex problem. To simplify the study of such reacting flows, two idealised modes of combustion are usually considered. In one mode the fuel and oxidiser are previously well mixed and this is referred to as *premixed* combustion, in the other the fuel and oxidiser are initially separate and combustion is termed *non-premixed*. States in between these two idealised modes are referred to as *partially premixed*; however, modelling partially premixed flows is challenging, and traditionally idealised flows are more often considered. The premixing of reactants produces fewer pollutants than when the reactants are not premixed. However, the separation of reactants allows easier control of the combustion process and is perceived as safer than premixing. For these reasons both regimes are employed in practical situations. We now consider these modes in greater depth.

#### 2.1.2.1 Non-premixed Combustion

In the non-premixed mode the initial separation of the fuel and oxidiser means that combustion can only occur where sufficient mixing has occurred. This mixing needs to happen on a molecular level and so diffusion—the process by which molecular *mixing* occurs—is very important, and non-premixed flames are also often referred to as diffusion flames. Non-premixed combustion is dependent on the rate of mixing of the fuel and oxidiser in the flow. Turbulence enhances mixing by *wrinkling* the interface of the fuel and air streams thereby increasing the gradients, and area, over which diffusion can occur. The mixing of two fluid streams can be represented by a scalar termed the *mixture fraction* which is unity in the fuel stream and zero in the oxidiser. In the simplest case the mixture

fraction $\xi$ of a species $i$ at some point in the flow field is given by eq. 2.5, with other formulations available in the literature for other cases; for example, the global reaction of a hydrocarbon in air. In eq. 2.5, $Z_B$ represents the proportion (usually by mass or mole) of a species $B$ in the current mixture, and the assumption of equal diffusivity (unity Lewis number) makes $\xi$ independent of the choice of $B$.

$$\xi_B = \frac{Z_B - Z_{B,ox}}{Z_{B,fuel} - Z_{B,ox}} \tag{2.5}$$

Mixing in turbulent flows can therefore be represented by the inclusion of a transport equation for the mixture fraction derived from eq. 2.4 and seen below in eq. 2.6.

$$\frac{\partial}{\partial t}(\rho\xi) + \frac{\partial}{\partial x_j}(\rho\xi u_j) = \frac{\partial}{\partial x_j}\left(\rho D_\xi \frac{\partial \xi}{\partial x_j}\right) \tag{2.6}$$

Several combustion models for non-premixed flows are based on eq. 2.6. The simplest model assumes that chemical reactions reach equilibrium very quickly and that no heat transfer occurs, making the chemical state of the flow dependent only on the mixture fraction field. The assumption of instant chemical reactions however, prohibits the study of effects such as extinction and pollutant formation which occur over finite timescales. More sophisticated models allow non-equilibrium or finite rate effects to be captured, and include Conditional Moment Closure (CMC), developed by Klimenko [86] and Bilger [20], and the Laminar Flamelet approaches developed by Williams [168] and Peters [121].

When seeking to include the effects of finite rate chemistry, the rate of mixing becomes very important and the scalar dissipation rate, $\chi$, which characterises the rate of mixing, is a central quantity in most modelling approaches. However, as this mixing occurs at the smallest scales, in most simulations $\chi$ itself has to be modelled; which, as noted by Kolla et al. [90], is still an "open question". The modelling task is made more difficult as the scalar dissipation rate is proportional to the square of the mixture fraction field gradient, as shown in eq. 2.7 (where $D$ is the molecular diffusivity), making it a very difficult quantity to measure in turbulent flames.

$$\chi = 2D(\nabla\xi)^2 \tag{2.7}$$

#### 2.1.2.2 Premixed Combustion

While non-premixed combustion is characterised by mixing, premixed combustion is characterised by the propagation of the flame through the mixed reactants. This propagation occurs towards the unburnt mixture at a flame speed that is dependent on the composition of the reactants. The flow conditions for a premixed flame therefore need

to be carefully selected to avoid undesired propagation of the flame towards the reactant source—termed *flashback*.

The wrinkling of the flame front caused by turbulence increases the active flame surface area from $A$ to $A_T$, with the result that the speed of flame propagation increases from the laminar value $S_L$ to the turbulent value $S_T$, which thus can be related by eq. 2.8. (A typical value for the laminar speed of stoichiometric methane/air combustion is 0.4 m/s). With too much turbulence, the flame, of finite rate chemistry, cannot produce hot burnt products fast enough to match the turbulent delivery of the unburnt reactants, in which case the flame can become unstable and extinguish.

$$\frac{s_T}{s_L} = \frac{A_T}{A} \tag{2.8}$$

The mixture of the unburnt reactants is defined by the *equivalence ratio* $\phi$ which relates the fuel-to-air ratio of the mixture $Y_{fuel}/Y_{ox}$ to the same ratio at stoichiometric conditions as given by eq. 2.9. Values of $\phi$ greater than unity denote a rich mixture and those less than unity a lean mixture.

$$\phi = \frac{Y_{fuel}/Y_{ox}}{(Y_{fuel}/Y_{ox})_{st}} \tag{2.9}$$

Just as the mixture fraction offers a simplified description of non-premixed combustion, similarly a scalar variable can be used to track the progress of premixed combustion. This scalar variable does not describe mixing but rather reaction progress by relation to burnt and unburnt quantities. This is usually called the reaction progress variable $c$ and can be defined by the temperatures of the burnt $T_b$ and unburnt $T_u$ gases as seen in eq. 2.10 (though other definitions are also used).

$$c = \frac{T - T_u}{T_b - T_u} \tag{2.10}$$

A transport equation can be defined for the reaction progress variable, as seen in eq. 2.11, which feeds an appropriate sub-model to define the chemical state of the flow. Several modelling approaches, such as the Flame Surface Density or Bray-Moss-Libby models, are based on the transport of this scalar quantity with particular focus on the closure of the chemical source term $\dot{\omega}_c$, which represents the rate at which (burnt) products are generated. Further details on combustion modelling as applied to LES are given later in section 4.1.

$$\frac{\partial(\rho c)}{\partial t} + \frac{\partial}{\partial x_j}(\rho c u_j) = \frac{\partial}{\partial x_j}\left(\rho D_c \frac{\partial c}{\partial x_j}\right) + \dot{\omega}_c \tag{2.11}$$

### 2.1.3 Reaction Kinetics and Mechanisms

Combustion is an exothermic chemical reaction between a fuel and an oxidiser. The product of this reaction depends on the fuel and oxidiser and an example reaction of

methane and oxygen is shown in eq. 2.12.

$$CH_4 + 2O_2 \longrightarrow CO_2 + 2H_2O \tag{2.12}$$

Such an equation describes the *global* reaction that occurs; however, in reality such a reaction would consist of many intermediate reactions, which given sufficient time approximately result in the global reaction. Although the global reaction is sufficient for many uses, when the flow and diffusion time scales approach the chemical time scales the rate (kinetics) of the reactions becomes important. Furthermore, it is in these intermediate reactions of finite time that pollutants and other species are formed, and if these are to be predicted, then the relevant *elemental* reactions, and their rates, need to be considered.

For a reaction of hydrogen in oxygen relatively few ($\approx 40$ [163]) elemental reactions are required. However, for larger and more complex hydrocarbons the number of reactions can exceed 1000. In such a case simply knowing the complete set of reactions and their rates is difficult enough, but their solution along with turbulent flows is usually too computationally expensive. Fortunately, for the prediction of particular species only a subset of the reaction mechanism may be required, and of this subset some reactions will be more significant than others. Therefore, for a given prediction objective, there is scope for reduction of the mechanism to a simpler, functional version, and this is an area of much research.

The rate of each elementary reaction can be determined from the species concentrations and the rate coefficient, $k_r$. The rate coefficients can depend strongly on the temperature $T$ and are usually approximated by eq. 2.13, which stems from the Arrhenius law, where $A$ and $b$ are coefficients and $E_a$ is the activation energy.

$$k_r = AT^b \exp\left(-\frac{E_a}{RT}\right) \tag{2.13}$$

Even with a reduced mechanism and suitable rate coefficients the solution of these reactions for turbulent reacting flows can still be difficult and hence modelling approaches are often employed, as described previously.

### 2.1.4 Chemiluminescence

Chemiluminescence is luminescence (the emission of light not resulting from heat) that occurs as a direct result of a chemical reaction. More specifically, this means that a chemical reaction produces a molecule in an excited state (i.e. where an electron is moved to a higher energy orbital than normal). This electronically excited state is unstable and at some point this molecule will decay to the more stable ground state by the emission of a photon of energy $h\nu$ (unless quenching occurs first), where $h$ is Planck's constant and

$v$ is the frequency of the light emitted. The photon energy is equivalent to the difference between the ground and excited states. Example transitions for commonly considered species are shown in table 2.1, taken from Gaydon [49]. This table includes $CH^*$ and $OH^*$ which are the two strongest chemiluminescent emitters in hydrocarbon flames.

| | **Transition** | **Wavelength** |
|---|---|---|
| $OH^*$ | $(A^2\Sigma - X^2\Pi)$ | 306.4 nm |
| $CH^*$ | $(A^2\Delta - X^2\Pi)$ | 431.3 nm |
| $C_2^*$ | $(A^3\Pi - X^3\Pi)$ | 516.5 nm |
| $HCO^*$ | $(A^2\Pi - X^2A')$ | 318.6 nm |
| $CH_2O^*$ | $(A^1A'' - X^1A_1)$ | 395.2 nm |

Table 2.1: Examples of commonly studied chemiluminescent transitions in combustion. The ground state is prefixed with X and the first excited state with A. The symbols $\Sigma, \Delta$ etc. represent the orbital angular momentum and the preceding superscript the spin.

The frequency $v$ of the emitted light depends on the specific *transition* that occurs for a particular molecule. Therefore, if the transition probability is known, and that collisional quenching can be accounted for (which is non-trivial), then the concentration of a particular excited molecule can be determined from measurements of the emitted light intensity. The ground state concentrations of the molecule may be obtained from the excited state concentrations using the Maxwell-Boltzmann distribution law [49]. However, this applies only in thermal equilibrium which may not be the case as many of the excited species of interest are radicals and are the result of non-equilibrium finite rate chemical reactions.

Chemiluminescence emission measurement, like other optical measurements, are affected by absorption and scattering of the emitted light in the flame before it reaches a detector. If a flame is optically thick then measuring reliable concentrations of a particular species becomes very difficult—particularly where the broad band emissions of soot mask chemiluminescence signals. For this reason lean premixed flames, which contain little or no soot, are most suitable for chemiluminescence measurements. However, optical measurements of non-premixed flames can also be performed by partial premixing of the fuel stream as demonstrated by Wang and Barlow [160]. Non-sooting flames can still be optically thick however, and result in self-absorption of the emitted light which reduces the measured signal. The spatial extent of the flame is important in this respect. For small flames, the emissions from some species, such as $CH^*$, have negligible self-absorption, while others such as $OH^*$ can be significantly affected.

## 2.2 Computed Tomography

The mathematical foundations of Computed Tomography (CT) were first laid down by Radon in 1917 however, it was not until the 1970s and the development of Hounsfield's EMI scanner[1] that interest in CT became more broadly subscribed. Since this time the basic premise of reconstruction from projections has not changed; however, the research field is still very active, for example as part of the more general *inverse problems* area of mathematics. Interest in CT continues in medical applications with its use in several modern imaging *modalities* such as Positron Emission Tomography (PET), Optical Coherence Tomography (OCT) and Single Photon Emission CT (SPECT), in addition to the ubiquitous X-Ray CT. (The reader should note that in medical applications *CT* is often used to refer specifically to X-Ray CT; in this work however, the more general meaning of CT is used: a discrete form of tomography using computer inversion algorithms.) Computed Tomography has also seen successful use in many other applications such as: Seismology, where it is used to determine the interior of the Earth; Non-Destructive Testing (NDT); the control and monitoring of industrial processes; and also in electron microscopy.

Computed Tomography (CT) is essentially an inverse problem, where a discrete estimate of a scalar field is obtained from the inversion of a number of integral measurements, of known geometries, taken at different angles through the original field. Or, to phrase more succinctly, it is the reconstruction of an object from its projections. With CT the reconstruction is always discrete and will always be an approximation of the continuous original field, although any error diminishes with increasing resolution. The resolution achievable using CT is strongly coupled to the number of angles at which measurements can be taken and generally higher resolutions require a greater number. Projection measurements can be obtained using many different measurement mediums. Indeed any medium is possible, provided the object is sufficiently transparent, or thin, to the chosen medium, and that the geometry of the projections are sufficiently known. In most applications the projection geometries are approximately rectilinear and characterised by a single angle as shown in fig. 2.1. However, provided they are known, any geometry is possible, and this allows strongly diffracting mediums such as ultrasound to be used.

This section provides a background to the theory of Computed Tomography. First the problem of reconstructing from projections is presented before continuing with a review of the common algorithms employed to perform reconstructions. Finally, the particular family of algorithms considered in this work—the Algebraic Reconstruction Technique—is considered in greater detail. For further information on Computed Tomography the *Principles of Computerized Tomographic Imaging* by Kak and Slaney [77] can be recom-

[1] For which he was awarded the Nobel prize in 1979 along with Cormack for work on the algorithms.

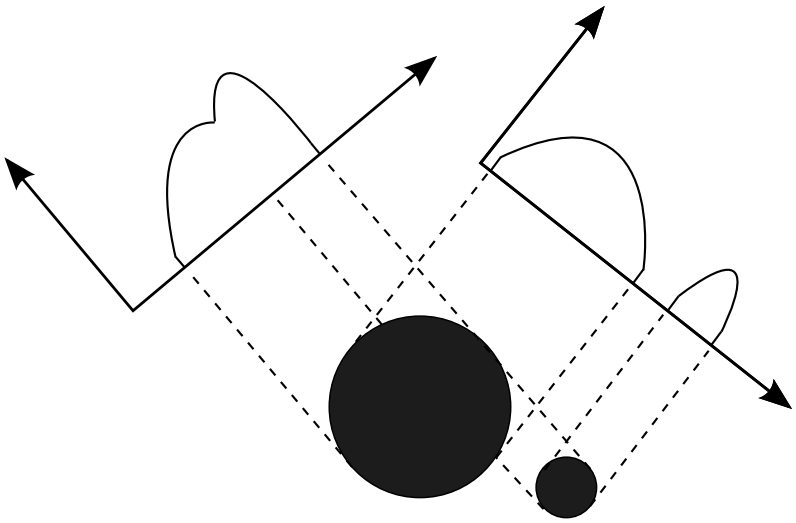


Figure 2.1: Example projections at two different angles for a two cylinder object.

mended along with books by Herman [67] and Barrett and Swindel [13].

### 2.2.1 Reconstructing from Projections

To present the problem of reconstructing from projections we begin by first defining the term projection. In this work a *projection* refers to a single integral, of index $p$, taken through a portion of the object domain, not the whole object domain. A *view* is the set of $N_p$ projections with common viewing angle $\theta$ that together span the whole of the object domain at that angle, as shown in fig. 2.2. A view is identified by the index $q$ and a total of $N_q$ views are used when reconstructing the *object* $f(x,y)$. Views can be either 1D or 2D (an image) and here we consider 1D views and 2D reconstructions for simplicity. (Note also that the term projections refers to all projections from all views.)

For simplicity, here we shall assume that the projections of a view are parallel line integrals. This is a reasonable representation of the propagation of X-Rays for example, and is a common basis for several algorithms. However, the reader should be aware that the problem can be cast as non-parallel and with other integral geometries. Please also note that the terms integral geometry and projection geometry are synonymous. Singular projections are often referred to as rays in the literature which implies line integrals. Later in this work projection geometries other than lines are used and so the term ray when used relates only to the orientation of a projection and not to its integral geometry. One final point on terminology, in much of the literature the discrete estimate of the object field reconstructed by the CT algorithm is termed the *Image*; however, this term makes less sense in the 3D case and can potentially be confused with 2D projections which are themselves images and so the term *reconstruction* is preferred in this work.

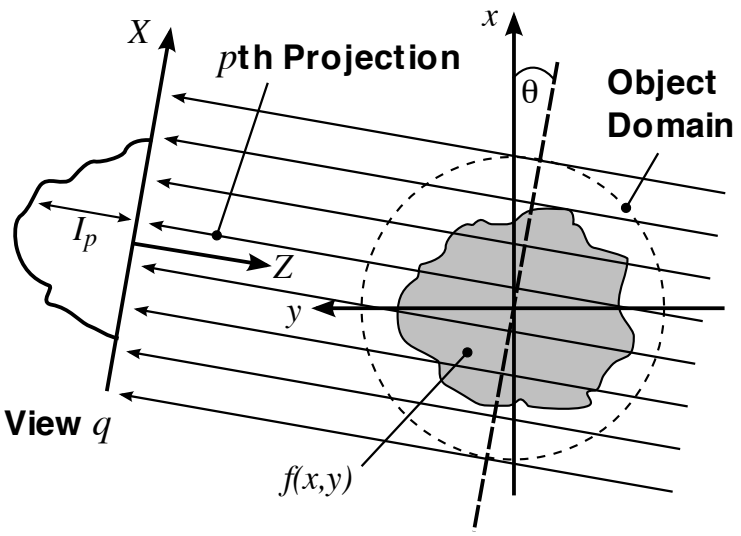



Figure 2.2: Example projections based on line integrals. Note how multiple projections at a common angle comprise a view.

For a generic 2D projection geometry the integral value for the projection $p$ of view $q$ is given by eq. 2.14.

$$I_{qp} = \int_{\theta,X} f(x,y)\, dA \tag{2.14}$$

If the view coordinates, $XZ$, are related to the object domain coordinates, $xy$, by a rotation of $\theta$, then for the case of parallel line integrals this equation can be rewritten using a delta function as seen in eq. 2.15.

$$\mathscr{R}_{qp}[f(x,y)] = I_{qp} = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y)\delta(x\cos\theta_{qp} + y\sin\theta_{qp} - X_{qp})\, dx\, dy \tag{2.15}$$

This is the Radon transform of the object $f(x,y)$ and the problem that is addressed by CT is the realisation of the inverse of the Radon transform $\mathscr{R}^{-1}$ of the measured projections $I_{qp}$.

#### 2.2.1.1 Fourier Slice Theorem

The inverse Radon transform can in theory can be achieved by using Fourier transforms. It can be shown that a slice of the 2D Fourier transform of an object $F(u,v)$, at an angle $\theta$, is given by the one dimensional Fourier transform of the view $S_\theta(U)$ (with parallel projections) at an angle $\theta$, as shown by eq. 2.16. This is called the Fourier slice theorem.

$$F(u,v)\Big|_{v=u\tan\theta} = S_\theta(U) \tag{2.16}$$

With an infinite number of views spaced equally over 180°, there are sufficient slices, $S(U)$, to fully populate the object in Fourier space $F(u,v)$. The 2D inverse Fourier transform of $F(u,v)$ then yields the object $f(x,y)$. However, as only a finite number of pro-

jections are possible, problems associated with the sparsity of the high frequency components in Fourier space, and the effect of noise, make this an impractical approach in most situations, and alternative algorithms are used.

#### 2.2.1.2 Emission and Transmission

Measurement of the projections can occur in two forms: emission or absorption (also referred to as transmission). For absorption, the attenuation of a known external source, for example a laser beam, is measured, and can in turn be related to the concentration of the absorbing field. For emission, the intensity of an internal source is measured, and can be either naturally occurring, as with chemiluminescence, or artificial as with PET. The main practical difference between these cases is that for absorption attenuation is desirable whereas for emission it is not, and needs to be accounted for.

Returning to our line integral example, attenuation can be accounted for in emission projections as seen in eq. 2.17.

$$I_{qp} = \int_{(\theta,X)\text{line}} f(x,y)\; e^{-\kappa \int c(x,y) dl_A}\; dZ \tag{2.17}$$

Here $\kappa$ is the absorption coefficient and $c(x,y)$ is the absorbing field. The extra integral spans along the line defined by $\theta, X$, but only from the point $xy$ to the detector (or the edge of the absorbing field if closer). We can see that if the emitting and absorbing field are related then the problem becomes non linear and cannot be solved by some algorithms. This issue is returned to later in section 5.1. However, for very *thin* objects, where $\kappa$, $c(x,y)$ and/or $l_A$ is small, attenuation can be discounted and for CT their is no mathematical difference between the two measurement methods.

### 2.2.2 An Overview of Algorithms

Numerous CT algorithms exist and can be classified in many ways, but usually algorithms are separated into analytical and iterative types. In general most analytical algorithms are based on the concept of *back projection*. Such algorithms generally require many low noise views for successful reconstruction [77] and they are generally less favoured in the physical sciences. These algorithms, in particular Filtered Back Projection (FBP), are applied most commonly to medical applications where static subjects and good access mean many measurements can be taken.

As can be expected, non-iterative analytical approaches are generally faster than iterative ones. However, iterative algorithms are generally more robust in the presence of noise, and for this reason are usually favoured in the physical sciences where post-processing time is less of a concern. With modern computing power however, the difference in the processing time is much reduced and iterative algorithms are seeing increasing

use even in medical applications, for example as with SPECT [26]. In the following sections we provide further detail on the main iterative algorithms that are relevant to this work. We begin however, by considering the ubiquitous FBP algorithm for comparative purposes.

#### 2.2.2.1 Back Projection

Back projection algorithms essentially work by successively projecting each measured view back across the discrete reconstruction domain. Because the spatial information is lost in the projection step there is no choice but to *smear* back the integral value. However, by smearing back views from many angles the spatial information and hence the object field can be recovered. To define this process mathematically we consider the Radon transform described earlier. This transform more generally is a *projection* operator for the particular case of line integrals and in turn has an associated *back projection*, as shown in eq. 2.18. This equation describes how a point $(x,y)$ in the back projection reconstruction $f_a(x,y)$ is given by the accumulation of the integral values of those projections that intersect with that point. This accumulation is performed over all $\theta$.

$$f_a(x,y) = \int_0^{\pi} I(Z,\theta)d\theta \tag{2.18}$$

Back projection alone however is not the inverse of the Radon transform/projection operator, as $f_a(x,y) \neq f(x,y)$. This is evidenced from the blurring seen in any reconstruction using only back projection, for example in [26], which results from the loss of spatial information in the projection step.

The FBP algorithm seeks to reduce this blurring in the reconstruction by filtering each view first before the back projection step. This is shown in eq. 2.19, where $\widehat{I}$ is the filtered version of $I$. Filtered Back Projection can actually be derived from the Fourier slice theorem by changing the inverse Fourier transform to polar co-ordinates and changing the limits of integration[2].

$$f_a(x,y) = \int_0^{\pi} \widehat{I}(Z,\theta)d\theta \tag{2.19}$$

The filter traditionally used is a ramp filter, which effectively weights the high frequency components in Fourier space and reduces the low frequency components that cause the blur. However, this also has the effect of enhancing the noise in a reconstruction and other filters are also commonly used. For example, the Hamming filter reduces both the lowest frequencies to reduce blurring and the highest to reduce noise. There is always a trade

[2]Indeed, Kak and Slaney [77] note that a wide variety of computer algorithms are obtainable by simple re-expression of this underlying theory.

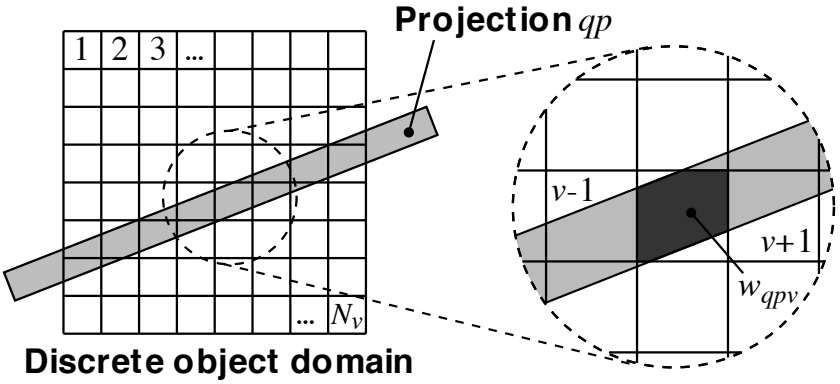



Figure 2.3: A pixels contribution $w_{qpv}$ to projection $qp$ is the intersection of the projection geometry and the pixel. The object domain has been discretised with $N_v$ pixels.

off between noise and contrast in the selection of an appropriate filter and generally some optimisation is required.

#### 2.2.2.2 Iterative Algorithms

Iterative algorithms are generally based on the re-expression of the reconstruction problem as a system of linear equations. This system is obtained from the discretisation of the projection operator, eq. 2.14, achieved by dividing the object domain into a grid of pixels of indices $i$ and $j$. To simplify the notation the pixel indices are vectorised and are represented by the single index $v$, of maximum value $N_v$, and which is a function of $i, j$. The projection operator in its discrete form is given in eq. 2.20.

$$I_{qp} = \sum_{v=1}^{N_v} w_{qpv} f_v \tag{2.20}$$

Here $w_{qpv}$ represents the contribution of the pixel $v$ to the $p$th projection of the $q$th view. The definition of the contribution depends on the type of integrals used. For line integrals, $w_{qpv}$ could be the $\Delta l$ of the line's intersection with the pixel. However, the discrete measurement of a view means each projection has an associated bin width, and 2D *strip* integrals are more appropriate. In this case $w_{qpv}$ is the intersection area (or volume in 3D) of the strip with the pixel $v$, as shown in fig. 2.3.

From eq. 2.20 we can then recast the reconstruction problem as a system of linear equations as shown in matrix form in eq. 2.21, where $W$ is referred to as the *projection matrix*. The projection matrix has as many rows as there are projections, and as many columns as their are unknown pixel values, and so can become quite large ($N_v \times N_p \times N_q$). Fortunately a single projection typically only covers a small portion of the object domain, and the columns of $W$ are quite sparse, allowing even large problems to be tackled.

$$\mathbf{I} = W\mathbf{f} \tag{2.21}$$

Equation 2.21 indicates that the object vector $\mathbf{f}$ can be found once the inverse of the projection matrix, $W^{-1}$, is known. In most situations the presence of noise means that $W^{-1}$ is ill conditioned and direct inversion is not possible. Furthermore in many practical situations only limited numbers of views, $N_q$, are available in which case the projection matrix is non-square and so has no inverse. In this case the system of equations is termed *under-determined* as the number of unknown pixel values, $N_v$ exceeds the number of known/measured projection values, $N_q \times N_p$, and so multiple solutions are possible. For these reasons iterative algorithms are usually employed that achieve an optimal solution to eq. 2.21 (e.g. maximum information entropy).

Most iterative algorithms conceptually follow the same approach: first an estimate of the object $f^{(0)}$ is made, then an error is generated by some comparison of the current estimate $f^{(h)}$ with the measured projections $I$, and then finally this error is used to correct the current estimate and give the next iterations estimate $f^{(h+1)}$. Iteration then continues until some convergence criteria is achieved. The algorithms differ mainly in how they generate and apply the correction, and, as with many iterative algorithms, the choice of the initial estimate and the convergence criteria is important. One added benefit of using such an iterative approach is that the reconstructions can be easily restricted to a range of values that is known *a priori*, for example non-negativity, by simply setting these values appropriately in between iterations.

The most widely used iterative algorithm, particularly in non-medical applications, is called Algebraic Reconstruction Technique (ART). The original ART algorithm was suggested by Gordon et al. [58], and has since been modified into many different forms. In the following section algorithms from the ART family are considered in greater detail, and for now we describe only the common additive version of ART.

Based on the *method of projections* by Kaczmarz [76], ART generates the error by comparing each measured projection value $I_{qp}$ with an equivalent projection taken through the current iteration's ($h$) estimate of the object $f_v^{(h)}$ using eq. 2.20. This error is normalised and immediately back projected into the reconstruction domain using $w_{qpv}$, as seen in eq. 2.22, and then the next projection is addressed. Once all views have been addressed then the next iteration begins. Generally a reasonable estimate of the object is produced after just one iteration, and higher levels of accuracy are obtainable with more iterations.

$$f_v^{(h+1)} = f_v^{(h)} + w_{qpv} \frac{I_{qp} - \sum_{v=1}^{N_v} w_{qpv} f_v^{(h)}}{\sum_{v=1}^{N_v} w_{qpv}^2} \tag{2.22}$$

Algorithms based on Expectation Maximisation (EM) are popular alternatives to those of the ART family, in particular for medical emission sensors such as SPECT, and have

been used in recent combustion emission studies [94, 73]. These methods consider the set of measured projections to be a single realisation of a stochastic process (incorporating noise) and determine a mean solution based on the maximisation of some criteria. Commonly the Maximum Likelihood criteria is used, in which case it is termed the MLEM algorithm. The MLEM algorithm for emission sensors in nuclear medicine was presented first by Shepp and Vardi [137] and is based on the measurement noise following Poisson statistics. MLEM is a two step process: in the expectation step, an expression for the likelihood of any solution based on the measurements is formed; in the maximisation step, the solution with the greatest likelihood is found. In eq. 2.23 is shown the MLEM algorithm presented in the simpler iterative form of Lange and Carson [91].

$$\bar{f}_v^{(h+1)} = \frac{\bar{f}_v^{(h)}}{\sum_{qp=1}^{N_q N_p} w_{qpv}} \sum_{qp=1}^{N_q N_p} \frac{I_{qp}}{\sum_{v=1}^{N_v} w_{qpv} \bar{f}_v^{(h)}} w_{qpv} \tag{2.23}$$

From eq. 2.23 we can see that the error is generated from the ratio of the measured projection to the current estimate's projection, rather than their difference as with ART. The other main difference to ART is that the correction is accumulated for all projections ($qp = 1 \rightarrow N_q N_p$) before being used to update the estimate. (Some ART variants also use such an approach as will be seen in the following section.) Strictly speaking, in MLEM, $w_{qpv}$ now represents the probability that the emission of the pixel $v$ will be measured by the detector $qp$ rather than its contribution to the projection $qp$, but these are functionally the same thing. The MLEM algorithm in eq. 2.23 was derived for photon emission measurements which feature Poisson distributed variations, and are therefore applicable to the measurement of chemiluminescence as has been demonstrated by Ishino and Ohiwa [73].

The principle advantages of both these iterative algorithms (MLEM and ART) is their ability to successfully reconstruct from limited and noisy projection data, and to not need an optimal filter to be defined as with FBP. Furthermore, both allow the incorporation of *a priori* information such as smoothness, which can be used to improve reconstructions or reduce the number of required views. Furthermore, unlike back projection algorithms, iterative algorithms such as ART and MLEM place no restrictions on the geometry of the projections nor their relative orientations (i.e. they can be non-parallel). ART and MLEM can be easily adapted to other geometries by an appropriate redefinition of the projection weighting factor $w_{qpv}$; neither the reconstruction problem nor the solution algorithms need to be changed.

A study by Herman and Meyer [68] found MLEM and ART to achieve similar reconstruction performance, though ART required an order of magnitude less convergence time. Also MLEM reconstructions from noisy data tend themselves to become noisier with greater numbers of iterations. This can be overcome to some degree by a including an additional criterion in the EM process that favours smoothness; although, the degree

of smoothness needs to be specified before. As noted by Verhoeven [152], EM methods are only expected to yield better results with the application of prior knowledge, and the simplicity and speed of ART means it is initially favoured in this work.

### 2.2.3 ART in Greater Detail

Algorithms from the ART family are commonly chosen in most non-medical applications, primarily because they can successfully reconstruct from limited numbers of noisy views and can easily incorporate prior information to improve reconstructions. ART can be thought of as a series expansion method: where a function—the object—is approximated by the summation of a series of basis functions. The standard basis is one of step functions that form an array of pixels; however, one of the principle ways of exploiting *a priori* information is by using a different basis to introduce smoothness, for example, to help reconstruction.

ART type algorithms are not limited to line projections and because of the finite bin width of projection measurements, generally the more exact 2D strips, as seen in fig. 2.3, are preferable. However, due to refraction, laser beam profiles etc., other integral geometries may better match the *true* volume of the integral measurements. When these *true* integration volumes/areas are known they can be easily incorporated into ART by appropriately adjusting $w_{qpv}$—although determining a better approximation than a line or strip is non-trivial. Nevertheless, it is true to say that the more accurately the projection geometry is described the better the reconstructions will be. For example, Andersen and Kak [6] have incorporated ray-tracing into ART algorithms to account for weak refraction.

More complex and *thicker* projection geometries take more time to determine and more computer memory to store. While this is less of an issue with modern computing resources, when ART was first implemented in the 1970's this needed to be addressed. The original ART implementation therefore, used a binary approximation of the projections, where when a pixels' centre was within the bounds of the strip integral a weight of unity was given, otherwise $w$ was zero. This leads to *salt and pepper* noise in the reconstruction, and later, improved reconstructions were obtained by using the length of intersection between the ray and the pixel as the weighting factor. Although this is still an approximation and results in some *reconstruction noise*. (Noise in the reconstruction resulting from projection inaccuracies, not noise in the projections). With modern computers however, there is little reason not to use the *exact* strip integrals as, even for high resolutions, the calculation of $W$ takes a reasonable time and it need only be performed once for a particular set-up.

To reduce the reconstruction noise as well as to improve convergence ART type algorithms often employ varying degrees of relaxation. The appropriate value to use depends

on a number of factors, including the projection noise, and a broad range of values have been employed. For example, in the work of Mishra et al. [108] relaxation values range from 0.001 to 1. However, as seen in [77, 68, 108], it is clear that stronger relaxation factors ($<< 1$) are required when the non-exact projection geometries (i.e. line integrals) are used, or when some other significant error is present such as truncated views. Where more *exact* integral geometries are used much higher relaxation factors are possible (Verhoeven [152] uses 0.5) indeed if they are necessary at all. Generally however, it is advantageous to still use weak relaxation with strip integrals due to unavoidable noise in the projections.

In this work *strip* integrals are used along with a relaxed version of ART shown in eq. 2.24, where $\beta$ is the relaxation factor. In this equation the correction for one view using ART is shown in vector notation, where $\mathbf{w}_{qp}$ is a row vector consisting of $w_{qpv}$ for all $v$, i.e. one row of $W$, and $\mathbf{a} \cdot \mathbf{b}$ denotes the dot-product of $\mathbf{a}$ and $\mathbf{b}$.

$$\mathbf{f}^{(h+1)} = \mathbf{f}^{(h)} + \beta \mathbf{w}_{qp} \frac{I_{qp} - \mathbf{w}_{qp} \cdot \mathbf{f}^{(h)}}{\mathbf{w}_{qp} \cdot \mathbf{w}_{qp}} \tag{2.24}$$

A popular variation on the standard additive form of ART is MART which applies its correction multiplicatively and can achieve better reconstructions than ART for objects with high gradients [152]. Unlike additive ART, which returns the minimum norm solution, MART has been shown to return the maximum entropy solution [152]. (Please note that in this work we refer to entropy in the context of information theory not thermodynamic entropy; see reference [75] for further information.) The MART algorithm considered in this work is that described by Verhoeven [152], shown in eq. 2.25, which necessarily has a relaxation factor added to improve stability. (Note that here the multiplication symbol $\times$ is an element by element multiplication not a matrix one.)

$$\mathbf{f}^{(h+1)} = \mathbf{f}^{(h)} \times \left( 1 - \frac{\beta \mathbf{w}_{qp}}{\mathbf{w}_{qp} \cdot \mathbf{w}_{qp}} \left( 1 - I_{qp} / \mathbf{w}_{qp} \cdot \mathbf{f}^{(h)} \right) \right) \tag{2.25}$$

As the correction for MART is applied multiplicatively, an advantage of MART is that once a value in the reconstruction is set to zero it remains so. This however, means that the initial estimate of $\mathbf{f}^{(0)} = 0$, commonly used for additive ART, is not suitable and some mean value derived from the projections is usually employed.

#### 2.2.3.1 Basis Functions

In casting the reconstruction problem as a system of linear equations, we discretised the object domain into an array of pixels each of uniform value. However, with ART we are free to choose a different basis, perhaps to exploit *a priori* information of the object

and obtain better reconstructions. Several studies have used a different basis, including cubic splines [62, 152]; spherically symmetric functions called *blobs* [104]; sinc functions [159]; and optimised functions from simulations [40]. In such a case, the object is approximated using $N_v$ basis functions, $b(x,y)$, as shown in eq. 2.26 and where $g_v$ is the coefficient of expansion.

$$f(x,y) \approx \hat{f}(x,y) = \sum_{v=1}^{N_v} g_v b_v(x,y) \tag{2.26}$$

We can then recast the reconstruction problem as a system of linear equations by substituting this equation into eq. 2.14.

$$I_{qp} = \sum_{v=1}^{N_v} g_v \int_{qp} b_v(x,y)dA = \sum_{v=1}^{N_v} g_v w_{qpv} \tag{2.27}$$

We can see that the only difference from eq. 2.20 is that $w_{qpv}$ is now expressed in the more general form of $\int_{qp} b_v(x,y)dA$, where the integral is over the intersection of the projection $qp$ with the $v^{\text{th}}$basis function. Returning to the pixel basis of $b(x,y) = 1$ inside the pixel and zero elsewhere, then $w_{qpv}$ is the inter-sectional area of the projection with the square pixel and $g_v$ is the average value of $f(x,y)$ in the $v^{\text{th}}$pixel. The change of basis does not affect the actual ART algorithm except that in this more general sense it now solves for $g_v$ rather than $f_v$.

#### 2.2.3.2 SART

In their ART implementation, called SART[3], Andersen and Kak [6] use a basis of bi-linear elements to better match the integral measurement of smooth fields, similarly to later work in section 5.2.2. The computational expense of using these bi-linear elements, however, prompted Andersen and Kak to approximate the integral by a finite sum of sample points. We note that, Kak and Slaney [77] also admit that exact strip integrals are still good representations of the integral. SART also *heuristically* applies a Hamming window to the reconstruction to emphasise corrections in the centre of the domain, though such heuristic filtering is less justifiable in the more general, non-medical case. Interestingly however, we note that this would have a similar effect as the optical projection adjustments made later in section 5.4.

The SART implementation is perhaps more noteworthy for claims of better noise performance: obtained by applying all the error from a single view at the same time, rather than ray-by-ray as with standard ART. This averaging reduces the oscillation of the error that occurs when two adjacent projections successively correct the same pixel. One can

[3]SART stems from a combination of the best of ART and another algorithm called Simultaneous Iterative Reconstruction Technique (SIRT)

obtain the noise performance of SART by the simple accumulation and application of the correction from standard ART for one view, $q$, as shown in eq. 2.28.

$$\mathbf{f}^{(h+1)} = \mathbf{f}^{(h)} + \frac{1}{\sum_{v=1}^{N_v} w_{qpv}} \sum_{p=1}^{N_p} \mathbf{w}_{qp} \frac{I_{qp} - \mathbf{w}_{qp} \cdot \mathbf{f}^{(h)}}{\sum_{v=1}^{N_v} w_{qpv}} \tag{2.28}$$

#### 2.2.3.3 Projection Ordering and Convergence Criteria

While both ART and MART are known to converge to a solution of the system of linear equations, the selection of an appropriate stopping point for iteration is non-trivial due to the presence of noise. A number of different convergence criteria have been suggested for use with ART and MART, and, in practice, are similarly implemented by testing if the change in some criteria $\mathsf{X}$ from one iteration to the next is less than some value $\Delta_c$ (i.e. $\mathsf{X}^{(h+1)} - \mathsf{X}^{(h)} < \Delta_c$). Perhaps the most obvious criteria is to minimise the error between the measured projections $\mathbf{I}$ and the equivalent projections of the current estimate, $W\mathbf{f}^{(\mathbf{h})}$. This however, is not recommended due to the presence of noise [26]. Gordon [58] has suggested the use of both the variance of $f^{(h)}$ and the entropy for ART convergence criteria and Mishra et al. [108] have used the change in $f$ itself. The convergence proof for under-determined systems of linear equations using the Kaczmarz method (on which ART is based) is given by Tanabe [147] and results in the minimisation of the norm criteria described by Kak and Slaney [77] as the minimisation of $|f^{(0)} - f^{(h)}|$. As this criteria approaches a minimum, $|f^{(h)} - f^{(h-1)}|$ also reduces and hence is consistent with the approach of Mishra et al. [108]. There is no single outstanding choice for a convergence criterion and, where the problem is well constrained, little practical difference would be observed between these methods. However, with noisy under-determined configurations the rate of convergence is slower, as the problem is less constrained, and here the choice of convergence criteria is more important. The issue of convergence for under-determined noisy configurations is returned to later in section 4.2.2.

Improvements to the rate of convergence of a reconstruction can be obtained [77, 60] by selecting the order in which each view is addressed, so that the correlation between subsequent views is minimised. Or to phrase it differently, to maximise the amount of new information being projected into the field. Though desirable, in practice full orthogonalisation is not possible and several different schemes have been suggested that seek to minimise the correlations of subsequent views. Guan and Gordon [60] have suggested a multi-level scheme, and Mueller et al [110] have devised a scheme they referred to as the *weighted distance scheme*. While both of these schemes yield significant improvements over the traditional sequential approach when only a few iterations are employed, the improvements are reduced with higher numbers of iterations and are negligible above approximately 10 iterations. Furthermore, both schemes offer only a marginal improvement

over random ordering of the views and, in some cases, are actually worse than the much simpler random approach. These tests are conducted using more than 30 views however, and with fewer views the random order approach could be expected to perform worse; although in such a case manual ordering is possible. Where multiple iterations are used the direct effect of any scheme on convergence is small, as reported by Verhoeven [152]. However, by reducing the correlation between views higher relaxation factors can be used and so convergence speed is improved.

## 2.3 Large Eddy Simulation of Combustion

Large Eddy Simulation is a numerical approach that delivers a time dependent, 3D solution to the set of equations governing turbulent reactive flows (described previously in section 2.1). Unlike DNS, where all scales must be resolved, LES considers only the larger energy containing scales and models the smaller (unresolved) scales allowing a much coarser grid to be used. The lower resolution requirements of LES mean that solutions can be obtained much faster than DNS, and simulations of much higher, practical Reynolds number flows are possible.

One of the principle aims of the current work is to develop a sensor that can provide validation data for combustion LES. In this section we provide an overview of the main principles of this advanced simulation approach and the common modelling approaches employed. For greater detail on the theory of LES the reader is referred to the books of Sagaut [132], Pope [128] and, for more attention to combustion, the book of Poinsot and Veynante [125] is recommended. In this work, LES is also later used to provide realistic scalar fields for the development and testing of the CTC sensor, and details of the solver 'FLOWSI' are also given throughout this section.

The reader should note that it is not the purpose of this work to investigate the predictive accuracy of combustion LES, but rather to provide fields of sufficient realism. To that end the proven LES solver of Forkel, Kempf, Stein and co-workers has been used, as it has a history of good predictions for a variety of flame configurations [45, 46, 82, 84, 143]. This allows simulations of previously studied flame configurations to be run without great consideration of the models or their parameters (as much as one can with such a powerful and yet relatively immature predictive technique).

### 2.3.1 The Filtered Mass and Momentum Equations

Large Eddy Simulation enables time-resolved solutions to be obtained by spatial filtering of the governing equations (eqs. 2.2, 2.3, and 2.4) rather than time averaging. While, in

principle it is possible to use any low-pass filter, in general a box/top-hat filter is used. Box filters can be conveniently implemented, when the governing equations are discretised using finite volumes, by using the grid to perform the filtering, and in this work such *implicit* filtering is used.

The filtered versions of the mass, momentum and mixture fraction equations are presented in eqs. 2.29, 2.30, and 2.33, where an over bar denotes filtering and a tilde represents Favre filtering. (Favre filtering simplifies the mathematical treatment of variable density flows.)

$$\frac{\partial \overline{\rho}}{\partial t} + \frac{\partial}{\partial x_i}(\overline{\rho}\widetilde{u_i}) = 0 \tag{2.29}$$

$$\frac{\partial}{\partial t}(\overline{\rho}\widetilde{u_i}) + \frac{\partial}{\partial x_j}(\overline{\rho}\widetilde{u_i}\widetilde{u_j}) = \frac{\partial}{\partial x_j}\left[\overline{\rho}\widetilde{\nu}\left(\frac{\partial \widetilde{u_j}}{\partial x_i} + \frac{\partial \widetilde{u_i}}{\partial x_j}\right) - \frac{2}{3}\overline{\rho}\widetilde{\nu}\frac{\partial \widetilde{u_k}}{\partial x_k}\delta_{ij} + \overline{\rho}\tau_{ij}^{\text{sgs}}\right] - \frac{\partial \overline{p}}{\partial x_i} + \overline{\rho}g_i \tag{2.30}$$

The filtering of eq. 2.3 yields a term, $\tau_{ij}^{\text{sgs}}$, called the sub-grid stresses which represents the effect of the small unresolved scales on the flow and requires modelling. A variety of models have been suggested and here we consider only the model employed in this work. The interested reader is referred to the book by Sagaut [132] for further detail.

In this work the sub-grid stresses are modelled using an eddy viscosity approach as follows. Firstly the assumption is made that the effective viscosity of the flow is the sum of the molecular viscosity and the turbulent viscosity that represents the effect of the small scale turbulence. Based on this, reformulation of eq. 2.30 gives the turbulent viscosity ($\nu_t$) as the only unclosed term. Closure is then obtained by using the Smagorinsky model [140] seen in eq. 2.31. The Smagorinsky model features a model constant $C_s$. Appropriate values for $C_s$ can vary even within a flow domain, and so the Germano procedure [50] is employed to dynamically determine the model constant.

$$\nu_t = (C_s\Delta)^2\left|\frac{1}{2}\left(\frac{\partial \widetilde{u_i}}{\partial x_j} + \frac{\partial \widetilde{u_j}}{\partial x_i}\right)\right| \tag{2.31}$$

### 2.3.2 Modelling Combustion in LES

To include the effects of chemical reactions in LES additional transport equations need to be solved as shown in section 2.1. As with the momentum equation, the filtering of these additional equations yields terms which for closure require modelling. Here we provide a brief overview of how this modelling is commonly performed in combustion LES before providing details of the specific approach used in the (non-premixed) LES performed in this work. The models are separated into the combustion modes to which they are usually applied; however, please note that several models are applicable to both modes as well as other simulation approaches such as RANS.

#### 2.3.2.1 Premixed Combustion

For premixed combustion several approaches are used to include combustion in an LES. These include: the *G-equation* method, which treats the flame as a surface and solves an additional scalar transport equation for the abstract scalar $G$; Linear Eddy Modelling, which describes mixing with a 1D sub-grid model; and a stochastic approach based on the transport of the PDF (Probability Density Function) of the scalars. These later two approaches can be applied to other modes of combustion and permit detailed chemistry to be used, but do require significant extra computational expense. Most common are models based on the transport of a reaction progress variable, $c$, as defined earlier in eq. 2.10, and here we consider these further to illustrate the modelling of premixed combustion in LES.

The transport equation for $c$ has been seen previously in eq. 2.11. Once appropriately filtered, this equation is formulated as shown in eq. 2.32 and all three terms of the right-hand side require modelling for closure.

$$\frac{\partial(\overline{\rho}\widetilde{c})}{\partial t} + \frac{\partial}{\partial x_j}(\overline{\rho}\widetilde{c}\widetilde{u_j}) = \frac{\partial}{\partial x_j}\left[\overline{\rho}(\widetilde{cu_j} - \widetilde{c}\widetilde{u_j})\right] + \frac{\partial}{\partial x_j}\left(\overline{\rho D_c \frac{\partial c}{\partial x_j}}\right) + \overline{\dot{\omega}_c} \tag{2.32}$$

Of these terms it is the third term, the filtered chemical source term $\overline{\dot{\omega}_c}$, that is of the most concern to modellers, and we consider the other terms only briefly. The first term is the sub-grid scalar flux term and can be adequately modelled in LES by a gradient diffusion assumption. The second term represents the molecular diffusion and can be separated into resolved and sub-grid components, with the later often neglected, though some approaches choose to model this term.

In the Eddy Break Up (EBU) model it is assumed that reaction kinetics are negligible and that Reynolds numbers are high. This allows the filtered chemical source term to be approximated from $\widetilde{c}$ and a characteristic turbulent time scale. The Bray-Moss-Libby (BML) analysis allows similar algebraic expressions to be formed from resolved quantities by assuming an infinitely thin flame, meaning $c$ can have only binary values. One modelling approach based on the BML analysis is the Flame Surface Density (FSD) approach which uses the level of turbulent flame wrinkling and a flame speed to determine the chemical source term. The Flame Surface density (of units $m^2/m^3$) represents the level of wrinkling of the flame and can be obtained algebraically or from an additional transport equation.

Typically in premixed combustion LES, the flame thickness is less than the grid resolution which can lead to numerical difficulties when large gradients exist across the flame front. One modelling approach overcomes this issue by artificially thickening the flame so it can be resolved, although this affects the turbulence chemistry interaction and some correction may be necessary.

#### 2.3.2.2 Non-Premixed Combustion

For LES of non-premixed combustion, models based on a conserved scalar that defines the state of mixing are popular, not least due to their relatively low computational overhead. These approaches assume that the chemistry is related to the mixing in the flow field, and hence all (dependent) scalars are functions of the conserved scalar. This scalar is typically called the mixture fraction and has been shown earlier in eq. 2.5. The filtered transport equation for $\xi$ is given in eq. 2.33.

$$\frac{\partial}{\partial t}(\overline{\rho}\widetilde{\xi}) + \frac{\partial}{\partial x_j}(\overline{\rho}\widetilde{\xi}\widetilde{u_j}) = \frac{\partial}{\partial x_j}\left(\overline{\rho}\widetilde{D_\xi}\frac{\partial\widetilde{\xi}}{\partial x_j} + \overline{\rho}\Xi_j^{\text{sgs}}\right) \tag{2.33}$$

The sub-grid mixture fraction $\Xi_j^{\text{sgs}}$, present in eq. 2.33, is modelled in this work using an *eddy diffusivity* approach that assumes that the effect of the small scale turbulence on mixing is similar to that of diffusion. This produces eq. 2.34 which allows closure of the filtered mixture fraction transport equation.

$$\Xi_j^{\text{sgs}} = D_{\xi,t}\frac{\partial\widetilde{\xi}}{\partial x_j} \tag{2.34}$$

The simplest form of the conserved scalar approach assumes equilibrium chemistry, making the dependent scalars a function of the mixture fraction alone. More sophisticated forms are dependant on the scalar dissipation rate as well as the mixture fraction, and allow non-equilibrium or finite rate effects to be captured. Conditional Moment Closure (CMC), developed by Klimenko [86] and Bilger [20] is one such approach which solves for conditional averages of the scalars on the basis that turbulent diffusion can be better modelled in mixture fraction space. The popular, Laminar Flamelet approach is applied in this work for LES of non-premixed flame configurations, and is explained further in the following section.

#### 2.3.2.3 Steady Laminar Flamelets

Flamelet approaches have been developed by Williams [168] and Peters [121] and are based on describing the turbulent flame as a series of small laminar flames located around the stoichiometric mixture fraction. Exploitation of the *steady flamelet* equation (eq. 2.35) then allows determination of the concentration of species $B$, and other dependent scalars, from the scalar dissipation $\chi$ and the mixture fraction. ($S_{X_B}$ in eq. 2.35 is a source term.)

$$-\rho\frac{\chi}{2}\frac{\partial^2 X_B}{\partial\xi^2} = S_{X_B} \tag{2.35}$$

In an LES of turbulent flows however, the sub-grid contributions of $\chi$ and $\xi$ are significant, and the sub-grid PDF of $\chi$ and $\xi$ must be modelled to account for the turbulence-chemistry interaction. For the scalar dissipation a simple Dirac function is assumed. For

the mixture fraction the sub-grid PDF is modelled as a $\beta$ function defined by the mixture fraction and its sub-grid variance.

The LES solver employed, therefore solves for mixture fraction which, together with a modelled sub-grid mixture fraction variance and scalar dissipation, is used to determine the dependent scalar concentrations from a *look-up* table of integrated values. Further detail on the Flamelet approach is given in Peters[122].

#### 2.3.2.4 Discretisation Scheme

In the LES solver used the filtered mixture fraction eq. 2.33 is discretised in space using a total variation diminishing scheme. For momentum a second order central difference scheme is used with oscillations limited by viscous stresses and pressure correction. Time integration is achieved using a 3-step Runge-Kutta scheme. The discrete equations are solved on a staggered cylindrical grid using the Finite Volume Method. In the type of implicit filtering used the width of the filter is determined by the size of the local finite volume.

# Chapter 3

# Computed Tomography of Chemiluminescence

Spatially and temporally resolved, whole volume, data is needed for improved validation of LES and to give further insight into turbulent combustion. In this work, a new measurement approach, called CTC, is developed that returns the 3D chemiluminescence intensity profile of a flame from multiple simultaneous images. In this chapter we define the generic CTC concept that is the basis of all the specific set-ups used in this work, and then through a review of previous work consider the potential capabilities of such a sensor. Following this our implementation of the CT algorithm is described.

## 3.1 The Sensor Concept

The CTC concept consists of two main stages: in the first, 2D projection measurements of a 3D chemiluminescent field are taken from multiple angles; in the second, a CT algorithm is used to reconstruct the 3D field from these 2D views. In this work several different CTC configurations have been investigated, and in this section we only define the two stages in general terms and leave more specific details to the relevant later sections.

The generic set-up is shown in fig. 3.1. Here the object flame is located at the centre of a circle of $N_q$ detectors which each provide a *view* (a set of $N_p$ projections) of the flame at a viewing angle of $\theta$. The specific number of views used is dependent on the application and defines the resolution of the reconstruction. (Except for the unlikely case where the angular sampling frequency $N_q$ is higher than the view sampling frequency $N_p$. This is further discussed in section 4.2). For stationary objects, as found in medical applications, measurements from many viewing angles can be taken by rotation of a single detector and very high resolutions, up to $N_p$, are possible. However, in this work resolution in time is desired and the measurements from different angles need to be taken synchronously,

meaning that multiple detectors are necessary.

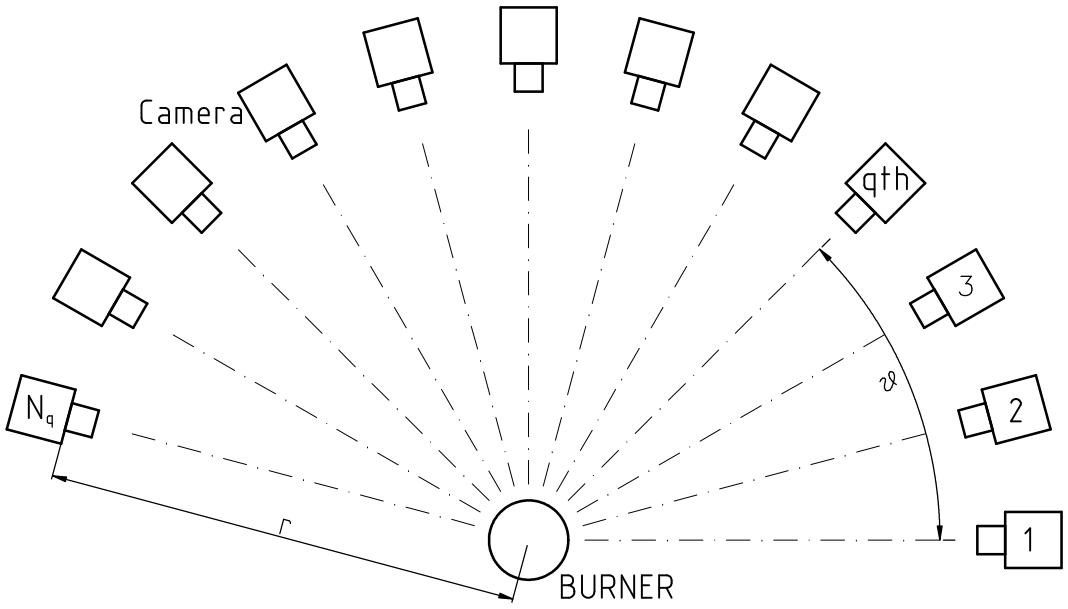



Figure 3.1: The generic CTC set-up. $N_q$ cameras are located around a burner at radius $r$. The angle of each camera is $\theta$.

The CTC sensor uses CCDs and lenses (a camera) to simultaneously obtain the projection measurements of the chemiluminescence intensity of a flame. A lens serves to increase the collection solid angle, and the light signal reaching the CCD, and therefore improves the Signal-to-Noise Ratio (SNR). In general each camera view has six degrees of freedom, although in the 2D case this is reduced to three: an $x$ and $y$ location, and a rotation $\theta$. For the best reconstructions, views should ideally be spaced evenly through 180° and at constant radius. Iterative algorithms such as ART do allow considerable flexibility in the view locations when needed, for example where optical access is limited, but with an associated reduction in the reconstruction accuracy. In this work we consider unconfined laboratory flames with generally good optical access and even spacing is usually used.

The CT algorithm requires that the location and orientation of all the projections be known to a sufficient precision, i.e the views need to be registered. Manual measurement of the camera locations is increasingly impractical with greater numbers of views, and instead an approach is used that is based on the comparison of images which feature common reference points. In this approach to view registration, once the cameras have been approximately located, images of a common calibration object can be taken and post processed to determine the view locations to higher precision. The disadvantages of this approach are that the precision is limited by the resolution of the camera images and a

suitable calibration object is needed for each parameter to be determined, though in the 2D case only the radius and angle $\theta$ are needed.

Standard CT algorithms use just 1D projection data, i.e. one co-planar row from the measurement images, to produce a 2D section of the object. The 3D reconstruction is then built up by stacking these sections. We note that this requires that all the cameras are located in the same plane and that out of plane effects are negligible. Iterative ART algorithms are used in this work because they perform well with the limited and noisy projection data likely to be encountered [4, 152]. An additional advantage of ART is that its simplicity makes it relatively easy to extend; for example, into 3D to permit arbitrary camera locations. Further details of the ART implementation are given later in this chapter.

## 3.2 Prior Art

Although most associated with medical diagnostics and X-ray attenuation in particular, Computed Tomography is a general technique that can be applied to any integral measurements. For this reason CT has been used successfully in many science and engineering applications; for example, Seismology, electron microscopy, and Non-Destructive Testing (NDT). Since its inception around 1970 Computed tomography has also seen use in general fluid dynamics, for example for visualisation purposes [71], and later in combustion applications. Notable applications in combustion include the work of Elsinga et al. [42], who have enhanced a Particle Image Velocimetry (PIV) system using CT to return a 3D velocity field, and Tse et al. [149] who investigate smouldering combustion using ultrasonic measurements.

In gaseous combustion, CT is most often coupled with optical line-of-sight measurements, though the application of Electrical Capacitance Tomography (ECT) to combustion [65, 167] is another notable exception to this. The use of optical CT in combustion can be split into three main groups: those using laser absorption, those based on the natural flame emissions, and those using refractive index measurements such as Schlieren or Interferometry images. Of these, the studies based on natural emissions are the most relevant to this work and are considered in greater detail in the following section. However, many issues, such as limited angles, are common to all approaches and so we begin with a brief review of the other measurement modalities.

Beginning with the work of Emmerman and co-workers [43, 133] laser absorption CT has been used in many studies [15, 16, 19, 31, 41, 55, 56, 106, 109, 115, 139, 138, 141, 170]. Laser absorption measurements allow absolute concentrations of specific species (e.g. $CH_4$ or CO) to be found for a whole section of a reacting flow. However, the

need for multiple laser sources has limited the number of angles from which projections can be measured, and restricted measurements to a single plane. (We note that less expensive diode lasers, as used by Wright et al. [170], would reduce such restrictions.) Early work [15, 43] used numerous mirrors or optical fibres to increase the numbers of views; however, these systems were relatively complex and were never implemented beyond the proof of principle stage. More recently, laser absorption CT has been performed using much lower numbers of views for the control and monitoring of combustion [31, 41, 109, 170] processes where pressure vessels limit optical access.

Tomographic reconstructions of Refractive Index (RI) measurements have been performed for flow visualisation purposes [71, 92], and have also been used to provide temperature and density information [1, 44, 88, 135]. Large numbers of views have been obtained by the rotation of steady flames, as performed by Agrawal et al. [1] and Feng et al. [44], or by using mirrors, as Snyder and Hesselink [141] and Schwarz [135] have implemented. A source is required for RI measurements, as with laser absorption, but unlike laser absorption the measured projections are usually 2D, and 3D reconstructions have been achieved via the stacking of sections [1, 44, 141, 135].

### 3.2.1 Focus on Emission Computed Tomography

Computed Tomography was first applied to chemiluminescence emission measurements by Hertz and Faris [70] and earlier Uchiyama et al. [150] used CT in their infra-red flame temperature sensor. Emission CT has since been mostly coupled to thermometry [25, 34, 61, 74, 94, 155] based on the infra-red emissions of soot or emissions from pervasive post-combustion species such as $H_2O$. The freedom from a source is advantageous where optical access is limited and many of these studies consider emission CT for the monitoring and control of furnaces. Edwards et al. [40] consider the control of a gas turbine combustor, and Leipertz et al. [93, 119] measure UV chemiluminescence emissions for industrial control, and identify several species as indicators of temperature and NO. Visible imaging of the flame structure is also performed for an SI engine by Philipp et al. [123].

Emission CT has seen less interest outside of control applications, and only Hertz and Faris [70] have used CT with a specific chemiluminescent transition, $CH^*$. (Although several studies [30, 35, 157] have used Abel transforms to localise measurements of $CH^*$, $OH^*$, and $C_2^*$ chemiluminescence for axisymmetric flames.) Quantitative species concentrations have been achieved by Hertz and Faris [70] albeit with a significant uncertainty of 40% resulting from errors in calibration. Calibration based on Rayleigh scattering has been performed by Walsh et al. [157] and DeLeo et al. [35] have used a deuterium lamp for absolute concentrations and report good agreement with simulations. Ishino

and Ohiwa [73] consider the 3D structure of a flickering flame using broad band (400 – 600 nm) chemiluminescent emissions, which includes $CH^*$, $HCO^*$, and $C_2^*$.

Where emission measurements are performed the attenuation of the emitted radiation within the flow domain (termed self-absorption) may be significant. Though studies concerned with visible emissions [70, 73] benefit from the optical thinness of most flames at these wavelengths. In the infra-red or ultra-violet ranges self-absorption can be significant and is an important consideration in optical thermometry studies, particularly those based on IR soot emissions; for example [33]. The issue of self-absorption is returned to in greater detail in section 5.1.

The prevalence of control applications in the literature has meant generally few viewing angles are considered and low resolutions achieved, particularly where time resolved measurements of turbulent flames are performed. The recent work of Ishino and Ohiwa [73] is the exception to this, using 40 simultaneous views of a (minimally) turbulent flame, and demonstrating the high quality of reconstructions that can be obtained when optical access is good and sufficient views can be measured.

### 3.2.2 Focus on the Algorithms

In the initial application of CT to optical combustion measurements, for example [43, 133, 141, 150], the standard back projection medical algorithms were used. However, these algorithms require many views to achieve good reconstructions leading to complex view measurement systems as presented by Beiting [14]. Iterative algorithms generally require less views than analytical algorithms and are the most commonly employed in the literature largely because of the prevalence of control applications and their associated optical access demands. The principle reason for the use of back projection algorithms is their speed; however, with modern computing resources, well implemented iterative algorithms require relatively little post processing time, and the flexibility and accuracy of these algorithms makes them preferable.

Iterative algorithms based on ART are used in most studies, for example [31, 55, 70, 42], and are found to perform well in the presence of noise and with irregular and non-parallel projections. The MLEM iterative algorithms are favoured in nuclear medicine emission studies such as SPECT [26] for their improved noise performance, and Ishino and Ohiwa [73] and Lim et al. [94] use such algorithms in their work, though the convergence of these algorithms is typically much slower than ART. Variations of ART have been investigated in several studies [153, 88, 108] using synthetic data referred to as *Phantoms*. These studies generally favour the multiplicative version of ART, although this is greatly influenced by the shape of the Phantom, and the continuous cosine and Gaussian Phantoms considered are less representative of chemiluminescence fields.

Another demonstrable advantage of ART is its adaptability. Several studies have tailored ART to specific applications; for example, Mishra et al. [108] have adapted MART for interferometric projections. The selection of different basis functions is another approach. While the Pixel/Voxel basis is most common, several studies use cubic splines [152, 88], blobs [104, 105], or more optimal bases [39] to utilise *a priori* information such as smoothness to improve reconstructions. The use of such bases may not be appropriate for chemiluminescence measurements as we can expect steep gradients in some configurations.

### 3.2.3 Resolutions achieved

The spatial resolution of any CT reconstruction is largely a function of the number of viewing angles at which projections can be taken, as described in many studies such as [70]. In the reviewed literature, a wide range of angles have been used, from 2 – 96 [15, 93], which reflects the wide range of applications of CT. For studies that consider industrial *in situ* sensor configurations [33, 93, 123], generally less than 12 viewing angles are used. This is partly due to difficulties with optical access, but also because often only low resolutions are needed in such applications. Although less common, several studies have practically demonstrated configurations that use high numbers of views [14, 73, 141] to achieve very high resolution reconstructions. All three main measurement modalities feature high and low numbers of views; however, generally higher resolution reconstructions are achieved where no strong laser source is required and optical access is good.

The resolved wavelength, $\lambda$, of selected studies is shown alongside the number of views used, $N_q$, in fig. 3.2. These studies are those for which it has been possible to determine the resolved wavelength of their reconstructions, and they have generally been estimated from twice the pixel resolutions (based on the Nyquist criterion of two points for resolution). For Emmerman et al. [43], Snyder and Hesselink [141], and Uchiyama et al. [150] this information was not available and so the values quoted in those studies are used without modification. Data from all measurement modalities is included, although only experimental results using CT algorithms have been shown; results from Phantom studies or those using Abel transforms are not included. Several studies, for example Beiting et al. [14], have defined configurations with a high number of views but only report results from a proof-of-principle study using just one rotated view of a steady flow. Ideally only studies with time resolved measurements would be shown; however, to provide more data other studies are also included.

The data in fig. 3.2 generally follows the expected trend of more views yielding higher resolutions. However, some scatter is seen and in some cases lower resolution is observed

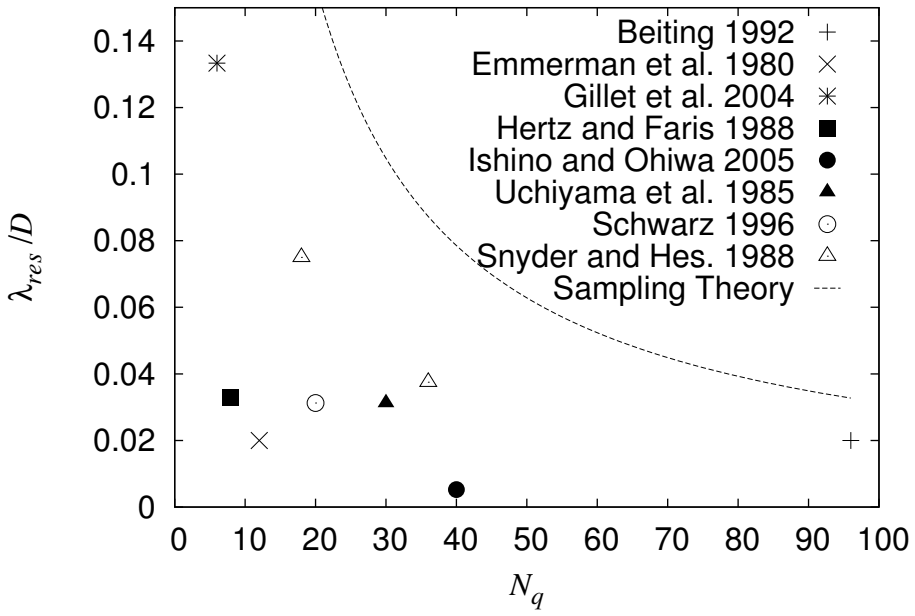



Figure 3.2: Resolutions achieved in other selected studies. Crosses represent laser absorption measurements, filled shapes are emission measurements, and unfilled shapes refractive index measurements. The resolved wavelength $\lambda_{res}$ has been normalised by the object domain diameter $D$ to allow comparisons. Also shown is a theoretical relationship obtained from sampling theory.

with more angles. This can, in part, be expected due to the different modalities and algorithms used; however, the lack of a common criteria for judging resolution in these limited angle cases is thought likely to be more significant. The resolution is often inferred from the size of the reconstruction grid which in most cases is given without justification. Gillet et al. [56] are however, the exception to this as they perform a complimentary Phantom study to determine an appropriate reconstruction resolution. Differences in the flow configurations may also contribute to the scatter; although, these are quite subtle, with all except Hertz and Faris [70] considering a circular jet of some sort.

The result for Emmerman et al. [43] appears anomalous as it has significantly better resolution than other studies using comparable numbers of angles. The significant noise seen in the reconstructions of Emmerman et al. [43] suggests that less stringent criteria are applied in this case. Interestingly, this resolution differs by an order of magnitude from a theoretical value derived in the same work (also given in [112]). This theoretical relationship between the number of angles and resolution is included for comparison in fig. 3.2, and much lower resolutions are predicted using this theory than are reported in all the literature. This adds to the ambiguity in the resolutions achieved and in section 4.3 we reconsider this issue in greater detail for these limited angle configurations.

Nevertheless, an initial estimate for the number of angles required to meet our objectives can be obtained using previous studies as a basis. An LES would typically resolve a laboratory scale flame with 64 – 128 grid points which implies that angles in the range of 20 to 50 would be required. Furthermore, studies such as that of Ishino and Ohiwa [73] and Synder and Hesselink [141] have demonstrated absolute resolutions of the order of 1 mm which are again representative of a typical LES. The numbers of angles required appears challenging at first. However, several studies have achieved such numbers already, and with the continued decline in the price of CCD cameras, deploying large numbers of cameras is entirely possible.

The temporal resolution achievable by any experimental technique is dependent on any factor that effects the Signal-to-Noise Ratio (SNR) of the measurements. (These factors are considered in greater detail in sections 4.2 and 6.1.2.1.) Thus the phenomena measured, and the detector used, are important, and the results of other modalities are not suitable indicators in this instance. Generally speaking stronger signals and more sensitive detectors with lower noise allow higher temporal resolutions. Measurements that incorporate a strong source can thus achieve very high resolutions as seen in the work of Beiting [14] who uses laser sources and claims a resolution of 200 ns. For chemiluminescence the signal strength is dependent on the chemical characteristics of the fuel and oxidiser, the excited species under consideration, and also the geometry of the flame. Hertz and Faris [70] consider a methane air flame slot flame and use intensified photodi-

ode arrays to achieve a temporal resolution of 16 ms (SNR of 25). Ishino and Ohiwa[73] achieve a comparable resolution of 1.2 ms using film for their study of a turbulent propane air flame. These resolutions would only be sufficient for *instantaneous* imaging (effectively where no blurring is observed) of only slightly turbulent flows. For more turbulent flames resolutions the order of 10 $\mu$s are likely to be necessary, though this is clearly dependent on the specific flow considered.

These studies however, use film and photodiodes for detection and generally a modern CCD could be expected to be more sensitive. Furthermore, signals could also be increased by using oxygen as the oxidiser or *brighter* fuels such as Acetylene, as also noted by Hertz and Faris [70]. Also, as spatial resolutions are limited by the number of angles, detector resolution can be forgone to allow larger and more sensitive CCD/CMOS pixels. The potential resolution for chemiluminescence is demonstrated by Bohm et al. [23] who achieve 10 $\mu$s exposures at repetition rates of 30 kHz with their high speed LaVision HSS5 CMOS camera. Repetition rates are dependent only on the hardware employed, although cameras of this specification would be prohibitively expensive for multiple deployment and would need to be used with mirrors or fibre optics to achieve sufficient views.

### 3.2.4 Conclusions

A review of the literature has shown that time resolved 3D turbulent flame structural information can be obtained from a sensor based on Computed Tomography of 2D Chemiluminescence images. One study [73] achieves this, though for only a weakly turbulent flame, broad band emissions and using a film based camera. To achieve resolutions in time and space that are comparable to those of turbulent LES, sufficient numbers of sufficiently sensitive cameras will need to be provided. Several studies have employed large numbers of detectors, and with the increasing capability and decreasing cost of suitable cameras, higher resolutions are feasible. Quantitative measurements of excited species concentrations have also been demonstrated, and used to assess reaction mechanisms and combustion simulations.

Configurations with limited numbers of angles pervade the literature and the reconstruction problem is typically under-determined. In such cases, ART type algorithms have demonstrated improved performance over back projection algorithms. Furthermore, the flexibility in projection spacing and the ability to reconstruct from even noisy datasets makes ART preferable to back projection in many cases. This flexibility, and the scope for tailoring the algorithms to the specifics of chemiluminescence measurements, means that ART and its multiplicative form MART are the main algorithms considered in this work. Further detail on the ART implementations follow this section.

## 3.3 The Basic ART Implementation

The two variants of the ART algorithms were implemented in a purpose written Fortran95 code using a pixel basis. The first algorithm used was the relaxed additive ART algorithm of Gordon et al. [59], referred to as ART, and the second, a multiplicative version of ART used by Verhoeven [152] and referred to as MART[1]. Both are defined in section 2.2.3. A third algorithm was also initially considered that mimics SART [6] by applying the correction from all projections of view simultaneously. However, preliminary testing showed that this approach took three orders of magnitude more iterations to converge[2] to a solution with equivalent accuracy to that obtained using ART or MART; for example 10000 iterations compared to 20 for ART and 12 for MART. For this reason it has not been investigated further in this work; however, the main reported advantage of simultaneously applying the correction for a whole angle is the improved noise performance and in future work where noise is significant its use should be considered. SART generates the correction the in the same way as the ART algorithm, and so the results of ART are representative of SART in the noise free case.

An important part of any ART implementation is the choice of projection geometries. In this work the *exact* strip integrals are used (except in sections 5.1 and 5.2.2), and the pixel weighting factor, $w_{qpv}$, is defined as the fractional intersection area of the pixel $v$ with the projection $p$ of view $q$. The calculation of the projection/pixel intersection areas is non-trivial and in this 2D case the approach of Deutsch [36] has been used which computes the exact intersection of an arbitrary polygon with a grid of pixel as follows.

First the integral geometry (polygon) is defined by a set of line vectors, in this case four as shown in fig. 3.3. The first vector (*AB* in in fig. 3.3) is traversed and the locations where the vector intersects the pixel edges are determined. These locations are each then localised to each pixel so that (0,0) is the pixel bottom left and (1,1) is the pixel top right. This gives $S_{in} = [x_{in}, y_{in}]$ and $S_{out} = [x_{out}, y_{out}]$ as seen in fig. 3.3. These local coordinates are then used to calculate, for each pixel, the area $a$ that is to the right of the vector using eq. 3.1 where $y_{max}$ is the maximum of $y_{in}$ and $y_{out}$, with $y_{min}$, $x_{max}$, and $x_{min}$ similarly defined. All pixels to the right of vector that are not intersected by the vector (*AB*) under consideration are assigned areas of unity.

$$a = 0.5(y_{max} + y_{min})(x_{max} - x_{min}) + y_{max}(1 - x_{min}) - y_{min}(1 - x_{min}) \tag{3.1}$$

The same process is then repeated for the remaining vectors and all the areas are then summed according to the rule that areas from vectors with negative gradients are de-

[1] Note, the MART algorithm was found numerically unstable without relaxation, as also reported by Verhoeven

[2] And two orders of magnitude more time.

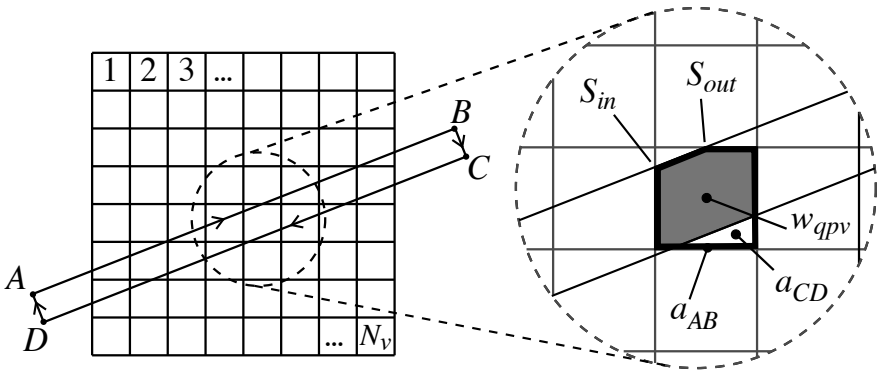



Figure 3.3: The computational method used to determine the 2D fractional area, $w_{qpv}$, of pixel $v$. The fractional area is equal to the area to the right of vector $AB$ ($a_{AB}$) minus (because of negative gradient) the area to the right of vector $CD$ ($a_{CD}$). The area to the right of a vector is determined from the local pixel intersection points $S_{in} = [x_{in}, y_{in}]$ and $S_{out} = [x_{in}, y_{in}]$ using eq. 3.1.

ducted; this summation provides the fractional areas.

The projection matrix $W$ is used to store all the weights of each pixel in the reconstruction, for each ray in the projection data set, resulting in a $N_p \times N_v$ matrix. With high resolution reconstructions using many projections this matrix can become very large, though in practice storing this matrix is not a problem for the 2D case. The effect on the speed of the algorithm is more problematic, because the number of operations required to perform the forward and back projections steps of ART increases greatly as $W$ increases in size. However, the projection matrix for a particular ray is very sparse and in this implementation we have taken advantage of this to greatly reduce the number of operations required on the forward and back projection steps. The significance of the speed-up achieved in doing this cannot be overstated. The use of sparse matrices has the added benefit of reducing the amount of computer memory required, and so also allows higher resolution reconstructions to be performed which is of particular importance when 3D problems are considered in later work.

This ART implementation allows the views to be addressed in any order. When the number of views are small ($< 10$) they are addressed in approximately the most orthogonal order; for example in this work when 10 views are used the order is 1, 5, 10, 3, 7, 6, 9, 4, 8, and 2. However when many views are used the orthogonal order is more difficult to determine and a simple random scheme is used which produces equivalent results as seen in Guan and Gordon [60]. The direct effect of projection ordering on the convergence speed is minimal in this work, because typically more than 10 iterations are required. However, reducing the correlation of the projections allows higher relaxation factors to

be used and so convergence speed is improved. A relaxation factor $\beta$ of 0.8 was used throughout this work and proved stable and of acceptable speed.

The reconstructions are judged to be converged once the absolute difference of the sum of the field, from one iteration to the next, is below the threshold $\Delta_c$ as seen in eq. 3.2. In the literature values for $\Delta_c$ are typically $10^{-3}$ – $10^{-6}$.

$$|f^h - f^{h-1}| \leq \Delta_c \times \beta \tag{3.2}$$

The use of variance criteria for ART was suggested by Gordon [58], although Verhoeven [152] notes that this does not predict the best stopping point and this was also confirmed in preliminary testing. The criterion used in this work has been used successfully by Mishra et al. [108] and has the advantage of being quicker to test as fewer operations are required, also this criterion is consistent with ART's proven minimisation of the norm of $\mathbf{f}$ [147].



(a) Phantom (b) ART 128 angles (c) MART 128 angles

Figure 3.4: Example reconstructions of the Flame D Phantom using ART and MART. The reconstructions used 128 evenly spaced angles with 128 projections per angle.

Example reconstructions of a section of *Flame D* using this described implementation are shown in figure 3.4. The section of Flame D is taken from an LES, described in greater detail later in section 4.1. These $128^2$ pixel reconstructions have used 128 angles each of 128 rays and took less than 10 seconds each to perform. As can be seen the reconstructions are practically identical to the Phantom, and demonstrate the correctness of the implementation.

# Chapter 4

# Testing Sensor Capability

This chapter is concerned with testing the capability of CTC to measure the chemiluminescent fields of non-axisymmetric turbulent flames. To assess whether CTC—using realistic numbers of views—can provide measurements of low enough error and high enough resolution for the validation of LES, two Phantom studies are performed. The first study looks at the accuracy achievable in the reconstructions with limited views, but unlike other studies, using realistic chemiluminescent datasets obtained from LES. The second study considers the resolutions that are achievable with limited angles, by using cosine Phantoms of known wavelengths. To clarify, we are concerned with the resolving power of ART, rather than the size of a pixel in the reconstruction grid. There is a degree of ambiguity in the literature regarding the resolution that ART actually achieves when using limited angles, and it is our aim to provide some clarity and also a basis for the estimation of the resolutions.

Following the Phantoms studies, the feasibility of the sensor concept is demonstrated using a single repositioned digital camera that provides experimental measurements of a non-axisymmetric steady Matrix Burner. We begin this chapter by describing the generation of the realistic Phantoms using LES adapted to predict the excited state species concentrations.

## 4.1 LES Phantoms

*Phantoms* are used in all CT development and are simply known fields through which the projection measurements can be simulated. The reconstruction obtained from these simulated measurements can then be compared to the original Phantom field and allow the error to be quantified absolutely. Furthermore, noise and other experimental phenomena can be incorporated into the simulated projections and studied. Phantoms are synthetically generated test cases and are often based on mathematical functions such as Gaussian

peaks [108], and cosines [152, 88]. Gillet et al. [56] have generated *realistic* Phantoms for their methane absorption measurements, and in this work we also use realistic Phantoms of excited state species, but taken from Large Eddy Simulations (LES) of turbulent flames. Such Phantoms provide a better basis for the assessment of the capabilities of CTC because they are considerably more representative of the fields likely to be encountered by the sensor. Furthermore, LES allows a variety of experimental configurations to be investigated, with not only varying flame geometries, but different chemiluminescent species as well.

Combustion simulations have been used previously for the development of a tomographic sensor by Torniainen et al. [148], who used LES to investigate the scalar field from a rectangular reacting jet. Later, Edwards et al. [40], from the same group, used a RANS simulation of a gas turbine combustor. In both these studies the LES/RANS data sets do serve as Phantoms, but they are primarily used to define optimal basis functions for tomographic reconstruction using a least squares approach. These optimal basis functions contain significant *a priori* information and dramatically reduce the number of views required. While acceptable for the control applications considered by Edwards et al. [40], using *a priori* information from simulations to reduce the number of required views is questionable where the measured data is later intended to validate LES as in this work.

Edwards [39] considers the emissions of excited $H_2O$ for Jet-A fuel ($C_{12}H_{23}$) and demonstrates how the measured signal intensity can be modelled for an optically thin ideal gas by a line integration of the excited molecule number density multiplied by three constants for that transition: an Einstein coefficient, Planck's constant, and the wavenumber. In this work a similar approach to simulating the emission measurements is used, whereby the intensity ($I_{qp}$) of a projection is given by the integral of the excited state concentration field $f(x,y)$ over the area $A$ which represents the projection geometry (in the 2D case). This is in turn approximated for the discrete LES Phantom fields by eq. 4.2 which is equivalent to the forward projection step of ART.

$$I_{qp} = \int_A f(x,y)\mathrm{d}A \tag{4.1}$$

$$I_{qp} \approx \sum_{v=1}^{N_v} w_{qpv} f_v \tag{4.2}$$

Some details of the LES solver used to generate the excited state Phantoms have been given previously in section 2.3. The implementation used is that of Stein et al. [143], Kempf et al. [82, 84], and Forkel et al. [45, 46] and has produced accurate results for several different flame configurations, both premixed and non-premixed, including single and opposed jet flames. The chemical state look-up tables used in the LES were updated to include $CH^*$ and $OH^*$ as well as other excited state species. Further detail of the

chemical modelling is given in the following section.

### 4.1.1 Simulating Excited State Species

To be able to include the excited state concentration fields in the LES it was necessary to update the chemical state look-up tables used in the LES to include these species. The look-up tables themselves are obtained from laminar solutions of a given chemical mechanism, and to be able to simulate the excited species suitable reactions need to be added to this mechanism. Some species and transitions that are commonly investigated have been given previously in table 2.1. Suitable reactions for these chemiluminescent species have been considered by Benvenutti et al. [17] for Ethanol ($C_2H_6O$) and are shown in table 4.1 where $hv$ is the energy of the emitted light and M represents quenching. The simulations performed by Benvenutti et al. using these reactions showed good agreement with their experimentally measured profiles.

| | | $A$ | $b$ | $E$ |
|---|---|---|---|---|
| a1. | $H + O_2 \longrightarrow O + OH^*$ | 9.75E+11 | 0.0 | 62178.2 |
| a2. | $CH + O_2 \longrightarrow CO + OH^*$ | 3.25E+11 | 0.0 | 0.0 |
| a3. | $OH^* + M \longrightarrow OH + M$ | 9.26E+11 | 0.0 | 0.0 |
| a4. | $OH^* \longrightarrow OH + hv$ | 1.00E+07 | 0.0 | 0.0 |
| a5. | $CH_2 + C \longrightarrow H_2 + C_2^*$ | 7.50E+11 | 0.0 | 0.0 |
| a6. | $C_2H + H \longrightarrow H_2 + C_2^*$ | 1.00E+10 | 0.0 | 0.0 |
| a7. | $C_2^* + M \longrightarrow C_2 + M$ | 2.04E+10 | 0.0 | 0.0 |
| a8. | $C_2^* \longrightarrow C_2 + hv$ | 1.00E+07 | 0.0 | 0.0 |
| a9. | $CH + O \longrightarrow CHO^*$ | 1.26E+12 | 0.0 | 0.0 |
| a10. | $CHO^* + M \longrightarrow CHO + M$ | 8.67E+11 | 0.0 | 0.0 |
| a11. | $CHO^* \longrightarrow CHO + hv$ | 3.33E+11 | 0.0 | 0.0 |
| a12. | $CH_3 + HO_2 \longrightarrow H_2O + CH_2O^*$ | 1.81E+11 | 0.0 | 0.0 |
| a13. | $CH_2O^* + M \longrightarrow CH_2O + M$ | 1.00E+10 | 0.0 | 0.0 |
| a14. | $CH_2O^* \longrightarrow CH_2O + hv$ | 2.26E+06 | 0.0 | 0.0 |
| a15. | $C_2H + O_2 \longrightarrow CO_2 + CH^*$ | 4.47E+13 | 0.0 | 104670 |
| a16. | $CH^* + M \longrightarrow CH + M$ | 6.50E+10 | 0.0 | 0.0 |
| a17. | $CH^* \longrightarrow CH + hv$ | 1.79E+07 | 0.0 | 0.0 |

Table 4.1: The reactions added to the mechanism of Lindstedt and coworkers [98, 130]. Units are J/mol and further details are given in Benvenutti et al. [17]

The reactions for the excited states were added to the reduced mechanism of Lindstedt and coworkers [98, 130] and solved directly in a 1D laminar opposed jet simulation by

Gkagkas [57]. For $OH^*$ the thermodynamic coefficients were taken directly from Burcat and Ruscic [29]. The coefficients for the remaining excited species were not available in this database and were provided by Robinson [131]; who estimated the thermodynamic coefficients by adding the energy of the emitted light, *hv*, to the enthalpy and Gibbs free energy of the ground state molecules. The simulation was performed for a partially premixed fuel stream of 25% methane and 75% air, by volume, that matched that used in the later considered flame configurations. The simulation was repeated over a range of strain rates from $10\,s^{-1}$ to extinction and used to construct the flamelet *look-up* tables.

In the RANS simulations used by Edwards [39], the ground state fields were found using equilibrium chemistry from which the excited state fields were determined using the Boltzmann distribution. In this work the excited states are considered as separate species and incorporated into the reaction mechanism used by the LES. This means that we obtain a *pseudo-instantaneous* excited state concentration field from the LES rather than the mean fields that would be given from RANS. This is an important distinction as the mean fields would be significantly smoother than the instantaneous fields, and be easier to reconstruct, thereby overestimating the capabilities of CTC performed at the high temporal resolutions intended.

### 4.1.2 LES Configurations and Results

Two different experimental flame configurations are investigated: a non-premixed jet flame and a non-premixed opposed jet flow. The jet flame considered is *Flame D*, which features in the TNF workshop's series of piloted jet flames [9]. This flame has an extensive body of experimental data having had scalar measurements performed at Sandia national Laboratories [10, 79, 11, 78] and velocity measurements at TU Darmstadt [134]. Flame D has fully developed turbulence ($Re = 22,400$) and features a degree of local extinction and re-ignition and as such has also been the subject of many numerical studies; for example, in references [95, 82, 113]. The Turbulent Opposed Jet (TOJ) configuration considered is that developed by Geyer et al. [51] and studied also at Sandia National Laboratories and more recently at Imperial College London. These flames are good test cases for this work because both are well understood and are considered benchmark cases for combustion LES. They provide two very different flame configurations; the jet flame is spread axially and has gradients in the radial direction that diminish with distance from the nozzle. For the TOJ a flat axially confined flame front is stabilised at the stagnation plane of the two jets and has the strongest gradients in the axial direction. Furthermore, while both are considered as non-premixed flames, the fuel stream in both cases is partially premixed to produce low soot flames that are suitable for optical measurements.

#### 4.1.2.1 Flame D

Flame D is a non-premixed piloted jet flame. The main jet has a diameter of 7.2 mm from which a partially premixed fuel stream consisting of 25% $CH_4$ and 75% air, by volume, exits at a bulk velocity of 49.6 m/s. The co-axial pilot has a diameter of 18.2 mm and an axial velocity of 11.4 m/s. The pilot burns a fuel equivalent to a mixture of $CH_4$ and air of 0.77 equivalence ratio. The whole burner is located in a co-flow of air at a velocity of 0.9 m/s.

The flame was simulated in a cylindrical domain extending 400 mm in the axial dimension and with a radius of 164 mm. The domain starts 10 mm upstream of the main nozzle exit, with the main jet and pilot nozzles represented as immersed boundaries: no flux is permitted across the cells concerned and a no-slip condition is enforced. The domain was discretised with finite volumes on a staggered cylindrical grid using $800 \times 32 \times 60$ cells in in axial, circumferential, and radial directions. The radii of the cells have variable spacing and are more refined at the centre of the domain where the highest gradients are seen. The simulations were performed using the inlet velocity conditions measured at TU Darmstadt by Schneider et al. [134] and with a zero-gradient condition at the outflow. Artificial turbulence was superimposed on the inlet velocity profiles using the procedure developed by Kempf et al. [83]. An example axial section of the instantaneous mixture fraction field is shown in fig. 4.1.

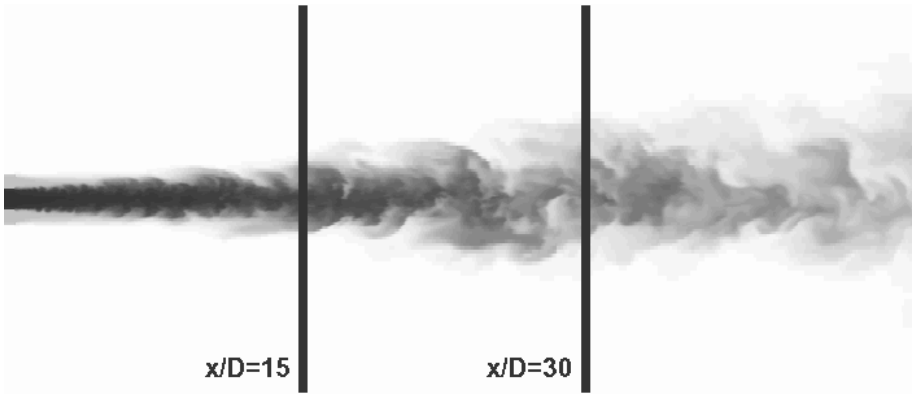



Figure 4.1: An axial section of the Flame D instantaneous mixture fraction field. Also shown are two downstream locations. D is the nozzle diameter of 7.2 mm.

To validate the simulation of the excited state scalar fields statistics were taken, once the simulation reached steady state, and compared with available experimental data. Unfortunately data for the excited states was not available for comparison; however, as we are concerned with obtaining representative fields only then comparison of other scalars will suffice. Figure 4.2 shows a comparison of the mixture fraction $\xi$ and the OH species mass fraction $Y_{OH}$, with the values measured at Sandia National Laboratories [10, 79]. Both the mean and RMS mixture fraction values show good agreement with the measured data

and are consistent with a simulation of the same flame by Kempf et al. [81]. The mean of $Y_{OH}$ peaks at a higher value than has been measured but the overall shape and magnitude, as well as the RMS, matches the experimental data reasonably well. These results give confidence that the excited state fields predicted in the LES are representative of those in the actual flame.

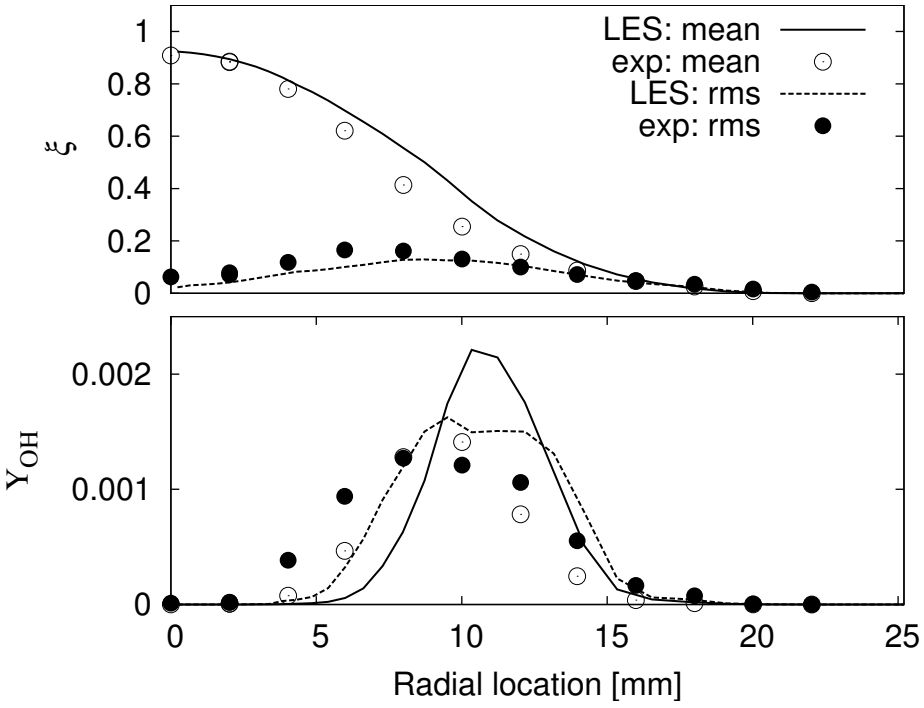



Figure 4.2: Radial profiles of mixture fraction $\xi$ and OH mass fraction $Y_{OH}$ taken from the LES alongside experimental data for an axial location 15D downstream of the nozzle.

#### 4.1.2.2 Turbulent Opposed Jet

The Turbulent Opposed Jet simulated in this work is that developed at TU Darmstadt [51] and consists of two vertically aligned and opposed nozzles of diameter 30 mm and separation 30 mm. The upper nozzle ejects air at a bulk velocity of 3.40 m/s, and the lower nozzle ejects the fuel at 3.61 m/s. The fuel is the same partially premixed methane/air mixture described previously for Flame D. The slight difference in the velocities of each jet stems from the momentum balance required for the stagnation plane to be equidistant from the nozzle exit. The nozzles are enclosed in a co-flow of nitrogen that prohibits the surrounding air from mixing with the jets. Perforated plates are located 50 mm upstream of either nozzle to generate turbulence ($Re = 6650$).

The opposed jet flame was simulated in a cylindrical domain extending 130 mm in the axial dimension and with a radius of 25 mm. The domain spans from one turbulence generating plate to the other and encompasses all of the nozzles and part of the co-flow.

The nozzles are again represented as immersed boundaries. The domain was discretised with finite volumes on a staggered cylindrical grid using $513 \times 64 \times 87$ cells in the axial, circumferential, and radial directions. (The cells were again refined towards the centre of the cylindrical domain.)

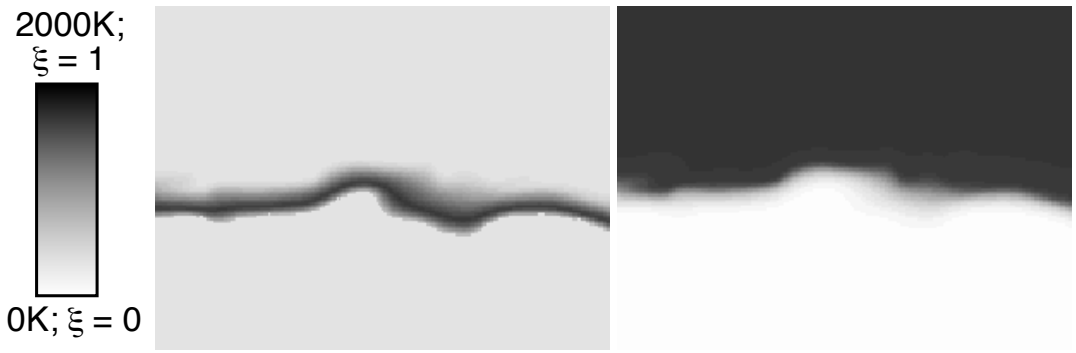



Figure 4.3: An axial section of the Turbulent Opposed Jet instantaneous temperature field (left) and mixture fraction (right). The top is the fuel side and these sections span 30 mm axially.

The simulations were run for sufficient time to get meaningful statistics, and example axial sections of the instantaneous mixture fraction and temperature fields are shown in fig. 4.3. The axial profiles for the mean and RMS mixture fraction and OH mass fraction are shown in Figure 4.4 and very good agreement is seen between the mixture fraction values measured by Geyer et al. [52] and those predicted by LES. Unfortunately no measurements of the OH mass fraction were available for comparison; however the magnitude and shape of the OH profile meets with expectations and given also the good agreement of the mixture fraction profile we can conclude that the LES results are sufficiently accurate for the needs of this work.

### 4.1.3 The Excited Species Phantoms

From each of the LES configurations a cross-section was extracted from the instantaneous field to serve as a Phantom. For the TOJ, the Phantom was taken from the $CH^*$ field close to the symmetry plane of the two jets. For Flame D, the Phantom was extracted from the $OH^*$ field at a location 30 nozzle diameters downstream of the jet exit, as indicated in fig. 4.1. The $CH^*$ and $OH^*$ species were selected as these are the most commonly studied in the literature, and so are likely to be the subject of future CTC measurements. Furthermore by selecting two different species we accentuate the differences between the fields because, as the previous chemistry simulations[1] have shown, $CH^*$ exists in a narrower mixture fraction range than $OH^*$, meaning higher gradients for $CH^*$ can be

[1] See section 4.1.1

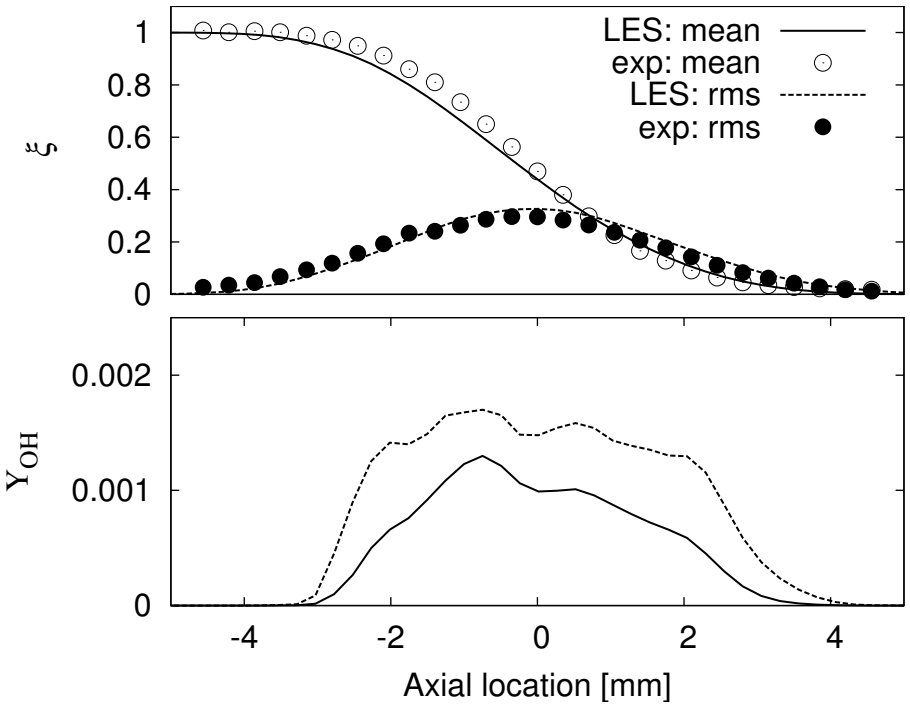


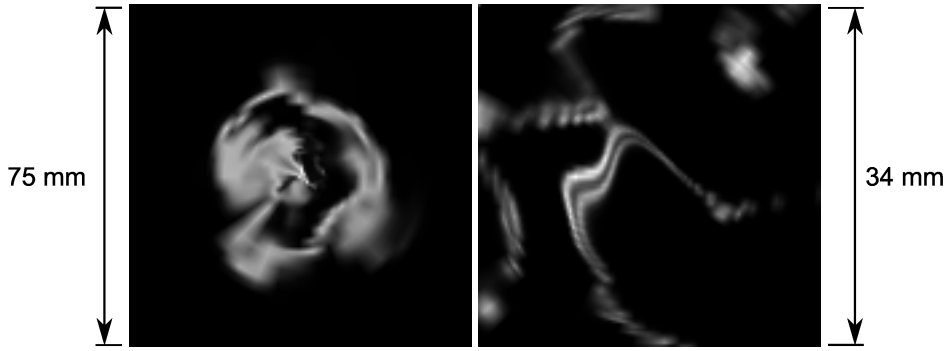

Figure 4.4: Axial profiles of mixture fraction $\xi$ and OH mass fraction $Y_{OH}$ taken from the LES alongside experimental data where available.

expected. Phantoms with different structures are desirable to better test the CTC method, and these Phantoms have been deliberately chosen to achieve this. The $CH^*$ opposed jet Phantom is sparsely populated and features reasonably high gradients in the measurement plane; the Flame D Phantom however, was extracted from a downstream location in the flame where the jet has spread and the field is more diffuse.

The radius of the LES domain extends far beyond the extents of the chemiluminescence fields and so the sections to be used as Phantoms were cropped to 75 mm by 75 mm and 34 mm by 34 mm for Flame D and the Opposed Jet respectively. These sections were then converted from the cylindrical LES grid to a $128 \times 128$ pixel Cartesian grid using bi-linear interpolation and the resulting Phantoms are shown in fig. 4.5. The cylindrical LES grid inherently suffers from reduced resolution at larger radii and artefacts of this can be seen towards the outer edges of each Phantom. These artefacts are accepted because it is desired to keep the resolution of the Cartesian Phantoms comparable to the resolutions of the LES at the centre of the cylindrical domain (this results from local grid refinement as previously discussed).

## 4.2 Reconstruction from limited angles

In combustion applications, and other non-medical applications, the number of viewing angles from which measurements can be taken is often limited. This can be the result of




(a) Flame D OH* (b) TOJ CH*

Figure 4.5: The 128 pixel, LES excited species Phantoms. For Flame D the section was taken 30 nozzle diameters downstream of the jet exit. For the Turbulent Opposed Jet (TOJ) the section was taken at the stagnation plane of the two jets.

limited access, for example due to an opaque pressure vessel, or as a result of equipment limitations. In this situation the reconstruction problem, as expressed as system of linear equations, is under-determined and the number of equations is less than the number of unknown values. In theory this can be resolved by reducing the reconstruction resolution, and so the number of unknown values. However, such a reduction in resolution can be avoided to some degree either by the use of *a priori* information or by accepting some error in the reconstruction. Limited angles are likely in this work as we desire high resolution reconstructions to permit useful comparisons to LES. The amount of error that one has to accept, or rather the number of angles needed for a certain error, is the subject of this section.

The problem of limited angles is somewhat offset in fluids applications by the smooth fields that are generally encountered as a result of diffusion effects. Such *a priori* information has often been used to improve limited angle reconstructions. Where a field is known to be axisymmetric, perhaps from averaging, the Abel transform can return the field using only a single projection. However when the field is non-axisymmetric series expansion methods allow significant *a priori* information to be exploited by the choice of different basis functions, seen for example in the work of Ravichandran and Gouldin [129] who use a Kaiser window type function as a basis to exploit *a priori* smoothness of the field. Even more *a priori* information is exploited by Torniainen et al. [148] and Edwards et al. [40] who use LES/RANS simulations to obtain optimal basis functions that minimise the number of unknowns. The approach of tailoring the basis functions to exploit specific *a priori* information can greatly improve results, but at the expense of generality, and at this stage of the current work only non-negativity of the scalar field is assumed *a priori*.

The effect of limited numbers of views on the reconstructions, using generic CT algo-

rithms, has been considered by Andersen [4], and for non-medical applications in several Phantom studies [152, 158, 159, 108] which all favour iterative series expansion algorithms over Fourier transform approaches. The work of Verhoeven [152] is the most referenced and determines the MART algorithm to have lowest errors without applying significant *a priori* information. However, ART is found to perform better than MART with smooth fields and with noisy projections. It is difficult to draw further conclusions from this study as it considers only five angles for a $25^2$ pixel reconstruction in its most relevant test case (a cosine-Gaussian Phantom with angles spread over $180^\circ$), and such an arrangement is only modestly under-determined by our considerations. (An LES section would typically consist of approximately $100^2$ pixels and we could expect to obtain in the region of 20 – 40 views). Mishra et al. [108] perform reconstructions of $100 \times 100$ pixels using 15 – 30 angles, but consider truncated views that are less likely to occur with open laboratory flames. Mishra et al. [108] do similarly favour MART, although they again note its sensitivity to noise. Furthermore, these studies have mainly used Gaussian and/or sinusoidal Phantoms that are much smoother than the chemiluminescent field of the TOJ Phantom for example. In the following section tests are conducted using a broader range of views and using realistic LES Phantoms to better assess the performance of ART and MART in limited views configurations.

### 4.2.1 LES Phantom Tests

The ability of the ART and MART algorithms to reconstruct an excited state species concentration field, from a limited number of viewing angles, is tested using realistic Phantoms from LES. To assess the reconstructions two global metrics are used: the correlation coefficient and the normalised RMS of the error field. The correlation coefficient, $r_{xy}$, seen in eq. 4.3, is a commonly used quantity when wanting to judge the correspondence of signal $x$ with signal $y$ (which are the Phantom and reconstruction in our case) and has been used in other Phantom studies [60]. The correlation of the signals is high when the shape (features) of the signals are similar, and a value of unity represents perfect correlation.

$$r_{xy} = \frac{\sigma_{xy}^2}{\sigma_{xx}^2 \sigma_{yy}^2} \tag{4.3}$$

$$\sigma_{xy}^2 = \frac{\sum (x-\bar{x})(y-\bar{y})}{N} \tag{4.4}$$

The value of $r_{xy}$ is independent of the mean signal amplitude and so tells us nothing of the correctness of the values themselves; therefore, the RMS of the error field (eq. 4.6) is also considered. This error is normalised by the maximum value of the Phantom field $f_p$ less its minimum value (which will usually be zero) to give $E_{RMS}$, and is analogous to a

noise-to-signal ratio. The RMS of the error field has also also been used in other Phantom studies [4, 152, 108].

$$e_{RMS} = \sqrt{\frac{\sum(f - f_p)^2}{N_v}} \tag{4.5}$$

$$E_{RMS} = \frac{e_{RMS}}{\max(f_p) - \min(f_p)} \tag{4.6}$$

Reconstructions of both Phantoms were performed for a range of viewing angles from 2 to 128 with $\Delta_c = 1.0 \times 10^{-3}$. For each reconstruction the viewing angles were spread evenly over 180°. The reconstruction grids were $128 \times 128$ to match those of the Phantoms and each viewing angle consisted of 128 parallel ray projections. Such a number of parallel projections is a conservative estimate of what could be achieved by using CCDs in the CTC sensor. The results for both algorithms can be seen in figs. 4.6 and 4.7 for numbers of angles ranging from 2 to 48. The results for 48 – 128 angles are all above 0.999 for $r_{xy}$ and have less than 1% RMS error and are not shown to improve the clarity at the lower angles.

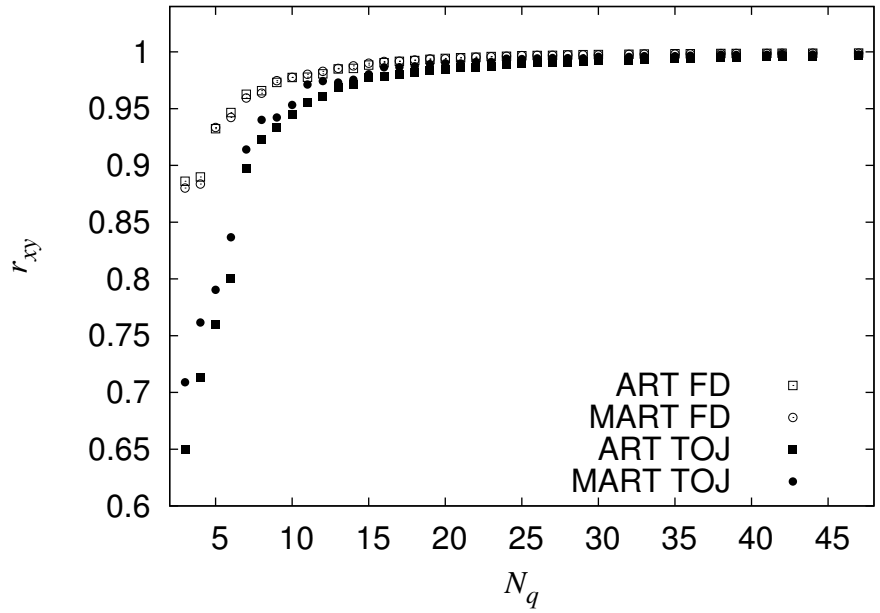



Figure 4.6: The correlation coefficient,$r_{xy}$ of the Phantom and reconstructions for different numbers of views, $N_q$. Results for both Phantoms using each algorithm are shown.

The results show that reconstructions capture the shape and the magnitude of the Phantoms well using either algorithm, and low errors can be obtained even when the number of angles is substantially less than the grid resolution. For these 128 pixel reconstructions an RMS error of just 2% is reported when using approximately 20 views. The same

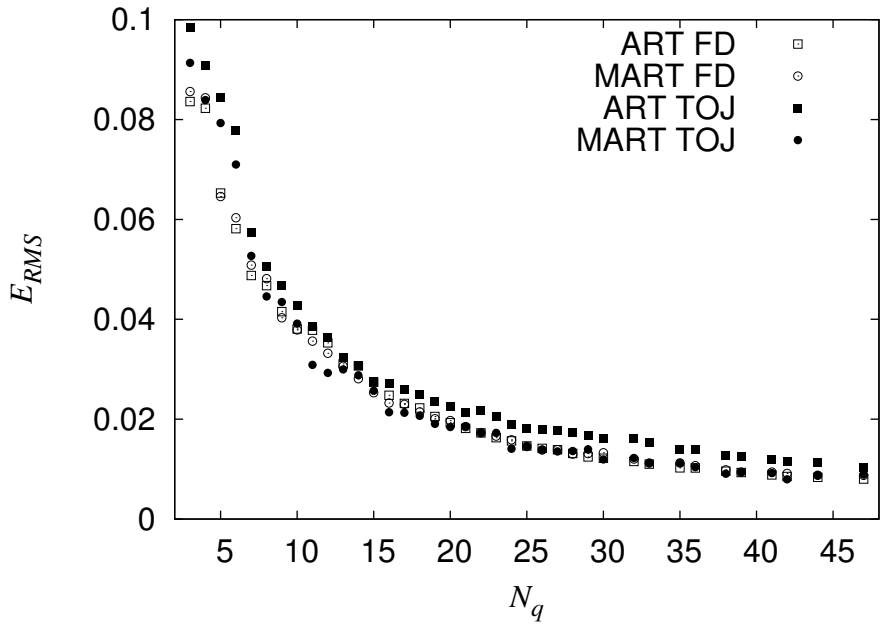
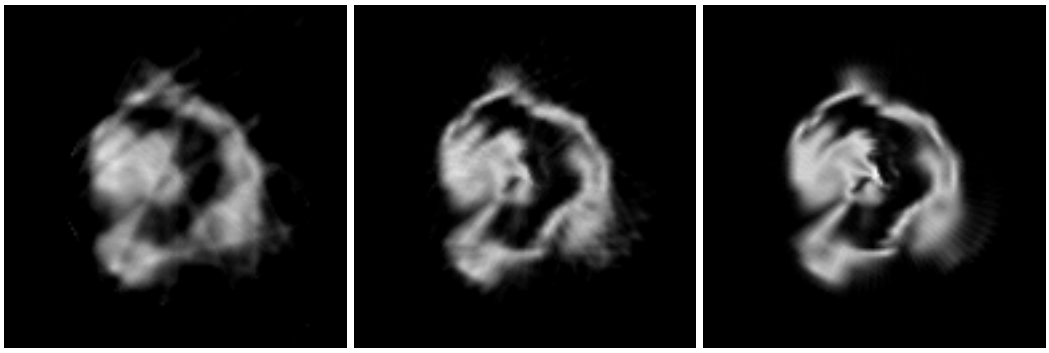


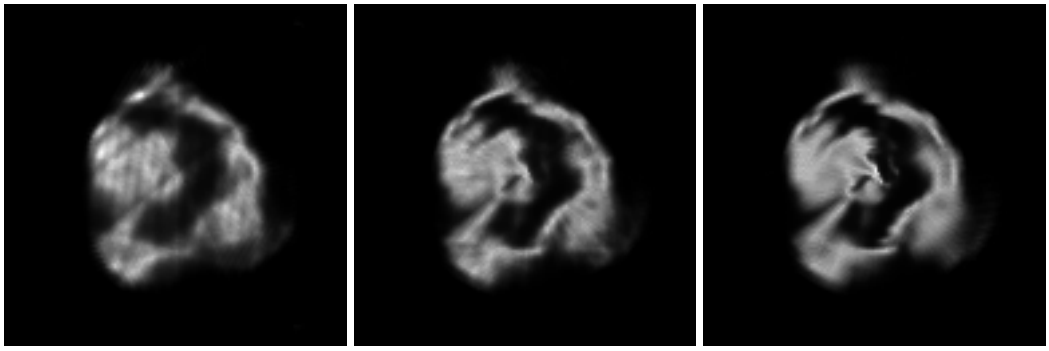

Figure 4.7: The normalised error RMS, $E_{RMS}$, of the reconstruction for different numbers of views, $N_q$. Results for both Phantoms using each algorithm are shown.

number of views also typically achieves correlation coefficients of 0.99 and suggests that good accuracy is achievable with modest numbers of views. These figures also show how the use of more angles is met with diminishing returns beyond approximately 15 – 20 angles, and this is also illustrated in figs. 4.8 and 4.9 which show reconstructions using 7, 20 and 47 views. With 7 views the low frequency features are captured but significant artefacts are observed, when 20 views are used these artefacts are removed and higher frequency features are captured; when 47 views are used however, only a relatively small improvement is seen.

For the TOJ Phantom the MART algorithm has consistently lower error than ART and this is quite significant below approximately 15 angles. However, for Flame D the ART algorithm actually has slightly lower error than the MART algorithm when using less than 10 angles. The small difference in the error of the ART and MART TOJ reconstructions could be a result of the convergence value $\Delta_c$. For the results shown here, both ART and MART required similar number of iterations to satisfy the convergence criteria; however, Verhoeven [152] found that the optimum number of iterations for ART was ten times that of MART. The TOJ Phantom results were repeated using a value of $\Delta_c$ of $1.0 \times 10^{-5}$ to investigate whether ART is converging prematurely in the TOJ case. Figure 4.10 shows that this lower value of $\Delta_c$ has lead to a reduction in the error of the TOJ reconstructions, but for both algorithms. The ART algorithm appears to have the larger reduction in error, and although the MART algorithm still has the lowest error, the margin between the two

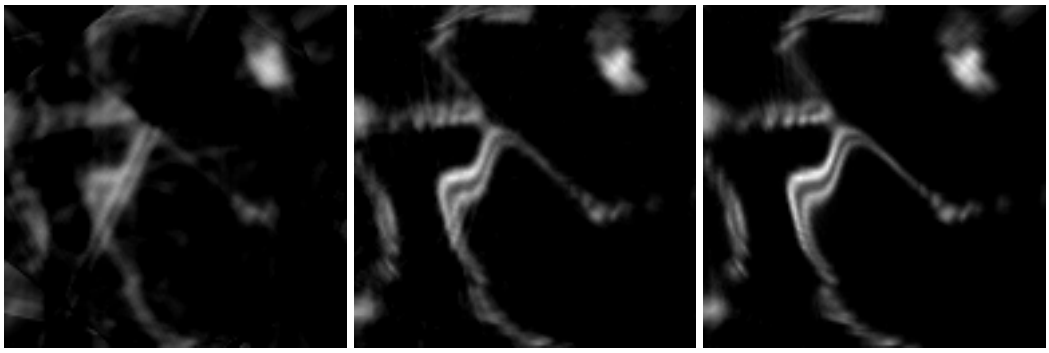


(a) ART 7 angles (b) ART 20 angles (c) ART 47 angles

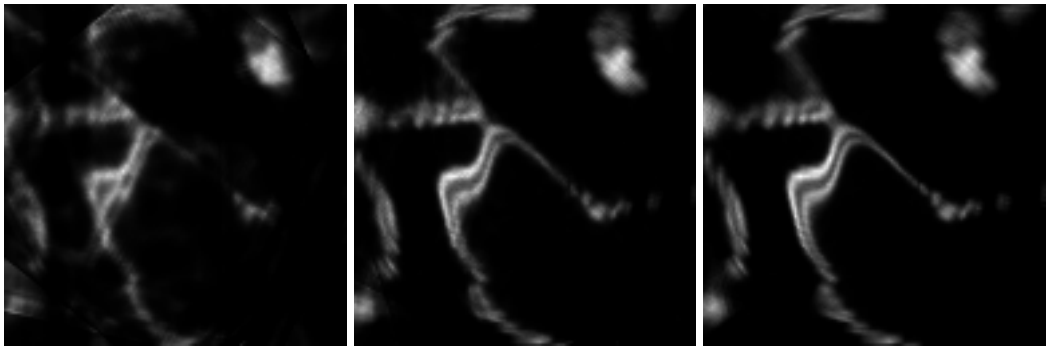


(d) MART 7 angles (e) MART 20 angles (f) MART 47 angles

Figure 4.8: Example ART and MART reconstructions of the Flame D Phantom, which is shown in fig. 4.5.



(a) ART 7 angles (b) ART 20 angles (c) ART 47 angles



(d) MART 7 angles (e) MART 20 angles (f) MART 47 angles

Figure 4.9: Example ART and MART reconstructions of the TOJ Phantom, which is shown in fig. 4.5.

is reduced, and could be expected to become negligible with smaller $\Delta_c$ (as is seen with the Flame D Phantom).

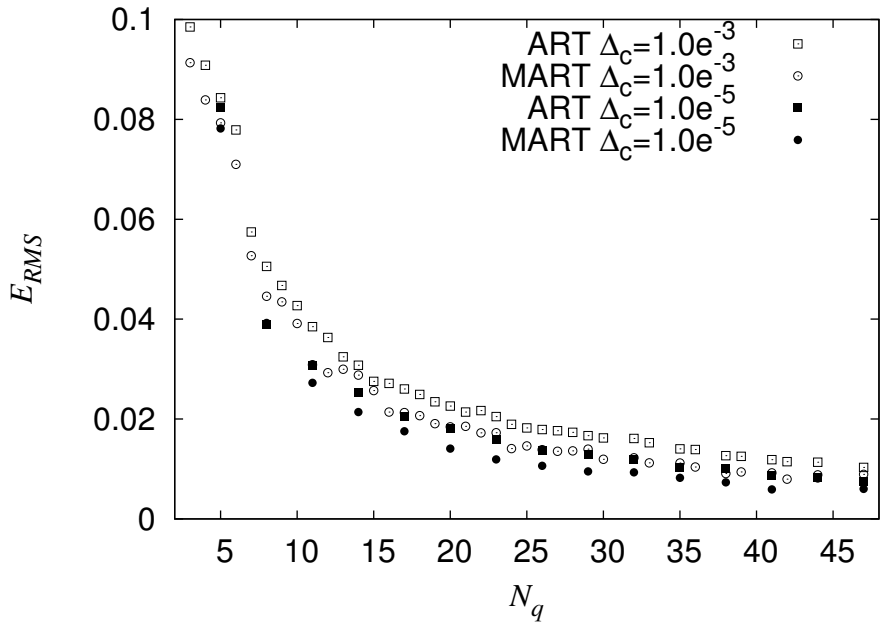



Figure 4.10: The normalised RMS error of reconstructions of the TOJ Phantom for two different values of $\Delta_c$.

In general, more under-determined configurations, i.e. fewer angles, require more iterations for convergence. This could explain the improvement seen in the TOJ reconstructions when using a smaller $\Delta_c$. The slight truncation of the views, which occurs as a result of the Phantom spanning outside the inner circle of the reconstruction in two corners, makes the reconstruction problem more under determined; hence the improvement with more iterations. This hypothesis is consistent with that observed by Mishra et al. [108] who considers severely truncated views and typically requires many more iterations than used in this work.

Overall results of this study agree with that of others [152, 108] and indicate that generally MART achieves the lowest error, except with smoother fields and few viewing angles where ART yields the lowest error. However, where more than approximately 10 angles are used, the difference between the errors is relatively small—even when different convergence criteria are used—and we can conclude that both ART and MART are capable of returning low error reconstructions in this range for both Phantoms. Numbers of viewing angles in the range of 20 – 30 would appear to be optimum for resolutions of $128 \times 128$ pixels as their reconstructions were free of significant artefacts and the addition of further views is met with diminishing returns in error reduction. The magnitude of the errors seen indicate that the number of views may not be the most dominant factor in the

experimental uncertainty, and other sources such as camera registration and calibration error, as well as noise, will be significant.

### 4.2.2 The Effect of Noise

Noise in experimental measurements is ever present and can only be minimised and not removed. The ART algorithms considered in this work are widely reported [56, 88, 108, 152] as performing well with noise in the measured views, hence their popularity. The iterative ART algorithms have been demonstrated in these previous studies to still converge to a solution in the presence of noise. Our main concern now is how noise in the measured views affects the reconstruction accuracy.

Noise in the measured views will affect the reconstruction as follows. Consider the fully-determined discrete reconstruction problem where sufficient angles and projections are taken such that $N_{qp} = N_v$. In such a case with no noise a single correct solution exists. If $N_{qp}$ is then reduced, the problem becomes under-determined and a number of solutions are possible, of which the correct solution is just one. With the addition of noise however, a correct solution may not be found and the iterative ART algorithms will oscillate in a solution space defined by the noisy views (as detailed in Kak and Slaney [77]). When the Signal-to-Noise Ratio (SNR) is high however, then this solution space is still small enough to return answers sufficiently close to the correct solution for most purposes. The effect of noise on convergence of ART type algorithms has been considered by Verhoeven [152] who declares ART to be the most stable in the presence of noise for the 5 view configuration considered. Mishra et al. [108] who consider more (truncated) views also find MART to be more sensitive than ART to noise in the projections. Given the relatively small difference in the RMS error and correlation coefficients seen in our earlier tests, ART in its additive form is preferred for its noise performance and is used in the remainder of this work.

#### 4.2.2.1 On Accuracy

It is desirable to know how the noise, or error, in the measured views translates to error in the reconstruction, as this would effect the time resolution of the sensor because low noise measurements typically require longer integration (exposure) times. Previous studies [152, 108] consider noise in the projections, but as previously mentioned not for the realistic Phantoms, or the likely number of (untruncated) views used in this work. Mishra et al. [108] state that noise is amplified in the reconstructions, but do not estimate by what factor; here we conduct a test to estimate this magnitude for CTC.

Using ART, reconstructions of the Flame D Phantom are performed with increasing

levels of noise added to the projections. The Flame D Phantom has been used because its reconstructions were the least sensitive to $\Delta_c$ as seen earlier; however, as noise can affect convergence as previously mentioned reconstructions were performed using values for $\Delta_c$ of $1.0 \times 10^{-3}$ and $1.0 \times 10^{-5}$. For the test 25 views, of 128 projections were used, and reconstructed at $128^2$ pixels. A noise signal $\varepsilon$ was added to the projection intensities $\mathbf{I}$ and was generated from a random number[2] signal scaled to have the correct variance $\sigma_\varepsilon^2$ for the desired SNR. The SNR is defined in eq. 4.7 where $\overline{I^2}$ is the mean projection intensity power.

$$\text{SNR} = \frac{\overline{I^2}}{\sigma_\varepsilon^2} \tag{4.7}$$

The RMS error ($E_{RMS}$) of the reconstructions is shown in fig. 4.11 plotted against the RMS error of the projections ($\varepsilon_{RMS}$) determined using eq. 4.5 but replacing $f - f_p$ with $\varepsilon$ and normalising by the maximum projection intensity $I_{\max}$. The SNR is also shown for its more intuitive meaning.

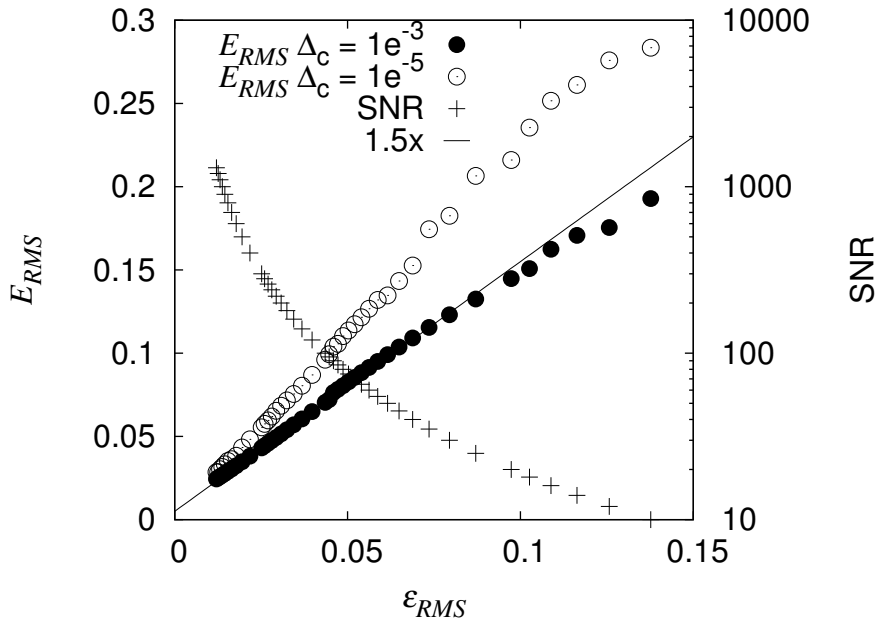



Figure 4.11: Projection noise amplification.

Figure 4.11 shows that projection noise can have a significant effect on the accuracy of the reconstruction of this Phantom. This figure also indicates that SNRs of over 1000 would be required to achieve an RMS error comparable to the 1.5% achieved in the no-noise case. Achieving such SNRs may require long exposure times; however, if we can accept a moderate error of 5% in the reconstruction then more achievable SNRs, below

[2]Using the Fortran95 intrinsic procedure

100, are possible. These results demonstrate an amplification of the error in the projections of approximately 1.5 in the reconstruction, when $\Delta_c = 1.0 \times 10^{-3}$. Comparison of the errors in fig. 4.11 with the earlier no-noise results indicates that the reconstruction accuracy is more sensitive to noise than to the number of angles. The error amplification also significantly increases when a smaller value of $\Delta_c$ and hence more iterations are used. This is the opposite of earlier no-noise results for Flame D which showed no significant difference in the RMS error using the same two values for $\Delta_c$. A similar result was observed by Verhoeven [152] where, beyond a certain number of iterations, ART and MART reconstructions increased in error, and this increase was made more severe by noise—though ART was less affected than MART. Such over-iteration causes significant error below an SNR of approximately 50, and where such SNRs are encountered reconsideration of the convergence criteria must be given.

## 4.3 Resolution Achievable with Limited Angles

In the previous section we considered the effect of limited numbers of angles on the accuracy of a reconstruction and provided estimates of the magnitude of the error. However, a successful reconstruction needs not only low magnitude error but to also faithfully reproduce the features of the object. While we have seen earlier that a high correlation also has a typically low error, this is not necessarily always the case. Consider the following; a reconstruction that matches the shape of the true field perfectly, but whose mean differs from the true field by say 50%, would resolve all the features of the field but have a high error. Conversely, a uniform grey level could have less error but would not resolve anything. In this section we consider the ability of ART to resolve the object *features* and seek to quantify this resolution for the limited angle case.

It is clear that the resolution achieved by any reconstruction is a function of the resolution in the views, and the number of views available. Each view provides resolution in a single direction in the reconstruction that is perpendicular to the view direction, and by having an infinite number of such views spaced equally through 180° resolution is achieved in all directions. However, where the reconstruction is a discrete estimate of the continuous object, as in all CT, only a finite number of discrete views is required.

Should we be free to choose as many equally spaced views as we want, then the reconstruction resolution would be limited only to that of the view imaging system, and the noise in the projections would be significant. However, the use of limited angles means that noise has a lesser effect on the resolution, when compared to the resolution reduction that results from the reduced angular sampling rate, as noted by Bertero and Boccicci [18]. While it is easy to see how the imaging resolution (i.e. number of pixels

in the CCD) determines the upper limit of the reconstruction resolution, it is less clear how the number of angles, or *angular resolution*, affects the reconstruction resolution. In this section we look more closely at the relationship between these two quantities, reconstruction resolution and number of angles, so that better estimates for the resolution achieved by the CTC sensor can be given.

### 4.3.1 Existing Theory

A theoretical relationship relating the number of equally spaced angles to resolution can be seen in eq. 4.8 [43, 112], where $f_s$ is the sampling frequency of the views of diameter $D$.

$$N_q = \pi D f_s = \pi k \tag{4.8}$$

Equation 4.8 returns the number of equally spaced views, $N_q$, required to achieve the same resolution in the reconstruction as in the views, and is equal to is $\pi$ times the view sampling wave-number. For example, if the views are sampled a modest 100 times (pixels) over their width the highest wave-number that can be captured is 50, assuming the Nyquist limit, which would thus require $>$150 angles!

If we assume that it is the number of angles that limits the resolution, as detailed earlier, then eq. 4.8 can be rearranged to give the resolved wavelength $\lambda_{res}$ for a given number of equally spaced angles, as seen in eq. 4.9.

$$\lambda_{res} = \frac{\pi D}{N_q} \tag{4.9}$$

Or in terms of wave-number, in eq. 4.10

$$k_{max} = \frac{N_q}{\pi} \tag{4.10}$$

Equation 4.9 has been plotted for a realistic range of angles in fig. 4.12. It is obvious from this figure that beyond a certain point the addition of further angles will be met with diminishing returns in improved resolution. This agrees with the intuitive assessment that, with greater numbers of angles, the decreasing spacing means, at some point, greater correlation between adjacent angles, and so less and less information is being added with each additional angle.

Equations. 4.8 – 4.10 are derived from the Fourier slice theorem (section 2.2.1.1) that also forms the basis of the back-projection family of algorithms and assumes line integrals. To confirm the applicability of eqs. 4.8 – 4.10 to ART using strip integrals, we can consider the following.

As seen previously in section 4.2, when the angular resolution and the view resolution are equivalent to the reconstruction grid resolution, the ART algorithm returns near perfect

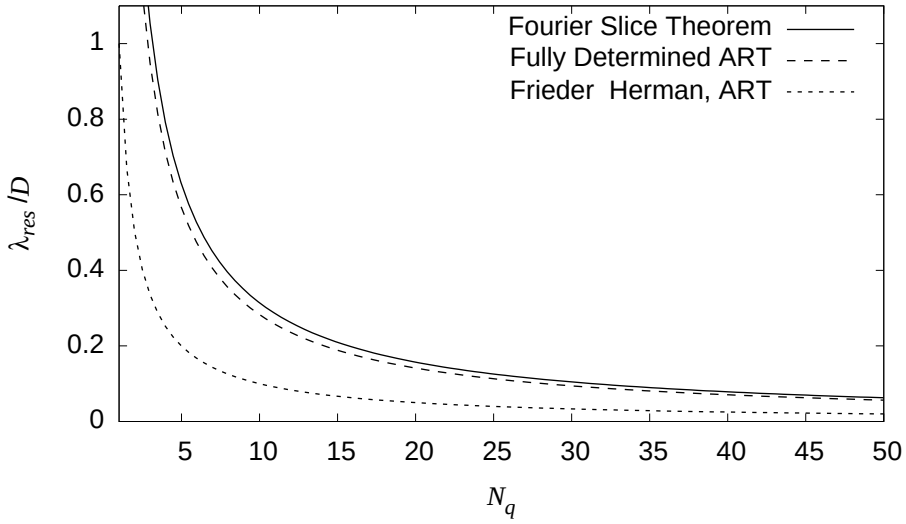



Figure 4.12: The theoretically achievable resolution for a given number of equally spaced angles, $N_q$. The derivation of the Fourier Slice Theorem relationship, $\lambda_{res} = D\pi/N_q$, can be found in references [112, 43]. The resolution, $\lambda_{res}$, has been normalised by the domain diameter/width $D$. Also shown is the resolution achieved by a fully determined ART reconstruction ($2\sqrt{2}/N_q$), and that predicted by Frieder and Herman [48].

reconstructions (in the noise free case). This implies that the ART reconstructions in these fully determined cases are resolved to the the highest frequency their grid can support: $f_{grid}/2$ assuming the well known Nyquist limit. For example, with $N$ views that are each sampled $N$ times and reconstructed on a grid of $N^2$ pixels, wave-numbers up to $N/2\sqrt{2}$ could be resolved (assuming the Nyquist limit along the grid diagonal). This approach shows only a small difference to the wave-number predicted by eq. 4.10 as can be seen in fig. 4.12. On this basis it appears that eqs. 4.8 – 4.10 are applicable to ART. However, in section 4.2 *near perfect* reconstructions are also achieved in significantly under-determined cases (where $N_q < N_{grid}$) implying resolutions much better than those predicted by the Fourier slice theorem.

Such an improvement for ART is supported by Frieder and Herman [48] who demonstrate that for relatively smooth fields, ART achieves the resolution given in eq. 4.11 and also shown in fig. 4.12.

$$\lambda_{res} = \frac{D}{N_q} \tag{4.11}$$

The resolution predicted by eq. 4.11 and eq. 4.9 differ by a factor of $\pi$; however, in their Phantom study Herman and Rowland [69] state that ART requires two to five times fewer angles than Fourier methods. (We note that the underlying assumption of smoothness in eq. 4.11 illustrates how the resolution achieved depends on the field to be reconstructed,

and agrees with that observed in section 4.2). It seems likely then that the resolution achieved by ART will be higher than that predicted by using the Fourier slice theorem approach, but by how much is dependent on the field to be reconstructed. To remove some of this ambiguity, in the next section, a Phantom study is performed using a series of different frequency cosine waves to attempt to better quantify the resolutions that can be achieved using ART, for the smoothly varying fields likely to be encountered by the CTC sensor.

### 4.3.2 Cosine Phantom Tests

The purpose of this test is to assess the resolution performance of ART in under-determined view configurations, (limited views), and so provide better estimates for the resolution achieved in this work. A similar study was performed by Herman and Rowland [69] but using Phantoms, such as bars, points, and checkerboards, that are not representative of the object fields likely to be encountered in this work.

The Phantoms now used are variations on a radially advancing cosine. Cosines were selected so that the Phantoms had known and selectable wavelengths, and they were chosen to be radially advancing so that they did not favour any particular direction. The set of cosine Phantoms used vary by their wavelength, sampling grid size, and the location of their centre. Example Phantoms showing the centre locations used are shown in fig. 4.13 and the parameters of all four variations are given in table 4.2.

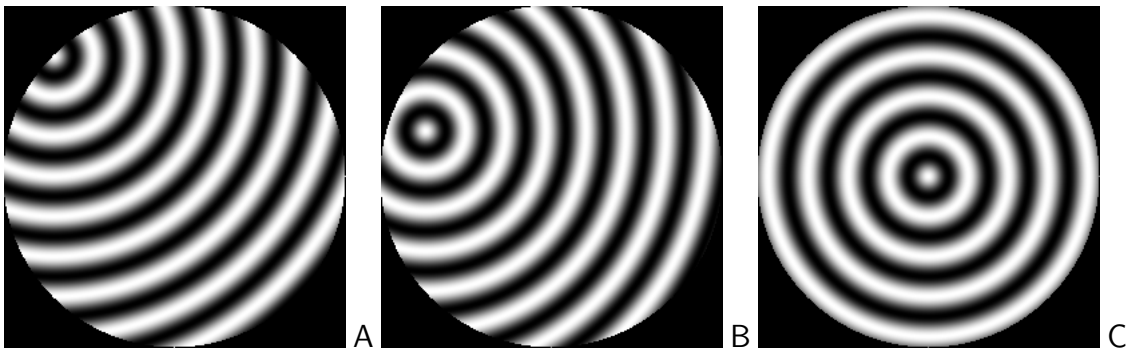


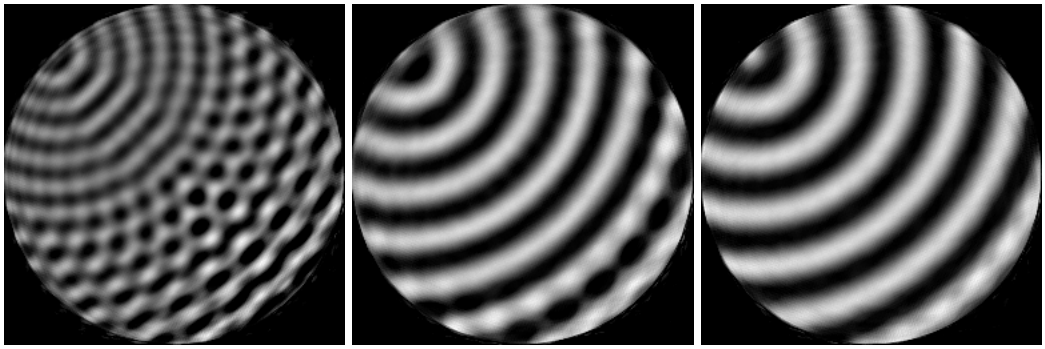

Figure 4.13: Examples Phantoms showing the different locations used for the wave centre. These examples have $256 \times 256$ pixels and wavelengths of 30 pixels.

For each centre location, Phantoms were generated for a range of wavelengths (5 – 80 pixels in steps of 5 pixels). These Phantoms were then reconstructed with a number of views ranging from 5 – 50 in steps of 5. All views were equally spaced through $\pi$ rads. The ART algorithm of eq. 2.24 was used for all reconstructions, and the reconstruction grid was chosen to be the same as the Phantom grid and the number of projections in the views.

| name | pixels | centre location | fig. 4.13 example |
|---|---|---|---|
| 256TL | 256 | 38 38 | A |
| 256OC | 256 | 34 94 | B |
| 256C | 256 | 128 128 | C |
| 114TL | 114 | 19 19 | - |

Table 4.2: Cosine Phantom Parameters.

The resolution achieved using a particular number of views, was determined by visually inspecting the reconstructions of a series of different wavelength Phantoms. The reconstructions were judged to have resolved a Phantom when they were free of strong aliasing and each radial band was visually distinct from its adjacent bands. An example of this process is shown in fig. 4.14.



$\lambda = 20$ ; not resolved $\lambda = 35$ ; not resolved $\lambda = 40$ ; resolved

Figure 4.14: An example of the resolution judging process, for $N_q = 20$ and Phantom family 256TL. At $\lambda = 20$ aliasing artefacts can be clearly observed, and at $\lambda = 35$ small artefacts are still seen between the outer rings. By $\lambda = 40$ the rings are distinct from each other and reconstruction is judged to be resolved at that wavelength.

To confirm the validity of the visual inspection process, the correlation coefficient of the Phantoms and their respective reconstructions was calculated for all values of $N_q$. As previously discussed when the correlation coefficient ($r_{xy}$ in eq. 4.3) approaches unity, the Phantom and its reconstruction are identical, and we can safely consider the reconstruction to have achieved the resolution in the Phantom. It is likely that the Phantom features will be resolved before $r_{xy} = 1$, as it isn't necessary for there to be perfect correspondence for the waves to be visually distinct. For any value of $N_q$, the correlation approaches 1 as $\lambda$ increases, levelling off at a certain point as can be seen in fig. 4.15 for $N_q = 20$. In all cases the levelling off point was consistent with the visually determined point in a range of approximately $\pm 5$ pixels of wavelength.

The resolved wavelength $\lambda_{res}$ is shown for each $N_q$ in fig. 4.16, for the Phantom types

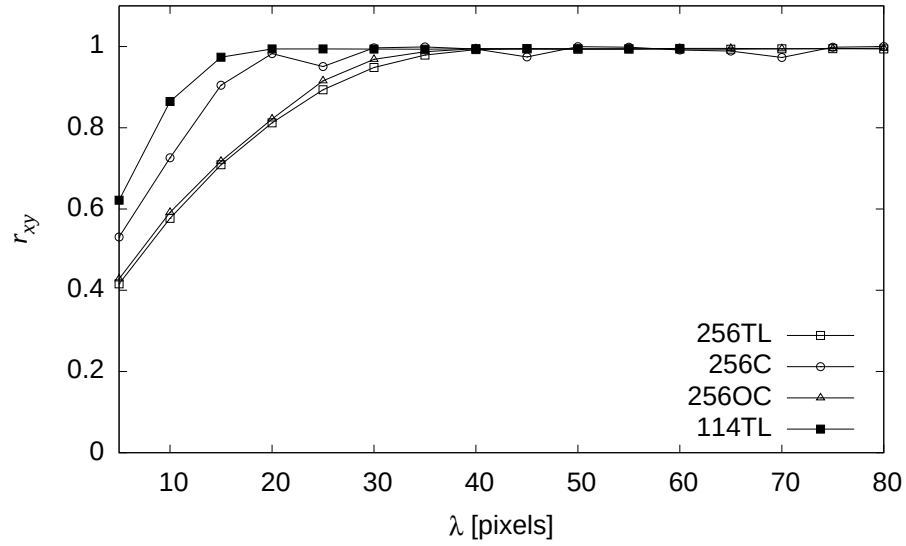



Figure 4.15: An example of the *levelling off* in the cross-correlation, $r_{xy}$, observed as $\lambda$ increases. The $N_q = 20$ case is shown for the Phantoms described in table 4.2.

listed in table 4.2. The resolved wavelengths were taken as the average of the upper and lower boundaries of the interval determined by the visual inspection process and normalised by the domain width $D$ (246 or 114 pixels). The interval boundaries are represented by the error bars.

The visually determined wavelengths show good agreement with those predicted by eq. 4.9 in all cases except Phantom type 256C, which has a wavelength roughly a third of that of the other types and shows close agreement with the prediction of Frieder and Herman [69]. At lower numbers of angles ($< 10$) the cosine results do deviate from those predicted using the Fourier slice theorem. However, this is consistent with the theory which only strictly holds for large $N_q$ [112]. The Phantom resolution, and hence the view sampling rate, doesn't appear to make any difference—as is expected for the case here where $N_q < k_{view}$.

The agreement with eq. 4.9 indicates that smoothness and non-negativity in the object has a limited effect on the resolution achievable with *standard* ART (i.e. without any modifications to exploit such *a priori* information). However, the difference observed between the 256C Phantom results and the theoretical relationship shows that the features of the field to be reconstructed can still affect the resolution that is achievable. One feature that distinguishes the 256C type Phantoms from the others is that they are axisymmetric; however, the ART algorithm used does not deliberately exploit any symmetry, and so any advantage is implicit. The 256C type Phantoms also differ from the 256TL and 256OC Phantoms in that the radial wave extends out to a smaller radius.

The reconstructions of the 256TL and 256OC Phantoms exhibit strong aliasing of the

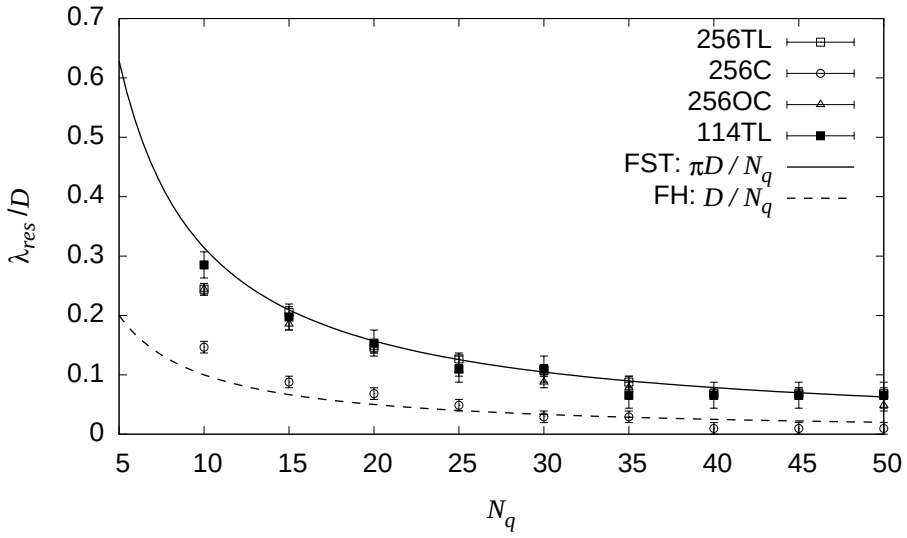



Figure 4.16: The normalised resolved wavelength achieved, $\lambda_{res}/D$, for different numbers of equally spaced views, $N_q$. The resolved wavelength is determined from ART reconstructions of radial cosine Phantoms. Shown are results for the Phantoms detailed in table 4.2. The continuous lines depict the theoretical result based on the Fourier Slice Theorem (FST) and the relation given by Frieder and Herman [69].

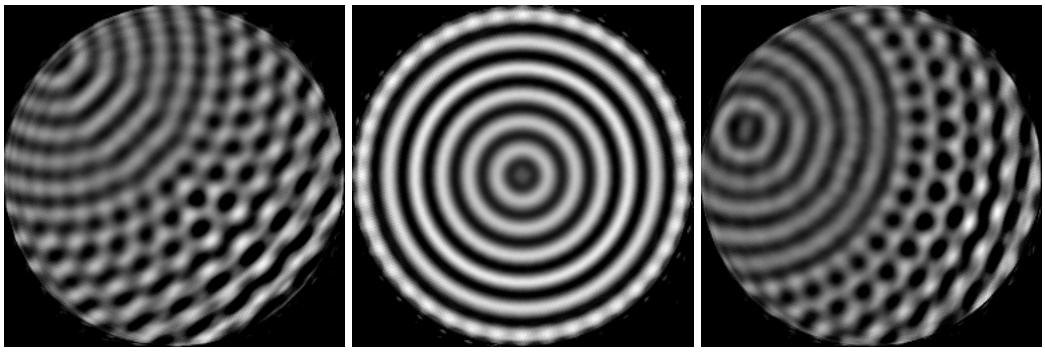
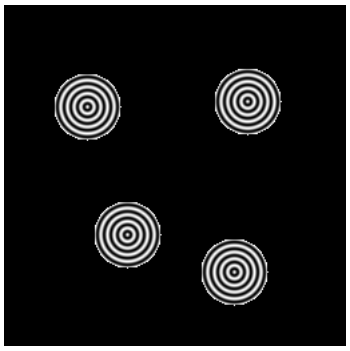



Figure 4.17: Example reconstructions of the three 256 pixel Phantom types, at $\lambda = 20$ and $N_q = 20$.

outer rings when the reconstructions of 256C show none—hence the better resolution. However, if we consider the examples seen in fig. 4.17; for the 256TL and 256OC cases the reconstructed wave is reasonably well resolved up to the first 6 peaks. This is the same number of peaks resolved in the 256C reconstruction, which has no more peaks to resolve. From this we can infer that it is not the symmetry of the Phantom that leads to better resolutions for the 256C case, but rather the more limited radius of the wave.

The resolution of the small radius peaks first, implies that where the high frequency waves are locally limited in their extent then better resolutions are be possible. To test this a new Phantom type has been generated that contains four radial cosine waves of limited radius, deliberately arranged non-symmetrically. An example of this Phantom can be seen in fig. 4.18.



256SU

Figure 4.18: Example of the 256SU Phantom. $\lambda = 5$.

Using the same method as described previously, the resolved wavelengths were determined for the same range of $N_q$. The results for the 256SU Phantom are shown alongside the earlier results in fig. 4.19. The 256SU Phantom results match well with that predicted by Frieder and Herman [69], and also appear to return a consistently lower $\lambda_{res}$ than the 256C Phantom. This supports the earlier point that it was the reduced radial extent and not the global symmetry that meant the 256C was resolved to smaller wavelengths than the other Phantom types.

This study has assumed equally spaced angles throughout; though, in some situations this may not be possible due to limited optical access. However, it is not possible to know the level of access in advance and so the effect of unequal spacing has not been considered here, but is considered in later work.

The selection of a radial cosine wave appears a comparatively stern test case for the ART algorithm, as the results here typically required more angles for $r_{xy} \rightarrow 1$ than was necessary for the Phantoms used in section 4.2. The agreement of the 256TL, 256OC, and 114TL Phantom results with eq. 4.9 implies that this equation is applicable to ART

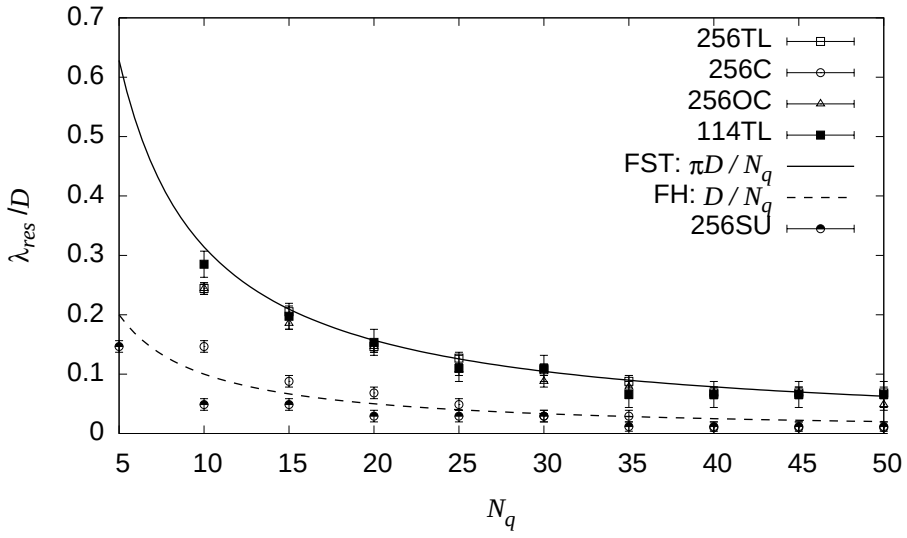



Figure 4.19: The normalised resolved wavelength, $\lambda_{res}/D$, for different numbers of equally spaced views, $N_q$. The resolved wavelength is determined from ART reconstructions of radial cosine Phantoms. Shown are results for the Phantoms detailed in table 4.2 and also for Phantom type 256SU. The continuous lines depict the theoretical result based on the Fourier Slice Theorem (FST) and the relation given by Frieder and Herman [69]

.

reconstructions. However, given the comparative difficulty of these cosine Phantoms, that others report much lower resolutions, and that the theory makes no assumptions regarding the nature of the field, eq. 4.9 can be considered applicable only as a limiting case for ART. Where prior knowledge of the field allows the radial extent of any waves to be considered limited, higher resolutions are achievable, as has been seen.

This is not an exhaustive study, but does strongly illustrate that the resolution achievable with a certain $N_q$ is dependent on the object to be reconstructed, although the results imply that smoothness alone does not appear to improve resolution. The selection of a different basis function for ART, such as blobs or cubic splines [152, 104, 105], could however, make better use of *a priori* smoothness information to improve resolution.

From the LES Phantoms seen so far, it appears that the radial/spatial extent of any waves is limited which helps to explain why their reconstructions in section 4.2 are so good. The chemiluminescent fields encountered in this work are likely to be similarly sparsely populated and waves will be more spatially limited than in most of the cosine Phantoms used here. Based on this a value for $\lambda_{res}$ of half that predicted by eq. 4.9 appears conservative and is consistent with Herman and Rowland [69]. This approach for determining resolution will be used as a working estimate, but will be reassessed on a case by case basis.

The use of *a priori* information to judge the spatial extent of any waves, and its the effect on resolution, is slightly subjective and should be used with caution. However, such an estimate for the reconstruction resolution is based on evidence and is reasonable given the observed difficulty in determining a representative resolution—as seen in this work, and by the lack of quoted figures in other work.

## 4.4 Proof of Concept Experiment

In this section the CTC concept is experimentally tested using $CH^*$ chemiluminescence intensity measurements of a Matrix Burner flame performed using a single consumer digital camera. The Matrix Burner produces a steady, but non-axisymmetric flame, that permits multiple views to be obtained from a single repositioned camera. The purpose of this test is primarily to provide a first test of the CTC concept, and expose the ART algorithm to actual experimental data. However, the success of a CTC sensor depends on being able to deploy sufficient cameras and so their cost is an important factor. Here we have deliberately selected a consumer digital camera, to assess what can be achieved with such low cost equipment and so provide a basis for the equipment selection of future multi-detector set-ups.

### 4.4.1 Experimental Set-Up

The study of temporally resolved, arbitrary turbulent flames is possible using the general set-up shown previously in fig. 3.1. In this section to test the CTC concept before expanding into the multi-detector set-up, a steady non-axisymmetric laminar burner has been selected that allows the multiple detectors of the general set-up to be replaced with a single repositioned camera as seen in fig. 4.20. The camera was mounted to a fixed height tripod so that the optical centre of the lens was at the same height as the burner. The camera is then rotated around the burner using an aluminium arm to maintain a radial distance of 300 mm and the alignment of the camera. The camera used was a Minolta Dimage A1 dating from 2004, which has a 2568 by 1928 pixel CCD, a 14 bit A/D converter and a 200 mm macro lens. This camera was chosen over other digital cameras because it has full manual settings and the direct CCD sensor readings are accessible using the *raw* image format, i.e. uncompressed and unaugmented.

The Matrix Burner consists of rows of two small concentric nozzles, of inner diameters 1 mm and 3 mm, organised in a cross pattern as depicted in fig. 4.21. The inner nozzle supplies methane and the outer oxygen, and the resulting flame actually consists of 21 small non-premixed laminar flames. The methane was supplied at a total (for all

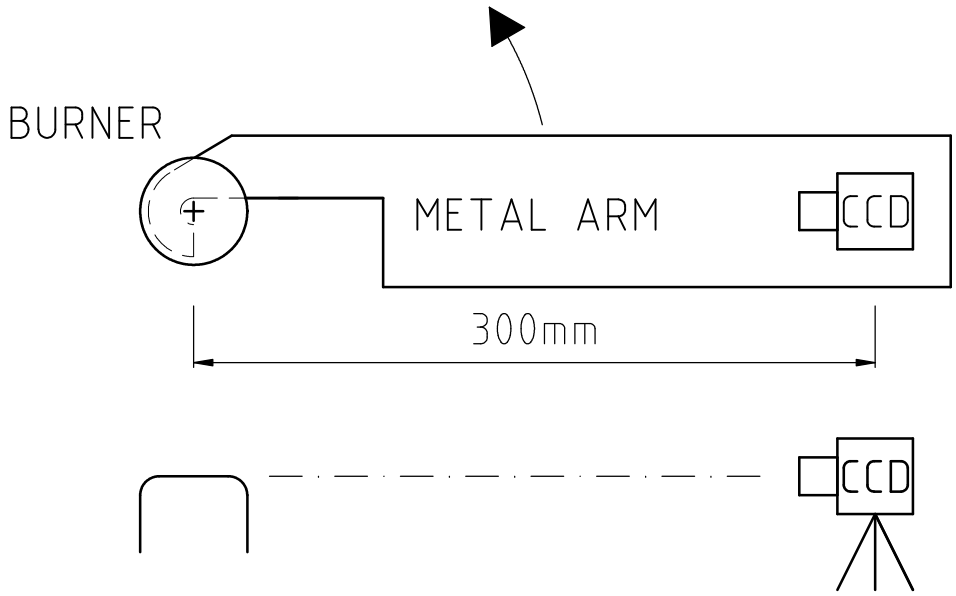



Figure 4.20: The specific CTC experimental set-up used in this study. The camera is repositioned to different angles to provide multiple views.

nozzles) of 2 l/min and oxygen at a total of 35 l/min. The flow rates were selected to provide a stable flame with low soot. The resulting flame is non-axisymmetric and provides up to ten flame fronts in one direction, similarly to the wrinkled flame fronts of turbulent flames, and so provides a good test of the CTC technique. Furthermore, this burner is good for validation purposes because the arrangement and shape of the flames is intuitively known beforehand; for example, we can expect the constituent flames to have a roughly circular section and any deviation from this constitutes an error.

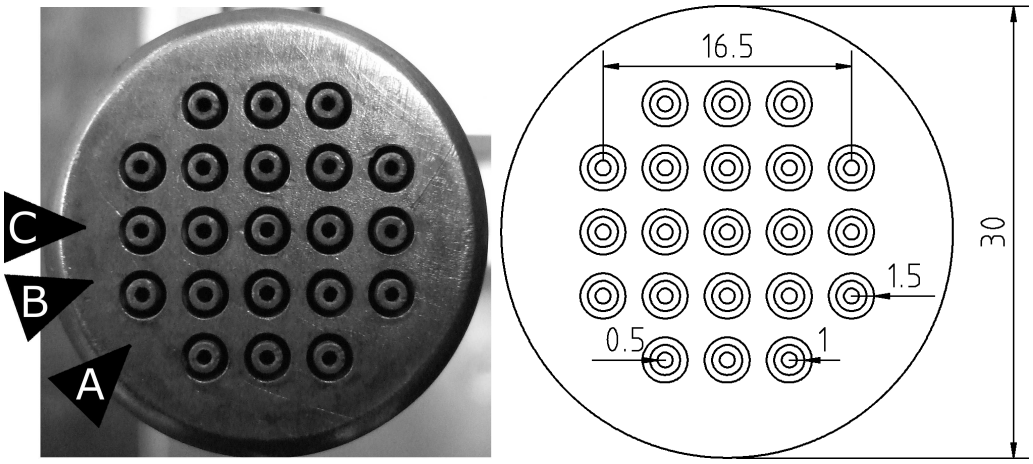



Figure 4.21: *Left:* Photograph of the Matrix Burner showing the angles of views 12 (A), 19 (B) and 24 (C) seen in fig. 4.22. *Right:* Drawing of the Matrix burner.

The radius at which the camera was located was limited to 300 mm by the size of the

available table and so an aperture value of $f/11$ was used to give a wider depth-of-field and an exposure time of 200 ms was in turn used to provide a good signal-to-noise ratio. To isolate the $CH^*$ emissions resulting from the $A^2\Delta - X^2\Pi$ transition, an interference band-pass filter of $430 \pm 10$ nm with 45% transmission at peak frequency was added in front of the lens. This transition is commonly used in chemiluminescence studies, for example [157, 70, 35], and is used here because its visible wavelength is compatible with the camera optics[3]. The camera, as with almost all color digital cameras, uses a Bayer filter and, to ensure that no interpolation of the signal occurs, the raw images from the camera are converted to pgm format using an open source program *dcraw* [32]. From this pgm file only the blue pixels have been used for the views and this results in a view projection resolution of 48 $\mu$m. To remove the unnecessary blank space all camera images were cropped to $450 \times 70$ pixels to just enclose the flame.

#### 4.4.1.1 View Registration

For reconstruction it was necessary for the centre of the cropped views to be co-located with the burner centre. However, a lack of stiffness in the camera tripod junction coupled with imperfections in the table surface meant that the burner centre did not always coincide with the centre of the camera image. (This offset was small enough not to alter the view angle significantly.) It was not possible to use a calibration object, to locate the centre of the burner in the views, because of the repositioning, and so images of the flame itself were used. The left and right flame boundaries were determined using *edge detection* and their mid-point was taken to be the burner centre. (As the flame centre is also the burner centre in this case)

Some rotation of the image about the camera's optical axis also occurred and was removed by an appropriate counter rotation of the image prior to cropping. The necessary rotation was determined from the angle between the image horizontal and the bottom of the flame (located using edge detection). While the approach used here proved sufficient for this proof of concept test, with hindsight improved view registration could have been achieved by illuminating a calibration object elsewhere in the flame image, perhaps the cylindrical burner casing itself.

### 4.4.2 Matrix Burner Results

Measurements of $CH^*$ chemiluminescent intensity were taken from 48 camera positions spaced equally over a total angle of 172° with an estimated positioning uncertainty of $\pm 0.4°$. A thick black cloth was placed behind the burner to minimise the background sig-

[3] Which filter out UV light.

nal. Example view measurements taken at three different angles can be seen in fig. 4.22. These views illustrate the information that is contained within the projections (and exploited by CT) as the effect of the obscured rows can be clearly identified.

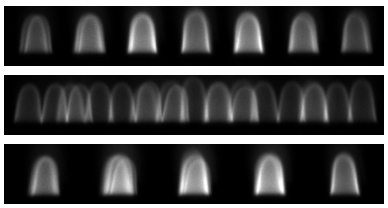
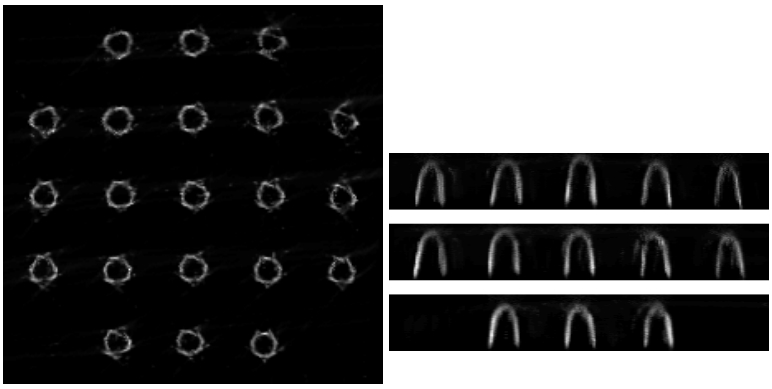


Figure 4.22: Example measured views. The angles of view 12 (top), 19 (middle), and 24 (bottom) can be seen in fig. 4.21.

Once aligned, co-planar rows were extracted from each 2D view to provide sets of the 1D *row* projections. These row projections are then used to reconstruct one section of the flame, and with the stacking of successive sections, the 3D flame chemiluminescence structure is built up. To reduce the computational time the row projections were filtered down from 450 to 225 values and a reconstruction grid of 230 by 230 pixels was used. Once stacked, this results in a 3D reconstruction of $230 \times 230 \times 60$ voxels (volume elements) each with a size of 96 $\mu$m in the camera plane and 48 $\mu$m normal to the plane, i.e. the stacking direction. The reconstructions were performed using an $\Delta_c$ value of $1.0 \times 10^{-4}$ and an example 2D reconstruction section is shown in fig. 4.23 alongside longitudinal sections of the 3D flame.



(a) Reconstruction/section at z=20 (b) Sections at y=115 (top), y=70 (middle), and y=22(bottom)

Figure 4.23: Sections from the reconstruction of the Matrix burner using 48 angles.

These sections show that all the constituent flames have been resolved and are located as would be expected. The flame fronts show a correct shape in the longitudinal *yz* plane and are generally circular in the *xy* measurement plane, although some deformation is seen in the outer flames particularly in the top right of fig. 4.23. This is thought to be due

to errors in alignment resulting from the camera repositioning, though it could also be a result of optical effects such as perspective distortion or limited depth-of-field. A volume rendering of the 3D reconstruction can be seen in fig. 4.24, where again good qualitative agreement is shown between the 3D reconstruction and a photograph taken from a similar angle. For example, in the photograph the central flame is noticeably higher and this is also captured in the 3D reconstruction.

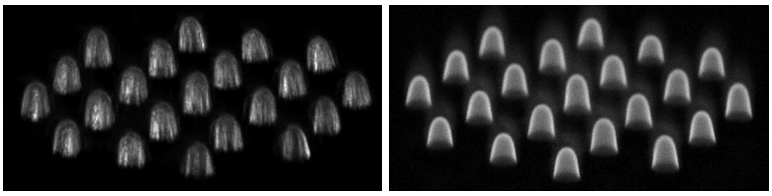


Figure 4.24: A volume rendering of the 3D reconstruction (left) of the $CH^*$ chemiluminescent light intensity alongside a picture (right) of the flame taken at a similar angle. The reconstruction used 48 angles with 225 projections per angle.

Based on our earlier analysis in section 4.3.2 we can estimate the resolution achieved in the reconstruction using 48 angles. Given that the spatial extent of the flames is similarly limited as the 256SU Phantom a value for $\lambda_{res}/D$ of approximately 0.01 is appropriate. For the domain width of 22 mm this gives an approximate resolution for the reconstructions of 220 $\mu$m. This seems a reasonable estimate given that the flame fronts have been resolved well and are just a few pixels wide. To consider the effect of the filtering on the resolution a 2D $460^2$ pixel reconstruction of a single section was performed for which the row data was not filtered. This 2D reconstruction is shown, alongside the equivalent section from the filtered reconstruction, in fig. 4.25 and we can see that there is little difference in the flames' thicknesses, indicating that they have indeed been well resolved on the coarser reconstruction grid. A small level of noise is evident in the higher resolution reconstruction, though it has not affected the resolution of the flames and has clearly been filtered out of the coarse grid reconstruction.

### 4.4.3 Conclusion

The CTC sensor has been implemented using a single repositioned consumer digital camera and used to measure the flame from a Matrix Burner. This simple implementation has demonstrated the ability of the CTC sensor to recover the 3D $CH^*$ chemiluminescence field of a flame from experimentally measured views. The reconstructions show good qualitative agreement with the observed shape of the flames and demonstrate the feasibility of CTC. The multiple flames fronts of the non-axisymmetric burner have been resolved indicating the suitability of CTC for use with turbulent flames.

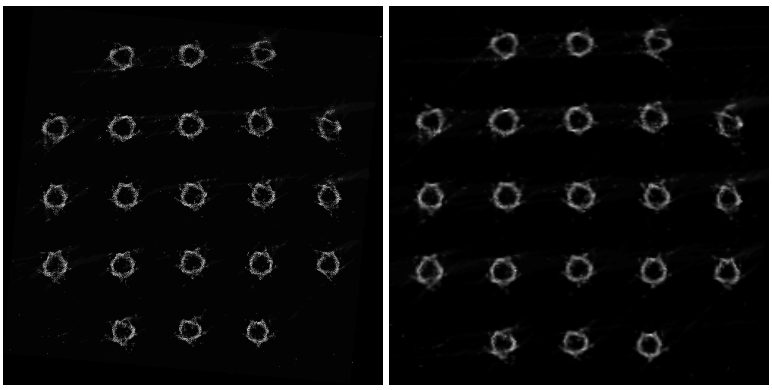


Figure 4.25: (Left) A *raw* 460 by 460 pixel section showing the $CH^*$ chemiluminescent light intensity. This section was reconstructed from 48 views each 450 projections wide. (Right) The same section reconstructed from views filtered to 225 projections wide.

The structures of the constituent flames have been more accurately captured in the centre of the image, with greater asymmetry observed away from the centre. This asymmetry indicates that either alignment errors or some misrepresentation of the projections has occurred; for example, due to perspective distortion caused by the relative proximity of the camera to the object flame. However, the results also demonstrate the robustness of the sensor as these good results have occurred despite the imprecision in the view angles and the camera location; thereby indicating a degree of insensitivity to such uncertainties.

The consumer camera employed was sufficient for the purposes of this test, however its use has highlighted some important issues regarding the use of off-the-shelf consumer cameras. Firstly, there is little potential for reliable synchronisation of such cameras; secondly, these cameras are sold as *black boxes* and it is very difficult to know the true parameters of the cameras, in particular the lenses, and what functions are automatically applied to the images. We can conclude that smaller simpler and better understood cameras would be desirable for a multi camera system. With the specific camera used in this work an additional problem is the size, which would make multiple deployment physically difficult. Significantly, this test has demonstrated that the CTC Sensor concept can work with less expensive, lower grade cameras.

# Chapter 5

# Improving ART for CTC

This chapter details several studies that aim to improve on the generic ART algorithm by tailoring it more specifically to the CTC sensor. Firstly the issue of self-absorption of the emitted light by the flame itself is addressed, before an improvement to the forward projection process for smooth fields is tested. Following this ART is extended into 3D to allow arbitrary camera locations and out-of-plane optical effects to be considered. Finally a new projection representation is presented that takes account of the optics of the CTC set-up.

## 5.1 Self absorption Study

When using measurements of emitted light, the optical thickness of the object flame needs to be considered. This is because attenuation within the object can produce significant reductions in the measured light intensities; in the worst case giving a zero signal. This *self-absorption* represents a potential source of error for CTC, which is of particular interest when quantitative species measurements are sought.

Self-absorption is usually significant for IR based temperature sensors and is accounted for in several studies [61, 119, 34, 94, 74]. Hall and Bonczyk [61] and Ito et al. [74] use complimentary absorption measurements to determine the appropriate correction but require additional sources to do so. Lim et al. [94] suggest a correction approach that does not require absorption measurements, but is applicable only to axisymmetric flames. Correia et al. [34] present a correction to the SART algorithm that is applicable to non-axisymmetric flames without additional absorption measurements. However, this uses a model for the absorption coefficient distribution applicable to soot only. Budinger and Gullberg [28] approach the problem of attenuation from a nuclear medicine stand point, and highlight it as the main difference between emission and transmission tomography. They provide a number of methods to address self-absorption in the nuclear medicine

case.

In previous work using emission CT for visible or UV wavelengths [40, 70, 157, 73], optically thin flames have been studied and self-absorption neglected. Hertz and Faris [70] deduce optical thinness for $CH^*$ emissions from absorption measurements, but do note that this is not the case for $OH^*$ and that each situation should be examined individually. Ishino and Ohiwa [73] also assume that absorption is weak in CH (as well as C2 and CHO). These studies have generally considered flames that are spatially limited in the measurement plane, making them optically thin [49]. Gaydon [49] notes however, that $OH^*$ emissions can suffer from self absorption even for small flames. In this section we consider how self-absorption could affect the CTC sensor, and how the assumption of optical thinness is affected by the flame configuration. A modification to the ART algorithm is then tested that corrects for self absorption of chemiluminescence, similarly to the corrections used in soot thermometry.

### 5.1.1 Investigating Error from Self absorption

In this study we aim to investigate self-absorption as a potential source of error for CTC by using realistic LES Phantoms. To better understand how the flame geometry affects self-absorption, we consider two flame configurations for their contrasting optical properties: a jet flame where the species are distributed longitudinally through the domain; and an opposed jet where the species are concentrated in the plane of measurement.

#### 5.1.1.1 Simulating Self-Absorption in the Measurements

The LES Phantom data was taken from the simulations of Flame D and the TOJ described previously in section 4.1. For Flame D sections were extracted from the axial locations shown in fig. 5.1 and for the TOJ the Phantom was taken from the main reaction zone at the mid-plane. The instantaneous sectional data from the cylindrical LES grid was linearly interpolated onto a $151^2$ pixel Cartesian grid, of 75 mm by 75 mm for Flame D and 34 mm by 34 mm for the TOJ.

In our previous LES Phantom studies the measurements have been simulated by a discrete integration of the LES data as given by eq. 4.2. To account for self-absorption, equation eq. 4.2 needs to be extended to include the attenuation of the emitted light as it travels to the sensor. The attenuation of monochromatic light passing through an absorbing medium can be modelled using the Beer-Lambert law shown in eq. 5.1. In this equation the attenuation of the light intensity $I_0$ is determined from the concentration of the absorbing species $c$, the distance $l_A$ travelled through the absorbing species and an

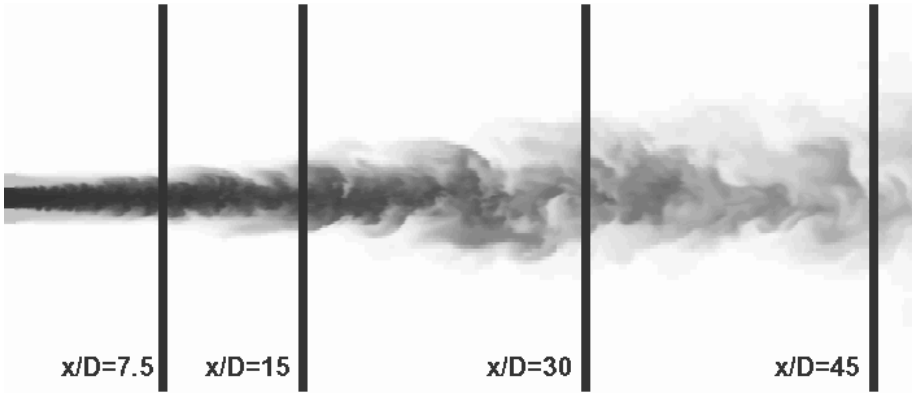



Figure 5.1: The locations of sections used as Phantom data for Flame D.

absorption coefficient $\kappa$.

$$\frac{I_0}{I_1} = e^{-\kappa l_A c} \tag{5.1}$$

For our discrete case, where $c$ varies throughout the object domain, we consider the properties within a pixel to be constant and so the attenuation of light $I_{ev}$ emitted from a pixel $v$ in the object domain is given by eq. 5.2.

$$\frac{I_{pv}}{I_{ev}} = e^{-\kappa \sum_{\upsilon=v}^{\text{edge}} l_{A\upsilon} c_\upsilon} \tag{5.2}$$

In eq. 5.2 the sum includes all the pixels that lie on the optical path between the considered pixel $v$ and the edge of the object domain in the direction of the sensor. Here we make the simplification that $\kappa$ is constant for a species, though it is temperature dependent, and consider only one absorbing species. The value of $l_{A\upsilon}$ is determined by calculating the intersection of a straight line, passing from the pixel of interest $v$ to the sensor, with the pixel $\upsilon$. This has the added convenience that we are also free to use such a path length as a value for $w_{qpv}$.

By using $I_{pv}/I_{ev}$ to additionally weight $w_{qpv}$ in eq. 4.2 and selecting suitable values for $c(i,j)$, we can simulate self-absorption in the measurements by using eq. 5.3. A similar formulation is seen in [28] and other studies, such as that by Correia et al. [34].

$$I_{qp} \approx \sum_{v=1}^{N_v} f_v w_{qpv} \frac{I_{pv}}{I_{ev}} \tag{5.3}$$

#### 5.1.1.2 LES Phantom Study

A Phantom study is performed using realistic $CH^*$ and $OH^*$ fields generated from LES. These species are considered the most likely to be used in a CTC sensor and we investigate self-absorption of their emissions as a potential source of error. The $CH^*$ emissions are

not usually found to suffer from self absorption, at least for small flames, but provides a base for comparison with $OH^*$. The error in the reconstruction is assessed primarily on the difference in its global standard deviation $E_\sigma$ when compared to the Phantom value as defined in eq. 5.4. Here $\sigma_p$ is the global standard deviation of the Phantom, and $\sigma_r$ of the reconstruction.

$$E_\sigma = \frac{|\sigma_P - \sigma_r|}{\sigma_P} \tag{5.4}$$

The ART algorithm is capable of returning reconstructions of very low error where a sufficient number of angles are used. In this section we concentrate on self-absorption and so consider only the case fully determined case. This means for the $64^2$ pixel reconstruction grids used, views were taken at 64 angles with 64 projections per view. Except where stated otherwise all reconstructions of Flame D are performed at axial location $x/D = 30$.

To simplify the modelling of self-absorption it is assumed that the absorbing field has the same geometry as the emitting field to be reconstructed. This simplification is reasonable given the similar spatial extent of the excited and ground state species in the flame, and that the ground state is the principal absorber of emissions from the excited state [116]. It is also assumed that the concentrations of the scalar absorbing field $c(i,j)$ are proportional to the scalar emitting field $f(i,j)$, which allows a nominal absorption coefficient $\alpha$ (eq. 5.5) to be defined that incorporates these concentration differences. By varying $\alpha$ over a range of values we can obtain a better understanding of how self-absorption effects reconstruction accuracy for each flame configuration, without the added difficulty of determining $\kappa$.

$$\alpha = \kappa \frac{c_v}{f_v} \tag{5.5}$$

Reconstructions were performed using ART with the projections subjected to increasing amounts of self-absorption (increasing $\alpha$). The error in the reconstructions of both Phantoms and species can be seen in fig. 5.2. This figure shows that, for both species, self-absorption increases the error in the TOJ reconstructions, at a lower value of $\alpha$ than in Flame D. For $OH^*$ the difference is relatively small, with both configurations reaching an arbitrary 5% error at $\alpha \approx 10^5$. However for $CH^*$ the difference is larger with 5% error occurring at $\alpha \approx 10^6$ for the TOJ and $\alpha \approx 10^8$ for Flame D. Interestingly, this difference is more than could be expected due to concentration differences alone. For both species the difference between the mean concentrations from each configuration is of the order of $10^1$. For example the TOJ mean $CH^*$ mass fraction is $0.25 \times 10^{-9}$ compared to $0.27 \times 10^{-10}$ for Flame D. This indicates that variation in spatial distribution must also have an effect on self-absorption to explain the difference in $\alpha$.

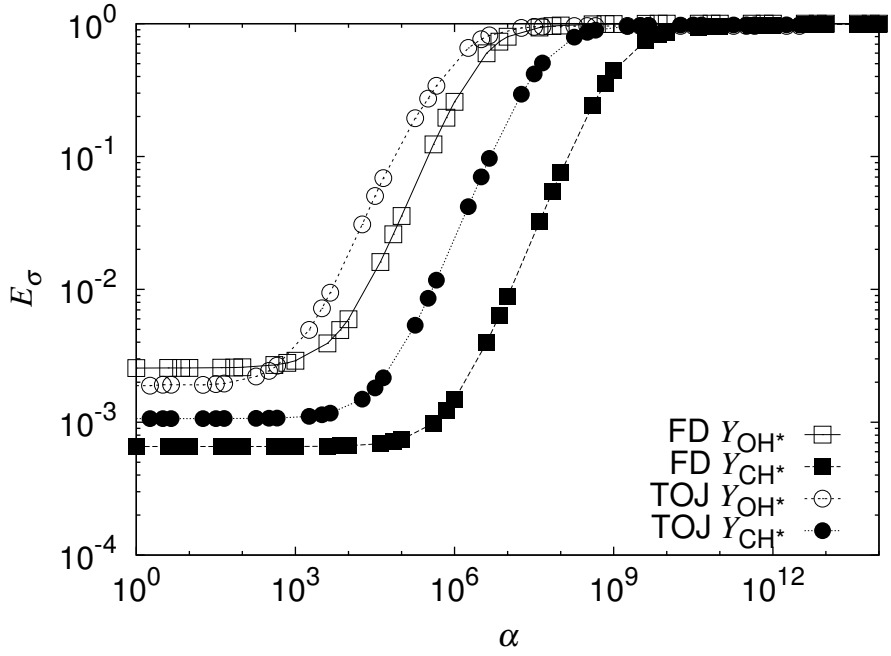



Figure 5.2: The normalised Standard Deviation error $E_\sigma$ in reconstructions using ART for increasing values of nominal absorption coefficient.

However variation in spatial distribution can also be seen within some configurations. For example in an unconfined jet such as Flame D, the reaction zone spreads out from the nozzle (see fig. 5.1) resulting in axial variation in the spatial distribution of the emitting species. Figure 5.3 compares the error the in reconstructions of Phantoms taken from the four different axial locations shown in fig. 5.1. (The difference in error seen initially is likely due to increased asymmetry and irregularity of the downstream Phantoms).

Figure 5.3 indicates that the effect of this variation on the level of self-absorption is minimal—for $OH^*$—as at all the locations the error begins to increase at approximately the same point. At first this seems to indicate that the spatial distribution has a limited effect on self absorption. However, the results seen in fig. 5.3 can be explained by the fact that as the jet spreads the concentrations reduce, as illustrated in fig. 5.1, and these two effects negate each other. This example reinforces the obvious dependence of self-absorption on both spatial and concentration factors.

It is likely that the main absorber of the light emitted by an excited state field is its ground state field [116], and it is desirable to know how the actual value of $\kappa$ for this species compares to the nominal values used so far. The excited state field is clearly a function of the ground state field leading to the strong spatial similarity seen in the LES data; for example, in fig. 5.4. This means that the supporting assumption of eq. 5.5 remains valid, and a reasonable estimate of the ground state $\kappa$ can, therefore, be obtained by simply scaling the excited state $\alpha$ by the ratio of the mean concentrations of the ground

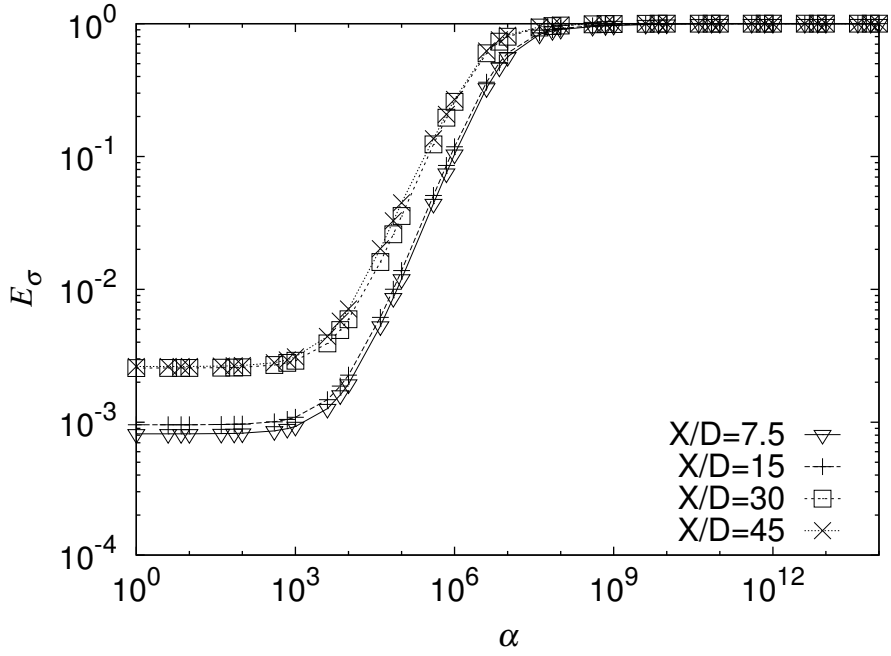


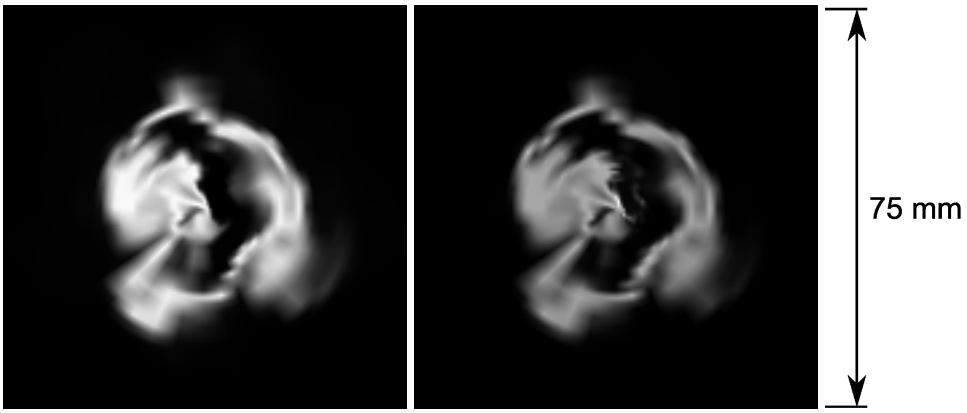

Figure 5.3: The normalised Standard Deviation error $E_\sigma$ in reconstructions using ART for increasing values of nominal absorption coefficient and self absorption. Results are shown for the $OH^*$ Phantoms taken from different axial locations in Flame D, as shown in fig. 5.1.

and excited states. In this study the concentrations are given by mass fractions and lengths in pixels (of Flame D for both configurations), hence $\kappa$ and $\alpha$ have units of pixels$^{-1}$.

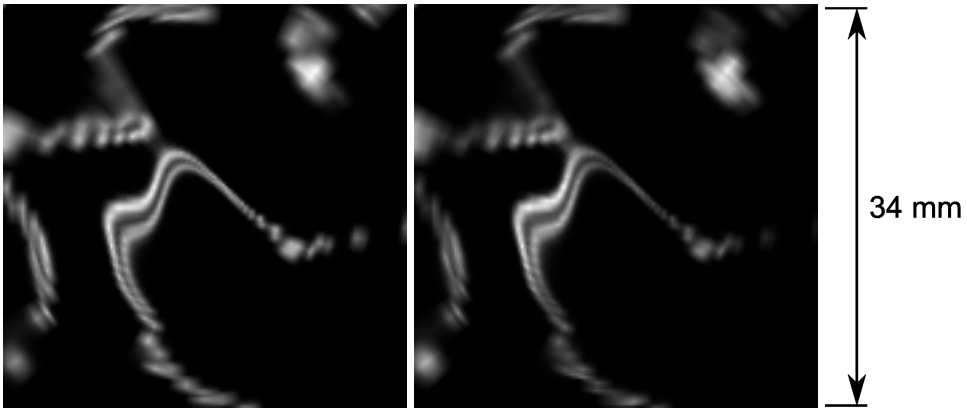
For Flame D and the TOJ, an arbitrary 5% error in the $OH^*$ reconstructions corresponds to an $\alpha$ of approximately $10^5$. From the LES data (for example, fig. 5.4) we can determine that the $OH^*$ and OH fields for Flame D and the TOJ differ in their mean concentrations by a factor of approximately $10^5$. Therefore, the OH absorption coefficient $\kappa_{\mathrm{OH}}$ that would result in 5% error is unity, according to eq. 5.5. The actual absorption coefficient of OH at 309 nm is approximately 180 cm$^{-1}$atm$^{-1}$ [136] at 2000 K based on partial pressures. To allow comparison, we can convert this actual value to our units of pixels and mass-fractions using eq. 5.6.

$$P_{\mathrm{OH}} = Y_{\mathrm{OH}} \frac{MW_{mix}}{MW_{\mathrm{OH}}} P_{total} \tag{5.6}$$

Both flames are unconfined, therefore $P_{total} \approx 1$atm, and from the earlier chemical state simulations (see section 4.1.1) $MW_{OH} = 17.008$ and $MW_{mix} \approx 28.571$. One Flame D pixel is 0.1125 cm, making the actual absorption coefficient of OH in our units as 35 pixels$^{-1}$. This value is much greater than the value of unity $\kappa_{\mathrm{OH}}$ that resulted in 5% error in the Phantom study, and suggests that greater than 10% error could occur. We can conclude then, that for both flame configurations self-absorption of the light emitted by $OH^*$ can cause significant error. By the same analysis, for $CH^*$, 5% error corresponds a ground




(a) FD: $\overline{Y}_{OH} = 0.55 \times 10^{-3}$ (b) FD: $\overline{Y}_{OH^*} = 0.55 \times 10^{-8}$




(c) TOJ: $\overline{Y}_{CH} = 0.16 \times 10^{-6}$ (d) TOJ: $\overline{Y}_{CH^*} = 0.25 \times 10^{-9}$

Figure 5.4: The mass fractions of the excited and ground state species from LES. Quoted below are the mean mass fractions $\overline{Y}$ for that species.

state $\kappa_{CH}$ of $1.6 \times 10^3$pixels$^{-1}$ for the TOJ and $8.4x10^3$ for Flame D (based on a difference ratio of $6.4 \times 10^2$). These are markedly higher than the values for OH and support the supposition that self-absorption is less significant for CH chemiluminescence. Unfortunately it was not possible to obtain an actual absorption coefficient for CH to compare against.

### 5.1.2 Modifying ART for Self Absorption

Having established that self-absorption can be a source of error for some flame configurations and species, we now investigate a modification of the ART algorithm that can reduce this error.

The Beer-Lambert law has been successfully used to incorporate self-absorption into the simulation of measurements using Phantoms. This simulation process is essentially the same as the forward projection step of ART, and so by performing the same weighting of $w_{qpv}$ self-absorption can be included in ART. A similar approach has been demonstrated by Correia et al. [34]. As the emitting and absorbing fields are likely to be correlated, non-linearity can be introduced into the reconstruction problem. Fortunately, as an iterative approach ART, is capable of solving non-linear reconstruction problems. Where previous studies have used absorption measurements or soot modelling to determine a suitable absorbing field $c(i, j)$, we assume that the spatial extent of the absorbing field is the same as the emitting field, as used for simulating the measurements. This model is necessary in the absence of more *a priori* information, and is justified based on the spatial similarity of the fields observed previously in the LES.

The Self-Absorption Modified ART algorithm (SMART) is thus formed by replacing the ART weighting factor $w_{qpv}$ with a new weight factor $w^A_{qpv}$ given by eq. 5.7. Here again the sum in the exponential term includes all the pixels that lie between the detector and the current pixel.

$$w^A_{qpv} = w_{qpv}\, e^{-\alpha \sum_{\upsilon=v}^{\text{edge}} w_{p\upsilon} f^h_\upsilon} \tag{5.7}$$

#### 5.1.2.1 Testing the ART modification

Example reconstructions of Flame D OH$^*$ and the Turbulent Opposed Jet OH$^*$ using SMART are shown in figs. 5.5 and 5.6 alongside the Phantoms and ART reconstructions for comparison. Both reconstructions were performed using projections that were subjected to simulated self-absorption with a nominal absorption coefficient of $10^5$. This value of $\alpha$ resulted in 5% error when unmodified ART was used, however both reconstructions using SMART have error of less than 0.25%.

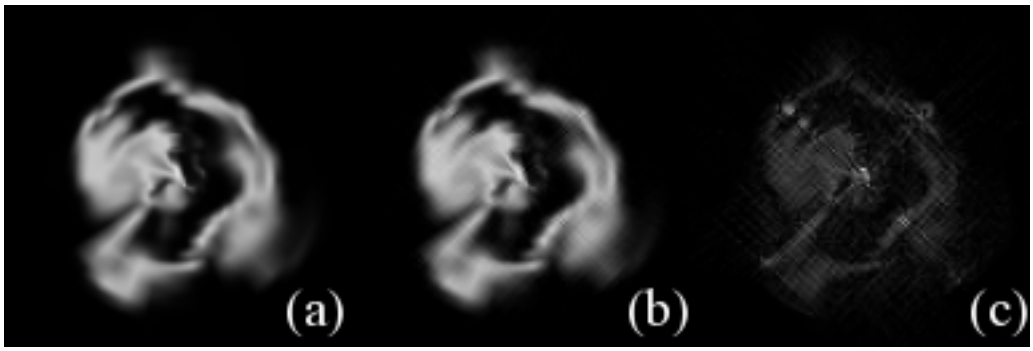



Figure 5.5: The Flame D $OH^*$ Phantom (a) and its reconstructions using SMART (b), and ART (c) when the projections are subjected to strong self-absorption ($\alpha = 10^5$). These example reconstructions were performed on $128^2$ pixel grids using 128 angles and 128 rays.

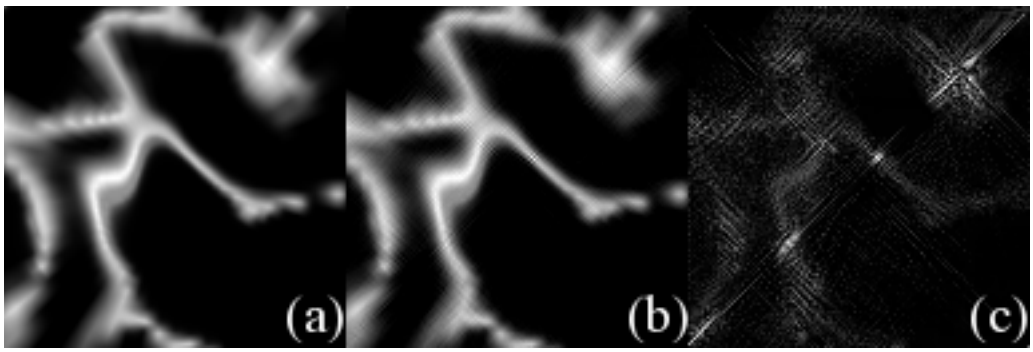



Figure 5.6: The TOJ $OH^*$ Phantom (a) and its reconstructions using SMART (b), and ART (c) when the projections are subjected to strong self-absorption ($\alpha = 10^5$). These example reconstructions were performed on $128^2$ pixel grids using 128 angles and 128 rays.

Figures 5.7 and 5.8 further confirm the ability of SMART to reduce the error resulting from self-absorption. In these figures, the error in reconstructions obtained using the SMART algorithm is shown alongside those of standard ART. For both species and flame configurations, SMART is seen to reduce the error by several orders of magnitude for the same level of self-absorption ($\alpha$) and allows reconstructions of low error to be obtained at higher values of $\alpha$ than is possible with the unmodified ART algorithm.

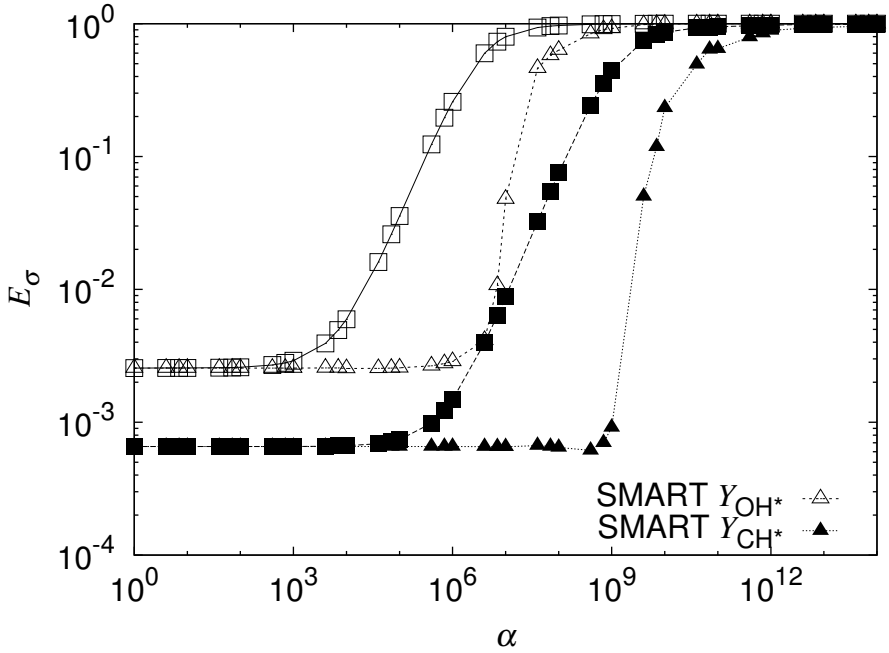



Figure 5.7: The normalised standard deviation error $E_\sigma$ in reconstructions using ART and SMART for increasing values of nominal absorption coefficient. Results are shown for the Flame D $OH^*$ and $CH^*$ Phantoms.

Figures 5.7 and 5.8 also clearly show that, as with ART, there is a point at which the SMART error increases significantly. This is however, several orders of magnitude above the equivalent $\alpha$ for ART and likely corresponds to the point at which the projections start attenuating to zero.

### 5.1.3 Conclusions

The effect of self-absorption on the projection measurements has been simulated by using the Beer-Lambert law. Reconstructions of $OH^*$ and $CH^*$ Phantoms were performed, using ART, for two contrasting flame configurations using increasingly attenuated projections. For the same absorption coefficient greater error was seen for the TOJ than for Flame D, highlighting the need to consider the spatial extent and concentrations of a flame when estimating its opacity, and not just the absorption coefficient.

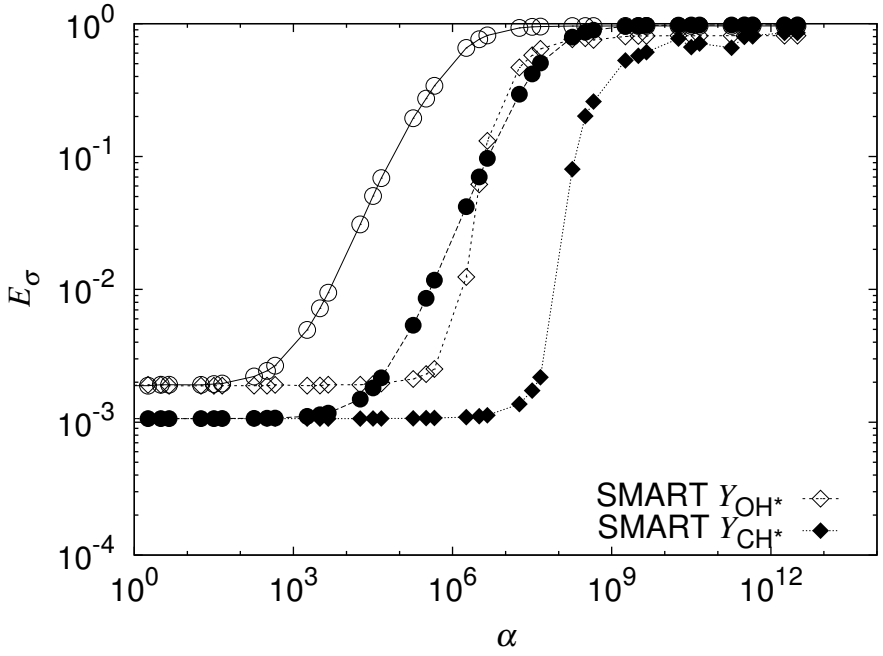



Figure 5.8: The normalised standard deviation error $E_\sigma$ in reconstructions using ART and SMART for increasing values of nominal absorption coefficient and self absorption. Results are shown for the TOJ $OH^*$ and $CH^*$ Phantoms.

The results show that for reconstructions of $OH^*$ self-absorption can cause significant error in the reconstruction when using unmodified ART, for both Flame D and the TOJ. Consideration of the actual ground state OH absorption coefficient given by Seitzman [136], suggests that this error could exceed 10%. For $CH^*$ the error becomes significant at much higher $\alpha$ than $OH^*$ supporting the supposition that self-absorption is negligible for this species.

A modification to ART that uses the Beer-Lambert law to weight $w_{qpv}$ has been presented. The algorithm, termed SMART, assumes that the absorbing field is proportional to the emitting field, and requires only the absorption coefficient to be known *a priori*. The results show that SMART can negate the error caused by self-absorption in the reconstruction up to the point where the measured projections are (possibly) attenuated to zero. The assumption of proportionality between the emitting and absorbing fields has been justified from the LES data for the case where the ground state is the significant absorbing field. However, this assumption is less valid where other significant absorbing species are identified, such as $H_2O$. In such a case, it may be possible that a modelling approach, such as laminar flamelets, could be used to relate the emitting and absorbing species and allow correction of self-absorption.

## 5.2 Enhancements to ART

### 5.2.1 Nested Grids

When using CTC to study transient events, measurements would need to be taken at enough points in time so that the event can be tracked. This can lead to very large 3D data sets and the convergence time of ART becomes significant. The use of sparse matrices has a dramatic effect on convergence times, and in this section we seek further efficiency increases through the use of nested grids.

The nested grid approach works by initially reconstructing a solution on a very coarse grid, and this coarse solution is then used as an initial estimate for solution on a finer grid. This is repeated until the desired resolution is reached, with transfer between grids achieved through linear interpolation. In general, a new projection matrix ($W$) would need to be determined for each intermediate grid. However, if these grids are chosen to be integer divisions of the ultimate reconstruction grid, then $W$ for the intermediate grids can be more quickly determined by simply reducing the projection matrix of the final grid.

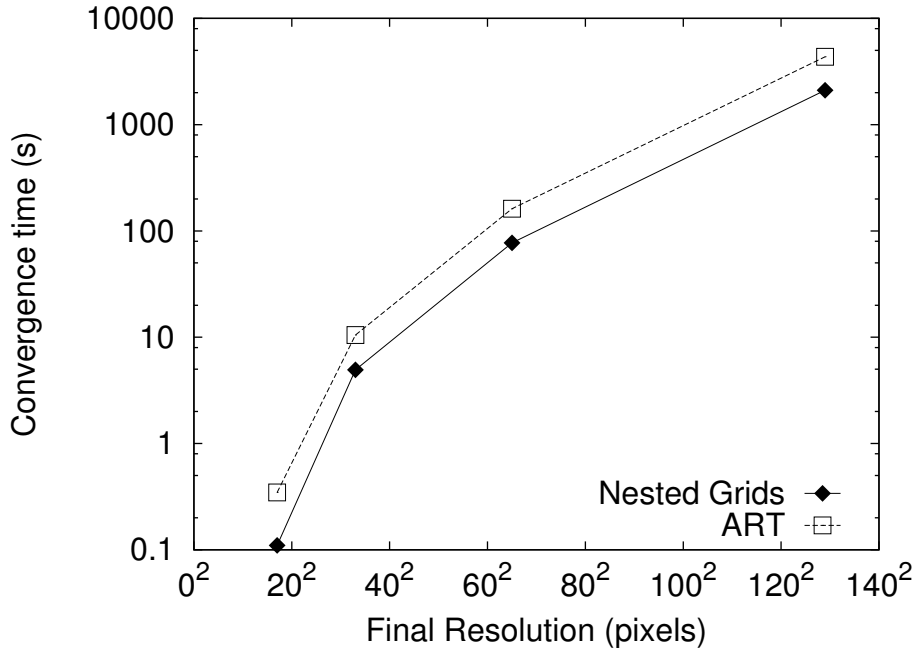



Figure 5.9: Convergence time comparison for ART and ART using nested grids, for different reconstruction grid sizes.

Reconstructions of a $128^2$ pixel Phantom were performed using the nested grid approach with intermediate grids of 16, 32, and 64 pixels (and without exploiting the sparse matrices). A comparison of convergence times using the standard ART algorithm and the nested grid approach can be seen in figure 5.9. This figure clearly indicates an improvement in the convergence time of approximately 50% and even greater savings could be

expected when 3D grids are used.

The use of nested grids has not been necessary in the rest of this work because the sparse matrices alone have improved the computation time sufficiently. Speed is not particularly important at this stage and maintaining freedom in the grid sizing is preferred. Where convergence speed is more important flexibility in the grid sizing can still be obtained using nested grids by pre-computing the intermediate $w_{qpv}$.

### 5.2.2 Linear Interpolation on the Forward Projection

The forward projection step of ART is essentially a numerical approximation of an integral taken through a (2D) field. The field is divided into pixels and within each pixel a constant value $f_v$ is assumed. The proportion of $f_v$ that contributes to the integral is given by $w_{qpv}$ and is usually taken as the intersection length or area of the projection and the pixel; however, we are free to choose another basis that may better approximate the true integral. When a 2D field is smooth and can be adequately sampled, then bi-linear interpolation can be used to determine any point in the field. A potentially more accurate value for the integral ($I_{qp}$) can thus be obtained by using interpolation to determine all the values that lie on or within a projection. In this section we investigate whether a bi-linear basis for the ART forward projection step offers improvements in the reconstruction accuracy of smooth fields.

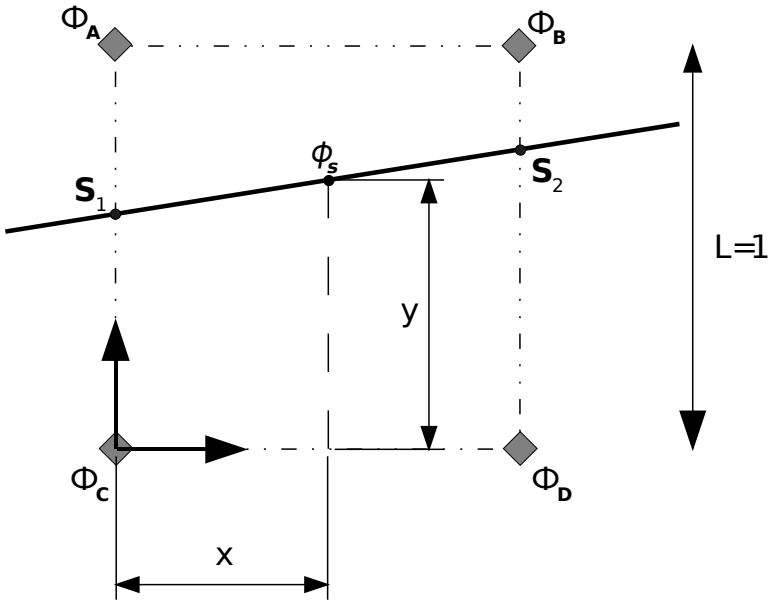



Figure 5.10: An interpolation cell. Diamonds represent object/Phantom sample points. $\mathbf{S}_1$ and $\mathbf{S}_2$ are vectors defined by local co-ordinates.

We consider the case where the projection is a line integral through the 2D field.

The value $\varphi_s$ of any point **S** on the line is obtained by bi-linear interpolation of the four surrounding sample points seen in fig. 5.10. The value $\varphi_s$ is thus given by eq. 5.8 where $T$ is the weighting of each sample point.

$$\varphi_{\mathbf{S}} = \Phi_{\mathbf{A}} T_A + \Phi_{\mathbf{B}} T_B + \Phi_{\mathbf{C}} T_C + \Phi_{\mathbf{D}} T_D \tag{5.8}$$

Within the interpolation cell defined by the surrounding samples the integral of $\varphi_s$ is given by eq. 5.9, where **C** and **Z** are given by eq. 5.10, and $m$ and $c$ are local constants that can be derived from $\mathbf{S}_1$ and $\mathbf{S}_2$.

$$\int_{\mathbf{S}_1}^{\mathbf{S}_2} \varphi_{\mathbf{S}} ds = \Phi \cdot \mathbf{T} \quad ; \quad \Phi = \begin{bmatrix} \Phi_A \\ \Phi_B \\ \Phi_C \\ \Phi_D \end{bmatrix} \quad ; \quad \mathbf{T} = \mathbf{CZ} \times \sqrt{1+m^2} \tag{5.9}$$

$$\mathbf{C} = \begin{bmatrix} c & \frac{m-c}{2} & -\frac{m}{3} \\ 0 & \frac{c}{2} & \frac{m}{3} \\ 1-c & \frac{c-m-1}{2} & \frac{m}{3} \\ 0 & \frac{1-c}{2} & -\frac{m}{3} \end{bmatrix} \quad ; \quad \mathbf{Z} = \begin{bmatrix} (S_{2x} - S_{1x}) \\ (S_{2x}^2 - S_{1x}^2) \\ (S_{2x}^3 - S_{1x}^3) \end{bmatrix} \tag{5.10}$$

For the case where the projection passes vertically through the cell **T** is given by equation (5.11)

$$\mathbf{T} = \begin{bmatrix} \frac{1-x}{2} \\ \frac{x}{2} \\ \frac{1-x}{2} \\ \frac{x}{2} \end{bmatrix} \tag{5.11}$$

Using eq. 5.9 we can recast the forward projection step of ART to include bi-linear interpolation as seen in eq. 5.12 where the sum is over all the interpolation cells that the $p^{\text{th}}$ projection passes through.

$$I_p = \sum_{\iota}^{\text{line}} \Phi_\iota \cdot \mathbf{T}_\iota \tag{5.12}$$

The use of bi-linear interpolation to improve the accuracy of the forward projection step in ART has been considered previously by Andersen and Kak [6] as part of their SART implementation. In SART the forward projection is re-expressed using bi-linear elements as basis functions instead of pixels and appears functionally the same as the approach described in this work. However, Andersen and Kak approximate the exact integral by sampling at a lower resolution to reduce computational cost.

Reconstructions of the jet Phantom seen in section 4.1.2 were performed using bi-linear interpolation on the forward projection (intART) and standard ART for a range of

resolutions. In each reconstruction the number of angles and projections were the same as the number of pixels in the reconstruction. The accuracy of the reconstructions is judged from the error in the RMS value of the reconstruction as compared to the Phantom. This error normalised by the Phantom RMS value and is shown in figure 5.11 for the range of resolutions considered.

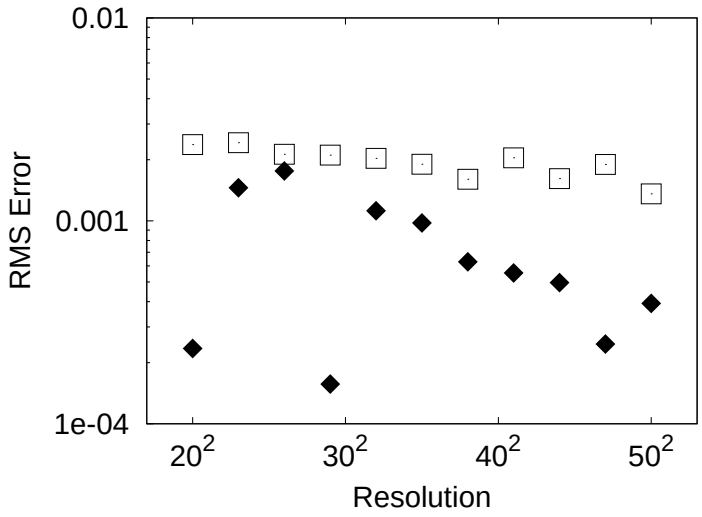



Figure 5.11: The error in the RMS of reconstructions using ART (squares) and intART algorithms (diamonds). All algorithms used the same convergence and relaxation parameters.

Figure 5.11 shows that the addition of interpolation to the ART algorithm has resulted in a reduction in the error of the reconstructions; thereby, indicating that the bi-linear approach to determining the projections is more accurate, as also claimed by Kak and Slaney [6]. However, the reconstructions using bi-linear interpolation took significantly longer to converge than standard ART, as four times as many operations are required in the 2D case to calculate each pixel's contribution to an integral. This modification has not been used elsewhere in this work because the modest improvements in accuracy do not justify the increase in complexity and computation time.

## 5.3 Extending ART to 3D

So far 3D datasets have been obtained by *stacking* a number of 2D reconstructions obtained from 1D views. However this approach greatly restricts the choice of camera locations as the sensors have to be located in the same plane. Where optical access is limited this could lead to a reduction in the number of angles that can be taken and so an undesirable reduction in the spatial resolution. Furthermore, the need to ensure that sensors are co-planar makes the physical set up of the experiment more difficult. By using 2D CT the

projections geometries are also restricted to 2D and any out-of-plane contribution made, for example due to the light collection optics, results in error. In medical applications, ART is already used with 3D *cone beam* projections [111], which reduce the patient radiation dosage. In this work to allow the maximum flexibility in sensor locations and to allow 3D projections to be used where appropriate, we seek to redefine ART in *true* 3D and avoid the stacking of 2D sections.

### 5.3.1 3D projections

We begin by defining the 3D coordinate system as shown in fig. 5.12. A significant difference to the 2D system is that the *view* is now a 2D set of projections, which all view the object from approximately the same 3D location and angle. A view can be thought of as the measurement image from one camera. Where previously $p$ was used to refer to a single projection, $p$ is now also synonymous with a single pixel of the 2D view with indices $m$ and $n$. The 3D coordinate system also introduces a new angle, $\gamma$, that is measured from the view normal direction $Z$, to the $xy$ plane of the object domain, and thus permits out of plane view locations. Where before the object domain was a 2D section divided into pixels it is now a 3D volume divided into voxels, and $v$ now refers to a voxel that has indices $i$, $j$, and $k$.

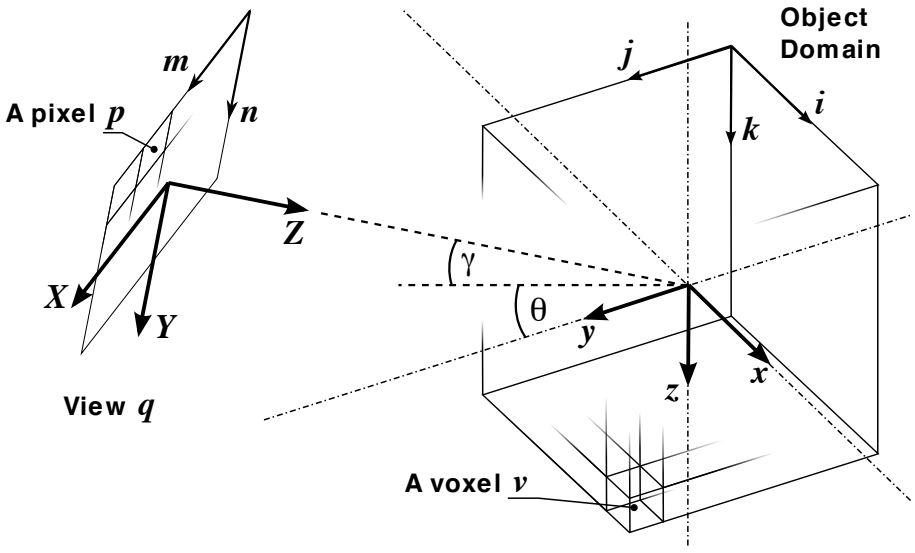



Figure 5.12: the 3D coordinate system.

The ART algorithm is based on projecting back an error generated by a comparison of the measured projection with an equivalent forward projection taken through the reconstruction domain. In extending ART to 3D this process is unchanged and all that is required is to re-express the forward and back projection steps in 3D. In 3D coordinates

additive ART is given by eq. 5.13, from which it can be seen that $w_{qpv}$ represents the forward and back projection steps.

$$f_v^{(h+1)} = f_v^{(h)} + \beta w_{qpv} \frac{I_{qp} - \sum_{v=1}^{N_v} w_{qpv} f_v^{(h)}}{\sum_{v=1}^{N_v} w_{qpv}^2} \tag{5.13}$$

Previously $w_{qpv}$ was taken as the fractional area of the projection-pixel intersection, and we now simply extend this to be the fractional volume of the intersection between a voxel and the 3D projection.

The use of 3D geometries permits out-of-plane contributions to the value of the projection to be included, and there is much scope for the development of specific (3D) geometries that better match the actual path of the light which reaches the sensor. At this stage we seek to define a simple 3D projection geometry that is consistent with that used so far in 2D; however, other 3D projection geometries are considered later in this work. For the 2D case the projections have been assumed to have some width and we can extend this into 3D by rotating these projections about their central axis to give a cylindrical projection geometry. Edwards [39] has also used a cylindrical approximation to simulate optical projection measurements, but has not, as far as this author is aware, employed it in CT. A perhaps more accurate choice would be to use an extended cuboid shape as this would more closely match the shape of the CCD pixel; however, the cylinder was chosen for its ease of implementation (when $\gamma > 0$) and is still a reasonable approximation to the pixel shape.

### 5.3.2 Notes on the Computer Implementation

The calculation of the intersection volumes of the 3D projections and the voxels is non-trivial and we describe the method used in this work for cylindrical projections. We begin by defining each cylindrical projection, $qp$, by three variables. The first is $\mathbf{C}_{qp}$ which is some point on the central axis of the projection, for example its centre in the object domain. The second is the unit vector $\mathbf{U}_{qp}$ which gives the direction of the ray. The third is the radius of the projection $r_{qp}$.

To determine the intersection a numerical approach is used whereby every voxel in the object domain is tested for intersection with every projection. To test if a particular voxel $v$ intersects the cylindrical projection, first the relative position vector $\mathbf{V}$ of the voxel is found by deducting $\mathbf{C}_{qp}$ from the centre of the voxel, as shown in eq. 5.14. (Note that

voxel units are used.)

$$\mathbf{V} = \begin{bmatrix} x_v \\ y_v \\ z_v \end{bmatrix} - \mathbf{C}_{qp} \tag{5.14}$$

Following this the radial component of $\mathbf{V}$ in the projection reference frame is determined from eq. 5.15.

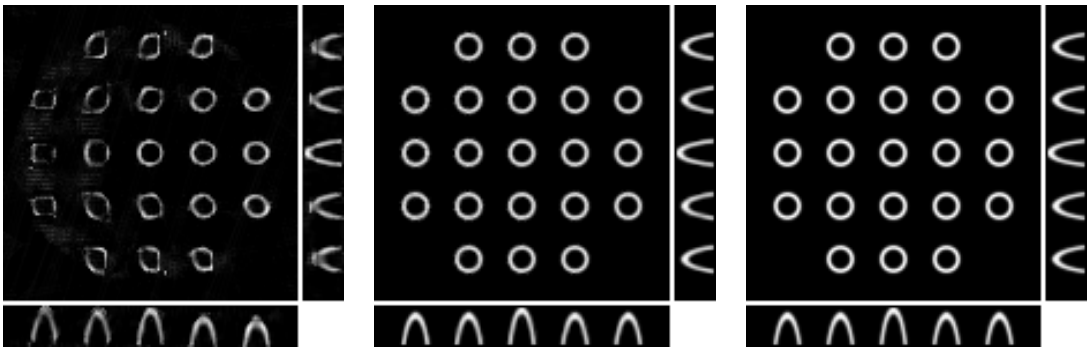
$$V_r^2 = |\mathbf{V}_{qp}| - \left(\mathbf{V} \cdot \mathbf{U}_{qp}\right)^2 \tag{5.15}$$

A voxel can then be said to intersect the cylindrical projection when $V_r < r_{qp}$. However, to avoid aliasing the condition $V_r < r_{qp} + 1$ is initially used, i.e. the projection radius plus one voxel width, to identify voxels which *may* intersect the projection. For these voxels a higher resolution sub-grid is then defined and then iterated over using the correct $V_r < r_{qp}$ condition. The sum of the values from the sub-grid is then used to approximate the intersection volume of the voxel with the cylindrical projection. In this work a sub-grid resolution of $10 \times 10 \times 10$ was found to be a good trade off between accuracy and computational cost.

The computational cost of this numerical approach to determine the intersectional volumes is acceptable because it only needs to be performed once for any CTC set-up. Also no analytical definition of the 3D intersection of a cylinder and a voxel could been found in the literature and the task of trying to derive one would not be for the faint of heart! Furthermore, this implementation has not focused on speed and if necessary it could be made more efficient by using integer arithmetic or analytically determining the set of possibly intersecting voxels. The method described here could also easily be computed in parallel.

### 5.3.3 Phantom Test

To test the 3D ART implementation, a Phantom shaped similarly to the Matrix Burner was synthetically generated on a grid of 128 by 128 by 18 voxels. The Phantom was then in turn used to generate 32 views which were spaced equally through 180 degrees about the $z$ axis. To make the views *more* 3D each was inclined by an angle of 5°to the $xy$ plane, i.e. $\gamma = 5°$for all views. For this set of views two reconstructions were performed. For one, the 3D nature of the projections was properly represented and $\gamma$ was set to the correct value; for the other, the views were considered to be co-planar and hence the value of $\gamma$ was incorrectly set to zero. Reconstructions were performed on the same grid as the Phantom with convergence achieved after 12 iterations for the inclined views, and 9 for the co-planar views. Sections of both reconstructions are shown in fig. 5.13 and a volume rendering of each is shown in 5.14, alongside the Phantom in both figures.



(a) Co-planar: $\gamma$ incorrect (b) Inclined: $\gamma$ correct (c) Phantom

Figure 5.13: Example sections from the co-planar and inclined reconstructions as well as the Phantom. Shown is a section in the $xy$ plane at $z = 16$ (middle) along with central sections in the $yz$ (right) and $xz$ (below) planes.

These figures show that for the inclined views the reconstruction is a near perfect match to the Phantom, and the validity of the 3D ART implementation is further confirmed by their cross-correlation value of 0.996. It is also interesting to note how well this reconstruction matches the Phantom, despite having used relatively few views. The reconstruction using co-planar views does not match the Phantom as well, and achieves a much lower cross-correlation value of 0.757. The co-planar reconstruction suffers from a degree of deformation, akin to stretching, of the circular flame shapes and overall the flames seem tilted in the $x$ direction. The difference between the inclined and co-planar results demonstrates the importance of correctly representing the path of the actual projections when seeking accurate reconstructions. However, in spite of the (deliberate) mis-registration of the viewing angle, the co-planar reconstruction has still captured the main features of the Phantom and resolved to some degree all of the constituent flames, thereby demonstrating the robustness of the ART algorithm. Furthermore, we have demonstrated a useful way of studying the error caused from mis-registration of the view location in the third dimension, and this is returned to later.

## 5.4 Accounting for Optical Effects

So far the optics of the cameras have been neglected, and all the projections have been assumed parallel and to project cylindrically from the CCD sensor. While these assumptions are clearly consistent with laser absorption measurements, they are also reasonable in our emission case when the object is sufficiently distant and the optical depth-of-field exceeds the object width. However, where the maximum temporal resolution is sought, such an optical arrangement may not be desirable as will be explained.



(a) Co-planar: $\gamma$ incorrect (b) Inclined: $\gamma$ correct (c) Phantom

Figure 5.14: Volume renderings of the co-planar and inclined reconstructions, and the Phantom, taken from a view 30°above the *xy* plane and at an angle of 30°to the *y* axis.

The depth-of-field of a camera depends on essentially two variables: the size of the opening and the distance of the camera from the object, with smaller apertures and a greater object distances giving a wider depth-of-field. However, a trade-off exists; because at greater distances and smaller apertures less light is able to reach the sensor as the solid angle is reduced and so weaker measurement signals are recorded. This would lead to a decrease in the temporal resolution achievable for a given camera, and furthermore, the increased object distance could lead to problems with packaging as well as a reduction in the spatial resolution.

We have seen previously that to avoid error the projections need to approximate the actual integral of the collected light. Therefore, as the camera optics serve to collect and direct light to the sensor, we can seek a better representation of the true integral, by using an appropriate optical analysis to determine $\mathbf{w_{qp}}$ for each projection. By including the effect of the camera optics we can remove the depth-of-field limitation and so avoid the earlier described trade-off and allow smaller object distances and wider apertures.

An alternative *optical* projection is now presented that is based on a weighted double cone shape, derived from geometric optics to represent a simplified version of the camera optics. Following this, a complimentary modification is described that allows non-parallel projections to be used to account for perspective effects. This modification and the new optical projection are then experimentally tested using views of insufficient depth-of-field obtained from the set-up seen previously in section 4.4.

### 5.4.1 Projections for Insufficient Depth-of-Field

The limitations to the depth-of-field are a result of the light collection optics (lenses) employed in the cameras. In this section an optical analysis is performed to determine an appropriate projection geometry that can account for this effect. A branch of optics

termed geometric optics provides a relatively simple but powerful way of describing the propagation of light through an optical system: ray-tracing. Ray-tracing can represent the depth-of-field effect, and is favoured here for its simplicity as a natural extension to the use of rays so far seen in the CT context. Only the main assumptions and the application of the ray-tracing are described here and for further information the book by Hecht [66] is recommended.

The use of ray-tracing to better define the optical projection geometry has been considered previously by Walsh et al. [156] who showed good agreement with measured data. However, these results are less informative as their measurements had a depth-of-field that exceeded the object dimensions. Also, this work considered stacked 2D Abel transforms, not 3D CT, and did not explicitly address depth-of-field as in this work. Andersen [3] has also applied geometric optics to improve reconstructions, for diffracting and refracting sources such as ultra-sound, but considered the forward projection only. Edwards [39] considered the collection volume of an optical emission sensor and noted that rather than a line integral, an optical sensor collects light from a finite volume of non-zero solid angle. Edwards however, went on to simplify the projection geometry of their fibre optic probe back to a line integral using various assumptions, most notably there being no variation in focus along the line-of-sight.

The assumption that the geometry of the optical elements (apertures in this case) is much larger than the wavelength of the light, allows the propagation of a spherical light wave to be represented by a number of rays that can be traced through the optical system. Using such ray-tracing means that wave effects such as diffraction are not accounted for. However, to maximise signals large apertures are used in this work, and the effect of diffraction is small. As a further simplifying assumption, the angles made by any traced rays with the optical axis are taken to be small, such that $tan(\theta) \approx \theta$. This implies that a wave passing through a lens remains spherical and converges to a *perfect* image at the image plane. This is termed the paraxial approximation, and is valid given that typically the lens to object distance $S_o$ would likely be much greater than the aperture diameter $f/N$, even in the limited depth-of-field case.

#### 5.4.1.1 Projection Geometry from Ray-Tracing

To further simplify the optical analysis the assumption is made that the system of lenses employed by the camera can be represented as a single thin lens that is co-located with the aperture. This is a somewhat necessary assumption as the specific details of a lens may not be known. This allows the thin lens equation to be used (eq. 5.16) that relates the distance from the lens to the object plane $S_o$, and the distance from the lens to the image

plane $S_i$, to the lens focal length $f$.

$$f^{-1} = S_o^{-1} + S_i^{-1} \tag{5.16}$$

It is also assumed that the lens responds uniformly to all wavelengths of light, which is reasonable given that the visible spectral range of a non-sooting flame is limited, especially where filters are employed. This allows the propagation of a point source of light in the object domain to be traced through the lens to the camera image plane as shown in 2D in fig. 5.15. This figure shows that a conical segment of the diverging spherical wave, originating from the point $v$, is refracted by the lens, and converged to a point $p$ on the image plane. (The relationship between $p$ and $v$ is defined by the thin lens equation.) The wave segment is represented by two limiting rays referred to as the *marginal* rays (the thick continuous lines in fig. 5.15) which define the extent (solid angle) of the wave captured by the circular aperture.

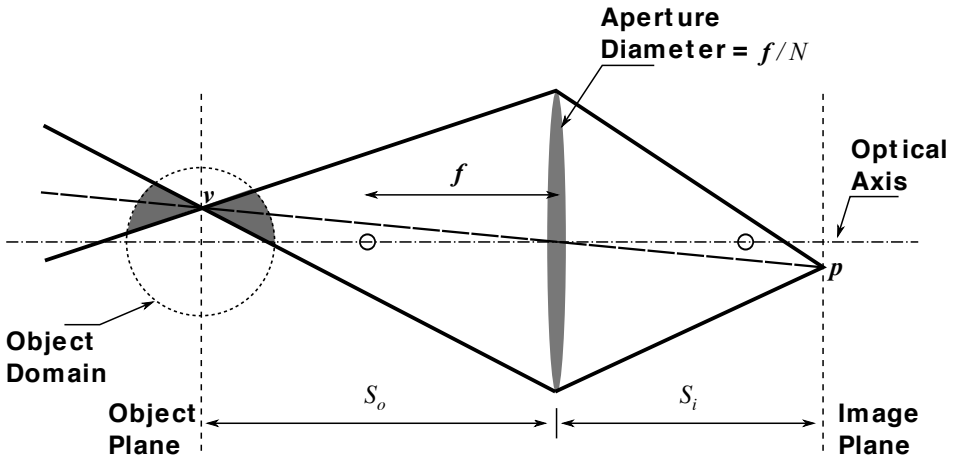



Figure 5.15: A diagram illustrating in 2D how ray-tracing describes the propagation of a spherical light wave diverging from point $v$ in the object domain, through a circular lens and aperture, and converging back to a point $p$ on the image plane. The double cone defined by the *marginal* rays (thick continuous lines) is formed about the central ray (dashed line). Note the diagram is not to scale, $f$ is the focal length of the lens.

In fig. 5.15 the central ray (the dashed line) is shown to illustrate how light propagation occurs at some angle to the optical axis for points located away from the optical axis. It follows then, that points located at different distances from the optical axis would have central rays that are not parallel. (Accounting for these non-parallel projections is returned

to later in section 5.4.2.) However, as the object distance $S_o$ increases, or the aperture decreases, the central ray becomes increasingly parallel to the optical axis, as do those of other points, and the assumption of parallel rays used previously becomes valid.

The ray-tracing analysis illustrates the trade-off between depth-of-field and SNR. A smaller aperture diameter increases the depth-of-field but means a smaller segment (or solid angle) of the spherical wave is collected, and so less energy focused to the image plane. Increasing the object distance has the same effect, though for most cases $S_o >> f$ and the diameter has the greatest influence.

One useful aspect of ray-tracing is that rays can be traced backwards as well as forwards. This means that a point on the image plane can be selected, say $p$, and by tracing back the rays from that point, the point of origin, $v$, can be found. This results in a double cone centred at $v$ that bounds a volume within which, it can be shown, a portion of any light emitted will arrive at $p$. If we restrict the emission of light to the object domain then the volume is given by the intersection of the domain and the double cone. This volume, shaded grey in fig. 5.15, therefore defines the volumetric extent of the projection of point $p$ in the object domain, and is the same as that used by Walsh et al. [156]. This double cone volume provides the geometry of the optical projection when it is measured at a point $p$. However, in this work an array of pixels (a CCD) is used to detect the light and the double cone needs to be extended to account for the finite area of the pixels, as seen in fig. 5.16. In this figure, rays have also been traced back from the pixel edges (the thick and dashed rays), and the double cone has expanded accordingly and no longer converges to a point in the object domain.

The pixels of most CCDs are usually quoted as being square or rectangular, though how much this relates to the ratio of the $X$ and $Y$ dimensions over the actual shape is unclear. Here we assume the pixel area is circular as this leads to a simpler circular section in the double cone which is important as a large number of the double cone projections will need to be numerically defined in the object domain. A circular pixel has the same dimensions in $X$ and $Y$ as a square pixel which itself is unlikely to be a *perfect* square, due to micro-lenses[1] for example. Furthermore, cylindrical ray projections have been successfully used earlier in this work.

To summarise: the expanded double cone in fig. 5.16 provides a reasonable approximation to the integral geometry of the projection for this single thin lens optical system. The double cone is defined by the object distance $S_o$, the aperture diameter $f/N$ and the magnification of lens. This geometry can then be used to define the voxels in the object domain from which the light, measured by a particular CCD pixel, originates. However, this geometric definition alone is not sufficient to describe the optical projection function,

[1] Used in interline CCDs to increase the fill factor

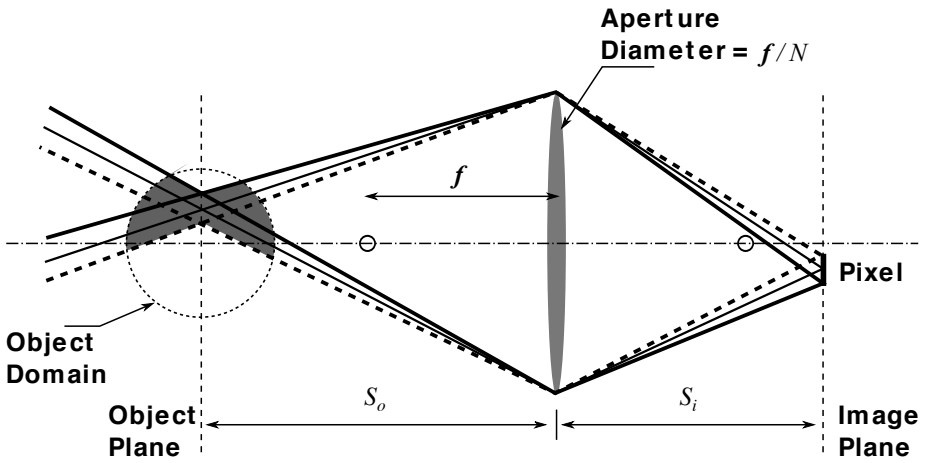



Figure 5.16: Rays traced back from the pixel limits (thick and dashed lines) expand the double cone of fig. 5.15 to the shaded area as shown.

i.e. $\mathbf{w_{qp}}$. This is because different voxels within the double cone geometry have differing portions of their emitted light reach the measuring pixel. This results from the inability of the lens to focus points away from the object plane to points on the image plane, and is related to the depth-of-field effect. The additional necessary weighting to account for this is dealt with in the following section.

#### 5.4.1.2 Blur Weighting of the Projection

For the cylindrical projections defined previously the lens optics can be ignored and all the light emitted within the ray geometry arrives at the same pixel, meaning that only the volumetric intersection of a voxel, $v$, with the cylinder is needed to define $w_{qpv}$. For the double cone projection however, the focal properties of the lens means that only a portion of the light emitted within the cone will arrive at the CCD pixel—giving rise to the familiar blurring of out of focus objects. This is closely related to the depth-of-field effect which we consider now, again using ray-tracing, to determine an appropriate weighting.

As seen, for a point source on the object plane the lens will focus the light to an equivalent point on the image plane neglecting other effects. If we now consider a point away from the object plane and nearer to the lens, the light is now focused to a point beyond the image plane, as illustrated by the thick continuous line in fig. 5.17. The sensor on the image plane interrupts the light wave before it is focused and its energy is spread over an area referred to as the blur circle (diameter $D_{blur}$). The same effect occurs for

points that are further away from the lens, only in this case the wave is focused to early, and diverges before reaching the image plane (illustrated by the thin continuous line in fig. 5.17).

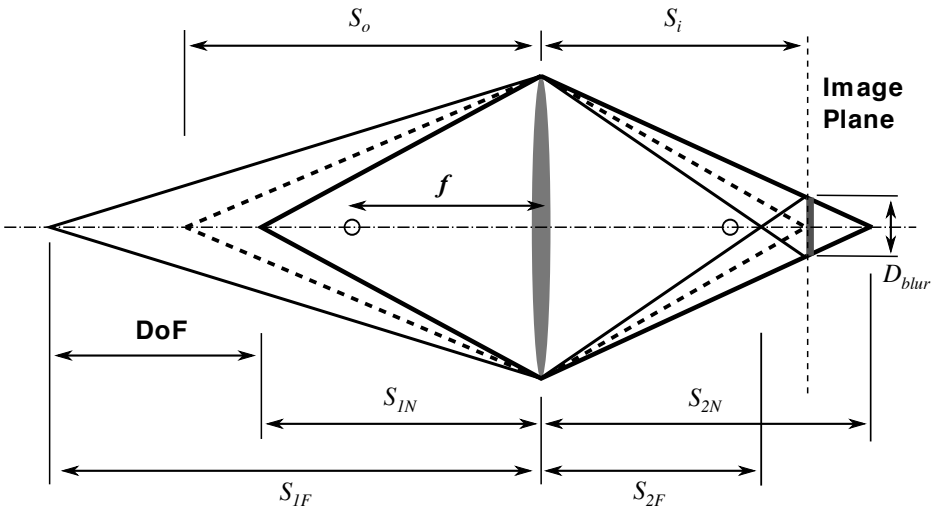



Figure 5.17: An illustration of the depth-of-field determined by ray-tracing. Points of light originating away from the object plane at $S_o$ are spread over a circle of diameter $D_{blur}$ on the image plane at $S_i$. The depth-of-field is determined from the near and far limits when $D_{blur}$ reaches the maximum acceptable size termed the circle-of-confusion.

Where $D_{blur}$ is greater than the pixel size then it is clear that some of the energy emitted will be spread on to other pixels causing blurring. Where $D_{blur}$ is less than the pixel size then all the energy will arrive at the pixel and the image will be sharp. The value of $D_{blur}$ for a point at $S_1 \neq S_o$ can be determined from similar triangles and the thin lens formula applied to the out of plane point (eq. 5.17).

$$f^{-1} = S_o^{-1} + S_i^{-1} = S_1^{-1} + S_2^{-1} \tag{5.17}$$

It follows then that the specification of an acceptable blur diameter, for example the pixel size, allows the near and far limits in the object domain to be determined, and within these limits the image will have acceptable sharpness. As seen in fig. 5.17 the difference between these near and far limits, $S_{1F} - S_{1N}$ is the depth-of-field (DoF). The area of acceptable blur is referred to commonly as the circle of confusion and is usually the order of the CCD pixel size. In this work we will use the CCD pixel size as an approximation though where pixel binning or down-sampling of the views occurs this is adjusted accordingly.

If the energy of a light wave is spread evenly over the blur circle of area $A_{blur}$, then the *portion* of the light measured by a pixel of area $A_p$ would be the ratio of these areas. This ratio thereby provides a suitable weighting for the double cone projection geometry, as defined in eq. 5.18, where $G_{qpv}$ is the fractional intersection volume between a voxel $v$, and the double cone representing the projection from pixel $p$.

$$w_{qpv} = G_{qpv}\frac{A_p}{A_{blur}} \tag{5.18}$$

The value of $A_p$ is known from the CCD specifications and any binning or downsampling performed. An expression for $D_{blur}$ has been derived from similar triangles (see fig. 5.17) as follows.

$$\frac{f}{NS_2} = \frac{D_{blur}}{(S_2 - S_i)} \tag{5.19}$$

$$D_{blur} = \frac{f(S_2 - S_i)}{NS_2} = \frac{f}{N}(1 - S_i S_2^{-1}) \tag{5.20}$$

Rearrangement of the thin lens formula (eq. 5.17) gives

$$1 - S_i S_2^{-1} = \frac{(S_1^{-1} - S_o^{-1})}{(f^{-1} - S_o^{-1})} \tag{5.21}$$

which allows $D_{blur}$ to be expressed using variables only from the object side as seen in eq. 5.22.

$$D_{blur} = \frac{f}{N}\frac{(S_1^{-1} - S_o^{-1})}{(f^{-1} - S_o^{-1})} \tag{5.22}$$

By taking the modulus of the numerator from eq. 5.22 we obtain eq. 5.23 which provides $D_{blur}$ for both near and far points.

$$D_{blur} = \frac{f}{N}\frac{|S_1^{-1} - S_o^{-1}|}{(f^{-1} - S_o^{-1})} \tag{5.23}$$

By using eq. 5.18 we implicitly assume that the pixel area is always enclosed by the blur circle. This would be true for points lying on or close to the *central* ray; however for other points the blur circle would only partially intersect the pixel area, as illustrated in fig. 5.18. In this case, using solely the pixel to blur circle area ratio would over estimate the amount of light reaching the pixel. Therefore, the proportion of the light reaching the pixel is now given by the ratio of the *intersection* area $A_i$ to the blur circle area $A_{blur}$ and eq. 5.18 can be adjusted accordingly. This gives eq. 5.24 as the new expression for the weighting factor.

$$w_{qpv} = G_{qpv}\frac{A_i}{A_{blur}} \tag{5.24}$$

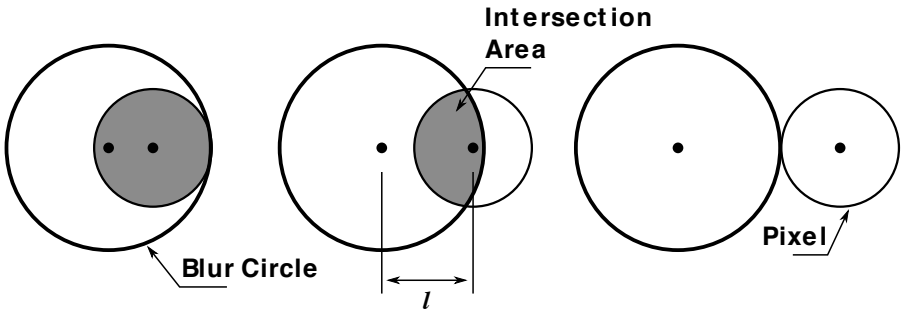



Figure 5.18: Diagram illustrating different intersection areas formed between a circular pixel and a blur circle whose centres are separated by a distance $l$.

Equation 5.24 however, requires that an expression for the intersection area be derived. We begin by defining three cases as illustrated in fig. 5.18. For the first case, the pixel is just wholly contained within the blur circle, and $A_i$ is equal to $A_p$; in the second case, the blur circle only intersects a portion of the pixel; in the final case the blur circle and pixel have, only just, no intersection. The first and final cases are the limiting cases, and between these the intersection area is considered to vary linearly with the separation of the two circle centres $l$. In the first limiting case, $l$ is equal to difference between the blur circle and pixel radii $r_{blur} - r_p$; in the final limiting case it is their sum. The intersection area $A_i$ is thereby given by eq. 5.25, which can be further simplified when units of pixels are used and $A_p$ is unity.

$$\frac{A_i}{A_p} = \begin{cases} 1 & (l <= r_{blur} - r_p) \\ \frac{r_{blur}+r_p-l}{2r_p} & (r_{blur} - r_p < l < r_{blur} + r_p) \\ 0 & (l >= r_{blur} + r_p) \end{cases} \tag{5.25}$$

This linear approximation has been used for simplicity because the exact equation, seen in eq. 5.26, involves a number of operations that would require a larger amount of computer processing time. This equation is plotted alongside the linear approximation in fig. 5.19 for the case where the blur radius is four times the pixel radius and is seen to be a reasonable representation of the exact equation.

$$\begin{aligned} A_i =& r_{blur}^2 cos^{-1}\left(\frac{d^2 + r_{blur}^2 - r_p^2}{2dr_{blur}}\right) + r_p^2 cos^{-1}\left(\frac{d^2 + r_p^2 - r_{blur}^2}{2dr_p}\right) \\ & - \frac{1}{2}\sqrt{(-d + r_{blur} + r_p)(d + r_{blur} - r_p)(d - r_{blur} + r_p)(d + r_{blur} + r_p)} \end{aligned} \tag{5.26}$$

To summarise: a new optical projection geometry and weighting scheme has been defined using ray-tracing. These optical projections can be used when the depth-of-field of the camera is limited and the previous assumption of parallel cylindrical rays is no

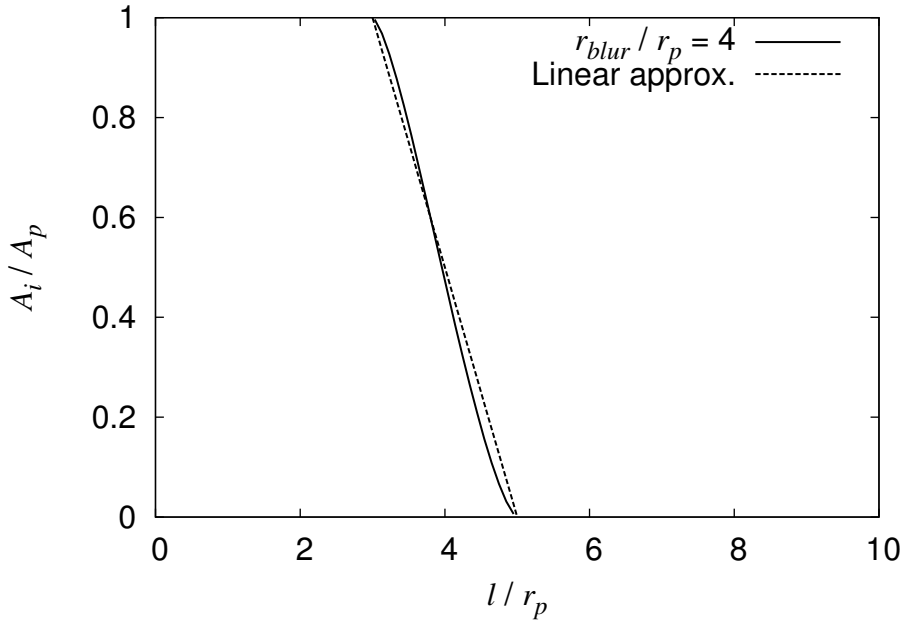



Figure 5.19: Comparison of the exact equation for the intersection area and the linear approximation used. The intersection area $A_i$ has been normalised by the pixel area and the separation $l$ of the centres of the two circles is shown as a multiple of $r_p$: the pixel radius.

longer valid. The value of $w_{qpv}$ is derived from the fractional volume of the intersection between the voxel, $v$, and a double cone, multiplied by the ratio of the blur circle area to the intersection area of the pixel and the blur circle. Walsh et al. [156] have considered a double cone projection geometry, but not applied to 3D CT and they do not apply a weighting based on the blur circle and pixel intersection ratio.

#### 5.4.1.3 Notes on the Computer Implementation

The determination of $G_{qpv}$ for the double cone geometry has been implemented using a small amendment to the approach seen in section 5.3. Here the centre of a voxel $v$, is tested using a variable radius derived from the internal angle of the double cone rather than a prescribed constant value as in the case of the cylindrical projections. The axial component of the voxel location vector V is obtained from its dot product with the double cone's direction vector U. This is used to determine the cone radius at that axial point, and compared to the radial component of V to test for intersection.

The radial component of V in metres is also used to approximate the magnitude of $l$ in eq. 5.25 by scaling by the transverse magnification $M_t$ given by eq. 5.27. (Note that the negative sign represents simply the inversion of the image and can be ignored). The error in this approximation is small when the object domain is small in comparison to $S_o$, as

$S_2 \approx S_i$, but should be reconsidered if this is not the case. The remaining value required to determine $D_{blur}$ is $S_1$ and is given simply by the addition of the magnitude of V to $S_o$.

$$M_t = -\frac{f}{f - S_o} \tag{5.27}$$

### 5.4.2 Non-Parallel Projections

As seen in fig. 5.15 the orientation of a double cone projection, with respect to the optical axis, can be represented by its central ray. The angle formed between the axis and the central ray increases as the central point of the double cone moves away from the optical axis along the object plane. Therefore, as a view consists of a range of projections spread across the object plane, each will form a different angle to the optical axis and they will be non-parallel, resulting in the familiar perspective distortion. When the object distance is large in comparison to the object diameter the variation in the angles is small and so the distortion is small and the projections can be considered parallel; however, these angular variations need to be considered where reduced values of $S_o$ are used, as is likely when depth-of-field is limited.

The angles that projections make with the optical axis depend only on the object's diameter $D_{obj}$ and the distance $S_o$ and are independent of the depth-of-field effect (though they are likely to coincide). Therefore this analysis can also be applied to cylindrical projections where appropriate; for example, perspective distortion can be seen with a pin-hole camera where no depth-of-field effect is possible.

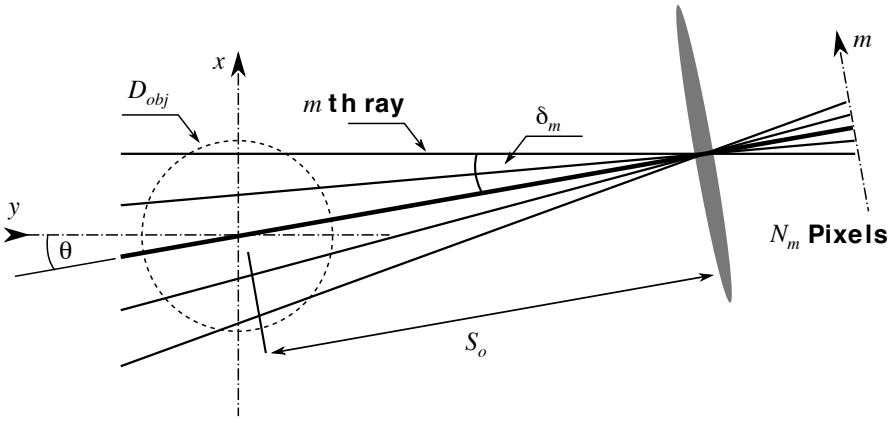



Figure 5.20: Illustration of the non-parallel light rays.

This angular variation within a view is illustrated for the $XZ$ plane in fig. 5.20 where projections are represented by their central rays. The angle of the $m^{\text{th}}$ ray is $\delta_m$ and is

taken relative to the central ray belonging to the projection located in the middle of the view (the thick line), and from which the views angles $\theta$ and $\gamma$ are also measured. As the view should extend to the width of the object, from fig. 5.20 we obtain eq. 5.28 as our expression for $\delta$ (where $\eta$ is the resolution of the view).

$$\delta_m = \tan^{-1}\left(\frac{\eta(m - N_m/2)}{S_o}\right) \tag{5.28}$$

$$\eta = \frac{D_{obj}}{N_m} \tag{5.29}$$

Assuming that the resolution in $Y$ is that of $X$, then the angle in $YZ$ is given by

$$\delta_n = \tan^{-1}\left(\frac{\eta(n - N_n/2)}{S_o}\right) \tag{5.30}$$

The ART algorithm from its beginnings has been capable of non-parallel projections and to account for these now is simply a case of adjusting the direction vector U of each projection by $\delta_m$ and $\delta_n$. The projection geometries in this case could also be altered so that they *fan out* or increase in radius towards the far side of the object domain. However, unless the object domain is very large, or the object distance is very small this would not result in a significant difference.

### 5.4.3 Testing Optical Projections using the Matrix Burner

To test the new optical projections the experiment of section 4.4 is repeated, but this time using the widest possible aperture setting of the camera to minimise the depth-of-field. Reconstructions are performed using four projection types. Reconstructions using parallel and non-parallel cylindrical projections are performed to assess the effect of non-parallel projections on perspective distortion. This reconstruction also provides a control for a reconstruction using optical projections. Reconstructions are also performed using only the geometry of the double cone, without weighting, to assess the importance of the blur intersection ratio $A_i/A_{blur}$. The optical and double cone projections both use non-parallel rays.

#### 5.4.3.1 Experimental Set-Up

The experimental set-up used is the same as that of section 4.4, with two differences. Previously the aperture value was chosen to ensure sufficient depth-of-field; however, this time we are seeking to minimise the depth-of-field and so use a much smaller aperture of $f/3.5$. The exposure time was accordingly reduced to 1250 $\mu$s to avoid over exposure. For these measurements the camera was also positioned on a 10 mm thick flat aluminium surface to improve the consistency of the camera $z$ location and to reduce the rotation of

the camera about its optical axis. This also made the repositioning of the camera easier and less prone to error. The camera was therefore raised above the burner by approximately 10 mm, and this was accounted for by setting $\gamma$ appropriately. To emphasise the effect of an insufficient depth-of-field a large number of angles were used; thereby minimising artefacts caused by limited numbers of views. A total of 48 views were taken with an angular spacing of $3.7 \pm 0.5^{\circ}$.

The camera employs a RGGB filter (see section 4.4) and only the blue component of the measured images were used giving an unfiltered view resolution of 45.2 $\mu$m. An object diameter, $D_{obj}$, of 22 mm and a domain height of 8 mm enclosed all the flames of the Matrix Burner and results in a pixel width in the raw unfiltered views of $N_{m_{raw}} = 487$ and $N_{n_{raw}} = 177$. Using the full resolution of these views was not possible due to the computer memory requirements for this number of views. The views were filtered and downsampled[2] to $N_m = 124$ and $N_n = 45$ giving a view resolution, $\eta$, of 177 $\mu$m. By filtering, the acceptable circle of confusion has effectively increased; thereby increasing the depth-of-field—although this remains less than the object width as will be shown. The increase in the depth-of-field, as a result of filtering, is evident from fig. 5.21 where only low levels of blurring are seen. However, these images still appear noticeably less sharp than those of the earlier proof-of-concept study, presented in section 4.4.

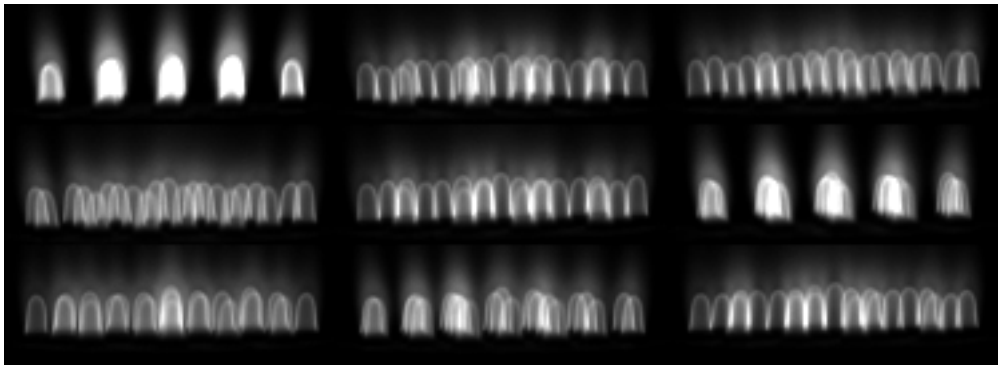


Figure 5.21: Example views taken with insufficient depth-of-field. Shown (from top left) are views: 1, 6, 11, 16, 21, 27, 33, 39 and 45.

#### 5.4.3.2 Estimating Projection Parameters

The optical and double cone projections are dependent on the parameters $f$ and $S_o$ and values for these have to be obtained. The focal length is reported by the camera itself to be 50.1 mm. However, $S_o$ has to be estimated because the double cone projections are based on a single thin lens, and the lens of the camera used is actually made up of a series of lenses (a literal black-box). Selecting a single point for the origin of the lens

[2]The filtered views still required over 1Gb of memory.

is therefore difficult; however, the approximate location of the CCD plane is known and coincides with the end of the aluminium arm at 300 mm. If we are consistent with the earlier model and assume that the actual lens is replaced by a single symmetric lens of the same focal length, then from the thin lens formula, eq. 5.16, the effective value of $S_o$ can be estimated. This assumption results in values of 236 mm for $S_o$ and 64 mm for $S_i$, and places the lens centre at approximately half the length of the actual barrel, which seems reasonable.

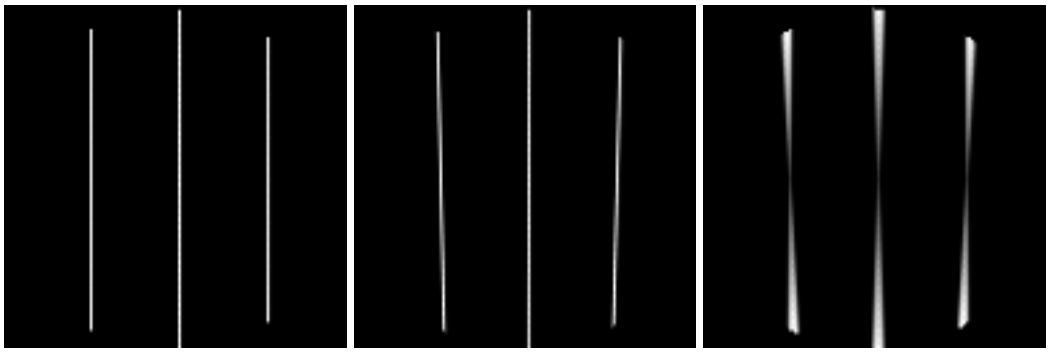
Using this value of $S_o$ we can then determine the appropriate value of $\gamma$. From a calibration image the cameras are centred 11mm above the burner and so $\gamma$ is equal to 2.67°. To determine the blur circle/pixel intersection the camera pixel diameter needs to be known. This can be determined from the actual camera pixel size or, because we select $N_m = N_i = N_j$, from the magnification of $\eta$ by eq. 5.27. However, even accounting for downsampling, the camera pixel size differs by 40% from that predicted by eq. 5.27. This is not unexpected as compound lenses of this kind can be quite complex and suggests that the image side of the actual lens is not that well represented by a single symmetric lens. The value of 47.8 $\mu$m determined from eq. 5.27 is used to be consistent with the single lens model. (The actual pixel size has however, been found to be consistent with the values determined using eq. 5.27 for the PicSight cameras used in later work.)

These estimated parameters result in a depth-of-field of approximately 6 mm—determined by equating $D_{blur}$ in eq. 5.27 to the estimated pixel size. This depth-of-field is significantly larger than desired but is still less than a third of the object size.

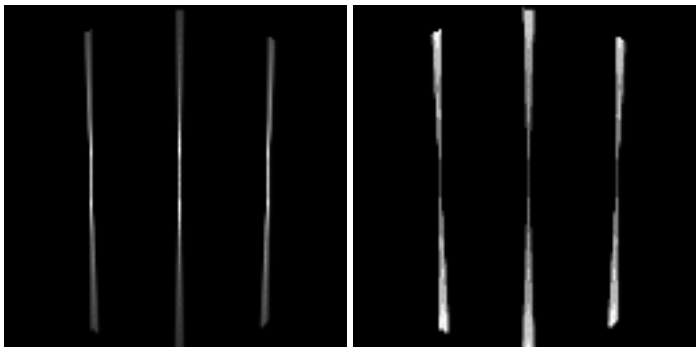
Examples of the four projection types used are shown in fig. 5.22. The first point to note is that the optical projection is markedly thinner than the double cone projection alone indicating a strong influence of the blur weighting. Also observable from this figure, is the similarity between the optical projection and the non-parallel cylindrical projection, which indicates that, for the optical parameters used, the depth-of-field effect is moderate. This figure also highlights the consistency of the optical projections whereby, as $S_o$ or the aperture increases, the double cone projection will converge to the cylindrical one. Nonetheless, these non-parallel cylindrical and optical projections do differ in their spatial extent as illustrated by the logarithm of the optical projection; however, this is offset by the blur intersection weighting that reduces $w_{qpv}$ away from the central ray of the optical projection.

#### 5.4.3.3 Results and Discussion

Reconstructions were performed for each projection type using a value for $\Delta_c$ of $1 \times 10^{-}3$. All the reconstructions converged after just 6 or 7 iterations, except for the unweighted double cone projection which required 19 iterations. A low number of iterations can be



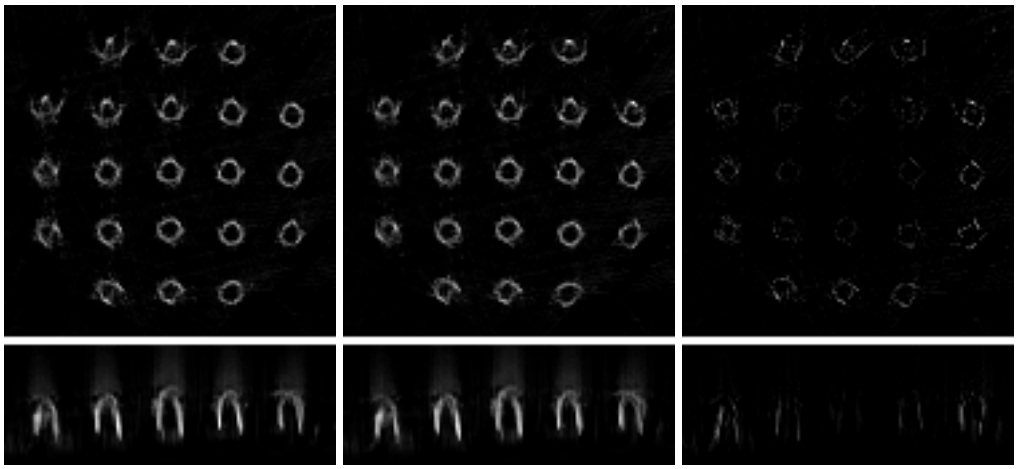
(a) Cylindrical: Parallel (b) Cylindrical: Non-Parallel (c) Double cone only



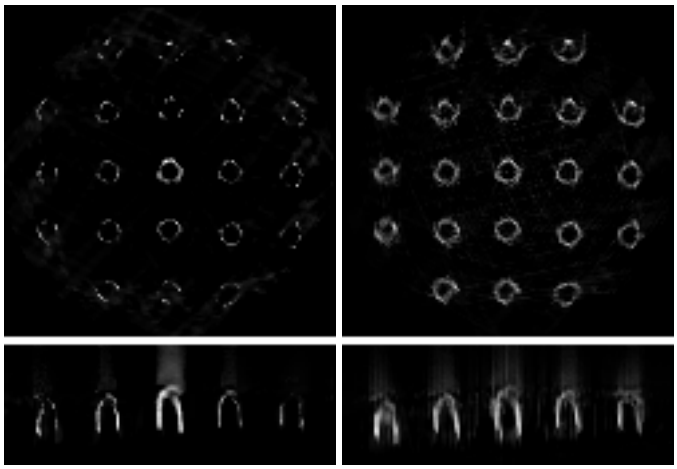
(d) Optical (e) Log. optical

Figure 5.22: Shown are examples of each of the four projection types used. Also shown is the logarithm of the optical projection (e). For each example three projections, $m = 32$, $m = 64$ and $m = 96$, from view 1 ($\theta = 0$) have been composited together.

expected due to the large number of angles used. Sections of each reconstruction are shown in fig. 5.23 and volume renderings are shown in figs. 5.24 and 5.25.



(a) Cylindrical: Parallel (b) Cylindrical: Non-Parallel (c) Difference: (b) - (a)



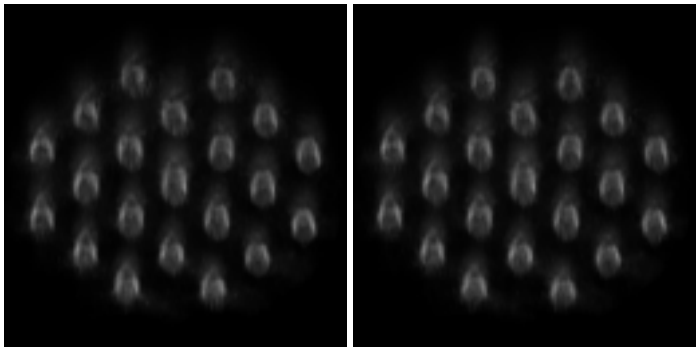
(d) Double cone only ($\times 4.2$) (e) Optical

Figure 5.23: Sections from the reconstructions of the Matrix Burner using the four different projections (a,b,d,e). Also shown in (c) is the difference of the sections in (a) and (b). The sections are taken from the *xy* plane at intercept $k = 29$ (upper) and in the *xz* plane at $j = 61$. The intensities in (d) have been scaled by 4.2 for clarity.

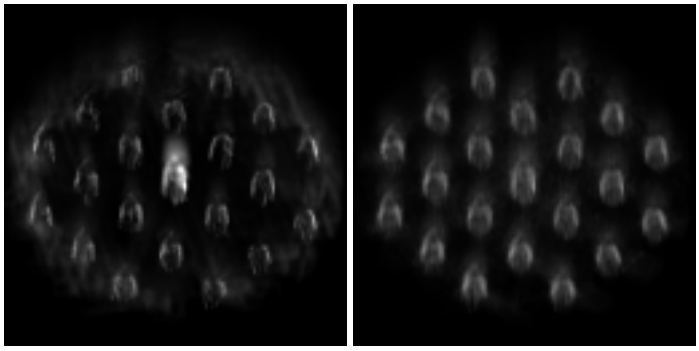
From these figures we can see that both types of cylindrical projections have produced reasonable reconstructions of the Matrix Burner, indicating that the depth-of-field effect is not strong in this case. However, the *difference* images show that the use of non-parallel rays has had a significant effect on the outer edges of the reconstruction, as can be expected, since this is where the greatest divergence from parallel would occur.

It is clear that there has been some improvement of the reconstruction from using non-parallel projections, as a significant reduction in the deformation of the outer flames located on the left of fig. 5.23 is seen, as compared to the reconstruction using parallel

cylindrical projections. This suggests that such deformation can be therefore be attributed to the perspective distortion effect. The occurrence of increased deformity on the image-left (negative $x$) side of the reconstructions is consistent with this, because, for most views, this side is the *far* side of the object domain as a result of the views being taken over 180°. The flames on this far side of the burner are slightly further away and so will be less resolved in the views; thereby explaining to the blurring and deformation seen when perspective distortion is not accounted for. The reduction in the perspective distortion from using non-parallel rays is clear when the images of fig. 5.23 and fig. 5.25 are viewed in quick sequence; the flame spacing is seen to decrease on the left side but increase on the right side, resulting in a generally more orthogonal flame distribution.



(a) Cylindrical: Parallel (b) Cylindrical: Non-Parallel

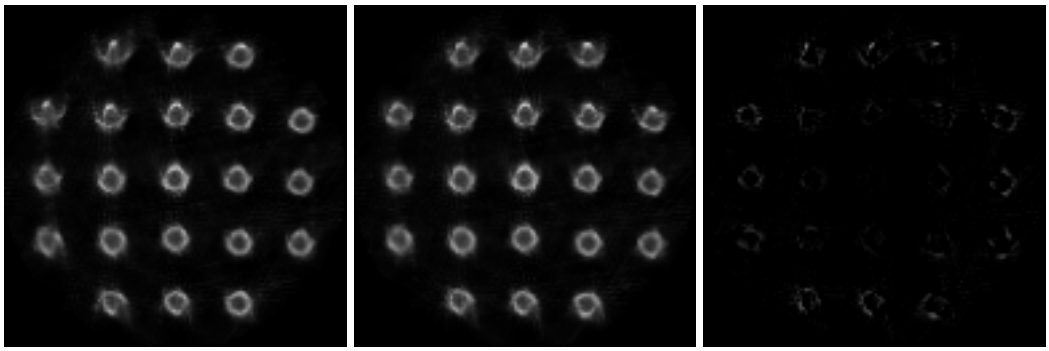


(c) Double cone only ($\times$7.3) (d) Optical

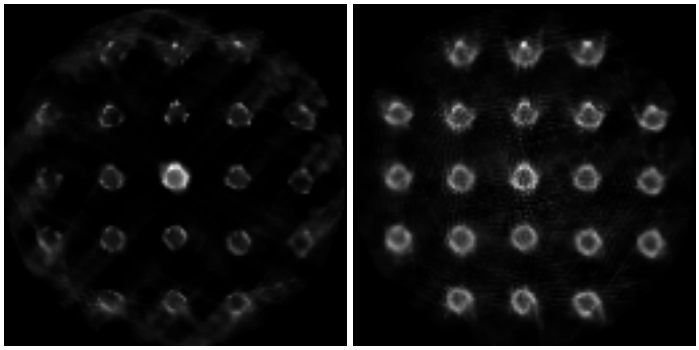
Figure 5.24: Volume renderings of the Matrix Burner reconstructions for each projection type. The viewer is at an angle of 53° to the $xy$ plane. The intensities in (c) have been scaled by 7.3 for clarity.

Although the reconstructions are generally improved by using non-parallel projections, these deformation artefacts do still exist and have actually increased slightly in the outer flames on the near (right) side. The marginal improvement at the near side when using parallel projections, suggests the object is at a greater distance than estimated (for the parallel case $S_o = \infty$). The deformation of the outer flames could also be affected, to some degree, by the unintended rotation of the views about their optical axis (although a

flat aluminium plate was used to minimised rotation). Such rotation would be greatest at the edges of the object domain and would mean that one view locates a flame higher than another view—as appears to occur in the *xz* sections in fig. 5.23.



(a) Cylindrical: Parallel (b) Cylindrical: Non-Parallel (c) [Difference: (b) - (a)



(d) Double cone only ($\times$3.9) (e) Optical

Figure 5.25: Volume renderings of the Matrix Burner reconstructions for each projection type. Also shown in (c) is the difference of the renderings (a) and (b). The viewer is perpendicular to the *xy* plane. The intensities in (d) have been scaled by 3.9 for clarity.

The optical projections have been used to successfully reconstruct the Matrix Burner and have produced a similar quality reconstruction to that using non-parallel cylindrical projections. This is somewhat expected due to the similarity of these projections, as seen earlier in fig. 5.22, which result from the optical parameters of this set-up. The optical projections produce a significantly better reconstruction however, than the unweighted double cone projections. The reconstructions for the unweighted double cone projections appear to significantly over emphasize the centre of the reconstruction and produce much greater blurring and deformation of the outer flames. The poor reconstruction performance of the unweighted double cone projections demonstrates the importance of weighting by the blur intersection area. This also illustrates the wider point that it is the *portion* of the emitted light, from a point within the projection geometry, that reaches the sensor which is important; not the projection geometry itself.

Although the test case has not proved as challenging as hoped, it has demonstrated that

the optical projections can be used to reconstruct the Matrix Burner to the same standard as the non-parallel cylindrical projections. This promotes confidence in the single lens model used, which has been tested here with a relatively complex, compound lens of unknown parameters. For the simpler and better characterised lenses that will be used in the future, the model can be expected to perform well.

### 5.4.4 Conclusion

A new projection type, termed an optical projection, has been derived by tracing rays from an image plane through a simple single symmetric lens. The new projection incorporates a weighting determined from the intersectional area of the blur circle and the image plane pixel. This projection type has been used to successfully reconstruct the Matrix Burner from experimentally measured views, with a depth-of-field deliberately less than a third of the object domain. The reconstruction was of a similar quality to those performed using non-parallel cylindrical views, indicating that where the depth-of-field is only moderately limited the cylindrical views remain a good representation.

The optical parameters of this set-up lead to strong similarity in the optical and non-parallel Cylindrical projections—explaining the correspondence of the reconstructions and indicating that this is a relatively easy test for the optical projections. Ideally measurements could have been taken with a yet smaller depth-of-field; however, the maximum aperture was already being used and memory limitations prevented the use of unfiltered views. The moderate depth-of-field did however, have the benefit of reducing the sensitivity of the optical projections to the value of $S_o$, which had to be estimated, somewhat unsatisfactorily, for the consumer camera used. From inspection of the reconstructions it appears that $S_o$ was under-estimated, and with a more limited depth-of-field this value would need to be more precisely known. The success of this test when using a large complex lens system indicates the general applicability of the single lens model, which will only be more valid for simpler lenses.

A modification that permits non-parallel projections to be used has been presented and tested using the Matrix Burner. The use of non-parallel rays has been seen to reduce the effect of perspective distortion in the reconstructions, and suggests that the deformation of the outer flames can be partly attributed to the use of parallel rays. The non-parallel rays did not completely eliminate the perspective distortion however, suggesting that more precise values for $\delta_m$, would produce better results.

By using these optical modifications it should be possible to position the cameras closer to objects and use wider apertures to maximise the measured signals, and in turn maximise the achievable temporal resolution of the CTC sensor.

# Chapter 6

# Experimental Results and Discussion

In this chapter we present and discuss results from several experimental investigations where the CTC sensor has been implemented, as ultimately intended, with multiple simultaneous views. This means time resolved reconstructions are achieved, unlike the earlier proof of concept test which had no time resolution. These results describe the first implementation of the CTC sensor with simultaneous views and using cameras for which large deployments are possible. Much attention is therefore given throughout this chapter to the practical aspects of the sensor deployment, in particular the camera registration.

The selection of the cameras is of critical importance and this chapter begins by defining a set of requirements that the camera selection was based upon. An industrial Machine Vision class camera that meets these specifications, and costs just £800, is presented and its noise performance tested. Sufficient funding was available to purchase 5 of these cameras and they are then used, along with a mirror system, in an implementation of the CTC sensor that provides 10 simultaneous co-planar views. This implementation is then used to take measurements of the Matrix Burner with time resolutions as low as 62 $\mu$s.

For successful CTC, knowing the precise view locations is of critical importance, and a Computer Vision (CV) approach to 3D camera registration is investigated, which scales to a large number of cameras, as well as providing the optical parameters necessary for optical projections to be used. The improved camera registration is then used to reconstruct the Matrix Burner, but this time operating in a new impinging configuration. Finally, to further demonstrate the applicability of CTC to turbulent flames, the 10 view set-up is used to take first measurements of a premixed Turbulent Opposed Jet flame.

## 6.1 Camera Selection

For time resolved results, measurements need to be performed simultaneously and multiple cameras are required. It is therefore critical to the success of a CTC sensor that

sufficient cameras are available to provide enough simultaneous views for the desired spatial resolution. As equipment cost is the main limit on the number of views, the use of commodity hardware, rather than low volume *scientific* grade cameras, is considered to allow larger numbers of views and high spatial resolutions. Cost, therefore, along with temporal resolution, are the main criteria on which equipment options are assessed. In this section we present the selection criteria and discuss how different types of camera meet these criteria. Following this, the details of the chosen camera model are presented and its noise performance is tested.

### 6.1.1 Selection Criteria

Generally, the the temporal resolution achievable by a camera determines its cost, and so the required specifications of the cameras are largely a result of the temporal resolution deemed necessary. For this sensor we wish to provide measurements of turbulent flames for comparison to Large Eddy Simulations. The turbulent flames typically considered (of which the TOJ and Flame D seen earlier are good examples) have characteristic flow speeds in the range of 3 – 100 m/s. LES of these flames typically have resolutions of the order of 1 mm, and so, for *instantaneous* measurement, temporal resolutions of 300 – 10 $\mu$s would be necessary. We set a desired temporal resolution of 100 $\mu$s, which covers most of this range and would still provide reasonable resolution for very high speed configurations.

The need for sufficient time resolution can be expanded to a more detailed specification. The camera must have a sufficiently sensitive, and low noise sensor, to produce measurements with usable Signal-to-Noise Ratios (SNR) at exposure times below 100 $\mu$s. The camera must also have a sufficient shutter, mechanical or electronic. The ability to synchronously take measurements with other cameras is essential. The physical size and construction is important to ensure that sufficient cameras can be packaged around the flame of interest. Other desirable specifications are a high dynamic range, low cost auxiliaries such as frame grabbers, and a high measurement repeat rate (frames-per-second) to allow transient phenomena to be tracked. The resolution of the sensor is of less importance due to the limitations imposed by limited angles and larger pixel sizes would be preferred for benefits of sensitivity. This also affords the use of CCD/CMOS *regions-of-interest* to decrease readout times and so increase the repetition rates.

Scientific grade cameras are available that meet the criteria, however the cost of such cameras prohibits their use in a current CTC sensor. Fortunately, because the ART algorithm is noise tolerant, we can afford lower grade cameras that do not feature expensive cooling for example. A consumer grade camera has been used previously in the proof-of-concept test, but, as detailed, achieving synchronisation would be very difficult, and their

black-box nature presents too many unknowns for their use in this application. Film has been used by Synder and Hesselink [141] and also more recently by Ishino and Ohiwa [73] for their bespoke 40 lens camera. However, film is generally less sensitive than a CCD[1] and suitable non-bespoke film cameras are not readily available; also digital images are more convenient.

Machine Vision cameras are employed in industrial settings in much greater volumes than scientific grade cameras. These *commodity* cameras are available in configurations that meet the time resolution specifications, and are affordable enough that 25 or more cameras could be deployed with a modest equipment budget. For these reasons Machine Vision cameras are used in this work, and the specific model, A Leutron PicSight P32M, is detailed and tested in the following section.

### 6.1.2 Chosen Camera: PicSight P32M

The camera model selected for use is a Leutron PicSight P32M. These monochrome cameras can be synchronously triggered using an external electronic signal and have a minimum global shutter time of 10 $\mu$s. They feature a 656 by 494 pixel Sony ICX414 interline CCD that allows faster image readout and is also used in higher grade cameras such as those of PCO. This 16 bit 1/2" CCD has large 9.9 $\mu$m pixels and uses micro-lenses to increase the fill factor and so the sensitivity of the CCD. The camera housing has dimensions of just $44.80 \times 31.40 \times 54.81$ mm, which facilitates packaging and mounting of the cameras, and also includes a removable UV filter which cuts off at 400 nm. The cameras are controlled via a High Speed USB interface, removing the need for expensive frame grabbers and additional power supplies, but still allowing up to 127 cameras to be controlled from one PC. Acquiring the images from multiple cameras at the same time with one PC could, however, overload the data-bus/CPU and lead to frame loss. Fortunately the cameras also include 16 MB of on-board memory where up to 50 full frames can be stored to avoid this problem. The cameras can write 62 full frames per second to the camera memory from which the controlling PC can later obtain the images without frame loss.

These cameras cost approximately £800 each, including lenses, making multiple deployment achievable within most equipment budgets. Sufficient funding for 5 cameras was available for the current project, with the intention of extending this by 20 cameras in later work. The PicSight cameras provide two options for synchronisation that both utilise an external electronic signal. The first, uses the signal directly and the width of the signal determines the exposure time; the second, isolates the external signal using

[1] Except perhaps low end consumer cameras.

an opto-coupler and, upon registering the leading edge, generates an internal signal of a predetermined length to start exposure. The first approach allows the minimum delay but requires additional electronics to control the pulse width. The second approach is simpler and requires only a simple on/off signal—though a small delay occurs. In this work the later approach is used for simplicity, because, as the cameras are all the same model, the delay between registering the signal and beginning the exposure is consistent. The PicSight cameras respond to a change in voltage between 5 and 24 V to trigger a frame acquisition and so a DC power source of 12 V was used and controlled using a push button switch. This approach should, in theory, allow a sequence of frames to be captured; however, the vendor supplied software was only capable of capturing a single synchronised frame[2] and a bespoke program using the vendor supplied libraries is required for sequences to be captured. The necessary resources were not available to produce such a program and so in this work single simultaneous frames are used.

#### 6.1.2.1 Testing Signal-to-Noise Ratio

The level of noise in the measured views can greatly affect the quality of the reconstruction. The quality of reconstruction that we are willing to accept ultimately defines our temporal resolution, because the noise level is related to the measurement exposure time. To better understand the noise performance of the cameras, images were taken of a test card consisting of seven bands of different grey levels, from black to white as depicted in fig. 6.1. For each of these bands of constant grey level, both the signal power and the variance of the noise signal can be reliably determined from the mean squared value and the variance within the band. These quantities then yield the SNR as defined in eq. 4.7.

Test images were taken in ambient lighting conditions using the maximum aperture of $f/1.6$ and the exposures listed in table 6.1. To achieve a broader range of signal strengths the gain was necessarily increased for shorter exposures. This would also be done when taking real measurements in order to make best use of the dynamic range of the camera. To assess the effect of the varying gain, test shots of 1000 $\mu$s were repeated for 3 different gain levels.

The SNRs achieved using the 3 different gain values are shown in fig. 6.2 against the normalised RMS values of the seven different signal levels. The RMS of the signal has been normalised by the maximum possible value for the test image, for example 255 for an 8 bit image, to indicate the portion of the dynamic range used. The SNRs are seen to vary significantly, particularly with greater signal levels, and higher SNRs are observed with lower gain.

[2]Due to a (vendor confirmed) bug in the image capture program. The software allows sequences of images to be captured when using the *software* trigger but not when using the external triggers.

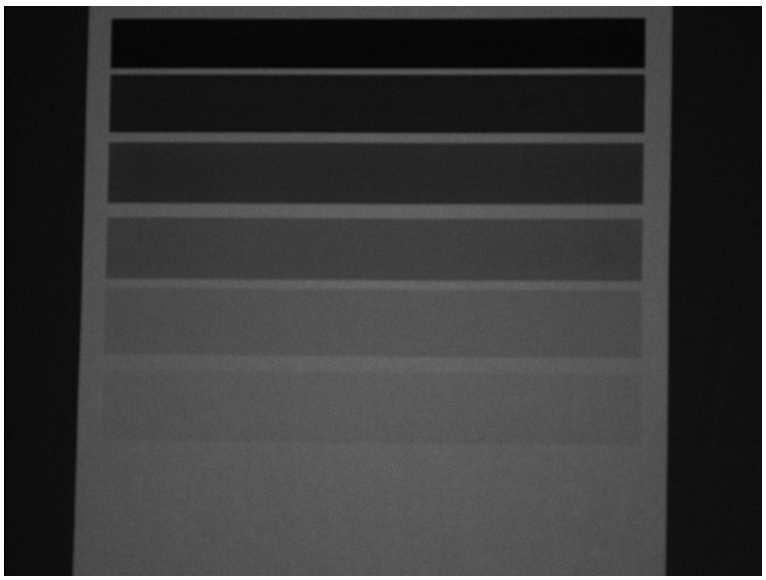


Figure 6.1: An example SNR test shot using an exposure of 500 $\mu$s and a gain of 30 dB.

| **Exposure**[ $\mu$s] | **Gain**[dB] |
|---|---|
| 2000 | 18 |
| 1000 | 18,24,30 |
| 500 | 30 |
| 250 | 33 |

Table 6.1: Test image settings

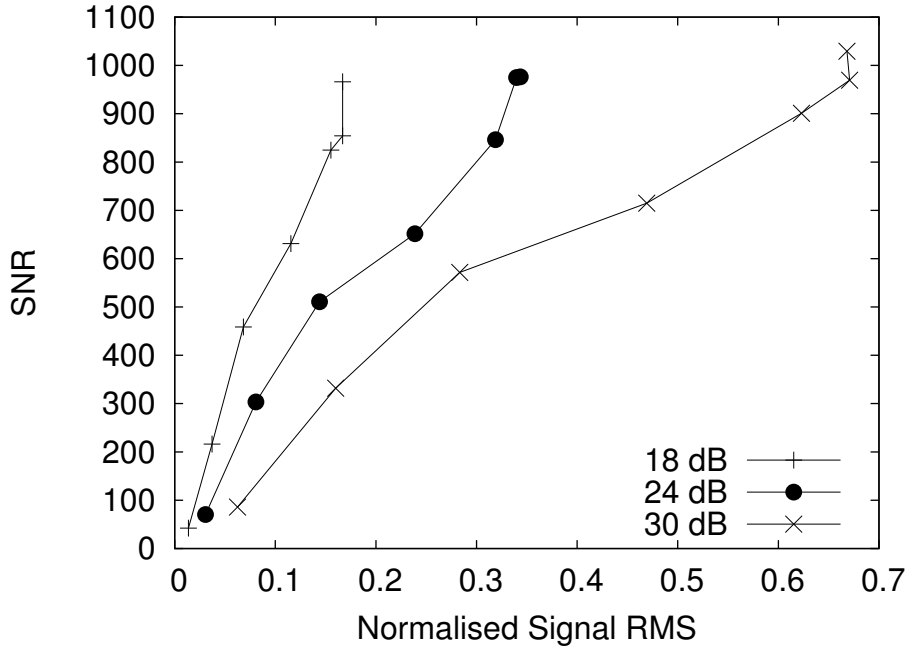



Figure 6.2: The variation of SNR with the normalised signal RMS for an exposure of 1 ms and different gain values.

The noise variance, $\sigma_n^2$, at different signal levels is shown in fig. 6.3 for the range of exposures considered and yields the expected result of greater noise with shorter exposures. Of more interest is the approximately linear dependence of the noise on the signal strength, which makes estimating the SNR achieved in the measured views more difficult. Despite this dependence the SNR still increases with signal level, as demonstrated in fig. 6.4, as the power of the signal increases with its square.

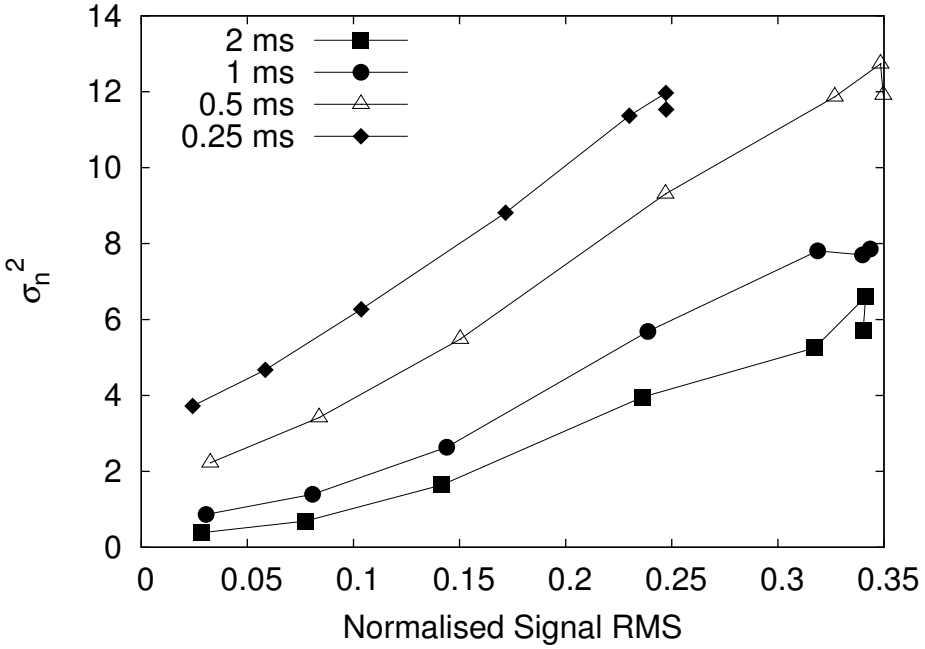



Figure 6.3: The variation of noise with the normalised signal RMS for four different exposure lengths. The noise is taken as the variance of the sample area.

These results show that high SNRs are achievable for the range of exposures considered; for example, an SNR of approximately 300 is achieved with an exposure of 250 $\mu$s and a signal just 25% of the possible dynamic range. The test images shown here used only ambient lighting and have relatively low signals. From our experience so far, images of a flame provide stronger signals, suggesting that lower gain values are possible,and even better SNRs can be expected with actual measurements.

#### 6.1.2.2 Lenses

The lenses chosen were 35 mm $\pm$5% Pentax C33500KP C-mount lenses. These lenses have a maximum aperture of $f/1.6$ and have sufficient resolution for Mega-Pixel cameras. When coupled with the 1/2" sensor of the PicSight cameras the lenses have a horizontal field of view of 10.76° which for the minimum working distance of 40 cm gives an object diameter of 7.7 cm. For the 656 by 494 pixel CCDs this translates to a maximum spatial

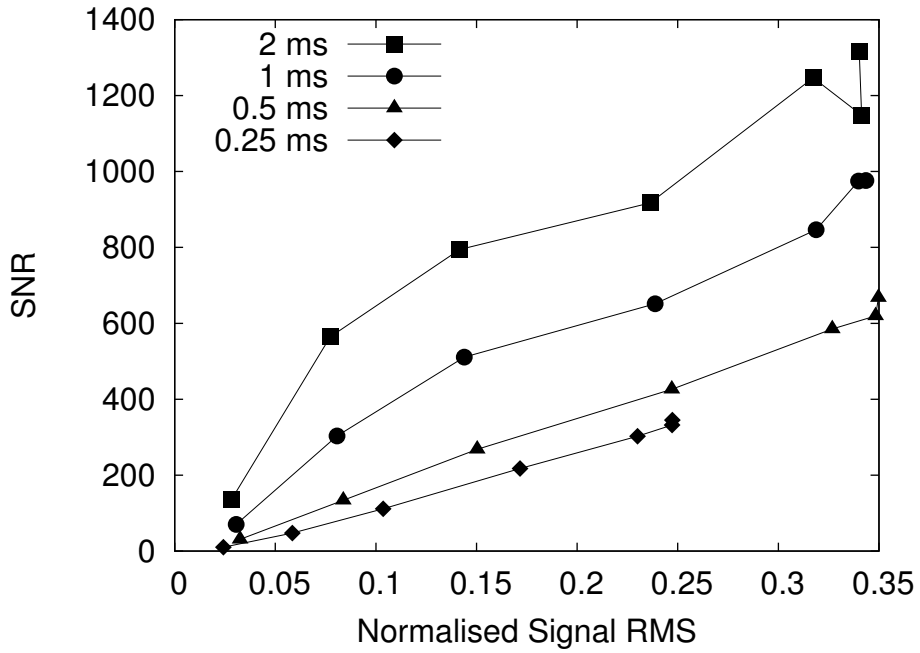



Figure 6.4: The variation of SNR with the normalised signal RMS for four different exposure lengths.

resolution in the views of 120 $\mu$m, which sufficiently exceeds both typical LES resolutions and what is likely achievable with limited angles.

These lenses were chosen to provide a large enough object distance to facilitate the camera packaging while still maintaining sufficient resolution in the images. The 7.7 cm object width is sufficient for many flame configurations, and for larger flames there is sufficient excess resolution to allow greater object distances if necessary. One criticism of these lenses is that they do not have discrete stops for the aperture setting, which means it is difficult to set this consistently across all cameras. However, this has not been an issue in this work because generally the maximum aperture setting has been used. One advantage of the lenses is that they are small and relatively simple, and as such are better approximated by the thin lens equation. For example, a test image of a 22 mm object was taken at a distance of 640 mm and the object was resolved by approximately 115 pixels. With the CCD pixel size of 9.9 $\mu$m this result is in a magnification of 0.05175. Using the single lens magnification equation (eq. 5.27) for the 35 mm lens gives an object distance of 641 mm, just 0.15% different from that measured.

## 6.2 Time Resolved Matrix Burner

This section shows the first results of the CTC sensor, implemented with ten simultaneous views obtained using mirrors and five machine vision cameras. The aim is to demonstrate

that the sensor using such commodity cameras can resolve the non-axisymmetric Matrix Burner chemiluminescence field to a high temporal resolution and so further demonstrate the applicability of the technique to turbulent flames. A resolution of less than 100 $\mu$s is sought based on our earlier requirements for *instantaneous* imaging. The sensitivity of CTC to camera registration error is also investigated using a Phantom generated to have a similar structure to the flame of the Matrix Burner.

We begin by detailing the specific experimental set up used to take the view measurements. The camera registration method is then given and the results of the Phantom study are presented. Finally the resulting reconstructed intensity fields for different exposure lengths are presented and discussed.

### 6.2.1 Specific CTC Set-Up

From the results of section 4.2, we have seen that significant artefacts and errors exist in reconstructions using less than 10 views, though these are likely to be less severe for the spatially limited Matrix Burner. To increase the numbers of views possible with the five available cameras a system of mirrors is employed. This mirror system, illustrated in fig. 6.5, results in each camera obtaining two views 18°apart; with five such systems 10 views are obtained spanning 162°. The mirrors used are general purpose grade, glass mounted, enhanced aluminium mirrors, with dimensions of 80 by 45 by 3 mm thick and with a reported flatness of $2\lambda$ over 25 mm. The angles, locations and widths of the mirrors were determined using ray-tracing to ensure that no vignetting[3] is caused by the finite width mirrors—for any light emitted in the 34 mm diameter object domain. The limiting rays are shown in fig. 6.5 by the dashed lines. A number of variations on this mirror arrangement are possible and this configuration was chosen because it is sufficiently compact to not interfere with adjacent views and has the shortest object distance possible, thereby maximising the view resolution.

The CTC experimental arrangement is shown in fig. 6.6. The burner is located at the centre of a semi-circle of five equally spaced cameras whose lenses are at a radius of 600 mm from the burner centre. The mirrors are located between the burner and cameras, and are fixed to individual mounts machined from a length of *angle* section aluminium. The mirrors, along with the cameras, are mounted to a single semi-annular 10 mm thick aluminium plate. Each camera is connected to the controlling PC and a common electronic external trigger to allow synchronous measurements to be taken. For these measurements, sufficient filters were not available and so the whole visible emission spectrum

[3] Vignetting describes when a *portion* of the light emitted from a point that could reach the sensor, is prevented from doing so by an obstruction or, for example, a finite width mirror. This results in a reduction in the intensity that would not otherwise occur.

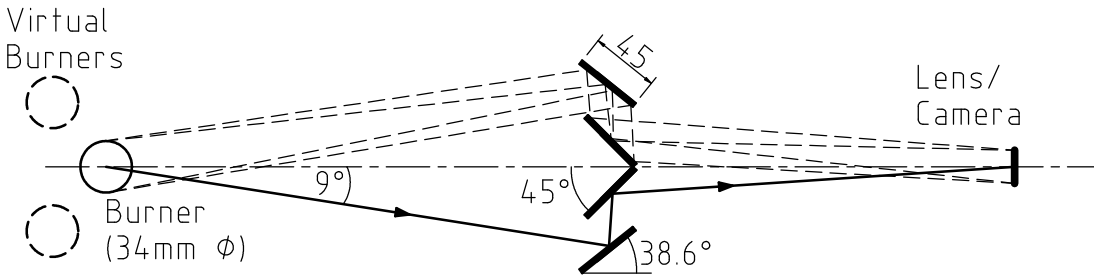



Figure 6.5: A scale diagram of the mirror arrangement, the lens is located 600 mm from the burner centre. The dashed lines show the marginal rays of points on the edge of the domain and define the required size of the mirrors. Also shown are the locations of the two *virtual* burners seen by the camera.

of the flame is imaged and stronger signals can be expected.

Although the burner diameter is 30 mm, the flames are confined to a smaller diameter of approximately 20 mm; an object domain of 22 mm has been used to reflect this. The optical distance from the object to the camera lens is lengthened to 637 mm by the mirrors. At this distance the views sample the object domain approximately 118 times across its width, leading to a view resolution of approximately 190 $\mu$m per pixel. Using eq. 5.23 and setting the acceptable size of $D_{blur}$ to 9 $\mu$m (slightly less than the CCD pixel size) the depth-of-field of the optical set-up is found to be 22 mm. As this is equal to the size of the object, cylindrical projections can be used.

### 6.2.2 Camera Registration

For successful reconstruction it is necessary to know the 3D location and orientation of each view. In this configuration the cameras are at known (designed) locations on the common mounting plate. However, the mirrors' actual locations differed slightly from that designed due to some small adjustment needed to correctly locate the burner in the image frame. A calibration object was therefore employed that allows the true angle $\theta$ of each view to be determined. This is thought the best approach generally for the CTC sensor as it removes the need for very precise location of the cameras and/or mirrors—which with greater numbers of views becomes increasingly impractical. The determination of the viewing angle will, however, be subject to some uncertainty and here we attempt to assess how this deviation could affect the reconstruction by using a 2D Phantom of similar shape to a section of the Matrix Burner. We begin first by describing the desired registration variables and how they are determined.

The cameras and mirrors are all mounted to a single annular plate and are considered

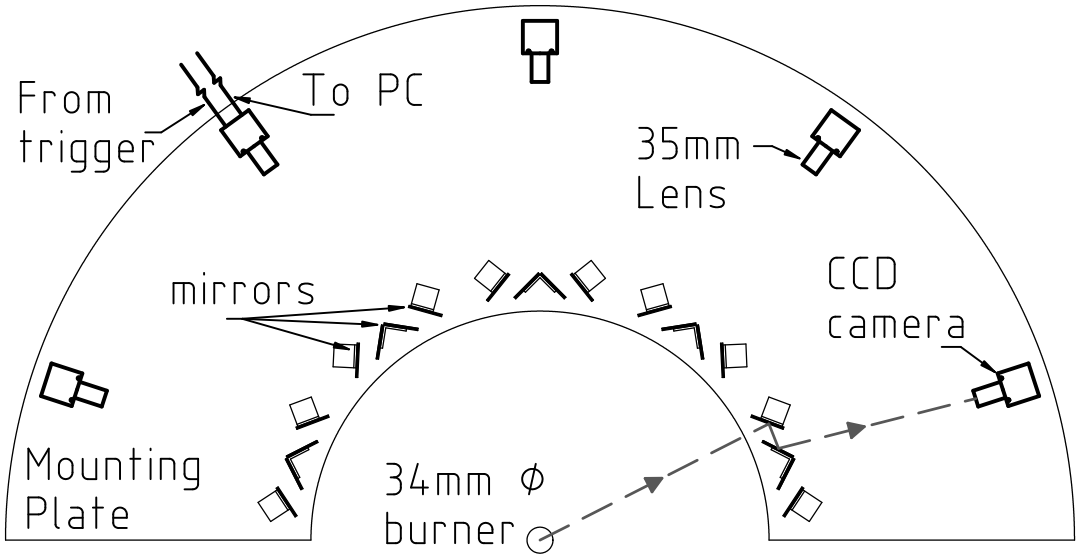



Figure 6.6: The CTC experimental set-up. The five cameras are positioned 36 degrees apart at a radius of 600 mm from the burner (measured to the lens). The mirrors are used to provide two views for each camera, with an example ray path shown by the dashed line. Each camera is connected to the controlling PC and an electronic external trigger source. All mirrors and cameras are mounted to a single 10 mm thick aluminium plate.

co-planar. The out-of-plane rotation at the cameras centre $\gamma_c$ is therefore assumed to be negligible. (Though $\gamma$ is adjusted over a small range to account for non-parallel projections.) As only small adjustments to the angles of the mirrors were necessary the radial location and hence the optical length $S_o$ of each view will not have changed significantly from that designed. Aside from $\theta$, the camera registration process needs to determine just two other quantities: the first is the the $X$ and $Y$ location of the central axis of the burner in each view, which serves as a common reference point; the second is the rotation of the view about the camera's optical axis.

Each view's rotation, angle ($\theta$) and reference point were determined using images of a cylindrical registration object. The object, shown in fig. 6.7, has a diameter of 43.5 mm and was designed to *just fit* over the Matrix Burner so that their central axes were co-located. The roll is determined by the angle between the image $Y$ axis and the vertical edge of the calibration object—found to high precision using a sub-pixel edge detection algorithm [24]. This rotation was removed by an appropriate counter rotation of the image. The central $Y$ reference was taken as the middle of the angular scale and the centre of the domain in $X$ is taken as the average equidistant point from the edges of the of the cylinder. Finally, the view angle $\theta$ is determined from the location of the central reference point on the marked angular scale. The scale has spacing of 10 degrees and linear interpolation was used to determine the values in between the marked points. The determined




Figure 6.7: Example registration image for view 10. The registration object sits atop the burner, so that their central axes are in line.

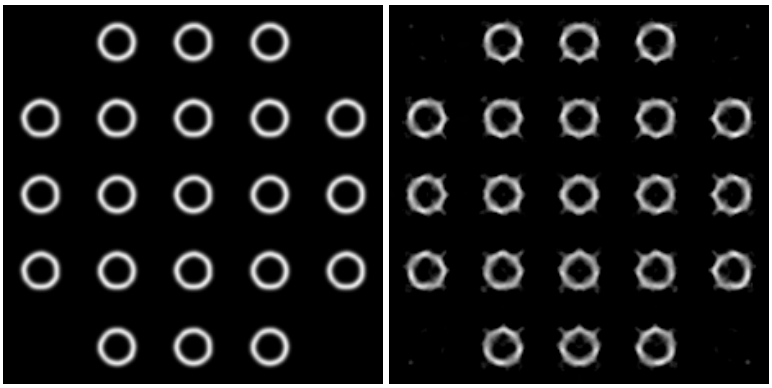
values can then be seen in table 6.2 and vary from the designed, ideally spaced angles by at most 0.55°.

Table 6.2: The calibrated viewing angle for each view shown alongside the ideally spaced (designed) angle. View 1 is used as the reference value.

| **View no.** | $\theta$ **[deg]** | **Ideal [deg]** |
|---|---|---|
| 1 | 0 | 0 |
| 2 | 18.55 | 18 |
| 3 | 36.27 | 36 |
| 4 | 54.41 | 54 |
| 5 | 72.27 | 72 |
| 6 | 89.88 | 90 |
| 7 | 107.80 | 108 |
| 8 | 126.00 | 126 |
| 9 | 144.15 | 144 |
| 10 | 162.38 | 162 |

To keep the precision high, sub-pixel edge detection was again used to locate points on the 10° scale. This implies an uncertainty of less than 0.5°, although manufacturing defects mean the error could be slightly higher. To test the effect of such error in $\theta$, simulated reconstructions were performed using a Phantom consisting of a series of smoothed rings (fig. 6.8) that is representative of a cross section of the Matrix Burner flame. Figure 6.8 includes a reconstruction that contains no error in the viewing angles. This reconstruction shows minimal artefacts and achieves good resolution of the constituent flame fronts in this ten view configuration.

Reconstructions of the Phantom were performed with random amounts of increasing



(a) Phantom (b) Reconstruction

Figure 6.8: The synthetic Phantom alongside a reconstruction with ten views and no error in the viewing angles.

error added to the viewing angles. The error for each view was obtained using a random number generator that was scaled to the desired peak value $\varepsilon_{peak}$. For each peak value four reconstructions were performed to give a better idea of the spread of the error induced in the reconstruction. As the main effect of misalignment will be on the structure—not the values—of the reconstructed flame, we again use the correlation coefficient $r_{xy}$ (eq. 4.3) to quantify the error in the reconstruction.

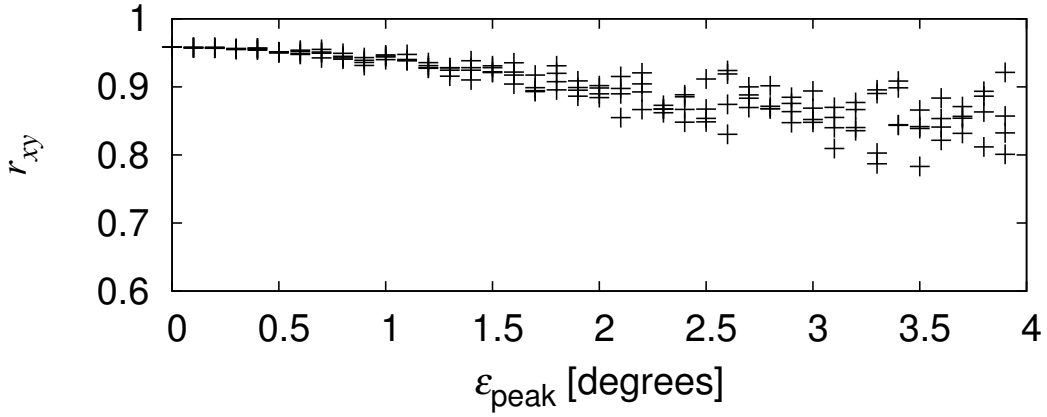
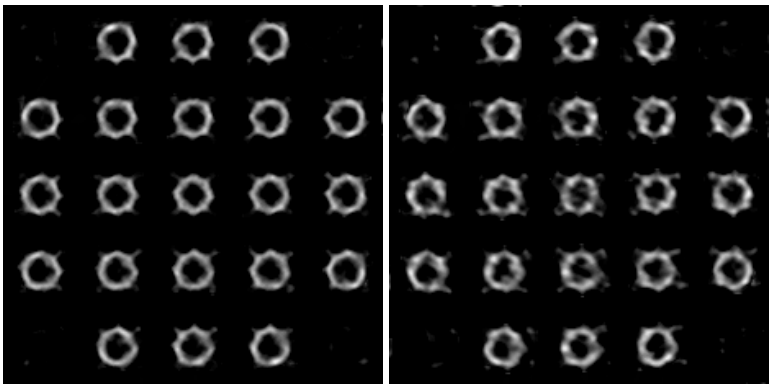



Figure 6.9: The correlation coefficient $r_{xy}$ of the Phantom and reconstructions subjected to increasing random error in the viewing angle.

The results of the test can be seen in fig. 6.9, where an almost negligible change in the correlation coefficient is seen below a peak error of 0.5°, and acceptable coefficients of more than 0.9 occur up to a peak error of approximately 1.5°. Inspection of the reconstructions themselves confirms this result, with significant deformation in the reconstruction only seen after a peak view error of approximately 2.3°, as demonstrated in fig. 6.10. These results indicate that the expected error of ±0.5° in the determined view-

ing angles will not have a significant effect on the accuracy of the reconstructions. One further observation is that the effect of error in the viewing angles is mostly distributed uniformly through the reconstruction, but is perhaps observed most strongly in the centre of the reconstruction—as demonstrated in fig. 6.10.



(a) $\varepsilon_{peak} = 1.0^\circ$ (b) $\varepsilon_{peak} = 2.3^\circ$

Figure 6.10: Reconstructions with 1° (*left*) and 2.3° (*right*) peak error in the view angle.

### 6.2.3 Results

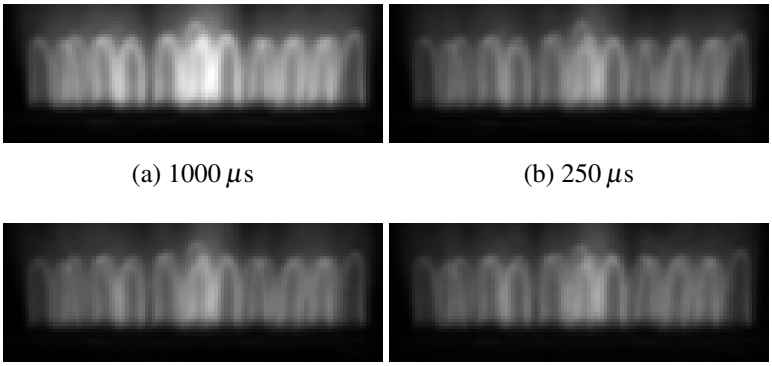
To test the CTC sensor's—and the P32M camera's—capability for high time resolutions, measurements were taken using the four different exposure times shown in table 6.3. Large Eddy Simulations would often have higher time resolutions than used in this test. However, this is usually for the benefit of numerical stability at localised points in the flow field, for example the jet exit of Flame D, and is not usually necessary to capture most features of the flow.

Table 6.3: Experimental Parameters

| **Exposure [$\mu$s]** | **Gain [dB]** | **Aperture [f/#]** | **Approx. SNR [dB]** | **$\Delta_c$ for conv.** |
|---|---|---|---|---|
| 1000 | 6 | 1.6 | 50 | $1.0 \times 10^{-3}$ |
| 250 | 15 | 1.6 | 40 | $1.0 \times 10^{-4}$ |
| 125 | 21 | 1.6 | 35 | $1.0 \times 10^{-5}$ |
| 62 | 27 | 1.6 | 30 | $1.0 \times 10^{-5}$ |

For this test an aperture setting of $f/1.6$ was used to maximise the measured chemiluminescence signal while still providing sufficient depth-of-field as previously detailed. A representative SNR of the views for each exposure setting was determined using eq. 4.7 and is shown in table 6.3. This value was determined from view 4 with the variance of the

noise signal approximated from a *black* region of the view. This approach would over-estimate the SNR as it does not account for the signal dependence of the noise. To counter this over-estimation, the black level variance is increased by a factor of 5. This factor is based on the noise variation seen in fig. 6.3 where an approximate five-fold increase is observed in $\sigma_n^2$, from a black (low) signal to a typical signal value of 0.3. This scaling is conservative given the much lower gain values used with these measurements. The lowest SNR is still found to be relatively high at 1000 (30 dB). Inspection of the views, e.g. fig. 6.11, confirms the high SNRs, as the noise level does not vary discernibly between different exposure times.



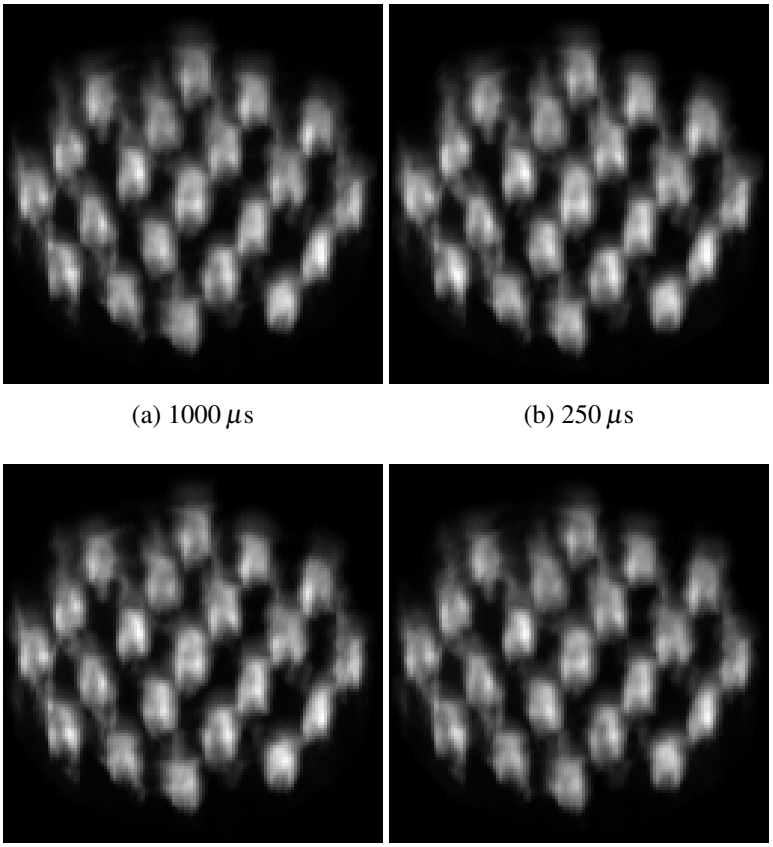
(a) 1000 $\mu$s (b) 250 $\mu$s

(c) 125 $\mu$s (d) 62 $\mu$s

Figure 6.11: The measured view 7 for each exposure used. The aperture is f/1.6. The images have been normalised by the maximum value of (a).

The views were cut down to 115 pixels in the view $X$ direction and 42 pixels in the $Y$ direction, resulting in a view resolution of 190 $\mu$m. The reconstructions were performed using the 3D ART algorithm with non-parallel cylindrical projections on a reconstruction grid of 115 by 115 by 42. The use of non-parallel projections means that adjustments to $\theta$ and $\gamma$ are made for individual projections to account for perspective distortion, though the effect on $\gamma$ is small because the domain is only 8 mm in $Y$. For an exposure of 1000 $\mu$s a value for $\Delta_c$ of $1.0 \times 10^{-3}$ was used; however, for the shorter exposures, and so lower SNRs, more iterations were required (as reported earlier, and in other work [108]) and lower values for $\Delta_c$ were necessary as shown in table 6.3.

Volume renderings of the reconstructions of each exposure can be seen in fig. 6.12. All reconstructions have resolved the individual flames well, as each is seen to be visibly distinct from the others. The reconstruction of the 62 $\mu$s exposure views has resolved the flame as effectively as the 1000 $\mu$s reconstruction, demonstrating that time resolutions of less than 100 $\mu$s are possible with this CTC set-up. The sections in figs. 6.13 and 6.14 do show that the internal cavity in the flames is not as well resolved as the Phantom reconstruction in fig. 6.8, though they are still identifiable. The Phantom however, does appear

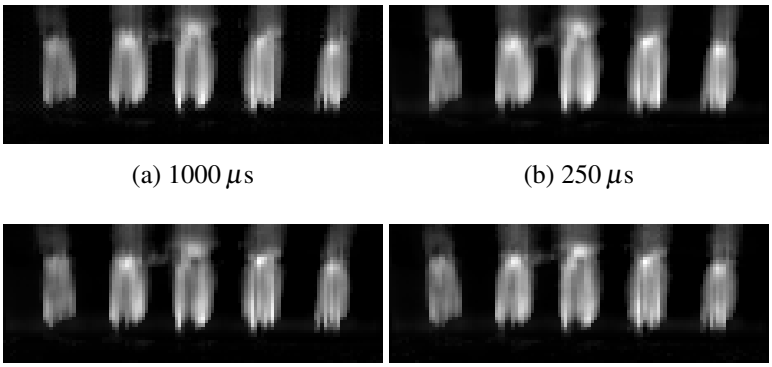


(a) 1000 $\mu$s (b) 250 $\mu$s

(c) 125 $\mu$s (d) 62 $\mu$s

Figure 6.12: Volume renderings of the reconstruction of each exposure. The viewing angle is 53 degrees from horizontal.

to have larger diameter constituent flames than those measured, which could explain some of the resolution difference because a larger cavity would be easier to resolve. From the analysis of section 4.3.2 we can estimate the resolved wavelength to be approximately 2 mm (assuming a radially limited flame as in section 4.4). This value for the wavelength agrees with that observed in fig. 6.14, where the flames are approximately 2 mm wide, and indicates that the Matrix Burner is less resolved than the previous Phantom study implies.



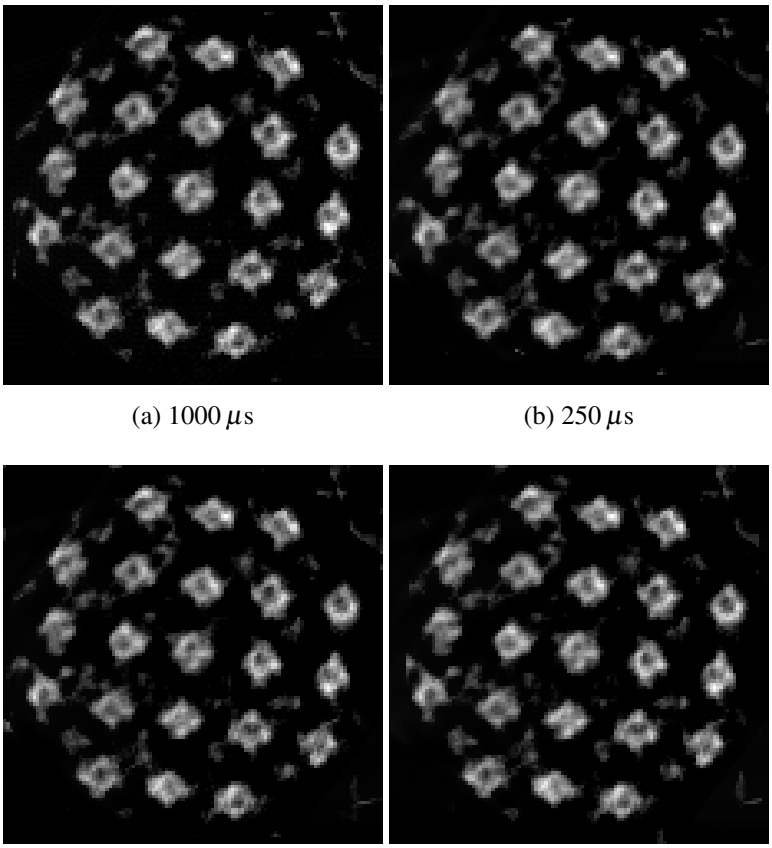
(a) 1000 $\mu$s (b) 250 $\mu$s

(c) 125 $\mu$s (d) 62 $\mu$s

Figure 6.13: Shown is a section of each of the reconstructions taken through the centre (plane normal [1 0.2 0]).

For all exposures the reconstructions are seen to suffer deformation similar to the *cross* shape observed in fig. 6.10. This deformation is more severe in the reconstructions of the actual flame than seen in the Phantom study—even considering the larger radius—suggesting a greater error in $\theta$ than previously determined. However, in the experimental reconstructions, greater deformation is seen in the flames at the edges of the reconstruction, while in the Phantom study the largest deformation was seen in the centre. Moreover, low intensity artefacts are also observed between the constituent flames that did not exist in the Phantom study. This suggests that other factors, in addition to error in $\theta$, are significant. The deformations also appear to be independent of SNR, and so we can conclude that it is some error in the projection geometries—other than error in $\theta$—that has lead to the deformations and artefacts.

The artefacts seen between the flames have also been observed in section 5.3.3, suggesting that the out-of-plane rotation ($\gamma$) of the views may have some effect. Some variation in the $Y$ location of the calibration object was observed between views, due to small errors in the mirror location. However, the sensitivity of the mirror arrangement meant the variation in $\gamma$ of each view was less than $1^\circ$: equivalent to just 2 cells over the reconstruction domain. To consider if such a small error can result in the deformations and artefacts seen in the experimental reconstructions, the reconstruction of the 1000 $\mu$s measurements



(a) 1000 $\mu$s (b) 250 $\mu$s

(c) 125 $\mu$s (d) 62 $\mu$s

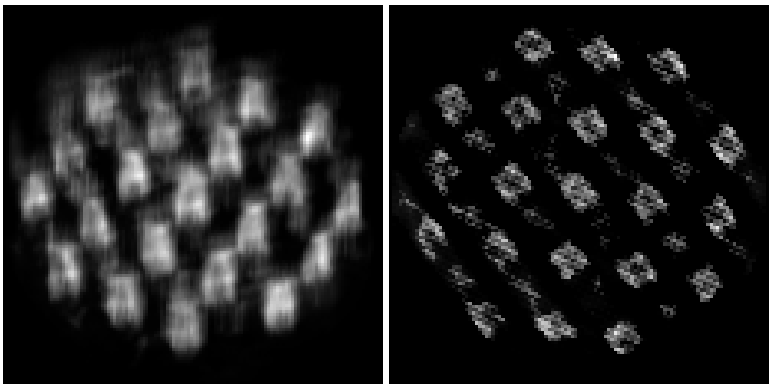
Figure 6.14: Shown is a section of each of the reconstructions at $z = 25$.

was repeated with adjustments made for $\gamma$.

Table 6.4: $\gamma$ estimates for each view

| **View No.** | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 |
|---|---|---|---|---|---|---|---|---|---|---|
| $\gamma$ [deg] | 0.77 | 0.58 | 0.32 | 0.54 | 0.86 | 0.70 | 0.68 | 0.86 | 1.27 | 1.14 |

The calibration object used was not designed to provide values for $\gamma$; however, estimates were obtained by considering the offset, $\Delta Y$, of the reconstruction domain centre from the camera image centre at $N_Y/2$. With the known object distance, $\gamma$ is then given by trigonometry and the resulting values are shown in table 6.4. A volume rendering and section of the reconstruction using these $\gamma$ values is shown in fig. 6.15. The deformations of the flames still exist, suggesting that error in $\gamma$ is not their cause. However, this new reconstruction appears sharper than the reconstruction where $\gamma$ was assumed zero and the artifacts between the flames appear reduced. This is thought to be due to better convergence in this configuration as fewer iterations were required; thereby indicating a more determined system of equations and hence that the projections are more accurately represented.



(a) Volume rendering viewed from 53°above the *xy* plane

(b) Section at $z = 25$

Figure 6.15: A volume rendering and section of the reconstruction of the 1000 $\mu$s exposure views using the $\gamma$ values in table 6.4.

The deformations could be a result of mis-location of the reconstruction central reference points in the views $XY$ frame, although this could be expected to have a uniform effect and not be larger towards the edges of the reconstructions as seen. Another possibility is the effect of lens distortion. The use of mirrors means that the flames are located towards the edges of the camera image (necessary to avoid vignetting) where lens distortion effects are generally greatest, and could be expected to cause greater error at the edges of the reconstructions. Lens distortion is considered further in section 6.3. Finally,

one possibility is that refraction within the flame (caused by the different gas optical densities) could steer the emitted light along a different path to that assumed by the projection geometries. This appears to be a reasonable explanation given that for this Matrix Burner any light usually has to traverse several flame fronts. However, if this was the case we could expect greater deformation in the centre, as the light from here has to cross the most flames. We could also expect less deformation at the sides nearest the cameras, as the light from there has the least flame fronts to cross. We note such errors have been considered by Andersen and Kak [5] who have developed an iterative correction approach using ray-tracing for use in Ultra-Sonic CT.

#### 6.2.3.1 Conclusions

The CTC set-up has been implemented using five machine vision cameras along with mirror systems to provide 10 simultaneous views of a Matrix Burner. Time resolved reconstructions have been performed in full 3D for views measured using 1000, 250, 125 and 62 $\mu$s exposures. No significant difference is observed for reconstructions of different exposure times and all have successfully resolved the individual flames of the Matrix Burner to a wavelength of 2 mm. This is consistent with that predicted for 10 angles using the analysis of section 4.3.2. The similarity in the results is due to the high Signal-to-Noise Ratios achieved for all exposures and demonstrates the machine vision cameras ability to provide high SNR measurements at temporal resolutions less than the target of 100 $\mu$s. Indeed the SNR of 1000 achieved for the shortest exposure of 62 $\mu$s suggests even shorter exposures are possible. The use of oxygen as the oxidiser has certainly contributed to the high SNRs; however the signals are so high that good SNRs could still be expected with air as the oxidiser.

A Phantom study has shown that the precision, $\pm 0.5^{\circ}$, obtained for the in-plane viewing angle $\theta$, results in almost negligible error in the reconstruction. The same study also suggests that resolutions higher than 2 mm are achievable for the Matrix Burner with 10 views. However, the actual reconstructions are clearly not this well resolved, because the internal cavity of the flames is not captured and significant deformation of the flame cross-sections is observed. This difference between Phantom and measured results is thought to be due, in part, to the slightly smaller diameter of the actual Matrix Burner flames. However, the severity of the deformations in the reconstructions using measured views, in addition to the small artefacts observed between the individual flames, and the increased deformity at the edges of the reconstruction, indicates that factors other than random error in $\theta$ may also be significant.

One possible factor investigated was error in the out-of-plane angle $\gamma$. This was initially assumed negligible due to the common mounts, but was tested by repeating the

reconstruction of the 1000 $\mu$s exposure views using estimated values for the $\gamma$. The angles involved were small and were found to have no effect on the deformations, though reconstructions had become sharper and required fewer iterations to converge. Error in the $XY$ location of the reconstruction centre, beam steering, and the effect of lens distortion have been suggested as other sources of error, though it is the author's belief that no single error results in these deformations and that it is likely to be the cumulative effect of several camera positioning errors. For this reason a more comprehensive approach to the 3D registration of the view locations and orientations is desirable.

## 6.3 Improving Camera Registration

For the successful application of CTC, the location and orientation of each camera needs to be known to a reasonable precision. Obtaining sufficient precision with manual measurements is quite challenging even for a single camera, and is increasingly impractical where multiple cameras are used. For this reason a calibration object is preferable, from which the locations and orientations can be obtained using Photogrammetry[4]. This approach was used successfully in section 6.2 to determine two rotations, including $\theta$, and two locations ($X$ and $Y$). However, the object used is specific to that co-planar (2D) case, as well as that burner and object distance. The object was not designed for the more general (3D) case of six degrees of freedom in the camera positioning (three rotations and three translations). To match the flexibility in sensor locations afforded by the 3D ART algorithm, we look to improve on the earlier camera registration techniques, and allow arbitrary 3D camera locations.

The ultimate aim of camera registration is to relate the coordinates of each camera's 2D image to those of the global 3D object domain. Such projective transformations are formalised in the mathematical discipline of *projective geometry*[5]. This formalism is employed to determine camera locations (and more) in Computer Vision (CV) applications, and we consider the techniques used in this discipline to improve our camera registration. Similar approaches are used by Ng and Zhang [114] for example, and also in commercial systems such as those of LaVision GmbH, as reported by Wieneke who used camera self calibration to improve Stereo PIV [164] and tomographic PIV [165]. These CV approaches still require a calibration object, but now, rather than specifically derived relations, the more general rules of projective geometry are applied.

In CV the projective transformation of a camera is defined by its *camera matrix*[6]

[4]Photogrammetry [107] is the measurement of objects from their images.

[5]The reader should not confuse this term with *projection geometry* defined previously.

[6]The is also called the Camera Projection Matrix, though we do not use this term to avoid confusion with the ART projection matrix.

and the process of determining this is termed *camera calibration*. The camera matrix incorporates a *camera model* and the internal parameters of this model, such as focal length, are estimated in the calibration process. With a suitable camera model, these parameters can be used to better define the non-parallel and optical projection geometries seen in section 5.4, and so potentially improve their reconstructions. The use of this CV approach offers further advantages. Firstly, by using automatic feature extraction algorithms, the process can be automated leading to faster calibrations while also reducing the scope for human error. Secondly, projective geometry is conceptually consistent with reconstructing from projections, indicating scope for improved implementations.

In this section we investigate the use of a standard Computer Vision approach to camera registration. Camera registration is performed using this approach for the five camera arrangement of the previous section (6.2) without the mirrors, and then compared to the values obtained using the cylindrical object used previously. The camera matrix also includes a model for lens distortion, and this is used to consider the potential effect of such distortion on reconstructions, as postulated in section 6.2. We begin by introducing the concept of camera calibration and the method used in this section.

### 6.3.1 View Registration using Camera Calibration

We begin here with a brief review of the main theoretical concepts of camera calibration. For greater detail the book of Hartley and Zisserman [64] is recommended.

Camera calibration aims to determine the transformation from a world coordinate system to the coordinates of the camera image plane. This is achieved using three sets of coordinates: world, camera, and image. The 3D world coordinates are transformed by rotation and translation to camera coordinates, which are then projected to the image plane. By choosing a common world coordinate system for calibration—in this case the object domain coordinates—the relative positions of all cameras can be determined.

In this work, the world coordinates correspond to the object coordinates $xyz$ and the camera coordinates are equivalent to the view coordinates $XYZ$. The image plane does not yet have a direct equivalent in this work, and we assign it coordinates $u$ and $v$. The transformation from world to camera coordinates is independent of the camera properties and is defined by *extrinsic* parameters. The projection from camera to image coordinates however, depends on the *intrinsic* parameters of the camera. While at first it may seem that only the extrinsic parameters are needed for our use, this is not the case, because we implicitly measure views in image coordinates and need to consider the intrinsic parameters also.

A camera is parametrised by a model, with different models incorporating different effects. The most common camera model used is the *pinhole* camera model [166], which

accounts for the effect of perspective. An alternative *thin lens* camera model exists, as used by Baba et al. [8], which also considers the focus and iris of the camera. For now, we consider the adequate depth-of-field case and the simpler pinhole model is sufficient (although future work may benefit from the thin lens model). Moreover, the pinhole model has the advantage of having already been implemented in the available camera calibration packages.

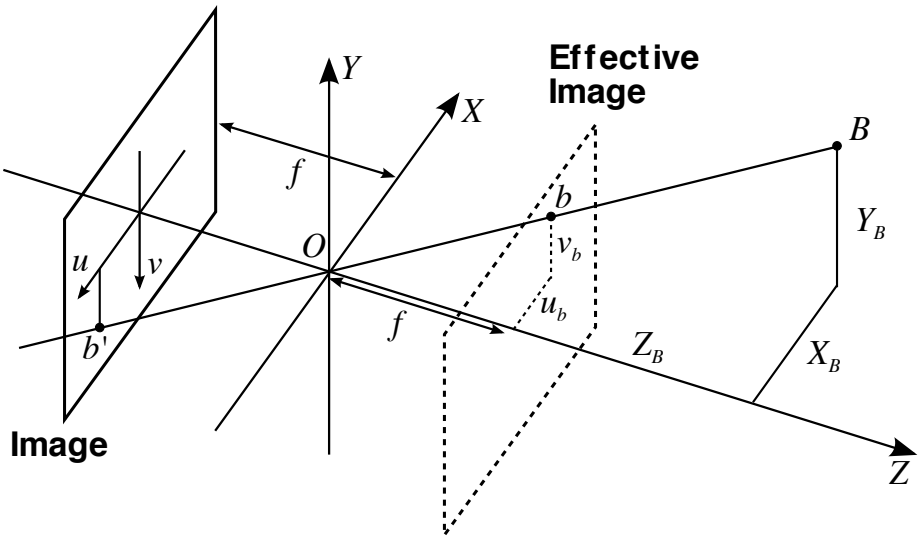



Figure 6.16: Pinhole camera model coordinates.

The pinhole coordinate system is shown in fig. 6.16. The *pinhole* is located at the origin of the view's $(XYZ)$ coordinate system, and the image $uv$ plane is parallel to the view $XY$ plane. From this figure we can derive the projection transformation. A point $B$ located at $(X_B, Y_B, Z_B)$ in the view reference frame is projected onto the image plane at a point $b'$ (or point $b$ on an equivalent image plane, as used in some of the literature) located at $(u_b, v_b)$. The image plane is a distance $f$, the focal length, from the view coordinates' origin, measured along the optical axis ($Z$ axis), and so $u$ and $v$ can be defined as follows in eq. 6.1.

$$\begin{bmatrix} u \\ v \end{bmatrix} = \frac{f}{Z} \begin{bmatrix} X \\ Y \end{bmatrix} \tag{6.1}$$

Homogeneous coordinates can be used to represent a point in projective space, and have several useful properties. Homogeneous coordinates are related to Cartesian coordinates by eq. 6.2 for 3D, where conventionally $A = 1$ but any non-zero number is allowable.

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix}_{euc} \Rightarrow \begin{bmatrix} xA \\ yA \\ zA \\ A \end{bmatrix}_{hom} \tag{6.2}$$

By expressing the *pinhole* projection from eq. 6.1 in homogeneous coordinates we obtain the linear expression in eq. 6.3, where $\sim$ denotes equality of any non-zero multiple and the matrix $\mathbf{C}$ is referred to as the *camera matrix*.

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} \sim \begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} \Rightarrow \mathbf{u} \sim \mathbf{C}\,\mathbf{X} \tag{6.3}$$

#### 6.3.1.1 Calibration Method

The camera calibration is performed using the approach of Zhang [172], as implemented in the popular MATLAB Camera Calibration Toolbox developed by Bouguet [24]. This implementation is based on a pinhole camera but extended as follows. The effect of lens distortion is included by an additional intermediate intrinsic transformation, shown in eq. 6.4.

$$\mathbf{u}_d = \left(1 + kc_1 r^2 + kc_2 r^4 + kc_5 r^6\right) \mathbf{u}_n + \begin{bmatrix} 2kc_3 u_n v_n + kc_4 (r^2 + 2u_n^2) \\ 2kc_3 (r^2 + 2v_n^2) + 2kc_4 u_n v_n \end{bmatrix} \tag{6.4}$$

Here $r$ is the radius of $\mathbf{u}_n$ which is the projected point given by eq. 6.1. The first term represents the radial distortion and the second term the tangential distortion. The coefficient $kc_5$ is only necessary for very distorted lenses and can usually be omitted. By expressing $\mathbf{u}_d$ as homogeneous coordinates we can determine the image coordinates from eq. 6.5.

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} \sim \begin{bmatrix} f & 0 & u_0 \\ 0 & f & v_0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u_d \\ v_d \\ 1 \end{bmatrix} \Rightarrow \mathbf{u} \sim \mathbf{KK}\,\mathbf{u}_d \tag{6.5}$$

This camera matrix **KK** has reduced by one dimension because the projection is previously performed by eq. 6.1 and the transformation is now in 2D only. Also components $u_0$ and $v_0$ have been added which represent any offset of the image coordinates origin from the optical axis, and are referred to as the *principle point*. Bouguet [24] notes that this point can be quite difficult to estimate and it is often preferable to set this as the centre point of the image.

The extrinsic parameters are unchanged and define a rotation and translation to map from world to camera coordinates as shown in eq. 6.6. By using this equation, along with eqs. 6.1, 6.4 and 6.5, we can transform from 3D world coordinates to 2D image coordinates.

$$\mathbf{X} = \mathbf{R}\,\mathbf{x} + \mathbf{T} \tag{6.6}$$

However, first we need to determine the 14 extrinsic and intrinsic parameters. This is achieved by using a calibration object, a checkerboard, to provide known values of $\mathbf{x}$ and

**u** from which the parameters can be obtained using an iterative technique. A gradient descent approach is used by the calibration toolbox, which minimises the error in re-projection.

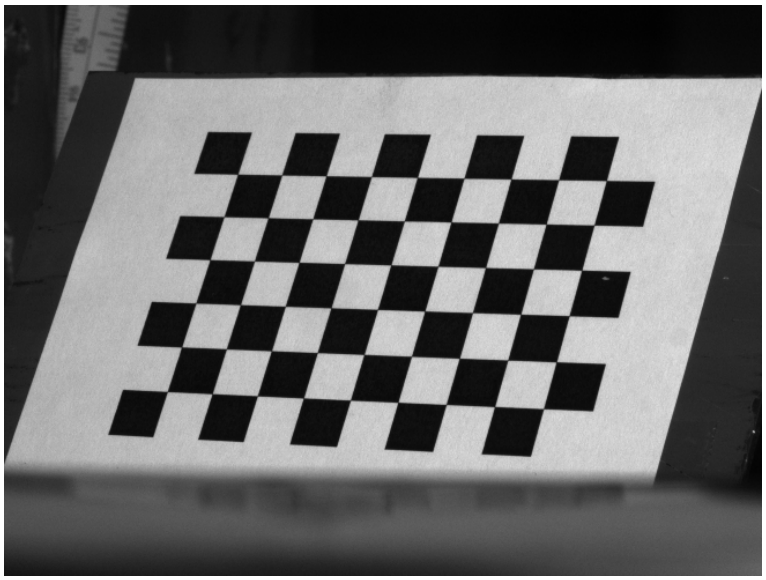


Figure 6.17: example calibration image.

The method used to perform the calibration of a camera consists of three steps. First, several images of the checkerboard are taken from different angles (an example image is shown in fig. 6.17); second, the corners of the checkerboard are located in each image; and third, the parameters are estimated. Each calibration image effectively defines a new origin for the world coordinates, with the *xy* plane the same as the checkerboard plane. The corner detection then provides the corresponding image coordinates and so we obtain several different sets of **u** and **x** values to be used for calibration. Ideally 15+ images should be taken and these should be spread through the imaging frame and set at different angles. The corner detection of so many images can take some time and so here we have used the automatic corner detection as provided by the Graphics and Media Lab at the Department of Computer Science, Moscow State University [154].

In this method we calibrate each camera separately and then, once all camera matrices are known, we determine the relative position of each by projecting back into a common world coordinate system, i.e. a common calibration image. It is also possible to jointly optimise the calibrations of multiple cameras by using images of the same object, and in theory this can yield lower errors in the estimated camera parameters [24]. In this work, joint optimisation has not been used because, in practise, taking sufficient suitable images of a common checkerboard is difficult due to the obtuse angles encountered; for example, between the first and last views. Svoboda et al. [146] overcome this problem by using a single point (a laser pointer) calibration object and achieve a re-projection error of just

one fifth of a pixel. This approach however, requires many synchronised frames (100+) that are not possible in the current set-up. (A *bug* exists in the control software provided by the vendor, as previously explained.) Nevertheless, the joint optimisation approach is promising and in future work, once the *bug* is fixed, it should be investigated.

### 6.3.2 Registration Approach comparisons

The CV approach was used to register the camera set-up of section 6.2, and the results are compared to those obtained using the previous cylindrical object method. However, the mirrors had to be removed because they resulted in parameters of very high uncertainty when the CV approach was used. These high uncertainties arose because the mirrors locate the checkerboard only in the two outer thirds of the image. This limits the range of poses of the checkerboard and so it is not possible to get a sufficient number of good calibration images.

The CV calibration was performed using an aperture of $f/8$ to ensure sufficient depth-of-field for the 80 mm wide checkerboard. The CV camera calibrations all achieved a re-projection error of less than 0.3 pixels, or 57 $\mu$m. A total of 36 different poses of the checkerboard were used. It was not possible to take images of all poses from each camera; however, all cameras had at least 19 calibration images. Pose five was visible from all cameras and was used to obtain their relative positions, which were then converted to object coordinates.

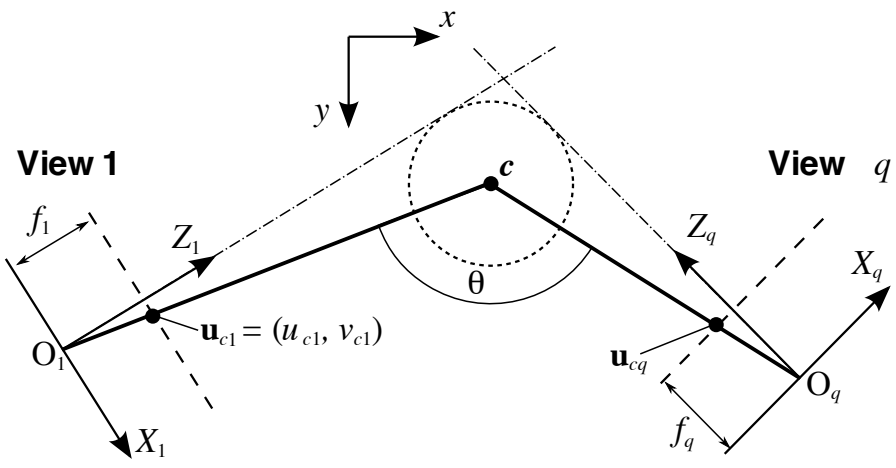



Figure 6.18: Diagram of the triangulation approach in the *xy* plane.

The conversion to object coordinates first requires that the location of the burner centre in pose 5 coordinates be found. The burner centre was located, as illustrated in fig. 6.18, from the intersection of the appropriate projections of each camera. The following procedure was used.

Table 6.5: Camera Calibration results. See the main text for discussion of the uncertainty.

| **View No.** | $f \pm 0.3$ [mm] | $S_{obj}$ [mm] | $\gamma$ [deg] | $\theta$ [deg] | $\theta_{\text{cyln}}$ [deg] |
|---|---|---|---|---|---|
| 1 | 36.6 | 642 | 1.002 | 18.000 | 18.000 |
| 2 | 36.6 | 642 | 0.657 | 54.035 | 54.154 |
| 3 | 36.5 | 642 | -0.047 | 90.056 | 90.153 |
| 4 | 36.4 | 647 | 0.580 | 126.053 | 126.031 |
| 5 | 36.4 | 646 | 0.500 | 161.828 | 161.921 |

**First** The centre of the burner $\mathbf{c}$ in each views image coordinates $\mathbf{u}_{cq}$ was determined using images of the cylindrical object seen in fig. 6.7. These images have already been corrected for rotation about $Z$ as previously described.

**Second** Then $\mathbf{u}_{cq}$ was converted using the camera matrix to camera coordinates $\mathbf{X}_{cq}$, and so used to define two arbitrary points ($Z$ =700 and 300) on the line between $\mathbf{u}_{cq}$ and $\mathbf{c}$ shown in fig. 6.18.

**Third** For each view these points were converted to the world coordinates of pose 5 using eq. 6.6, and then used to determine the direction vector $\hat{\mathbf{v}}_{cq}$ of each line.

**Fourth** The location of the burner centre in $xyz$ was then taken as the mean intersection point of all the $\hat{\mathbf{v}}_{cq}$ lines. This was found simply from the cross-product of the vectors.

The resulting offset from the centre from pose 5 is shown below.

$$\mathbf{c} = \begin{bmatrix} x \\ y \\ z \end{bmatrix}_{p5} = \begin{bmatrix} 23.8 \\ -3.9 \\ 7.5 \end{bmatrix} \pm \begin{bmatrix} 0.17 \\ 0.12 \\ 0.02 \end{bmatrix} \text{ mm}$$

For $\theta$, only the relative angle of each camera is necessary and this can be obtained from the dot product of a reference line (i.e. view 1, $\hat{\mathbf{v}}_{c1}$) with the lines of the remaining views ($\hat{\mathbf{v}}_{cq}$). The angle $\gamma$, however, needs to be measured from the $xy$ plane. The normal of this plane, the $z$ axis, is determined using the four step procedure described above, but with $\mathbf{c}$ chosen at a greater $z$ value $z_{up}$. This new centre point $\mathbf{c}_{up}$ together with the original point are then used to determine the direction vector of the $z$ axis. The dot product of the $z$ axis with the view lines' direction vectors, ($\hat{\mathbf{v}}_{cq}$), thereby yields $\gamma$ as per eq. 6.7.

$$\gamma_q = \frac{\pi}{2} - \hat{\mathbf{v}}_{cq} \cdot \vec{z} \tag{6.7}$$

The results of the CV registration of each view are shown in table 6.5, alongside the calibrated focal length. The value of the in-plane angle, obtained using the cylindrical object registration approach, $\theta_{cyln}$, is also shown for comparison. The values for $\theta$, obtained

using the CV approach are very close to those predicted using the cylindrical object alone, with a maximum of 0.12° difference indicating low errors in both approaches. For the CV approach, the uncertainty in $\theta$ is estimated to be $\pm 0.035°$, with the exception of view 3 which has an estimated uncertainty of $\pm 0.500°$. These uncertainties are obtained by propagating the error in $\mathbf{u}_c$ (0.5 pixel), the re-projection error (0.3 pixel), and the extrinsic parameter error (typically less than 0.2%) through the calculation. Except for view 3 these errors are very low, and demonstrate the accuracy achievable with this CV approach. The higher error in view 3 occurred because pose 5 was almost perpendicular to this camera. This was necessary to ensure that cameras 1 and 5 could each obtain a sufficient view of the checkerboard of pose 5, but resulted in a much higher error for the extrinsic rotation parameters of view 3. Nevertheless, the uncertainty for the angle in view 3 is still equal to that using the cylinder alone and would be sufficient for the purposes of CTC as already seen. Furthermore, this problem could be eliminated in future applications by choosing a better, non planar, calibration object, as used by Svoboda et al. [146].

The error of $\gamma$ will be higher than $\theta$ as this relies on the determination of the object $z$ axis. However, this can be expected to be a similar order and these values have sufficient precision for our use. The values of $\gamma$ are relatively small, although a difference greater than 1° is observed between view 1 and 3. Similar differences have been tested in section 6.2 and found to have only a small effect. In this section we have not considered the rotation of the cameras about $Z$ because the CT implementation does not include this rotation. The rotation has been measured and removed via a counter rotation, as described previously in section 6.2. However, the CV camera calibration does allow for the rotation about $Z$ to be determined and it could be included in the CT implementation if necessary.

The camera calibration provides an estimate for the focal length of the lens and also allows the object distance to be determined. This is a significant advantage when optical/ non-parallel projections are used as these important parameters are not otherwise necessarily easy to determine. For example, the focal lengths quoted by the manufacturer are not very precise ($\pm 5\%$) and measuring to the internal iris of a lens is difficult. The figures shown in table 6.5 support these examples as the focal length is approximately 1.5 mm greater than the manufacturer quotes (but within the quoted uncertainty), and the object distance is approximately 42 mm longer than the measurement to the front of the lens. While this difference in object distance at first seems large, it is valid as the distance from the front of the lens to the CCD is approximately 53 mm and the iris is located towards the rear of the lens assembly.

#### 6.3.2.1 Conclusions

We have performed the registration of the cameras using a Computer Vision approach for the arrangement of section 6.2 but without the mirrors. The mirrors have necessarily been removed to allow adequate calibration images to be taken, leaving only five views. Reconstructions of the Matrix Burner have not been shown because, with just five views, significant aliasing artefacts occur that prohibit further analysis. The CV registration approach has been compared to the bespoke registration approach used in section 6.2 and both are found to produce very similar values for $\theta$. However, the CV approach achieves much lower uncertainties, is generally applicable, and is employed successfully in other optical measurement techniques. Perhaps more importantly for CTC, it can also fully locate the views in 3D (6 degrees of freedom) as well as provide useful intrinsic camera information, such as the focal length. For these reasons the CV approach should be preferred in future work. The quality of calibration is dependent on the quality of the calibration images and we have found that obtaining such images of the checkerboard is non-trivial. Consideration should be given to a more appropriate calibration object, such as a 3D arrangement of point sources, and also to the multiple optimisation of cameras to improve the calibration process.

### 6.3.3 Lens Distortion

In section 6.2 it was suggested that lens distortion may contribute to the deformations present in the reconstructions of the Matrix Burner. The CV camera calibration provides an estimate of the parameters *kc*, which are used to model the distortion of the camera lens. We can therefore use this model to assess the extent of the lens distortion, and its potential effect on a reconstruction.

The calibrated lens distortion of camera 1 is shown in vector form in fig. 6.19. This shows, as expected, the greatest distortion at the image corners and very low distortion in the centre. The distortion in the top corners of the image (positive $v$) is greater than in the lower corners (negative $v$). This is the result of the tangential component of distortion, which for this case is always in the positive $v$ direction. Overall however, the distortion of the Pentax lenses appears low for this aperture, and this is also illustrated qualitatively in the accompanying calibration image in fig. 6.19, which has been undistorted using eq. 6.4.

Lens distortion affects reconstructions by changing the individual viewing angle, $\delta_{mn}$ (see fig. 5.20) of each projection in a view, and so results in error if not accounted for. We can quantify the angular change, by comparing the angle predicted using the distortion model, with the angle predicted without the model, i.e. with the pinhole camera model alone. We define this change as $\Delta_{mn}$, which is the difference between $\delta_{mn}$ and $\delta^D_{mn}$;

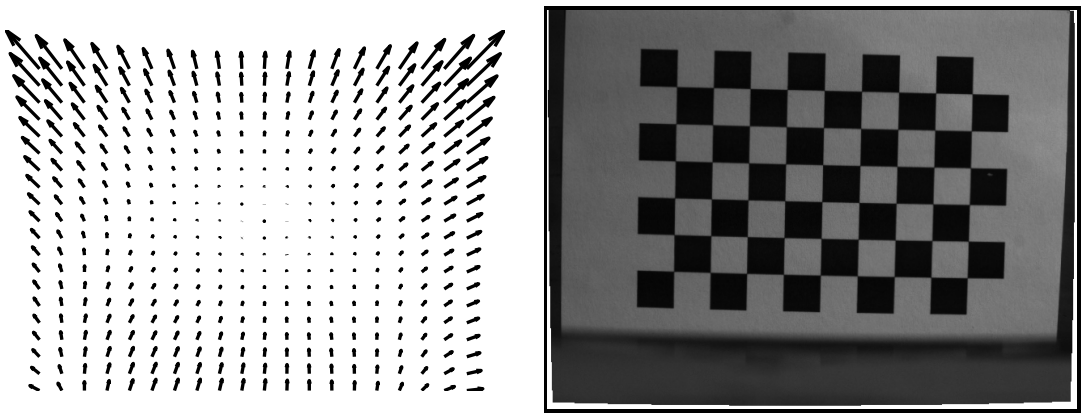


Figure 6.19: A vector plot of the distortion of camera 1. The vectors have been scaled by 2 for clarity. Also shown is a calibration image taken with camera 1 that has been appropriately undistorted. Please note that the vector plot is shown inverted.

where $D$ denotes lens distortion adjusted. This change is plotted in fig. 6.20, for rows and columns in the centre and at the edges of the image, using the calibrated parameters of camera 1.

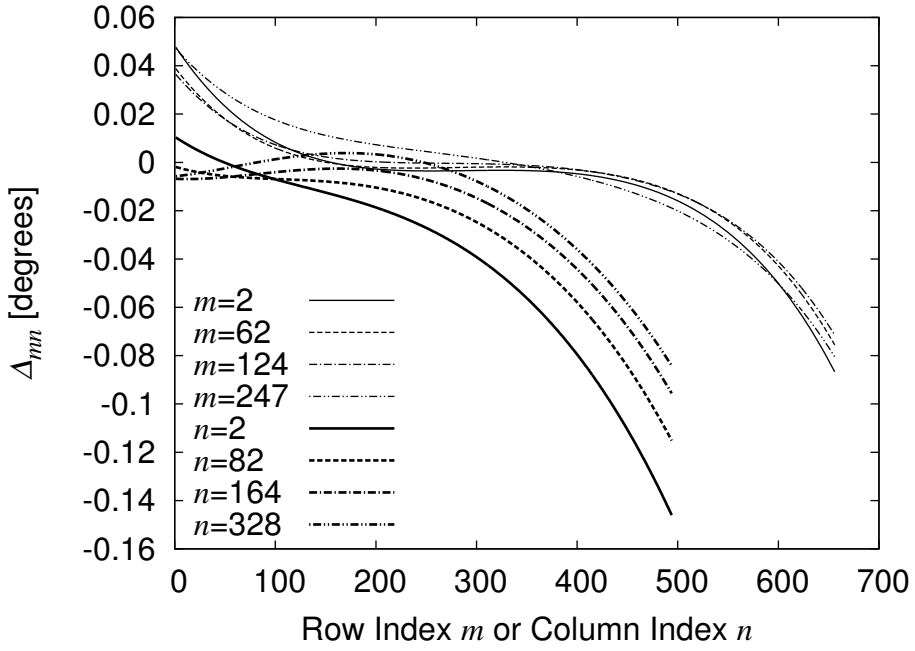



Figure 6.20: The difference in the viewing angle of each pixel, with and without distortion for camera 1.

This figure confirms that the angles in the image centre change little ($< 0.01^\circ$), and a larger change is observed at the image edges and, in particular, the corners. In section 6.2, the two views measured with camera 1 are vertically located in the centre of the camera image. However, these views are horizontally located at the edges (between $n$=41 – 155 and $n$=501 – 615) and are placed in areas of higher distortion, with resulting $\Delta_{mn}$ values

of approximately 0.02 – 0.01° and -0.02 – -0.06°. While these errors appear small, for the object distance of approximately 680 mm, 0.06° is equivalent to almost 4 voxels width in the object domain—the thickness of a single flame in the Matrix Burner.

Previously in section 6.2, much larger random errors in $\theta$ (up to $\pm 1°$) have been observed, and have had little affect on the reconstruction accuracy. These previous results suggest that the small errors in $\theta$ caused by lens distortion are negligible. However, there are two important differences between these cases. Firstly, $\theta$ is the general angle of the whole view and any error effects each projection equally. Secondly, the error in $\theta$ was random, with lens distortion the error is more systematic. These differences suggest that uncorrected lens distortion could lead to error in the reconstructions, and could also be a significant contributor to the flame shape deformations observed in section 6.2.

#### 6.3.3.1 Correcting Lens Distortion

In theory, the distortion coefficients determined using the CV approach with no mirrors are also applicable to the mirrored views of section 6.2; because it is the lens that is calibrated and this is common to both arrangements (also the aperture size does not affect the distortion [66]). The measured views of section 6.2 could therefore be corrected and so allow improved reconstructions to be obtained. Unfortunately, with the available set of calibration images for cameras 2 – 5, it has not been possible to obtain sufficiently reliable values for the distortion coefficients *kc* and so this has not been performed. The reason cameras 2 – 5 have much higher uncertainties in *kc* than camera 1 is not fully understood, but seems to be related to the prediction of the principle point. In preliminary calibrations of all cameras, the predictions for $u_0$ and $v_0$ were far from the image centre, and had high uncertainties and counter-intuitive distortion fields. For these reasons the calibrations were performed with the principle point fixed to the centre point of the image (328,247), as recommended by Bouguet [24]. Fixing the principle point slightly reduced errors in other parameters, but lead to an increase in the error for *kc* in all but camera 1 (which also produced a more intuitive distortion field: low in the centre and higher at the edges). These differences are repeatable and suggest that the calibration images for camera 1 are superior somehow to those of the other cameras. We note that in the previous sections studying CV camera calibration, the relevant areas have been in the centre of the images and so their results would not be adversely affected by uncertainty in *kc* as this area has very low distortion.

The potential for reliable models for the lens distortion with which to improve reconstructions makes this CV approach attractive; particularly where views stretch the width of the imaging domain. However, to realise this potential in future work, improvements to the procedure used here are required—for example a better calibration object—and also

further consideration needs to be given to the principle point estimation. Nevertheless, the results of camera 1 indicate that a low error model of the lens distortion is possible with the right calibration images, and such a model would open up interesting possibilities for CTC. For example, *fish eye* lenses could be used, which allow very wide angles of view, and would be useful for *in situ* sensors.

## 6.4 Impinging Matrix Burner

So far the Matrix Burner has been run using low flow rates, to provide a known, non-axisymmetric arrangement of small distinct laminar flames. However, the Matrix Burner is commonly used for heating purposes, and is run with high flow rates which lead to less defined constituent flames that spread as they impinge upon a surface. In this section we perform CTC of the Matrix Burner running in this mode to expose the sensor to a different flame shape (see fig. 6.21). With the mirror system employed in section 6.2 we are restricted to a domain diameter of 34 mm, which is not sufficient for the impinging flame of approximately 70 mm diameter. With the mirrors removed the possible domain size increases to approx 120 mm and we can consider the whole impinging flame. By removing the mirrors we can take only five views of the flame and so will not be able to capture the high frequency detail. For example, near the burner itself, significant aliasing can be expected. However, where the flow impinges, much lower spatial frequencies can be expected that should be captured with five views.

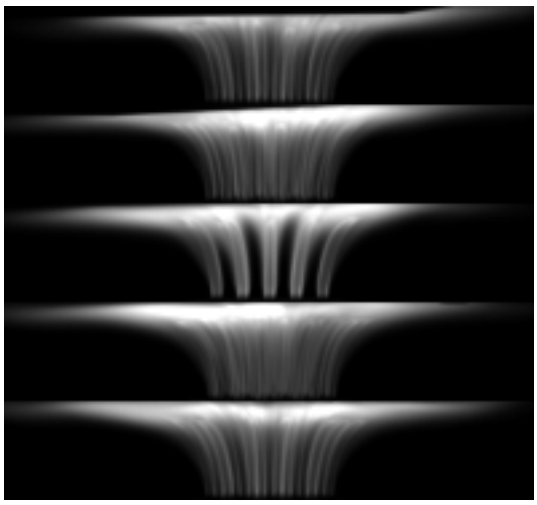
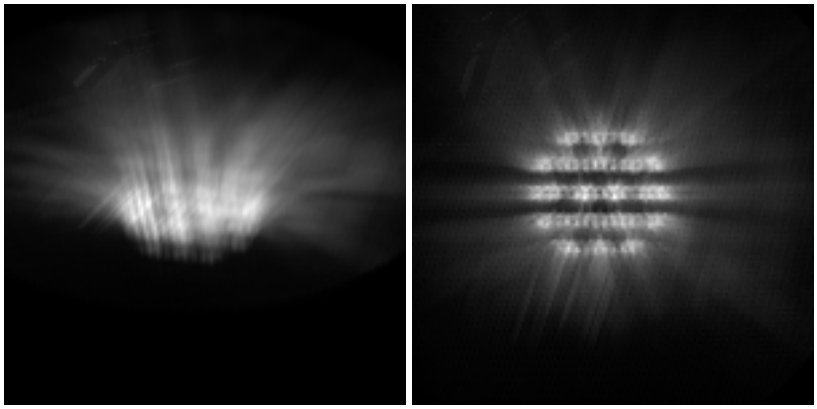


Figure 6.21: The 5 measured views of the impinging Matrix Burner. View 1 is shown at the top and view 5, at the bottom.

The impinging flow was created by suspending a thick rectangular plate of high-temperature metal alloy into the Matrix Burner flame. The five views of the impinging flame were measured using the camera set-up shown previously in fig. 6.6, but without

the mirrors. This set-up locates the views in the centre of the camera images, so distortion is low, and also enables the values for $\theta$ and $\gamma$ to be taken from the CV registration in section 6.3. The views span 80 by 15 mm in the object domain, with a pixel resolution of 456 by 86. However, to reduce the computational cost, the views have been filtered and down-sampled to 198 by 37 pixels. An exposure of 2.5 ms was used along with an aperture of $f/8$ to give sufficient depth-of-field. The resulting views are shown in fig. 6.21. The reconstruction was performed in 3D using cylindrical parallel projections on a grid of $N_i = 198, N_j = 198, N_k = 37$ and using the ART algorithm with $\Delta_c = 10^{-4}$.

Volume renderings of the reconstruction of the light intensity are presented in fig. 6.22, and show that the global shape of the impinging flame has been captured. The individual flames are seen to spread away from the burner centre line as they rise, and spread into a flat rectangular secondary reaction zone that matches the dimensions of the metal plate. It is interesting to note the low intensity bands that pass through the middle of the flame, seen most evidently in fig. 6.22b, which are clearly the result of the advantageous angle of view 3. Indeed, if a view perpendicular to view 3 were also measured, we could have expected to resolve the lower individual flames more clearly (albeit at the expense of global resolution, as five views cannot be evenly spaced *and* perpendicular). In this work, we aim to avoid reliance on such prior knowledge in order to maintain generality. However, the use of such *a priori* information is sensible in specific cases—in particular where the number of views is very limited.

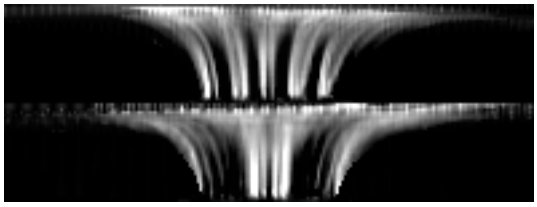


(a) $\theta = 29°, \gamma = 36°$ (b) $\theta = 18°, \gamma = 90°$

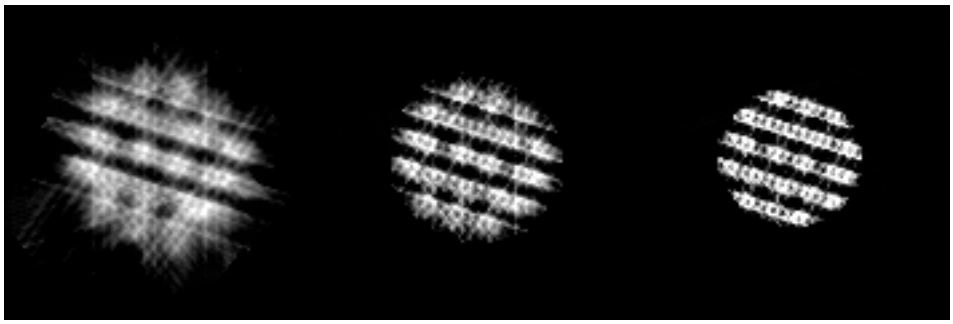
Figure 6.22: Volume renderings of the Impinging flow reconstructions. The viewing angles of each rendering are shown below the images.

The individual flames have not been well resolved, as was expected, and this is evident in fig. 6.22b. This is further confirmed by reconstruction sections shown in fig. 6.23, where the 21 constituent flames can hardly be identified due to the severe aliasing that occurs. It would be appropriate therefore to filter the reconstructions to lower resolution.

However, this case presents an interesting problem as the impinging flame effectively presents two domain widths: 22 mm at the burner exit and 70 mm at the impinging point. In such a case of variable domain widths the absolute resolution also varies through the domain, as per eq. 4.9, and variable filtering could then be justified.



(a) top: $x = 99$; bottom: $y = 99$



(b) $k = 15, 22, 30$

Figure 6.23: Sections of the impinging flame reconstruction. The section indices are shown below the images.

As almost all practical combustion systems are confined, the interaction of flames with surfaces is of great practical interest, for example when considering heat transfer. Here we have demonstrated CTC for a flow impinging on a surface and successfully recovered the low spatial frequency shape of the flame. The five views we have been able to obtain are not sufficient to resolve the higher spatial frequency features of this flame, but we reiterate that the CTC sensor is planned to obtain 25 or more views in future work, and would achieve much higher resolutions.

## 6.5 Premixed Turbulent Opposed Jet

The Matrix Burner has proved to be a very good test case for CTC for reasons already given; however, while it has some of the properties of turbulent flames (several flame fronts, non-axisymmetric), it is nonetheless an arrangement of steady laminar flames. The CTC sensor is intended primarily as a tool for the study of turbulent combustion and so in this section we consider the application of CTC to a turbulent flame: a Turbulent Opposed Jet (TOJ). As CTC has only so far been tested with non-premixed flames we consider the TOJ in the premixed mode—the mode CTC will perhaps most likely be used with. A

similar turbulent flame configuration was the subject of an LES in section 4.1.2 and used for Phantom studies; however, here we consider this configuration experimentally and in the premixed mode. Turbulent Opposed Jets in both modes have been the subject of several experimental studies [51, 52, 97, 99, 103, 102] and have also recently been studied with LES by Kempf [85] and Stein [142]. The TOJ configuration is therefore a relevant test case for CTC and has the advantage of a spatially compact flame and also good—though not perfect—optical access.

## 6.5.1 Experimental Set-Up

### 6.5.1.1 Premixed Turbulent Opposed Jet

The TOJ burner was originally designed and manufactured at TU Darmstadt and studied in the non-premixed mode by Geyer et al. [51, 52]. This burner has been installed at Imperial College London and modified to burn in the premixed regime by P. Geipel and R. P. Lindstedt in as yet unpublished work. The burner is depicted in fig. 6.24 and consists of two opposed circular nozzles of 30 mm diameter ($D$) located on their common centre line with a separation of $1D$. Each nozzle provides a premixed fuel/air stream and is concentric with a larger diameter nozzle which provides a stabilising co-flow of air. Both fuel/air streams have the same equivalence ratio and speed; therefore, locating the stagnation point an equal distance from each nozzle. The two streams stabilise a twin flame brush which is centred at the stagnation point. (This contrasts the non-premixed configuration where only one flame is observed.) Turbulence is generated upstream of each nozzle by a turbulence generating plate (TGP) consisting of a brass sheet perforated with 4 mm diameter holes.

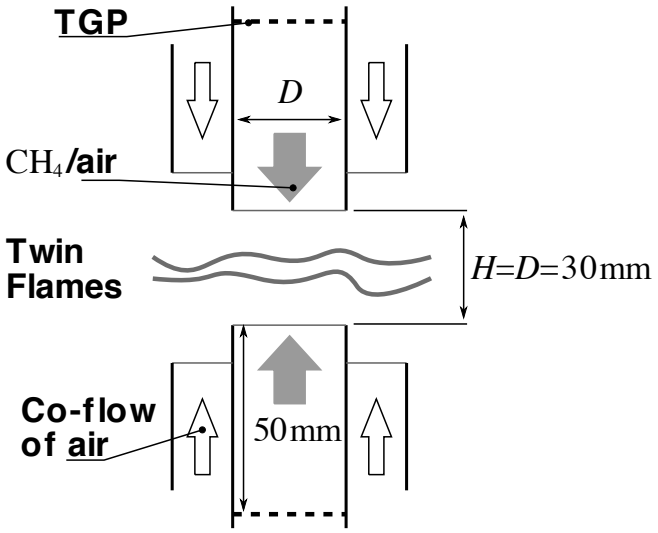



Figure 6.24: Diagram of the Turbulent Opposed Jet.

For these experiments the fuel used was methane and the oxidiser was air. The mixtures were set using mass flow controllers by Bronckhorst and only lean mixtures have been considered so no soot is present in the flames. To observe the differences in the measured intensity signals, three different stoichiometries were used: 0.7, 0.8, and 0.9. A jet bulk velocity of 3.6 m/s was used giving a Reynolds number of approximately 6000.

#### 6.5.1.2 CTC set-up

The TOJ flames considered have a chemiluminescent profile of approximately 100–120 mm in diameter, which is much larger than the allowable object width of the mirrored set-up in section 6.2. The removal of the mirrors would provide a sufficient field of view; however, as seen with the Impinging Matrix burner, using only five views leads to significant artefacts in the reconstruction. We therefore employ a new mirror arrangement that allows ten views of the TOJ flame to be obtained. The cameras and the triggering mechanism remain the same as in section 6.2.

The new mirror system is shown in fig. 6.25 and as previously is used to provide 2 views per camera, but this time for an object domain 120 mm in diameter. As before the 2 views are ideally spaced at 18° apart, and the mirror locations and angles have been obtained using ray-tracing to ensure that no vignetting occurs. Because of the wider domain the cameras and mirrors have had to be located further from the burner to ensure a sufficient horizontal field of view, and the front of the lenses are now at a radius of 1674 mm. The outer mirrors are positioned approximately 200 mm from the camera's optical axis and result in an effective object distance of 2 m. However, because we use only 10 views there is ultimately no resolution penalty associated with this large object distance, as the view resolutions still exceed that achievable in the reconstruction. Because of the large radius a single mount was no longer practical and each of the 5 camera/mirror systems had its own mount. This 500 mm wide mount was machined from 10 mm aluminium and the technical drawing is shown in the appendix.

To obtain a reconstruction of consistent resolution views should be ideally spaced through 180°. Such spacing however, was not possible for these TOJ measurements because the supports of the Burner would significantly obscure views 3, 4, 7, and 8 as illustrated in fig. 6.26. Therefore, a non-ideal view arrangement had to be chosen as shown in fig. 6.27. This arrangement was selected as a compromise between having the least correlation between views, and having as few views obscured as possible. This results in a large spacing of 51° between views 4 and 5, which we accept, as spatial resolution is not the main priority of this test. In this set-up we also have to endure some truncation of view 7 and, as indicated in fig. 6.27, this view reliably spans only 86 mm of the 120 mm object domain. Here we use the term *reliably* to describe the view area that is neither

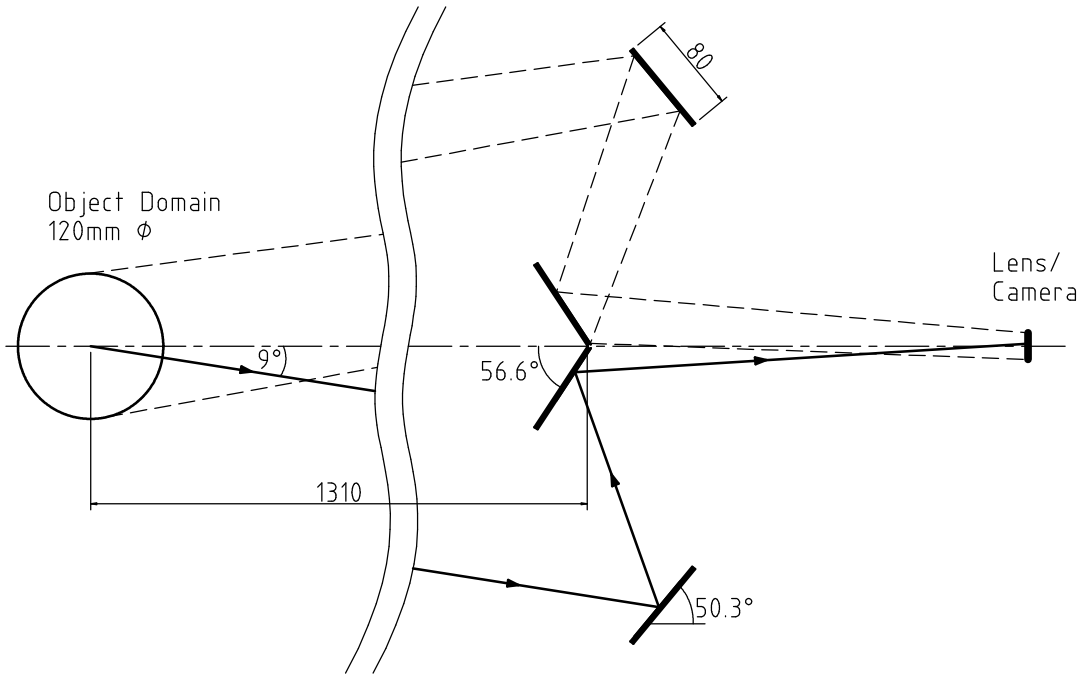



Figure 6.25: A scale diagram of the TOJ mirror arrangement, the lens is located 1674 mm from the burner centre. The dashed lines show the marginal rays of points on the edge of the domain and define the required size of the mirrors. The diagram has been truncated to fit, as indicated by the parallel curved lines.

fully obscured, nor partially obscured (i.e. subject to vignetting) by the burner support; although because of the large object distance there is only a small distance between these two cases. While not ideal, the ART algorithm does allow truncated views as demonstrated by Mishra et al. [108] and we can expect to pay only a small penalty for this single truncated view.

The 5 mounts were arranged around the burner as shown in fig. 6.27; however, practical difficulties in positioning the mounts meant that some views (7 and 8) deviated significantly from the positions shown in this figure. The mounts of views 1 – 6 were fixed to tables and could be located with a sufficient precision of approximately $\pm 10$ mm. For views 7 – 10 however, it was not possible to use tables due to space constraints and mounts had to be supported by tripods[7]. The correct location of the tripods was much more difficult and lead to some deviation from the intended angles for these views. This is discussed further in the following camera registration section.

Each mount was located in the same plane with an estimated precision of $\pm 10$ mm and levelled using a spirit-level. For view 10 this positioned the burner centre at the top of the image frame, where the measurements would have been subjected to vignetting. This mis-

[7]The tripods were good quality Manfrotto models capable of supporting a much larger mass than the aluminium mounts.

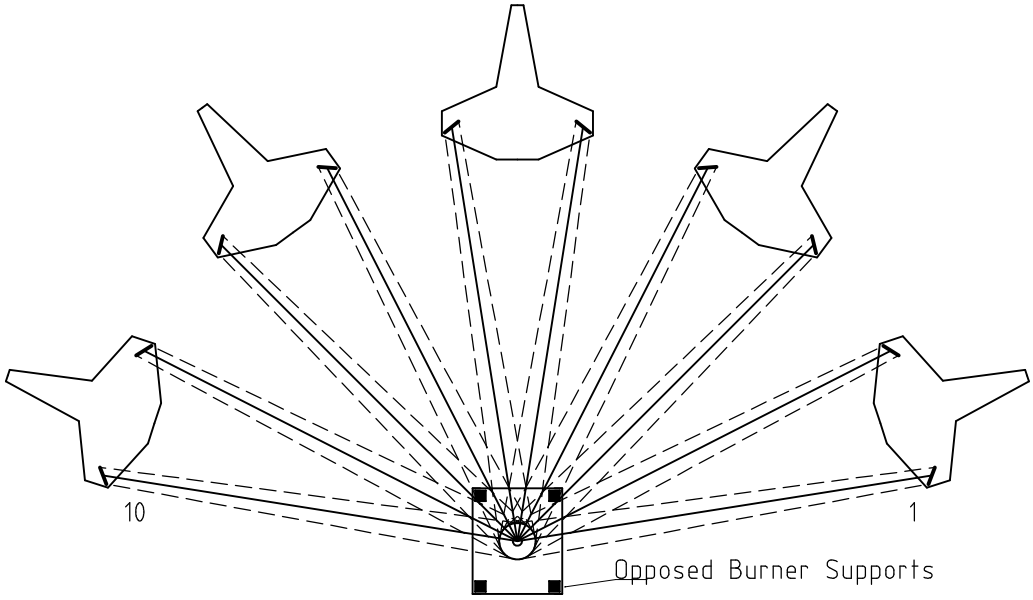



Figure 6.26: Scale drawing of the ideal view arrangement around the TOJ. The dashed lines represent the limiting marginal rays and the continuous line the central ray of the view. Only the rays from the outer mirrors are shown. This ideal CTC set-up is prevented by the upright supports of the burner, which obscure views 3, 4, 7, and 8.

location of the object centre was due to small errors in the mirror/camera locations on the mount which were amplified by the relatively large object distance for this set-up of 2 m. The attachments of the mirrors to the mounts did not allow sufficient adjustment and so to correctly locate the object in the image frame the whole mount was lowered by 50 mm. This means that views 9 and 10 are located in a plane approximately 1° from that of the other views. However, because such a deviation has been seen to have a relatively small effect, to make registration simpler and quicker we assume the views are still co-planar and accept this as a source of error. The degree of error introduced by the non-planar views 9 and 10 is estimated in the later section 6.5.2. With hindsight, by rotating rather than lowering the whole mount views 9 and 10 could have remained co-planar.

#### 6.5.1.3 Camera registration

With the assumption of co-planar mounts, the parameters to be determined for each view are $\theta$ and the object domain centre (chosen as the stagnation point) in the view $X$ and $Y$ coordinates. The use of mirrors prevents registration using the Computer Vision approach and so a cylindrical calibration object marked with an angular scale was again employed to determine the view angles, $\theta_q$. A cylindrical object was used successfully in the co-planar Matrix Burner set-up of section 6.2. However, because the object distance was much greater in this set-up a much larger cylinder was used to ensure sufficient resolution

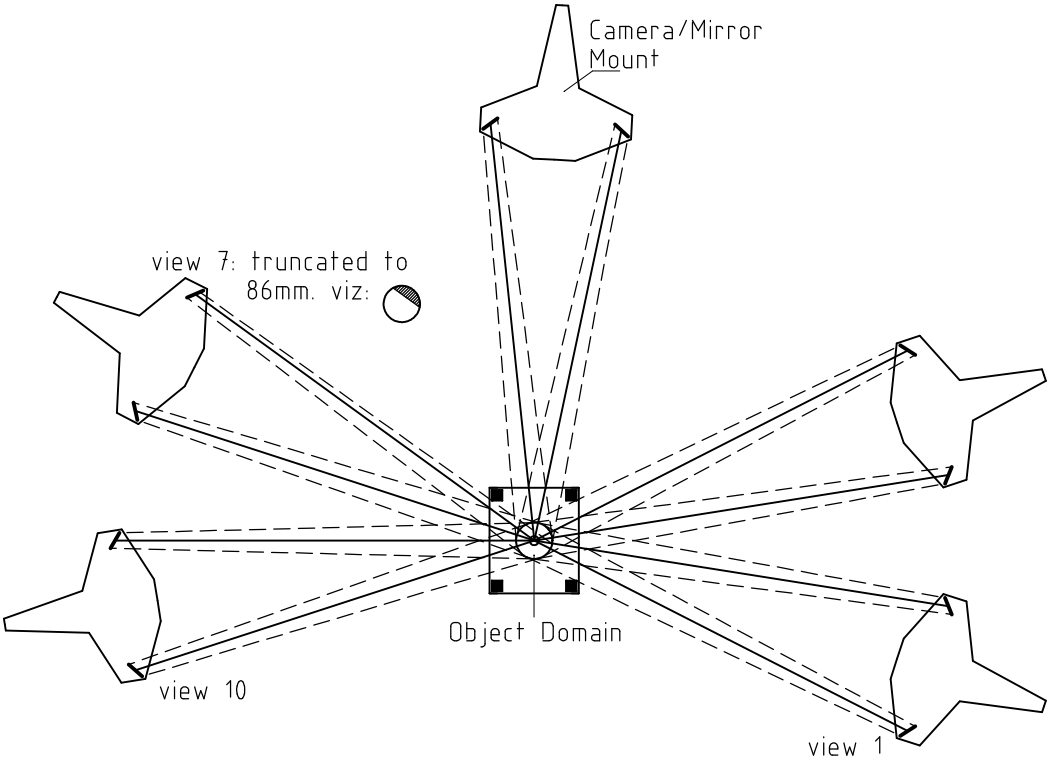
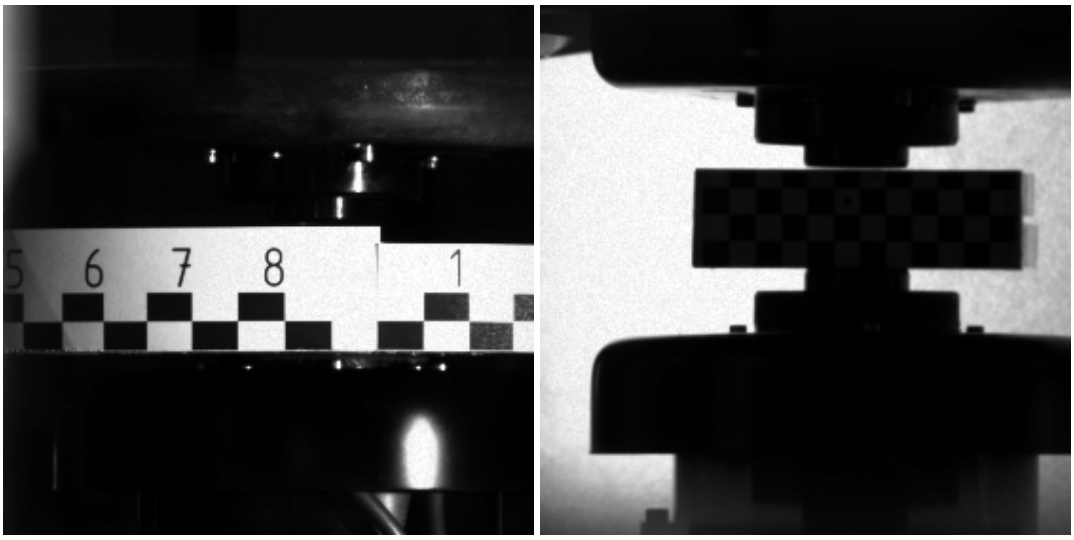



Figure 6.27: Scale drawing of the *intended* view arrangement around the TOJ (note, the arrangement is not symmetrical). The dashed lines represent the limiting marginal rays and the continuous line the central ray of the view. Only the rays from the outer mirrors are shown.

of the angular scale. The circular object was machined from 20 mm thick Nylon sheet to a diameter of 300 mm—these dimensions were limited by the supporting structure. The centre of the object was also machined away to a inner diameter of 62 mm so that it would *just fit* over a co-flow nozzle, and ensure the object was concentric with the burner nozzles. As previously a scale was affixed around the circumference, with marks every 5° as shown in fig. 6.28a. Because the width of the angular calibration object extends beyond the imaging area an additional calibration object was required to locate the object domain centre in each view ($XY$). This object was a simple rectangular sheet attached to a 30 mm diameter cylinder that fits within the lower nozzle and so locates the object between the two nozzles, as shown in fig. 6.28b. Images of both calibration objects were recorded for all views.




(a) View 1 angular registration (b) View 2 centre registration

Figure 6.28: Example registration images of the TOJ set-up.

Prior to registration the images of both calibration objects were first corrected for rotation about the optical axis. The appropriate counter rotation was determined from the edges of the upper and lower co-flow nozzles in the centre-registration images (e.g. fig. 6.28b). The burner centre line is determined from the mean of the nozzle edges' $X$ locations, and the rotation is then taken to be the angle between this line and the image $Y$ axis. Once the image is corrected the nozzle edges are again used to locate the object centre in the $X$ direction. The object centre $Y$ location was determined from the mean of the top and bottom edges of the centre calibration object. All edges were obtained using sub-pixel detection [24]. Finally the view angle $\theta$ is determined from the location of the object centre in $X$ on the marked angular scale. Linear interpolation was used to obtain values in between the 5° divisions. The determined values can then be seen in table 6.6.

Table 6.6: The registered relative viewing angle for each view alongside the intended angles shown in fig. 6.27. View 1 is used as the reference value.

| **View no.** | $\theta$ **[deg]** | **Intended [deg]** |
|---|---|---|
| 1 | 0 | 0 |
| 2 | 17.45 | 18 |
| 3 | 35.92 | 36 |
| 4 | 53.50 | 54 |
| 5 | 107.43 | 107 |
| 6 | 124.27 | 125 |
| 7 | 174.27 | 171 |
| 8 | 191.62 | 189 |
| 9 | 207.70 | 207 |
| 10 | 226.84 | 225 |

In general the registered relative viewing angles are close to the intended angles with only views 7, 8 and 10 differing by more than 0.75°. For views 8 and 10 this results only in a higher correlation with views 2 and 4 respectively as they now differ by only approximately 6° rather than the intended 9° from these views. With the non-equal spacing already in use, this increased correlation is not expected to significantly affect the reconstructions. For view 7 the correlation with view 1 is also similarly increased, but the increase in the relative angle has also decreased the truncation of this view by the burner support.

The truncation of view 7 has been further decreased by a *global* rotation of all views relative to the TOJ supporting frame (anticlockwise in fig. 6.27). The angular calibration object is designed to return only the relative angles of each view as this is all that is required by the CT algorithm. However, we can estimate the global difference between the actual and intended angles from the location of the supports in the calibration images. These indicate that the actual views differ globally by 2–3° from the intended views. While this rotation has decreased the truncation of view 7, view 4 is now slightly obscured by a burner support. Similarly view 3 is slightly truncated by a misalignment. Fortunately the truncation of views 3 and 4 is small and the decrease in the view widths has been determined from the calibration images to be 5 mm and 4 mm respectively. For view 7 the view width has been determined as 108 mm rather than the intended value of 86 mm.

### 6.5.2 Phantom Studies

In sections 4.2 and 4.3 we considered the effect of limited but ideally spaced angles on the accuracy and resolution of the CTC sensor. In the experimental set-up of this section, it has not been possible to equally space the views through 180° and our earlier analysis is less applicable. To better understand the accuracy and resolution of the TOJ reconstructions, we therefore perform a Phantom study using the specific, unequal experimental viewing angles defined in table 6.6. For the remainder of this section (6.5) we refer to these angles as the *experimental* angles.

To obtain a representative Phantom for this premixed TOJ set-up we again use LES. The implementation used earlier to provide the partially-premixed Phantoms was not set-up to simulate premixed flames and so data from a different implementation has been used. The premixed TOJ has been simulated by Stein [142] for methane/air mixtures using the *PsiPhi* code of Kempf, and it is from this LES that our Phantom is taken. This LES dataset has the advantage of a constantly spaced Cartesian grid: meaning the Phantom is free from the cylindrical grid artefacts seen in the earlier LES Phantoms. The simulation used an algebraic Flame Surface Density (FSD) combustion model, but does not predict any chemical species. However, due to the proximity of Chemiluminescence to the flame front we can obtain a sufficient representation from the FSD itself ($\Sigma$).

The Phantom is taken from a single time step of the LES which has a grid of $180 \times 180 \times 512$ cells. This grid spans from TGP to TGP and we consider only the central 39 cells in $k$ that encompass the twin flames. With this grid the LES has a spatial resolution of $260\,\mu$m; however, the view resolution is $\approx 580\,\mu$m and so we filter and resample the LES data to match this resolution using a combined *voxel* mixing approach[8]. We note that without the mirrors view resolutions of $120\,\mu$m are possible which exceed that of this highly resolved LES. The down-sampling results in a Phantom grid resolution of $N_i = N_j = 83$ and $N_k = 18$. As we are now concerned with full 3D reconstructions we in turn use all three dimensions of the LES data for the Phantom, rather than only extracting 2D sections as previously. This allows, for example, errors in $\gamma$ to be considered. Views of the 3D Phantom from the experimental angles were generated using parallel cylindrical projections and are shown in fig. 6.29, where the twin flame brushes of the premixed TOJ are observed.

[8] In pixel/voxel mixing the value for the new grid is simply the weighted average of the cells of the old grid that are encompassed.

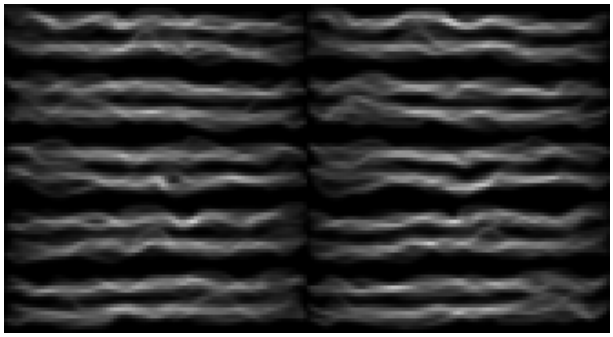


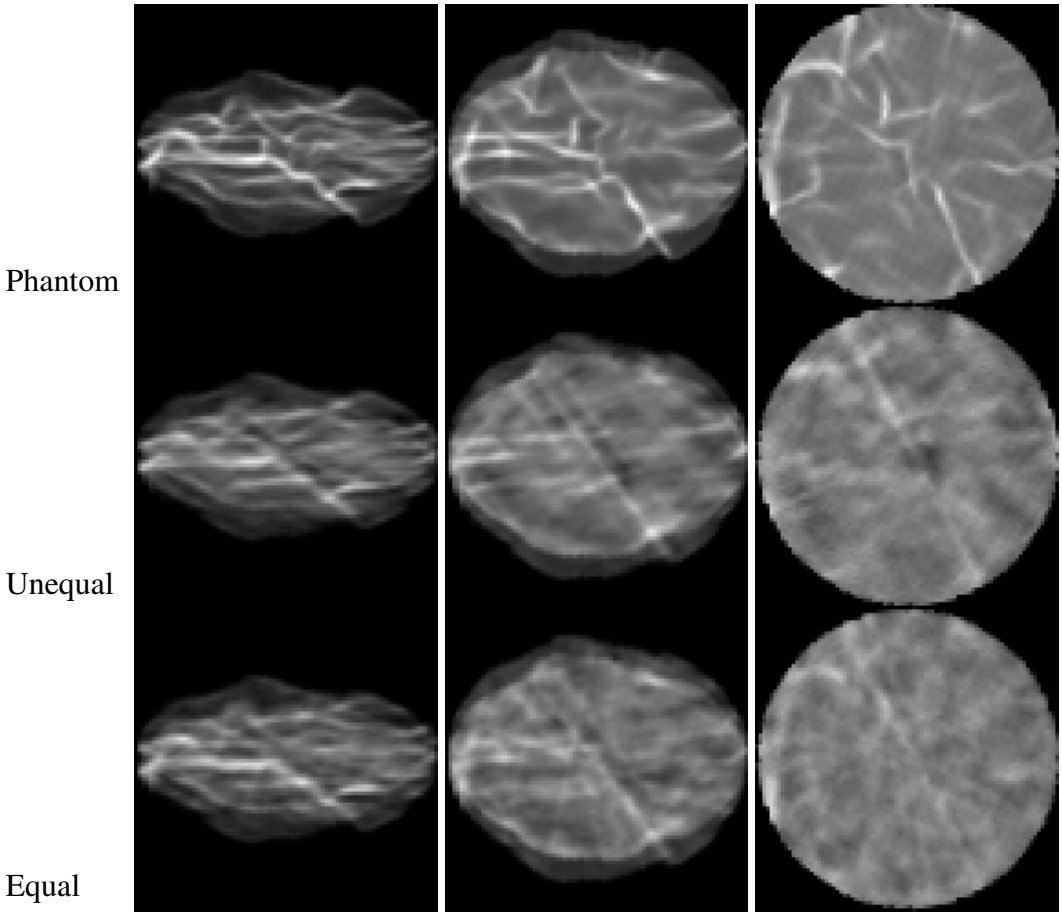

Figure 6.29: Views of the 3D FSD Phantom from the registered angles in table 6.6. View geometries have been scaled by two.

#### 6.5.2.1 Unequal Angles

Reconstructions of the 3D TOJ Phantom were performed using the ART algorithm with parallel cylindrical projections and the experimental angles of table 6.6. For comparison, a reconstruction using equally spaced angles was also performed. For both cases a $\Delta_c$ of $10^{-3}$ was used, which resulted in convergence after 14 iterations for the experimental angle case and 13 for the equal angle case. Volume renderings of the Phantom and resulting reconstructions are shown in fig. 6.30 and selected sections in fig. 6.31. These figures indicate that both arrangements are capable of capturing the largest features of the Phantom and both also achieve a high resolution in $z$, as expected for this co-planar set-up. However, though neither resolves the smallest in-plane features (again as expected), the equally spaced angle reconstruction provides a better—and more consistent—resolution of the larger Phantom features as particularly seen in fig. 6.31a.

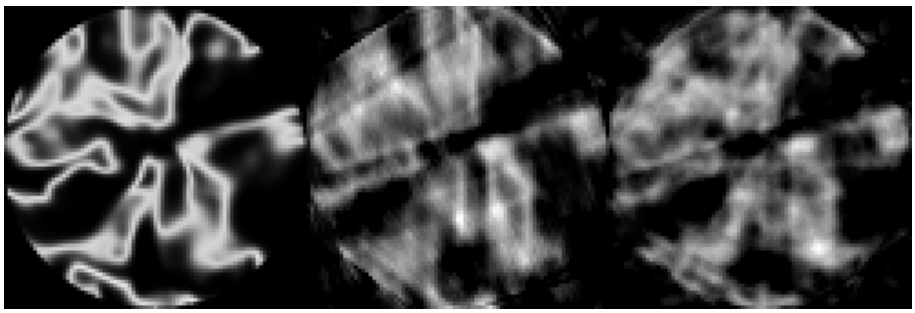
The improved resolution of the equal angle case is also evidenced by comparison of the reconstructions with the Phantom using the error measures described in section 4.2. The reconstruction using the experimental angles achieves a correlation coefficient of 0.90 and a normalised RMS error of 0.108, whereas the equal angle reconstruction achieves more favourable values of $r_{xy} = 0.93$ and $E_{RMS} = 0.092$. These results confirm that using the non-ideal experimental angles leads to an increase in the reconstruction error; however, the magnitude of this increase is moderate—an increase in the RMS error of 1.6% would be acceptable in many applications. On the other hand, the decrease observed in the correlation coefficient when using experimental angles, suggests a significant reduction in resolution, as this parameter has previously (section 4.3) shown to be a good indicator of resolution.

Knowing that the unequal angles used in the experiment will reduce the resolution of the reconstructions, it is desirable to know by how much. To quantify the effect on resolution we perform a cosine Phantom analysis (seen previously in section 4.3) using

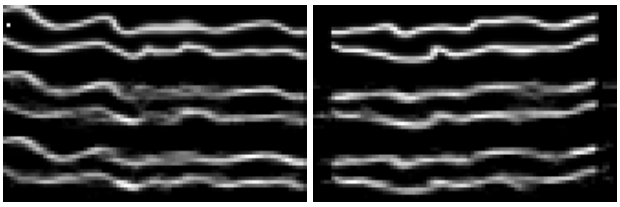



(a) $\theta$=80°, $\gamma$=22° (b) $\theta$=80°, $\gamma$=45° (c) $\theta$=80°, $\gamma$=90°

Figure 6.30: Volume renderings of the FSD Phantom (top) and its reconstruction using 10 unequally spaced views (middle), and 10 equally spaced views (bottom).



(a) $k$=12



(b) $y$=42; $x$=62

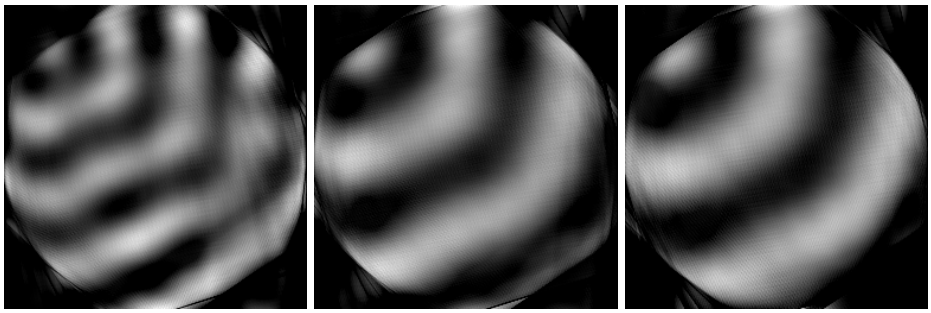
Figure 6.31: Selected sections of the FSD Phantom (top and left) and its reconstruction using 10 unequally spaced views (middle), and 10 equally spaced views (bottom and right).

the experimental angles of this specific set-up. The $z$ resolution is unaffected by the co-planar angular spacing and so we again consider only 2D Phantoms, specifically the 256TL Phantom type from section 4.3.

Using experimental spacing the resolved wavelength is 90 – 100 pixels (as demonstrated in fig. 6.32), which is significantly greater than the $\lambda_{res}$ of 60 – 65 obtained previously using 10 equally spaced angles. Figure 6.32a suggests that the difference in resolved wavelengths is due to a kind of localised aliasing, as significant local blurring and deformation of the wavefronts is observed in places, while elsewhere on the same wave apparently better resolution is achieved. This observation supports the intuitive assessment that non-uniform angular spacing results in non-uniform reconstructions.

Although rigorous and generally applicable, as seen previously, this cosine analysis is quite a stern test and returns only the limiting, or *worst-case*, prediction for the resolution. Where objects are more limited in their radial extent[9] higher resolution is achieved, and this explains why we have already seen much better resolutions using 10 *equally* spaced angles, than the wavelength of 60 pixels predicted by the cosine analysis. To improve our prediction of the resolution we need to apply prior knowledge of the object to be reconstructed. As a representative field, the TOJ Phantom provides a degree of *a priori* information that we can perhaps exploit. We have seen previously in section 4.3 that as its correlation coefficient approaches 1, a reconstruction achieves the resolution of the

[9] see section 4.3.



(a) $\lambda$ = 60; not resolved (b) $\lambda$ = 90; not resolved (c) $\lambda$ = 100; resolved

Figure 6.32: Reconstructions of the 2D cosine Phantoms using the experimental unequal viewing angles. Three different Phantom wavelengths are shown which illustrate how the resolved wavelength is determined.

Phantom. It therefore follows that we can estimate the resolution of a reconstruction by comparison to successively lower resolution Phantoms. This is conceptually similar to the cosine analysis, but rather than using cosines of known wavelengths we use realistic Phantoms that are filtered to known wavelengths. And, rather than visual confirmation, we use the criteria $r_{xy} \rightarrow 1$ to determine resolution.

To perform this test we have reconstructed a section of the 3D TOJ Phantom at $z = 5$, using both experimental and equal angles. We then apply a low pass filter (a sinc function) to each reconstruction and the Phantom, and so remove the high spatial frequency components. The correlation of the Phantom and the reconstructions is then recorded before they are again filtered but this time with a longer cut of wavelength, $\lambda_{co}$. The results of this process are shown in fig. 6.33, where we see that with equal angle spacing $r_{xy}$ *levels-off* substantially earlier than when the experimental spacing is used. In section 4.3 we saw that perfect correlation was not necessary for resolution, and observed that the cosine waves were resolved typically 0.02 normalised wavelengths ahead of the levelling-off point. We can therefore estimate the normalised resolved wavelength, $\lambda_{co}/D$, to be approximately 0.13 for the equal angles and 0.25 for the experimental angles.

This equal angle resolution of 7.7 waves-per-domain (wavenumber) is greater than the 4 waves predicted by the cosine analysis and is in better agreement with the 10 waves predicted using the relation suggested Freider and Herman [48]. This gives confidence in the applicability of this approach.

For the experimental angles the resolution of 4 waves across the domain is also an improvement on the cosine analysis but is still quite a significant reduction from the equal angle case, as is demonstrated in fig. 6.34. Although from inspection of the reconstructions we can see that some areas are better resolved than this value suggests, it is also evident that in other areas the prediction of 4 waves is a reasonable representation of

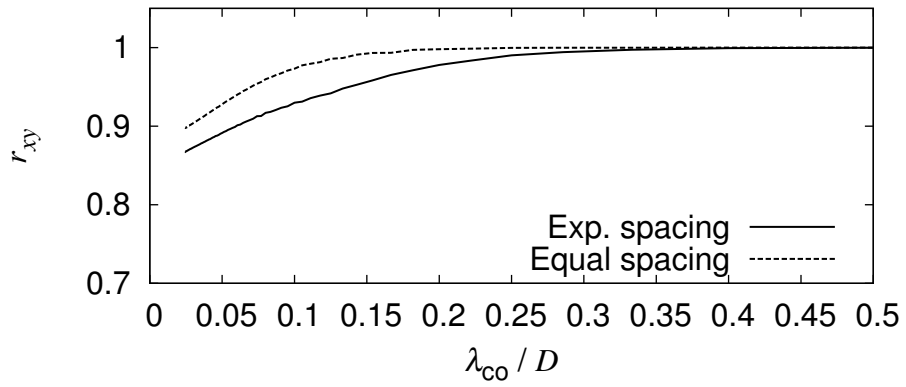



Figure 6.33: The correlation between filtered reconstructions and the similarly filter TOJ 2D Phantom. The cut-off wavelength, $\lambda_{co}$, has been normalised by the domain diameter $D$.

the resolution. Because of the non-uniform nature of the reconstructions using the experimental angles, we can therefore only be sure of resolving wavelengths greater than 30 mm for our experimental object domain of 120 mm. However, it may be possible in future work to identify the areas of lower and higher resolution from the viewing angles alone, and perhaps even provide a *resolution map*, so that we do not need to forgo these better resolved areas when having to use non-equal angular spacing.

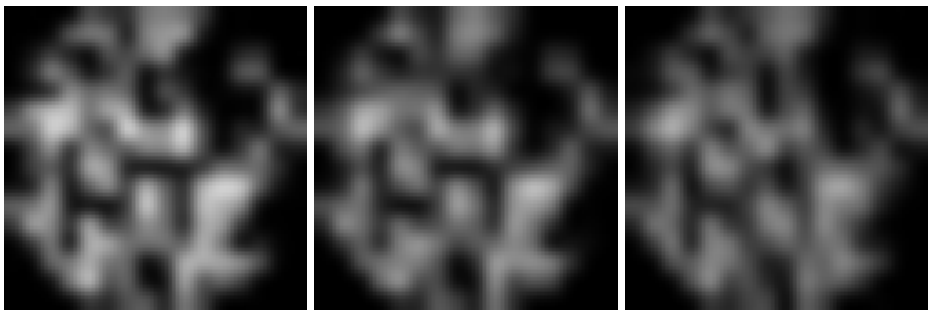


Figure 6.34: The 2D TOJ Phantom ($k = 5$) and its reconstructions using equal angles (middle) and experimental angles (right). All have had wavelengths less than $0.13/D$ filtered out.

#### 6.5.2.2 Systematic $\gamma$ Error

In correcting the in-frame location of the burner in view 10 we set this view and the co-mounted view 9 in a slightly different plane to the remaining views. This difference is relatively small—estimated from the object distance to be approximately 1°—and to facilitate camera registration[10] all views are assumed co-planar. By making this assumption

[10] necessary due to experimental time constraints.

we therefore introduce a systematic error, and here we use the 3D TOJ Phantom to assess the effect on the reconstruction that this error in $\gamma$ will have.

To simulate this systematic error, views 1 – 8 of the 3D TOJ Phantom were obtained for the co-planar experimental angles, as shown previously in fig. 6.29. The same approach was used for views 9 and 10 except that here $\gamma$ was set to $1^\circ$ rather than zero. The reconstruction was then performed assuming all these views were co-planar ($\gamma = 0$). When compared to the 3D Phantom this reconstruction achieves an RMS error of 0.118 and a correlation coefficient of 0.89; earlier without systematic error values of $r_{xy} = 0.90$ and $E_{RMS} = 0.108$ were achieved. This difference is consistent with that observed in other registration error studies performed in this work. It appears then that the expected increase in error by assuming co-planar views is moderate; and, from the sections shown in fig. 6.35, that the effect on the resolution is limited, though some additional *streaking* artefacts are observed. While we can accept such error here, we reiterate that should $\gamma$ be known then this error could be entirely avoided.

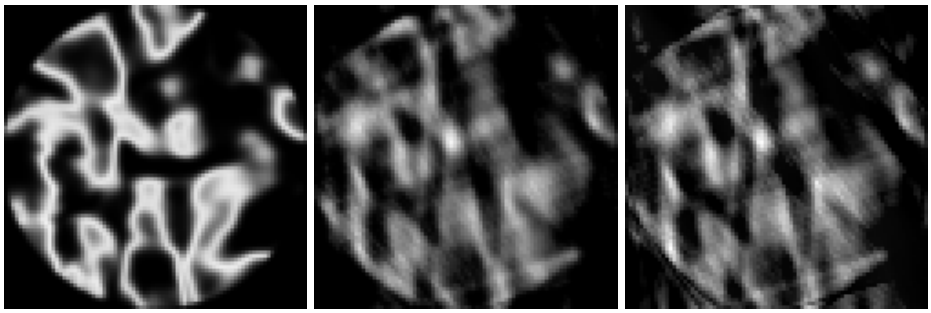


Figure 6.35: Sections of the 3D TOJ Phantom and its reconstructions using the experimental angles (middle) and with $1^\circ$systematic error in views 9 an 10 (right). The sections are at $k = 5$.

#### 6.5.2.3 When to Filter

In the case of limited viewing angles, we know the resolved wavelength of the reconstruction is typically much less than the resolution of the views. It is therefore reasonable to use a coarser grid for reconstruction, i.e. one with two or more points per resolved wavelength. For example, if twenty views sample an arbitrary object 100 times each, using our earlier cosine analysis we could expect a minimum of approximately 6 waves to be resolved across the object domain and a reconstruction grid of $24 \times 24$ could be justifiably used assuming 4 points per wave. In such a case the problem of reconstruction would become more determined and we could perhaps expect better reconstructions. However, we have seen that ART is capable of good reconstructions even in very under-determined conditions, and so far we have performed reconstructions on grids that are sampled an

*equivalent* amount as the views—in our example a grid of $100 \times 100$ would be used. This is because we wish to exploit the maximum amount of information contained in the views, and by using a coarser grid there would be some implicit filtering of the views and so a loss of information. Furthermore, this filtering may or may not be sufficient to prevent aliasing that results when trying to fit higher frequency information onto a lower frequency grid. This could, however, be avoided by explicitly filtering the views before, though again still with loss of information.

From a development point of view it is also better to avoid filtering as this can mask informative deformations and artefacts. Also, as we have seen, the resolved wavelength depends not only on the number of views but also on the object itself and so to apply an appropriate cut-off filter we need some *a priori* knowledge of the object to be reconstructed. For simple familiar cases such as jet flames, such prior knowledge is intuitively available; however, for new complex flame geometries it may not be and we could therefore justify only using the limiting case for ART, and could potentially remove much valid information from our reconstructions. By delaying any filtering until after reconstruction, we can incorporate the new information from the reconstruction to determine a more appropriate filtering wavelength.

These are compelling reasons to avoid filtering prior to reconstruction. Nevertheless as stated using a coarser grid with implicit or explicit filtering does make the problem more determined and so potentially more accurate. To clarify the difference, if any, between filtering *before* or *after* reconstruction we conduct a test using a 2D section of the TOJ Phantom ($z = 5$ shown in fig. 6.35). We define three cases: filtering the reconstruction, i.e. *after*; filtering the views prior to reconstruction, i.e. *before*; and implicitly filtering by reconstructing directly onto a coarser grid. For the first case the reconstruction is performed on a $83 \times 83$ grid and then wavenumbers above a certain cut-off are filtered using a sinc function. For the second case the same grid and filter are used. For the third case the views are not filtered and the reconstruction is performed on a coarser grid that applies equivalent filtering based on the Nyquist criterion. For example, for a cut-off wavenumber of 20 a grid of $40 \times 40$ would be used. These cases are performed for a range of cut-off wavenumbers and compared to the Phantom filtered with the same sinc function. The coarser gird is up-sampled using linear interpolation to match the grid of the filtered Phantom. The results are shown in fig. 6.36.

For highest correlation with the Phantom it appears to be best to reconstruct at the full view sampling resolution, and then filter if desired. At higher cut-off wavenumbers this improvement is only small, but with strong filtering this difference is more pronounced. While the reconstructions filtered *before* and *after* both see improving correlation with more filtering, the coarser grid solutions actually become worse. Interestingly the reduc-

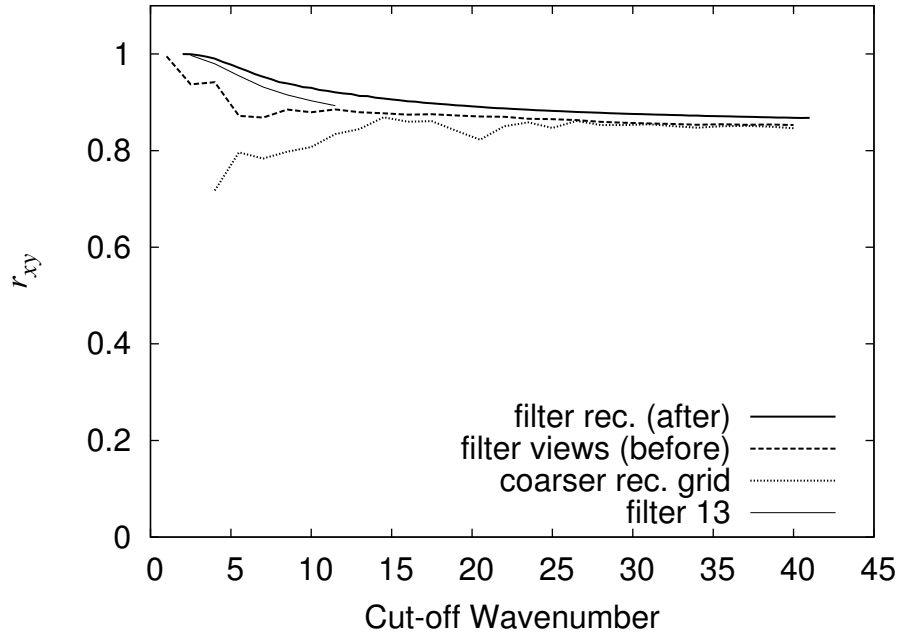



Figure 6.36: The correlation of the Phantom and its reconstruction using different filtering approaches at a range of cut-off wavenumbers.

tion in $r_{xy}$ begins at a cut-off wavenumber of 14 which coincides with when the problem is adequately determined, i.e. 830 unknowns to 830 equations. The convergence of coarser grids after this point was erratic and could result from the over-determined conditions. Reconstructions with less filtering or finer grids have shown more consistent convergence; this could be because the problems are less determined and so change less between iterations making catching the correct convergence point easier.

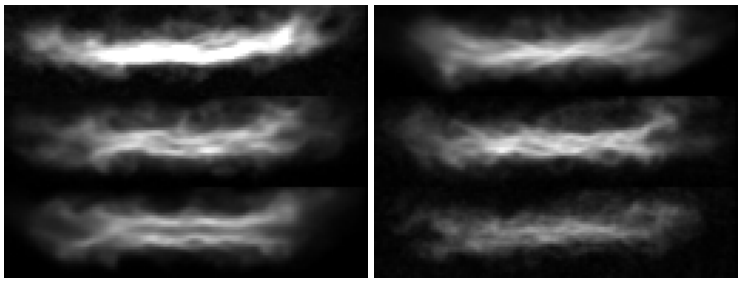
It is possible that this employed test may favour filtering reconstructions afterwards, as the Phantom is similarly filtered. To try to avoid this possible bias the coarse grid reconstructions, and those filtered before, also had a *neutral* filter applied, i.e. one at the same cut-off wavenumber of the Phantom, but this only reduced correlation and so was not used. To match the grid of the coarse grid solution, resampling of the Phantom, the *before* reconstruction, and the *after* reconstruction was also tried but did not result in any significant change. The chosen filter may also have an effect on the results, particularly at low cut-off wavenumbers and it is possible that a more appropriate filter could be found in future work.

These results support reconstruction with unfiltered views in this noise free 2D case; however, with limited computational resources and also noise in the views some filtering and down-sampling of the views could be desirable. These results show that this is possible without too much penalty provided the filtering isn't too strong. This is demonstrated in fig. 6.36 where we see a reconstruction filtered before, at a cut-off wavenumber

of 13, and then also filtered afterwards, which eventually achieves the same correlation as filtering the full view resolution reconstruction. In such a case it is better to filter and down-sample the views before rather than using a coarser grid, as the reduction in the number of projections would increase the speed of and stability of convergence.

### 6.5.3 Experimental Results

Using the PicSight camera/mirror system, arranged as per fig. 6.27, views of the premixed methane/air TOJ were measured. To give a range of SNRs, views were measured at three different equivalence ratios, 0.7, 0.8, 0.9, and four different exposure lengths: 2000 $\mu$s, 1000 $\mu$s, 500 $\mu$s, 250 $\mu$s. For each measured view a background image was also taken immediately after the measurement once the flame was extinguished, and this was deducted from the view prior to reconstruction. The views were cut to $\pm 60$ mm in $x$ and $\pm 15$ mm in $y$ from the registered burner centre to span the full object width and the full area between the opposed nozzles. To reduce the computational expense of reconstruction, each view was filtered and down-sampled to 128 pixels in $X$ and 32 in $Y$. An aperture of $f/1.6$ was used for all measurements, which provides sufficient depth-of-field for the down-sampled views. As an example, view 6 of each $\phi$ and exposure setting is shown in fig. 6.37. These views show that for the flow and mixing conditions used, except for an equivalence ratio of 0.9, the two flame brushes are barely distinguishable from each other.



(a) From top, $\phi = 0.7, 0.8, 0.9$ (b) From top, exposure = 2000, 500, 250 $\mu$s

Figure 6.37: View 6 of the TOJ at different exposures and equivalence ratios. Note each view has a different gain: (a) gain = 30, 24, 18 dB (exposure = 1000 $\mu$s); (b) gain = 18, 30, 33 dB ($\phi = 0.8$).

It is evident from fig. 6.37 that the chemiluminescence of the flame does not extend over the full width of the 120 mm object domain and so the slight truncation of views 3 and 4 should not affect the reconstruction. This is also suggested by fig. 6.38 which shows that the truncation of views 3 and 4 occurs only in the background noise region. For view 7 the truncation is more significant, and the portion of the flame intensity signal

beyond $x = -49$ mm is lost, though it appears only a weak signal occurs beyond this point and we can expect a limited effect. Incorporating truncated views into ART is simple, because a view does not have to span the whole object domain (though it is preferable). We can therefore avoid the uncertainty of extrapolating values, as the truncated regions can simply be ignored; for example, for view 7 we need consider only those projections whose $n$ index is in the range 15 – 128.

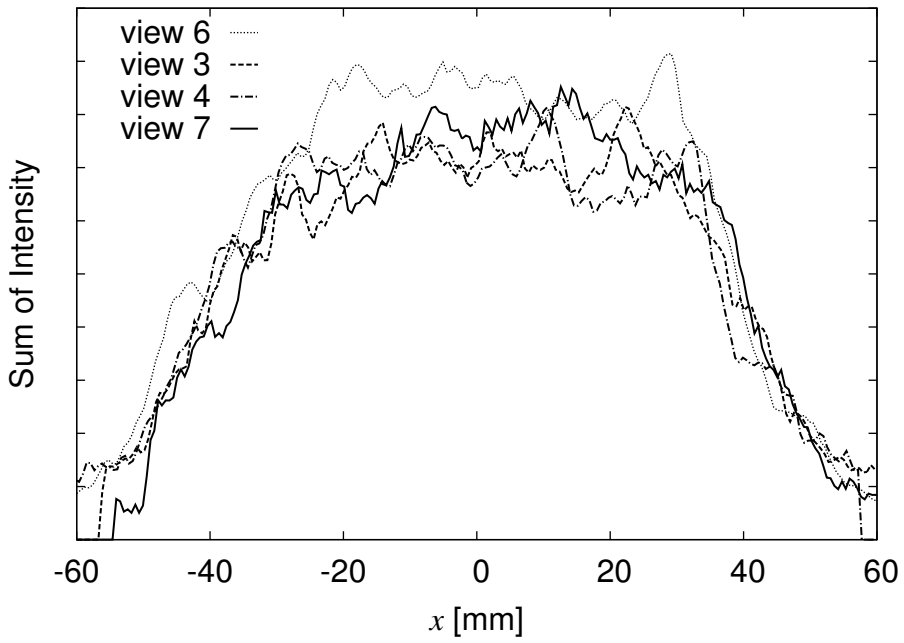



Figure 6.38: The column sum of the measured intensity ($\sum_n I_{mn}$) for the 3 truncated views: 3, 4, and 7. Also shown is an untruncated view 6 for comparison.

The experimental parameters for each set of measurements are summarised in table 6.7. The exposure and gain settings are the same as used in the SNR test of section 6.1.2.1 and so from the results of that section the SNRs of the measurements of this section have been estimated. This is achieved by determining the normalised signal RMS value of the view and then simply interpolating from figs. 6.2 and 6.4. A representative RMS value is taken from view 5 for each set of measurements, as little variation was observed between the views. The SNRs in section 6.1.2.1 were determined from test shots of bands of continuous grey; however, the actual flame measurements have varying signals and contain large *non-flame* areas. To to account for these differences, when calculating the representative RMS signal value of the measurements, only values greater than 25% of the peak signal value were used. The deduction of the background signal in this case would mean greater noise than measured in section 6.1.2.1, but only by a small amount for these relatively high SNRs. The measured views show the expected behaviour of having higher SNR for mixtures closer to stoichiometry, though interestingly the increase from

0.7 to 0.8 is quite small.

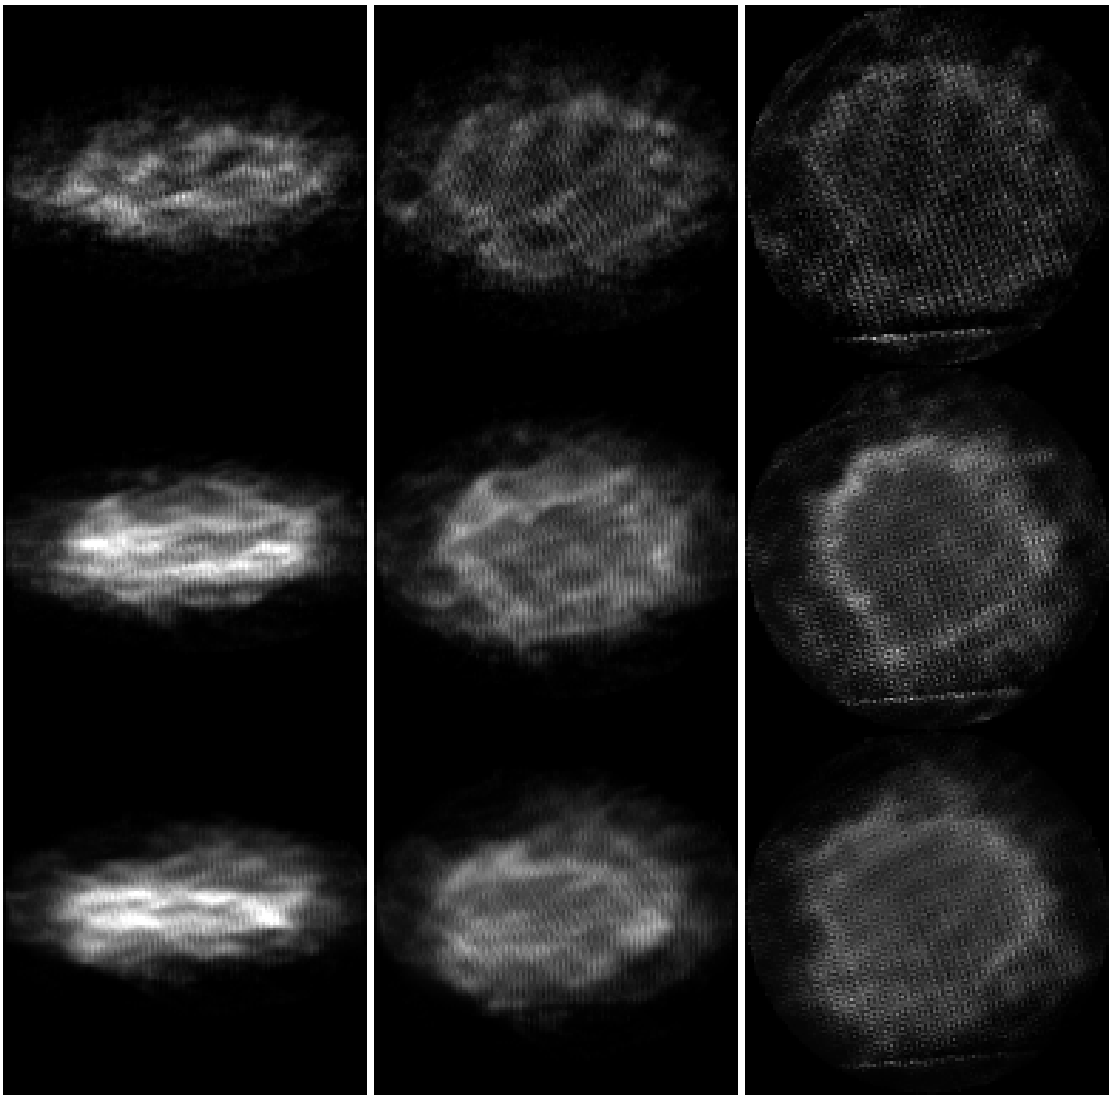
Table 6.7: Experimental Parameters

| $\phi$ | **Exposure [$\mu$s]** | **Gain [dB]** | $f/\#$ | **SNR [dB]** | $\Delta_c$ **for conv.** |
|---|---|---|---|---|---|
| 0.9 | 1000 | 18 | 1.6 | 27.3 | $10^{-5}$ |
| 0.8 | 1000 | 24 | 1.6 | 25.7 | $10^{-5}$ |
| 0.7 | 1000 | 30 | 1.6 | 25.5 | $10^{-5}$ |
| 0.8 | 2000 | 18 | 1.6 | 28.2 | $10^{-5}$ |
| 0.8 | 500 | 30 | 1.6 | 22.4 | $10^{-5}$ |
| 0.8 | 250 | 33 | 1.6 | 17.8 | $10^{-5}$ |

Each set of views was reconstructed on a grid of $N_i = N_j = 128$ and $N_k = 32$ cells using the ART algorithm and cylindrical parallel projections. A value for $\Delta_c$ of $10^{-5}$ was used for all reconstructions as this was found to yield the best results. This is consistent with the earlier Matrix Burner reconstructions where an $\Delta_c$ of $10^{-5}$ was best for SNRs under 35 dB. Indeed one might expect $10^{-6}$ to be preferable for the lower view SNRs here; however, this value was tried and not found to improve the reconstructions. In figs. 6.39 – 6.42 volume renderings and sections of the TOJ reconstructions are shown.

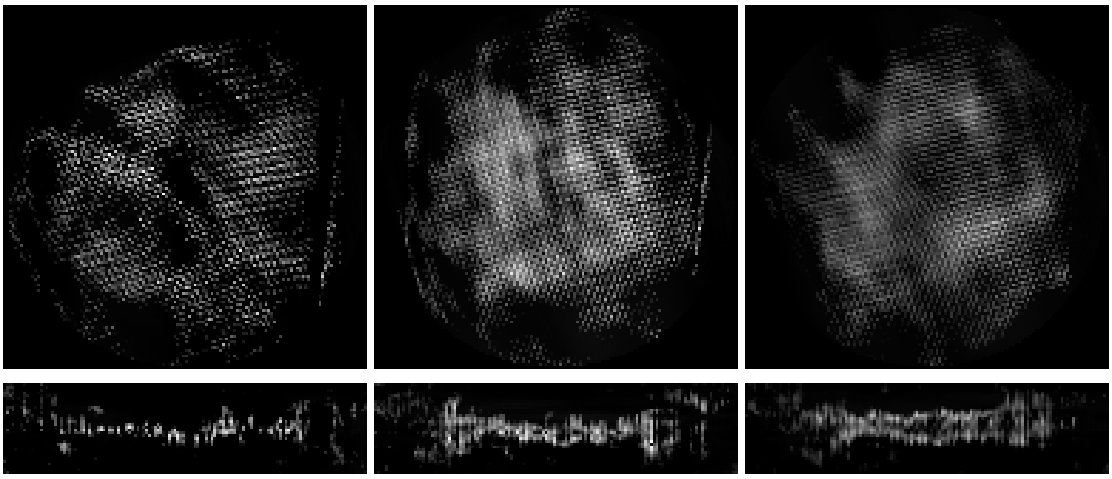
The TOJ reconstructions have all captured the larger features of the flame quite well, and all renderings show good qualitative agreement with the flame observed during measurements. For example, the higher intensity ring towards the edge of the circular flame when viewed from above, as seen in fig. 6.39, was also observed in the actual flame. This corona is related to the familiar curling of the flame at its edges which we again observe with the reconstructions, for example fig. 6.42. We can also see that as with the actual TOJ flame the reconstruction is relatively flat and smooth near the burner centre line and becomes more chaotic as it spreads to the edges. The reconstruction of the 0.7 equivalence ratio flame appears to have captured some local extinction, as dark holes in the flame surface are observed, whereas for the measurements with $\phi = 0.8$ and 0.9 the flame is unbroken in this region. The reconstruction at $\phi = 0.7$ also shows a much thinner flame brush than for $\phi = 0.8$, and for 0.9 the flames are visibly distinct as previously observed. The 0.7 reconstruction is also noisier than the other reconstructions as expected due to its lower SNR. (We remind the reader that the different gains mean the intensities are not directly comparable.)

In all reconstructions a thick line like artefact can be observed at one side of the flame, for example as seen in figs. 6.39c and 6.41c. This artefact corresponds to the truncation point of view 7. The precise reason for the artefact is unknown, though as it disappears when the full width of view 7 is used (i.e. $n$ spans 1 – 128), it is almost certainly



(a) $\gamma = 20^\circ$ (b) $\gamma = 40^\circ$ (c) $\gamma = 89^\circ$

Figure 6.39: Volume renderings of the TOJ at different equivalence ratios. From top: $\phi = 0.7$, $\phi = 0.8$; $\phi = 0.9$.



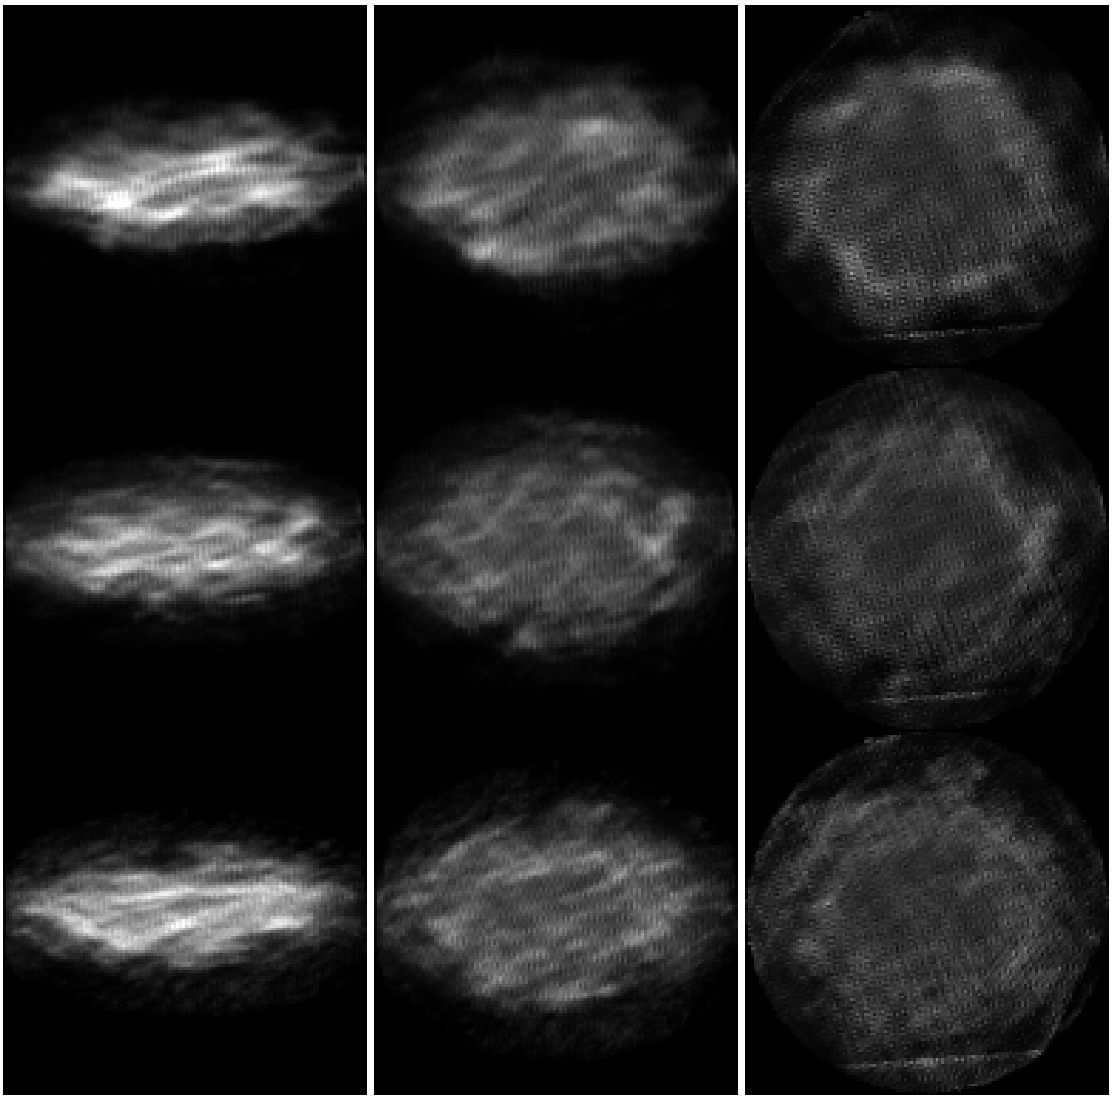
(a) $\phi = 0.7$ (b) $\phi = 0.8$ (c) $\phi = 0.9$

Figure 6.40: Selected mid-sections (upper, $z = 16$; lower, $y = 64$) of the TOJ reconstruction for different equivalence ratios.

related to the view truncation. It is thought likely that this line artefact is the result of the locally reduced determinacy[11] in the truncated region, as this region effectively has 1 less view than the rest of the object domain. This locally reduced determinacy then allows something akin to a *gathering* of *surplus* intensity from the more determined regions. The very edges and corners of the reconstructions have similarly reduced determinacy, and such gathering is frequently observed in these areas; for example, at the edges of the 250 $\mu$s reconstruction seen in fig. 6.41c. In this case we could expect that with more views and more even spacing any truncated regions would be better determined and any line artefacts would be reduced. In spite of the artefact, truncating the view is still preferable to using estimated intensity values for the obscured projections because that approach affects the whole reconstruction rather than just the obscured region.

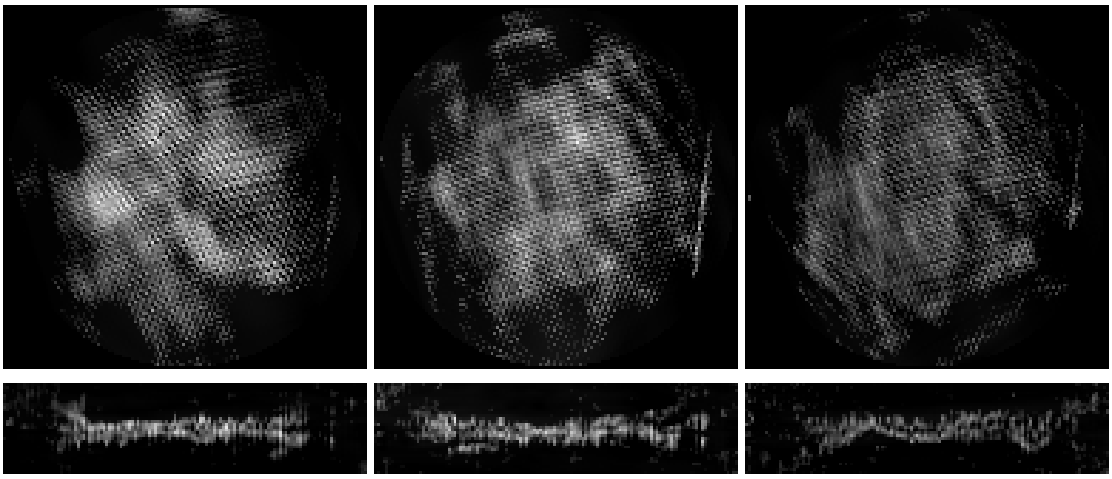
All reconstructions feature small noise like artefacts that are discernible in figs. 6.39 – 6.42. This is not random noise however, because on close inspection regular patterns are observed. These patterns are quite like the *salt-and-pepper* noise discussed previously in section 2.2.3, which implies that they result from inaccuracies in the projection geometry. We know the cylindrical projection geometry used here to be a simplification and, as previously mentioned, by using a more exact extruded cuboid type projection geometry (or better) we could expect to remove this salt-and-pepper noise. We note that this kind of noise has not been observed in the Phantom studies because the same projection geometry is used to both simulate the views and perform the reconstruction. Salt-and-pepper noise was also not seen in the previous Matrix Burner results which used the same projection geometry; this is possibly because the small spatial extent of these flames makes it less

[11] As in a system of linear equations, as explained in section 2.2.2.2.



(a) $\gamma = 20^\circ$ (b) $\gamma = 40^\circ$ (c) $\gamma = 89^\circ$

Figure 6.41: Volume renderings of the TOJ for different exposure times. From top the exposures are: 2000, 500, 250 $\mu$s.



(a) 2000 $\mu$s (b) 500 $\mu$s (c) 250 $\mu$s

Figure 6.42: Selected mid-sections (upper, $z = 16$; lower, $y = 64$) of the TOJ reconstruction for different view exposures.

obvious. However, re-inspection of fig. 6.15b suggests that some salt-and-pepper noise is present in that $\gamma$ corrected reconstruction. For these TOJ reconstructions salt-and-pepper noise has only been observed after many iterations, and was not present when an $\Delta_c$ of $10^{-3}$ or $10^{-4}$ was used. This suggests that *salt and pepper* noise may have been observed in the Matrix Burner results if a lower $\Delta_c$ and so more iterations had been used. In these TOJ results the salt-and-pepper noise is small and has had very little effect on the quality of the reconstructions, with the exception perhaps of the $\phi = 0.7$ reconstruction which has a thin and sparse flame. Indeed, salt-and-pepper noise aside, the SNR of the reconstructions themselves appear quite high, because the reconstruction from the lowest SNR views (60 for 250 $\mu$s) is not discernibly more noisy than the reconstruction using the highest SNR views (537 for 2000 $\mu$s). This suggests that the filtering of the views (performed to reduce the computational expense) has removed much of the noise. It is possible therefore, that reconstructions of the TOJ using even shorter exposures can be obtained by greater filtering/pixel-binning[12] of the views. Or even without any filtering if we can accept an increase in the reconstruction error.

### 6.5.4 Conclusions

We have successfully reconstructed a turbulent flame—a premixed Turbulent Opposed Jet—using 10 unequally spaced views, obtained simultaneously, from 5 PicSight cameras and a system of mirrors.

Unequal angular spacing was necessary to avoid the burner supporting structure severely

[12]Combining the results of 2 or more adjacent pixels. This can even be done *on-chip*.

obscuring 4 views. The effect of unequal angles on the reconstruction was investigated using a realistic LES Phantom generated by Stein [142]. From this Phantom study we found that this particular non-uniform view spacing leads to a non-uniform reconstruction and a larger resolved wavelength: 25% of the object domain width (120 mm). However, due to the non-uniformity some regions of the reconstructions were as well resolved as the equal angle case, which uniformly resolved wavelengths 13% of the object domain. The normalised RMS error was more moderately affected by the unequal spacing: increasing by just 1.6% of the Phantom maximum value.

The particular view arrangement used was chosen to minimise the correlation between views for the mirror system employed and so obtain the most uniform reconstruction possible. It is possible that a more optimum arrangement exists, particularly if mirrors are not used as this affords much greater flexibility in the view locations. In the future, the optimum arrangement could be relatively easily identified using a Phantom study, which could also provide a *map* that identifies areas of higher and lower resolution/error—perhaps using a checkerboard Phantom. It would also be possible to redesign the TOJ burner support structure to allow equally spaced views to be obtained by moving the supports to a larger radius for example.

It was not possible to locate the mount of views 9 and 10 in the same plane as the remaining mounts, and by assuming co-planar views, a systematic error was introduced. An LES Phantom study showed this $1^\circ$ error in $\gamma$ results only in a moderate increase in the normalised RMS error of 0.01, and has only a minimal effect on the resolution. The registration approach used returned only $\theta$—hence the co-planar assumption. However, we note that with an alternate method, such as the CV approach, $\gamma$ can be determined and this source of error removed.

We have also used the LES Phantom to determine whether it is better to reconstruct using the full resolution of the measured views, and so exploit the maximum amount of information from the views, or to implicitly/explicitly filter and downsample the views to make the problem more determined. The study has revealed that strictly speaking it is better to reconstruct using the full view resolution and then filter to a lower resolution if desired. However, with mild filtering of the views the difference is small and so can be preferred for the improvements in reconstruction speed and the reduction in view noise.

Views of the TOJ have been measured using a range of exposures and 3 different equivalence ratios, and reconstructed using ART. The shape of the TOJ has been well captured in all reconstructions and several familiar features can be identified; for example, the reconstruction of the 0.7 equivalence ratio flame has captured areas of apparent local extinction. All reconstructions have displayed a small amount of *salt and pepper* noise that result from inaccuracies in the cylindrical projection geometries. This was not

observed in the earlier Matrix Burner reconstructions because the flame shape was not amiable to its detection. Its effect here is small however and with more precise geometries, for example an extruded cuboid, we could expect this noise to disappear.

The SNR of the views has been estimated from the test shots of section 6.1.2.1, and are found to range from 60 – 537, which is much less than obtained with the methane oxygen Matrix Burner flame. Nevertheless the effect of measurement noise was not really discernible in the reconstructions because the views were filtered and resampled from approximately 210 to 128 pixels wide. This filtering effectively increases the SNR of the views and explains why even for the shortest exposure (250 $\mu$s) the reconstruction contains little apparent noise. This suggests that with pixel-binning much lower exposures are possible, as the SNR is seen to approximately halve with a similar halving of the exposure time—suggesting an exposure of 60 $\mu$s would have an SNR of approximately 15 without any binning or filtering. Even shorter exposure times would be possible with *brighter* higher hydrocarbon fuels such as propane or ethylene. Indeed test images of the TOJ burning ethylene have been taken with exposures as low as 31 $\mu$s and achieved good SNR without any binning. We note however, that the TOJ flame provides a much stronger signal than many other flames by virtue of its orientation with respect to the cameras. Nevertheless these SNRs achieved do further suggest that very high time resolution are possible for chemiluminescence measurements of turbulent flames using the PicSight cameras.

These results have demonstrated the robustness of ART and the CTC technique, because despite using only ten poorly spaced, mirror-doubled, views—of which 5 views had either systematic location error or are obscured to some degree—the reconstructions are remarkably good.

# Chapter 7

# A Novel Application for CTC

So far we have focused on the development of the CTC sensor; in this chapter we turn our attention to a new potential application of the CTC sensor: the measurement of the flame front local strain rate. This is based on a hypothesised relationship between the integral of chemiluminescence and the strain rate in a non-premixed flame. Because the strain rate can be related to the scalar dissipation rate by making some simplifying assumptions, we can also provide estimates of this difficult to measure quantity.

We begin by presenting the speculated relationship between strain rate and local integral chemiluminescence and how this can be coupled to CTC. Following this, evidence is presented from simulation and experimentation that supports, but does not prove, the existence of such a relationship.

## 7.1 The Measurement of Flame Local Strain Rate

Local strain is an important feature of turbulent flows, particularly with respect to extinction. Furthermore, the strain rate can be related to the scalar dissipation rate [163] which is a central quantity in many non-premixed combustion modelling approaches that incorporate finite rate chemistry, e.g. Laminar Flamelets [168, 121], Conditional Moment Closure [86, 20] and Transported PDF [126, 127] methods. The scalar dissipation rate has been studied experimentally [27, 79, 78, 144, 145, 51]; however, its robust measurement in turbulent reacting flows has proven very difficult [21, 53, 162, 151]. This is because the scalar gradients on which $\chi$ depends (see eq. 2.7) occur at the smallest spatial and temporal scales of turbulence, and require very high resolutions to be properly captured [21, 162]. For example, Karpetis and Barlow [79], and Geyer et al. [51] quote resolutions of 300 $\mu$m and 350 $\mu$m respectively in their scalar dissipation measurements. The small probe volumes required mean weak signals are encountered and noise is a considerable problem. Measurement is made even more difficult by the 3D nature

of turbulent flames, which means all three components of the gradient vector must be considered simultaneously.

Here we investigate a relationship between the strain rate and the local *integral* measurement of chemiluminescence across the flame normal, $I_\omega$. This relation can then be used to estimate the local, conditional average of the scalar dissipation rate close to the stoichiometric mixture fraction[1]. We define $I_\omega$ in eq. 7.1 where $N_B$ is the number of light emitting molecules of species $B$ and $\omega$ is the flame normal coordinate.

$$I_\omega = \int N_B \, d\omega \tag{7.1}$$

By considering an integral measurement we avoid the very high resolutions that are necessary to resolve the scalar gradients. By using CTC, the local, 3D $I_\omega$ can be measured for arbitrary flame configurations, for which the direct measurement would not be possible (see fig. 7.1).

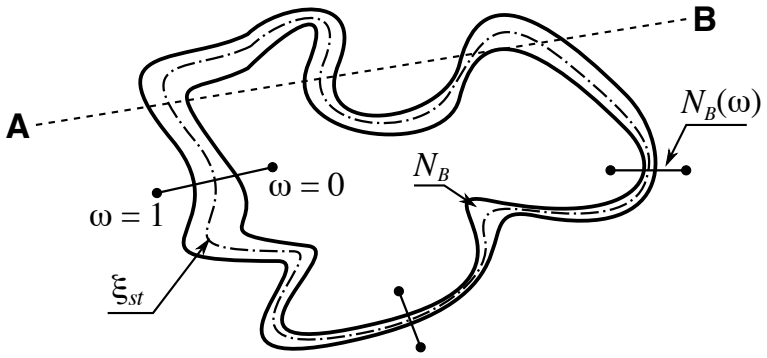



Figure 7.1: Illustration of how the local flame normal integral of chemiluminescence can be obtained once the spatial distribution is known. The spatial distribution of chemiluminescence can be provided for arbitrary flames by CTC. The line **AB** illustrates how line-of-sight measurements alone cannot always provide the local flame normal integral of chemiluminescence. The flame normal direction is $\omega$, $\xi_{st}$ is the stoichiometric mixture fraction, and $N_B$ is the number of emitting molecules of species B.

It is interesting to note that for small flames ($< 20$ mm) and with enough views (30+), the CTC resolution could exceed the $300\,\mu$m [79, 51] used in previous studies to resolve the scalar gradients. Although in CTC we measure only a single dependent scalar, whose gradient is not equal to the scalar/mixture fraction gradient, it is true that the chemiluminescence gradients are to some degree dependent on the mixture fraction gradients. If a suitable relationship could be found—perhaps based on the laminar flamelet assumption—then we could directly estimate $\chi$ from CTC measurements. The direct res-

[1]For most applications the scalar dissipation rate close to stoichiometry would be of most interest

olution of the scalar gradients is not, however, the focus of this work and we continue, considering the integral relationship only.

The proposed relationship between $I_\omega$ and the local strain rate is based on the concept of decomposing turbulent flames into a series of laminar flames: the commonly employed laminar flamelet model (see section 2.3.2.3). By measuring $I_\omega$ for a series of strained laminar reference flames, for example in a opposed jet configuration, we can obtain a *look-up* table, in terms of $N_B$ from which the local strain in the turbulent flame can be estimated. Turbulent flames are unsteady, and the strained reference flames should also, strictly speaking, be unsteady. However, steady laminar flamelets may provide a sufficient, and simpler, representation. The strained reference flame used in this work is a Laminar Opposed Jet flame (LOJ). In such a flow, the bulk strain $a_b$ is directly related to $\chi$ because the transport of the scalar is determined by the velocity field and the strain is, hence, proportional to $\nabla\xi$. Therefore, assuming laminar flamelets also allows us to estimate $\chi$ in turbulent flames.

Other studies support a (non-integral) relationship between chemiluminescence and strain, such as that of DeLeo et al. [35] who show an increase in peak $CH^*$ concentration from $20\,s^{-1}$ to $40\,s^{-1}$, though the authors note uncertainty in this result. Hardalupas and Orain [63] also report some variation between non-integral $CH^*$ and strain in their premixed laminar opposed jet, and also highlight the strong dependence on strain of the $C_2^*/OH^*$ and $C_2^*/CH^*$ ratios reported by Angermeier [7]. In their recent chemical simulations Panoutsos et al. [120] also report a dependence on strain of $CH^*$ and $OH^*$ chemiluminescence for a non-premixed methane-air flame.

To reiterate, using CTC we can obtain the local, integral measurements of chemiluminescence, in 3D, for turbulent flames. From these measurements we can potentially determine the local flame strain rate by making the laminar flamelet assumption. From the strain rate, the average local scalar dissipation rate can be estimated, and because chemiluminescence occurs in a narrow mixture fraction range, this average is conditioned on the flame front and centred on the stoichiometric mixture fraction $\xi_{st}$.

### 7.1.1 Laminar Opposed Jet Simulation

The relationship between the flame normal integral of chemiluminescence and the strain rate has been investigated using data from 1D strained steady laminar flame simulations performed by Gkagkas [57]. These non-premixed simulations were performed using the *in-house* code of Lindstedt and co-workers and are the same as those used previously in section 4.1 to provide the excited state chemistry *look-up* tables. The simulations were available for a range of strain rates, and provide the flame normal concentration profiles of several excited state species in terms of their mole-fraction, $X_B$ (where $B$ is the species).

The amount of emitted light is a function of the mole fraction and therefore, by integrating these profiles over $\omega$ the relationship between the strain rate and $I_\omega$ can be assessed.

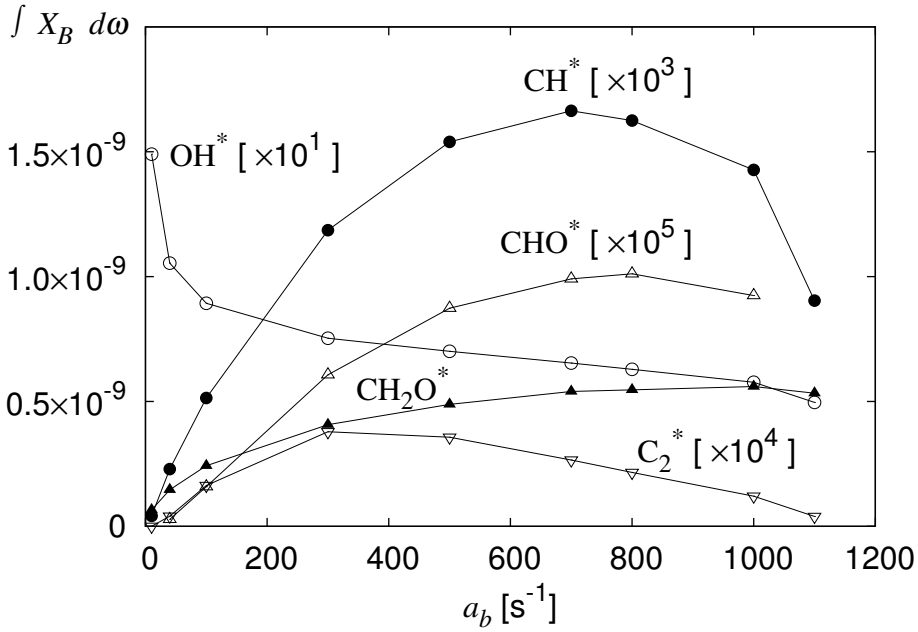



Figure 7.2: The integral of the mole fraction $X_B$ across the flame normal $\omega$ plotted against the bulk strain rate $a_b$ for the excited species $B$: $OH^*$, $CH^*$, $CHO^*$, $CH_2O^*$, $C_2^*$. The data results from the 1D solution off a Laminar Opposed Jet by Gkagkas [57] presented previously in section 4.1.

The flame normal integral $\int_0^1 X_B d\omega$ of mole fraction is shown in fig. 7.2, at different bulk strain rates, for the excited species in table 2.1. For all species the integral varies with bulk strain. For $OH^*$ this relationship appears monotonic and the integral decreases with increasing strain, this contrasts the $CH^*$, $C_2^*$, $CHO^*$ and $CH_2O^*$ profiles which increase initially but are not strictly monotonic. Monotonicity is not essential where more than one species can be measured or prior knowledge of the strain rate range can be employed. The magnitudes of the $OH^*$ and $CH_2O^*$ integrals in fig. 7.2 are substantially higher than the other species, and so could be better selections for high temporal resolution measurements.

The difference of the $OH^*$ profile in comparison to the other species, prompts suspicion of this result. We would expect that an increase in scalar dissipation rate leads to (at least initially before extinction) an increase in the reaction rate and higher chemiluminescence. While observed with the other species, this is not the case for $OH^*$. As noted by Docquier and Candel [37], the formation and destruction reactions for many excited radicals are not yet completely understood. Therefore, we can expect a degree of uncertainty in these simulated concentrations. Also the kinetic parameters for reactions

involving $OH^*$, $C_2^*$, $CHO^*$ and $CH_2O^*$, which were taken from the study of Benvenutti et al. [17], had to be estimated by those authors as they not available in the literature. This prompts most confidence in the $CH^*$ profile.

### 7.1.2 Laminar Opposed Jet Experiment

The relationship between the flame normal integral of chemiluminescence and the scalar dissipation rate is further tested using experimental measurements of a non-premixed Laminar Opposed Jet.

The nozzle set-up is that used by Luff et al. [99] and Lindstedt et al. [96], with the turbulence generating plates removed and a nitrogen co-flow added to allow laminar diffusion flames to be studied. The pipes leading to the 25 mm diameter ($D$) nozzles are shaped to maintain a laminar profile at relatively high Reynolds numbers. The fuel stream consists of a partially premixed mixture of 25% natural gas and 75% air by volume. A single cooled 12 bit CCD of 1376 by 1040 pixels from a LaVision Flowmaster system was fitted with a 50 mm lens and a interference band-pass filter of 45% efficiency at wavelength $430\pm10$ nm to capture the light emitted by $CH^*$ in the $A^2\Delta - X^2\Pi$ transition. As illustrated in fig. 7.3, the camera was mounted at an angle ($\gamma$) of 19° to the horizontal, which was as close to perpendicular as was possible without the nozzle obscuring the flame. Results at different radial locations ($\delta$) show the radial profile to be sufficiently homogeneous that the measurements could be simply scaled by $\sin(\gamma)$ to get the equivalent flame normal integral. It would be possible to use CTC to obtain $I_\omega$, but at the time of this experiment, the CTC set-up was not available.

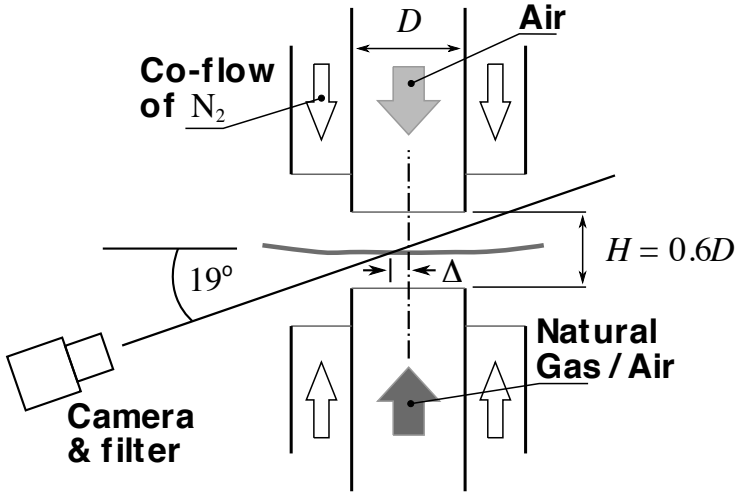



Figure 7.3: A schematic drawing of the experimental set-up used to measure the integral chemiluminescence of the Laminar Opposed Jet.

The integral of the emitted light was measured at different bulk strain rates $a_b = (u_{ox} -$

$u_f)/H$, where $H$ is the nozzle separation of $0.6D$, and $u_{ox}$ and $u_f$ are the magnitude of the velocities in the oxidiser and fuel stream respectively (measured at the jet exit). The momentum fluxes of the fuel and oxidiser streams were balanced, and the velocity range of 1.01 to 2.16 $\mathrm{ms}^{-1}$ in the fuel stream corresponds to a range of 78 to 168 $\mathrm{s}^{-1}$ in the bulk strain (and Reynolds numbers ranging 1577 to 3380 in the oxidiser stream). Ideally measurements at higher strain rates would have been made, but it was not possible to reduce $H$ further without compromising optical access, and higher velocities would have lead to turbulent flows. Measurements were taken at an exposure time of 30 ms, and the results presented are averaged over thirty exposures. In this set-up the CCD provided integral line-of-sight intensity values, which were spatially averaged over an interrogation window of 0.64 by 0.16 mm in the CCD plane, and are presented in fig. 7.4.

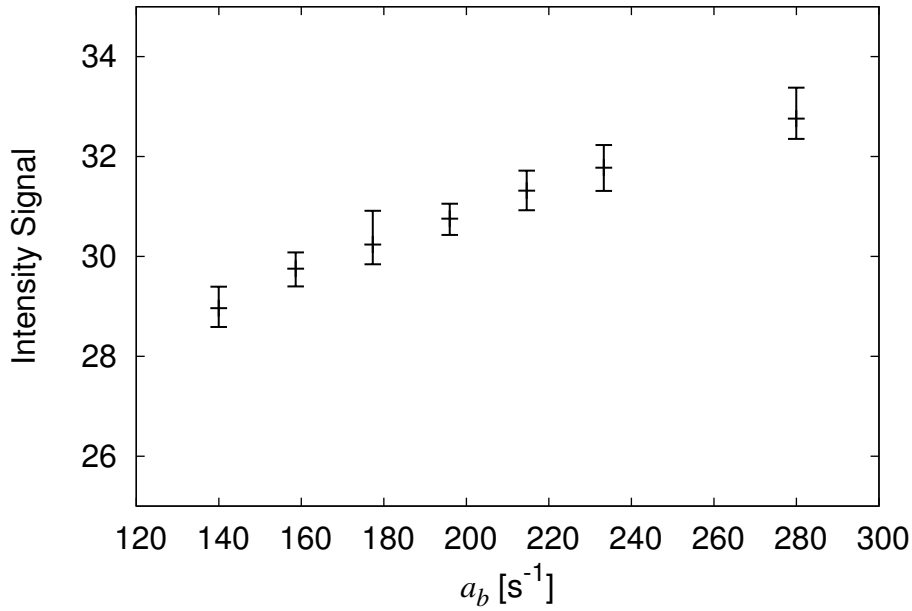



Figure 7.4: The measured intensity of the 430 nm wavelength light, emitted by decaying $CH^*$ radicals, against the bulk strain rate. Intensity values were taken 2 mm from the centre line of the Laminar Opposed Jet Burner. The mean of 30 samples is shown, with the (error) bars representing the maximum and minimum of the 30 samples.

Figure 7.4 shows an increase in the integral of the $CH^*$ emitted light with increasing bulk strain rate, as was also seen in the simulation (fig. 7.2). However, there appears to be a significant difference in the gradient of the simulated and measured results. The measured intensity varies by approximately 12%, whereas the integral of the mole fraction from the simulation data varied by almost 200% over the same range. For an optically thin flame, the light intensity, as measured here, is proportional to the number of *emitting* excited molecules, whereas our simulation returns the excited state concentrations whether

emitting or not. Because some of the excited state molecules return to the ground state by quenching [157], not emission, we can thus expect some differences in our intensity measurements and the simulation concentrations. Furthermore, the simulation considered methane and the measurements were made for natural gas. Nevertheless the difference in gradient is quite large and suggests further work is needed to fully understand the link between the integral of chemiluminescence and strain rate.

### 7.1.3 Discussion and Conclusions

The measurement of the 3D, local, flame normal strain rate, and by extension the spatially averaged scalar dissipation rate, is an exciting potential application for CTC given the importance of these parameters and the difficulty in measuring them using other techniques. The ability to measure these quantities is based on a postulated relationship between the flame normal integral of chemiluminescence and the strain rate. Such a relationship has been investigated using mole-fraction data from a 1D LOJ simulation, and chemiluminescence intensity measurements of a LOJ. The measured results have shown that the integral of chemiluminescence varies with strain rate, and the simulation has shown that the integral of excited species mole fraction also varies with strain rate.

The simulation results suggest, that the relationship may not be monotonic (though for $OH^*$ it may be). However, monotonicity is not essential when two species can be measured simultaneously. This could be practically achieved without requiring any more cameras by using a system of mirrors as used previously in chapter 6. By inserting an appropriate filter between the mirrors, for example using 431 nm for the left sub-system and 516 nm for the right sub-system, a single camera could image two species: the left image would be $CH^*$ and the right $C_2^*$ in our example. Furthermore, such a system could be more compact, and so better for packaging, than seen previously, as the angular spacing of the mirrors would not need to be the same as the overall view spacing.

The results shown here are for steady laminar flames. Turbulent flames are unsteady and so using a relation based on the steady laminar flamelets assumption will always be an approximation. Obtaining a suitable relationship based on unsteady flamelets would be more complicated, indeed if it is even possible. However, we note that with CTC and appropriate cameras, it is possible to take several measurements in quick succession and so, potentially, capture the time dependence of $I_\omega$. Steady laminar flamelets are also used successfully in simulations, but whether they can provide a sufficient approximation to allow the measurement of strain in more general turbulent flows is unclear, and requires further work. The disparity between the gradients of the simulated and measured results also suggests more testing is required to prove that usable strain-integral chemiluminescence relationship exists.

# Chapter 8

# Conclusion

This chapter first summarises the work undertaken in this project, before the main conclusions are reviewed with respect to the principle issues of this work; for example, the resolution achievable with CTC. Finally a list of specific recommendations is given to aid future work.

## 8.1 Project Summary

Time resolved, 3D experimental data of turbulent flames is required to both further the fundamental understanding of combustion and to support the development of advanced simulation techniques. This project has been concerned with the development of a new measurement approach called CTC that returns the 3D chemiluminescence intensity profile of a (non-sooting) turbulent flame, and has the potential to also provide 3D measurements of excited species concentrations, heat release rate, equivalence ratio, and strain rate. In this work we have taken measurements using industrial machine vision type cameras which are affordable enough for multiple deployment; this avoids the resolution limitations traditionally imposed by the cost of specialist scientific cameras.

The CTC sensor is based on Computed Tomography and two algorithms from the ART family (additive ART and MART) have been implemented in a custom Fortran95 code. The ART algorithms performance in under-determined configurations has been tested using realistic *Phantoms* derived from LES—updated to include prediction of $CH^*$ and $OH^*$. The algorithms were able to achieve low error reconstructions even in very under-determined noisy configurations; for example, only 20 – 30 angles were necessary for the $128 \times 128$ pixel reconstructions to be free of significant artefacts. Views with SNRs of over 1000 produced negligible increase in the reconstruction error (compared to the no-noise case), and even with an SNR of 100 only a 3% increase in error was observed. The additive form of ART was preferred in the remainder of the work for its

improved noise performance.

A study has also been performed using cosine Phantoms of different wavelengths to better understand—for ART—the spatial resolution achievable with limited numbers of views. The resolutions determined using this cosine assessment agree with a theoretical relationship derived using the Fourier slice theorem, except for the less radially extensive Phantoms which achieve higher resolutions. A proof-of-concept experiment has also been conducted using a single repositioned camera which provides 48 views of a steady, but non-axisymmetric, Matrix Burner flame. The reconstructions have resolved each of the flame fronts of the burner and, despite some deformation of the constituent flames at the edges, overall the reconstruction shows strong qualitative agreement with the observed flame.

Several modifications to the representation of the projections in ART have been investigated. Firstly, having established that self-absorption can cause error in $OH^*$ measurements, a modification that incorporates the Beer-Lambert law into ART has been studied using LES Phantoms, and found to successfully correct for self-absorption (for the case where the ground state field is the main absorbing field). Redefining the projections in 3D has permitted arbitrary camera locations and allows out-of-plane contributions to the projections to be included. The camera optics can result in out-of-plane contributions to the measured intensities, and we have used a geometric optics analysis to define a more accurate weighted double-cone projection geometry. This *optical* geometry permits shorter object distances and wider apertures to be used and so stronger signals and shorter time resolutions. The *optical* projection geometries are tested using measurements of the Matrix Burner made with the repositioned digital camera set to the maximum aperture of $f/3.5$ to minimise the depth-of-field. While the non-parallel rays successfully corrected for perspective distortion, little difference was observed between the cylindrical and *optical* projection geometries as the depth-of-field was not limited enough in this set-up, and also the consumer camera was not well represented by the single lens assumption of the *optical* projections.

The CTC sensor has been implemented using 5 PicSight P32M cameras and a mirror system to provide 10 simultaneous views and allow time resolved measurements. The selected cameras cost £800 each and so can be deployed in much greater numbers in future work. Using this set-up the Matrix Burner flame has been measured for a range of exposures with the shortest ($62\,\mu$s) achieving an SNR of approximately 1000. Each 3D reconstruction has captured the individual flames of the Matrix Burner, indicating a resolved wavelength of 2 mm. A study performed using a synthetic Matrix Burner Phantom has shown the $\pm 0.5^\circ$ uncertainty in the viewing angles has little effect on the reconstruction accuracy. As with the proof-of-concept experiment deformation of the outer flames

was observed, and a Phantom study has shown systematic uncertainty in $\gamma$ (the out-of-plane angle) not to be the cause. It has been concluded the deformations result from inaccuracies in the projection representation; possibly for example, due to beam steering or lens distortion; however, systematic registration error is also thought to significantly contribute. Registration using a Computer Vision (CV) approach has, been investigated and found to be at least as accurate as the previous bespoke methods, while also providing the 3D location and orientation of the cameras as well as internal parameters such as focal length, and lens distortion. Analysis of the distortion model suggests a systematic error exists when using the mirrors that could contribute to the deformations observed previously. The camera registration was necessarily performed without the mirrors, because the restricted volume they allow to be imaged was not sufficient for reliable calibration.

Measurements of a pre-mixed turbulent opposed jet have been performed using the 10 view, mirrored set-up but adjusted to account for the wider flame. A combination of obscuring supports and practical location difficulties meant the ideal equal angular spacing of views could not be achieved. A study using premixed LES FSD Phantoms provided by Stein [142] has shown that the unequal spacing of views results in non-uniform resolution, and an increase in the minimum resolved wavelength from of approximately $0.13D$ to $0.25D$. Measurements of the TOJ were taken for a range of exposures as low as 250 $\mu$s which achieved an SNR of approximately 60; although the reconstructions contained little apparent noise due to the filtering of views. The larger features of the TOJ flames were captured in all the reconstructions, including extinction regions observed in the $\phi = 0.7$ case. All reconstructions featured mild *salt and pepper* noise, which is thought to be the result of the inaccuracy of the cylindrical projection geometry.

Finally, a new potential application for CTC has been presented, whereby the reconstructed 3D chemiluminescence field in turn provides the 3D, local, flame front averaged, scalar dissipation rate. This is based on a postulated relationship between the integral of chemiluminescence across the flame front normal and the local strain rate. Analysis of the LOJ simulation performed by Gkagkas [57]—which features excited state chemistry—suggests that such a relationship exists, but is not monotonic except for $OH^*$. Experimental measurements of a laminar opposed jet showed a relatively linear relationship between the integral of $CH^*$ chemiluminescence and strain, though the range of strain was quite limited. The disagreement in the gradients of the measured and simulated results, and the narrow range of parameters considered, suggests more work is needed to establish this link.

## 8.2 Review of Main Conclusions

The Computed Tomography of Chemiluminescence sensor has been successfully implemented using commodity CCD cameras, and used to obtain the 3D, time resolved, chemiluminescence profile of a turbulent opposed jet flame. Previous work by Ishino and Ohiwa [73] has achieved, time resolved, 3D reconstructions of flame chemiluminescence using bespoke film cameras, but only for a weakly turbulent flame and by stacking 2D sections. As far as the author is aware this is the only work that has considered high resolution 3D (without stacking) profiles of chemiluminescence, for turbulent flames, using an implementation based on multiple affordable CCDs. Sufficient resources were available for five cameras, and the time resolved CTC implementation uses only 10 viewing angles—resulting in spatial resolutions that are lower than desirable for the study of turbulent flames. The resolution required however, does depend on the application and for qualitative LES validation even these low resolution results would be useful, as the results have captured the larger scales of the TOJ flame and several familiar features can be identified; for example, areas of apparent local extinction in the $\phi = 0.7$ reconstruction. The TOJ results also demonstrate the robustness of ART and CTC, as the reconstructions still resolve the main features of the flame despite using only 10 poorly spaced views.

Nevertheless, CTC is intended as a high resolution sensor and so perhaps most importantly the Matrix Burner and TOJ results have demonstrated that affordable commodity cameras can be used in time resolved CTC. The view resolution achieved with the PicSight cameras (120 $\mu$s) exceeded that of even the the highly resolved LES of Stein [142] (260 $\mu$s) and without mirrors would be even higher. In future work, it is simply a matter of scaling up the experiment to obtain the higher resolutions, because both the LES Phantom studies and the 48-view Matrix Burner results have shown, with more viewing angles high spatial resolutions—exceeding that of LES—are possible.

To better understand the relationship between the numbers of angles and the resolution achievable with ART we have conducted an analysis using cosine Phantoms. Generally this analysis agrees with the theoretical relationship obtained using the Fourier slice theorem; except for those Phantoms we have termed *radially limited* (Phantoms 256OC and 256SU) which have resolved much shorter wavelengths than that predicted by the theory. The higher resolutions of the *radially limited* Phantoms agree with the relation given by Frieder and Herman [69], the resolutions obtained in previous work [14, 56, 141, 150], and with the Matrix Burner reconstructions which resolve a wavelength of $0.1D$ using 10 angles and $0.01D$ using 48. We have concluded that the theoretical relationship is applicable to ART, but only as a limiting case, and where the field to be reconstructed is more *radially limited* higher resolutions are achieved. For LES Phantoms the chemilumi-

nescence fields were quite sparse, which explains why their reconstructions with limited angles were so good; however these fields are representative of those likely to be encountered by a CTC sensor and so we expect to achieve at least double the resolution predicted by the Fourier slice theory in future work. It may in fact even be possible to obtain an engineering model to determine this improvement, perhaps based on the second moment of area of the field or similar. The theoretical relationship assumes equal angular spacing of the views; however, where this is not the case then non-uniform resolution in the reconstruction can be expected, as has been demonstrated from the cosine analysis of the TOJ measurement angles.

Based on these experimental and Phantom studies we have estimated the number of views that would be required to achieve a resolution comparable to our development target: LES. Typically an LES would use a grid of 100 – 200 cells across the domain width $D$, implying a resolved wavelength range of $0.02D$ – $0.01D$. To achieve such a resolution with CTC then approximately 30 to 40 viewing angles would be required. While this seems quite a high number, 40 of the PicSight P32M cameras used in this work would cost just £40000, which is comparable to the cost of a single high speed scientific camera. Furthermore, slightly lower resolution information, say $0.04D$, would still be very useful for LES validation and a CTC sensor using 20 – 30 cameras could achieve this.

The PicSight cameras used have also been shown to achieve very short temporal resolutions. Successful reconstructions of the the Matrix Burner were performed using views with exposures as short as $62\,\mu$s—which is below the $100\,\mu$s target set for the measurement of turbulent flames. However, the SNR of the views in this case was 1000, and so exposures closer to $10\,\mu$s are possible for this flame. The use of oxygen as the oxidiser has certainly contributed to the high SNRs of the Matrix Burner measurements, and for the TOJ lower SNRs are observed. For example, with $250\,\mu$s exposures an SNR of approximately 60 is obtained (and by extrapolation an equivalent exposure of $60\,\mu$s would have an SNR of approximately 15). Nevertheless, for the TOJ the effect of measurement noise was not discernible in the reconstructions due to the filtering of the views.

Typically the reconstruction problem is under-determined and there is usually excess spatial resolution in the views. In this work we have filtered and resampled the views to remove this excess and so reduce the datasets and speed up processing; however, this has the added benefit of significantly increasing the view SNR. Much higher SNR measurements of the TOJ could, therefore, have been obtained by the removal of the mirrors; because this would have allowed higher resolution views and so more binning/filtering. Given that with an SNR of 100 an ART reconstruction has just 5% RMS error, we can then suppose that with pixel-binning and removal of the mirrors (which reflect only 80% of the incident light) it would have been possible to perform measurements of the TOJ

using exposures well below our target of 100 $\mu$s and have still obtained sufficient SNR for reconstruction.

The measurements using the PicSight cameras have considered the visible wavelength chemiluminescence; however, the chemical simulations suggest that the Ultra Violet (UV) wavelength $OH^*$ emissions would be stronger and could allow even shorter exposure times. Shorter still exposure times are also possible by using oxygen as the oxidiser and *brighter* fuels; for example, (mirrored) test images of the TOJ burning ethylene achieved good SNR for exposures as low as 31 $\mu$s.

The PicSight P32M cameras are capable of measurements at repetition rates of 62 Hz, but only single measurements have been performed in this work due to a bug in the vendor supplied camera operation program. (In future work a new control program can be written to avoid this problem; unfortunately there was insufficient time to do so in this work.) To track small scale turbulent events repetition rates of the order of many kHz would be required; however, the PicSight P32M cameras will be capable of measuring larger scale transient events, for example low frequency acoustic oscillations. While cameras of high enough speed to track turbulent events are available, they are perhaps too costly to be deployed in a current CTC sensor, although with the continual progress in CCD/CMOS technology a future CTC sensor may achieve such high repetition rates. For example, a recent consumer camera, the Casio EX-F1, is capable of 1200 frames-per-second.

The CTC development undertaken in this work has considered LES validation as its target application; however, the data provided by CTC will also be more generally useful in the study of combustion. CTC can be used to study important combustion parameters such as equivalence ratio and heat release rate in 3D and can be used to validate chemical reaction mechanisms that incorporate excited state species. Indeed the recent resurgence of interest in chemiluminescence and its modelling [120, 117] suggests that new uses for chemiluminescence data will be found.

In this work we have considered a new potential use for the 3D chemiluminescence data provided by CTC: the measurement of the local flame-front strain rate from the integral of chemiluminescence across the flame normal. The link between integral chemiluminescence and strain has been investigated using both simulation [57] and measurement of a laminar opposed jet flame. Both results suggest a relationship between strain and integral chemiluminescence exists, although the difference observed in the simulated and measured $CH^*$ gradients suggests more work is needed to properly define this relationship. In particular a reassessment of the excited species simulation data is needed as this was originally intended only to allow representative excited state species concentrations to be obtained by LES and was not extensively validated. Questions also remain regarding whether the results from steady laminar flames are a sufficient approximation to unsteady

turbulent flames. However, the steady *flamelet* modelling approach is already used in turbulent combustion, and strain rate measurements, even with large uncertainties, would be useful.

The Matrix Burner has proved to be a good test case for the development of CTC, as it is non-axisymmetric and has multiple flame fronts that are similar to a wrinkled turbulent flame. Perhaps most usefully for this development, the Matrix Burner has an intuitively predictable flame shape that permits qualitative assessment of the reconstruction. The reconstructions of the Matrix Burner have generally resolved the constituent flames well.

However, the sections of some of the outer flames are not as *circular* as would be expected—suggesting that some error is present in the reconstruction process. Several possible causes for the deformation have been investigated, although none have proved to be a dominant cause of the deformation. The deformations appear greater in the reconstructions using ten views than those using 48 views, indicating the number of views plays a role. However, reconstructions of a Matrix Burner Phantom show less deformation than the experimental reconstructions, and this implies that it is some mis-representation of the projections, not the limited numbers of angles, that causes the deformations. Several of the studies in this work have considered random error in the viewing angle $\theta$, and systematic error in the out-of-plane angle $\gamma$, as possible sources of the outer flame deformations. However, neither error has been found to be a significant contributor. The use of non-parallel, i.e. perspective corrected, projections has been found to reduce, but not remove, the deformations. It is thought that lens distortion could also contribute, because the greatest deformation occurs at the edges of the reconstructions. It is the authors belief that the perspective distortion and the cumulative effect of view registration error—in all 6 degrees of freedom—are main causes of the flame deformation and a more comprehensive view registration approach has been deemed necessary.

A Computer Vision (CV) registration approach has been demonstrated which provides the view locations in 3D along with estimates for the camera parameters. This more general CV approach obtained values for $\theta$ similar to that of the previous bespoke approach, but with lower uncertainties, and also fully locates the view in 3D as well as providing estimates for camera parameters such as the focal length. The CV registration approach also estimates parameters for the lens distortion, and analysis of the model shows that distortion can affect the viewing angle of the outer pixels of an image by as much as $0.15^\circ$. The mirroring system locates the views in these outer pixels, suggesting that lens distortion could contribute to the flame deformations where mirrors have been used. We note that the distortion model provided by the CV registration can be used to correct this source of error.

In this work a number of improvements to the standard ART algorithm have been

presented and tested. Self-absorption has been investigated using realistic LES Phantoms and, based on the absorption coefficient derived by Seitzman [136], the results suggest that for $OH^*$ emissions this could be a significant source of error—though probably not for $CH^*$ as also reported in other studies [70, 73]. A modification to the ART algorithm (SMART) based on the Beer-Lambert law has been tested and found to correct the error caused by even very strong self-absorption. This modification assumes that the emitting and absorbing fields are proportional, which is reasonable when the ground state field is the main absorbing field (This assumption could be easily tested in the future by using the ground state data from the LES to simulate such self absorption.) The success of the SMART algorithm has also shown how ART can be used to solve even non-linear reconstruction problems.

The extension into 3D allows out-of-plane effects to be considered. A new projection geometry has been presented that better represents the camera optics that give rise to the out-of-plane contributions to the projections. These new *optical* projections allow shorter object distances and wider apertures to be used that in turn mean stronger measured signals and so higher time resolutions. The new optical projections have been derived from the ray-tracing, and include a necessary weighting to account for the intersection of the CCD pixel and the image blur circle. The optical projections have been used to reconstruct the Matrix Burner, and although the reconstructions are good, they are not noticeably better than those performed using cylindrical projections. This is because the depth-of-field of the measurements was larger than intended, making the cylindrical projections a good approximation. For measurements with a shorter object distance and a wider aperture a greater improvement from using the optical projections could be expected. The optical projections were also hindered by the consumer camera used in this experiment, which did not have well defined optical parameters, and was not well represented by the single-lens model assumed by the optical projections. The optical projections would benefit, therefore, from the CV registration approach, which provides reliable estimates for the focal length and object distance, and there are synergistic advantages to their combined use. The PicSight cameras used in later work have been seen to be well represented by the single-lens assumption, and their combination with optical projections and CV registration, could be expected to achieve even more accurate and even lower time-resolution reconstructions, than achieved in this work.

## 8.3 Recommendations

This project is part of an ongoing development of the CTC sensor. In this section we first present four specific recommendations for future work that are judged to be the most im-

portant to the success of CTC. Following this more general directions for future research are considered.

### 8.3.1 Specific Recommendations

This work has demonstrated that CTC can be implemented with inexpensive PicSight P32M cameras, and that they can be used to measure a turbulent flame. It has also been separately demonstrated that with sufficient numbers of views (20 – 40) resolutions comparable to LES are achievable. However, high resolution measurements of turbulent flames have not been demonstrated. Future work should, therefore, deploy more cameras to produce the high resolution measurements of turbulent flames that CTC has been shown capable of providing.

The low numbers of views and the limited flame volumes have meant the view datasets have been small and have required only moderate computing resources. By deploying substantially more views, the convergence time and, more importantly, the memory requirements of ART will increase substantially, and could *artificially* limit the resolution. Fortunately the ART algorithm has considerable scope for parallelisation—for example the order in which each view is addressed is relatively unimportant—and this should be taken advantage of in future work to allow greater numbers of views and high resolutions reconstructions to be performed.

The third recommendation is to avoid the use of mirrors. Mirrors have had to be used in this work to prevent significant aliasing in the reconstructions that occurs when using only 5 views (as seen in the Impinging flame reconstruction). However, we have seen that the mirrors reduce the measured signal by their imperfect reflection, by forcing longer object distances, and by reducing the resolution of the views and the potential for binning/filtering to improve SNR. Furthermore, the reduction in the resolution means the mirrors also increase the uncertainty in registration, as well as making the positioning and packaging of the cameras much more difficult. (We note that much of the difficulties had in measuring the TOJ resulted from using mirrors.) Finally using mirrors prevents the use of the superior CV approach to view registration.

While a well designed bespoke approach to registration can achieve as good registration of a few parameters, the CV approach provides a way to register all six degrees of freedom and its use is strongly recommended. Furthermore, the CV approach also provides the camera parameters necessary for the use of the *optical* projections and perspective correction. There is also much scope for improved implementations; for example, the camera model from CV registration can be used as an input to ART. In this work we have calibrated each camera individually using a flat plate; however, improvements may be obtained by optimising for all cameras and using a more suitable calibration object as

suggested by Svoboda et al. [146].

### 8.3.2 Areas for Further Research

The 3D reconstructions of the TOJ have shown a form of structured noise that is similar to the *salt-and-pepper* noise observed when using inexact 2D integral geometries with ART. To remove this *salt-and-pepper* noise, future work could consider replacing the approximate cylindrical/double-cone(optical) projections with extruded cubes/double-square pyramids, that better match the actual projection from a square CCD pixel. These improved geometries could be implemented using the same approach as the cylindrical/optical projections, but testing a relative $X$ and $Y$ coordinate rather than a single radius. It may also be possible to derive better still, bespoke, projection geometries from the CV registration data by considering more comprehensive camera models such as that used by Baba et al. [8].

In addition to more accurate projection geometries, the reconstructions could also be improved by considering new algorithms and also new basis functions such as *blobs* [105]. ART has been favoured in this work for its adaptability and noise performance however, optimisation methods such as MLEM (Maximum Likelihood-Expectation Maximisation) and MAP (Maximum A Posteriori) are often favoured in Nuclear Medicine applications and have been successfully used by Ishino and Ohiwa [73] for chemiluminescence measurements. The stochastic approach of these methods appears complimentary to the photon emission phenomena of chemiluminescence measurements and therefore may offer improvements over ART.

The realistic LES Phantoms used in this work have greatly aided the development of CTC, for example, by allowing the study of self-absorption. However, while the excited species predictions were reasonable, recent work, for example by Panoutsos et al. [120], suggests that improvements to the excited state reaction mechanisms used in this work are possible. The LES Phantoms could also be further improved by predicting the emitted light, as well as the excited species concentrations, i.e. predict $hv$ as well as $Y_{\mathrm{CH}^*}$, as in CTC it is the light that is measured not the concentration. This would pre-account for quenching and so make comparison with experimental data easier. Such LES data would also allow the study of other optical effects; for example, beam steering.

Large Eddy Simulation has been used in this project both as a development aid and target. While the 3D LES data can be easily used to simulate CTC measurements, the use of 3D data for LES validation is quite new. Consideration of appropriate new validation methods (3D spatial correlations for example) is therefore needed to fully utilise the validative potential of the 3D data provided by CTC. The comparison of two different instantaneous realisations may be of limited benefit and so attention should also be given

to obtaining correlated sequential CTC measurements to track transient events.

High repetition rates are necessary to track small scale turbulent events, but currently the high speed cameras needed to make such measurements are too costly for use in a CTC sensor. However, as typically in CTC the cameras will have excess resolution (pixels), it may be possible to employ optics to perform *spatial multiplexing* that allows all views to be measured by just a few cameras. For example fibre optic bundles, as used by Wright et al. [170], could be used to direct light from the viewing point to a high speed camera. The light from different view points could be directed at different areas on the CCD/CMOS and so provide simultaneous view measurements. Such an implementation could achieve both high spatial resolutions and high repetition rates, and be affordably implemented with current technology.

# Bibliography


[1] A. K. Agrawal, N. Butuk, S. R. Gollahalli, and D. Griffin. Three-dimensional rainbow schlieren tomography of a temperature field in gas flows. *Applied Optics*, 37(3):479–485, 1998.

[2] P. G. Aleiferis, Y. Hardalupas, A. M. K. P. Taylor, K. Ishii, and Y. Urata. Flame chemiluminescence studies of cyclic combustion variations and air-to-fuel ratio of the reacting mixture in a lean-burn stratified-charge spark-ignition engine. *Combustion and Flame*, 136(1-2):72–90, Jan. 2004.

[3] A. H. Andersen. Tomography transform and inverse in geometrical optics. *Journal of the Optical Society of America A*, 4(8):1385–1395, 1987.

[4] A. H. Andersen. Algebraic reconstruction in CT from limited views. *Medical Imaging, IEEE Transactions on*, 8(1):50–55, 1989.

[5] A. H. Andersen and A. C. Kak. Digital ray tracing in two-dimensional refractive fields. *Journal of the Acoustical Society of America*, 72(5):1593–1606, Nov. 1982.

[6] A. H. Andersen and A. C. Kak. Simultaneous Algebraic Reconstruction Technique (SART): A superior implementation of the ART algorithm. *Ultrasonic Imaging*, 6(1):81–94, Jan. 1984.

[7] S. Angermeier. PhD thesis, Imperial College London, London UK, 1998.

[8] M. Baba, N. Asada, A. Oda, and T. Migita. A thin lens based camera model for depth estimation from defocus and translation by zooming. In *Proceedings of the 15 th International Conference on Vision Interface, Calgary, Canada*, pages –, 2002.

[9] R. S. Barlow, editor. *International Workshop on Measurement and Computation of Turbulent Nonpremixed Flames*. Sandia National Laboratories, 1996 – 2008.

[10] R. S. Barlow and J. H. Frank. Effects of turbulence on species mass fractions in methane/air jet flames. *Proceedings of the Combustion Institute*, 27:1087–1095, 1998.

[11] R. S. Barlow, J. H. Frank, A. N. Karpetis, and J.-Y. Chen. Piloted methane/air jet flames: Transport effects and aspects of scalar structure. *Combustion and Flame*, 143(4):433–449, 2005.

[12] R. S. Barlow, G.-H. Wang, P. Anselmo-Filho, M. S. Sweeney, and S. Hochgreb. Application of Raman/Rayleigh/LIF diagnostics in turbulent stratified flames. *Proceedings of the Combustion Institute*, 32(1):945–953, 2009.

[13] H. H. Barrett and W. Swindell. *Radiological Imaging: The Theory of Image Formation, Detection and Processing*, volume 2. Academic Press, 1981.

[14] E. Beiting. Fast optical absorption tomography. *Optics Letters*, 16(16):1280–1282, 1991.

[15] E. Beiting. Fiber-optic fan-beam absorption tomography. *Applied Optics*, 31(9):1328–1343, 1992.

[16] K. Bennett and R. L. Byer. Optical tomography: experimental verification of noise theory. *Optics Letters*, 9(7):270–272, July 1984.

[17] L. H. Benvenutti, C. S. T. Marques, and C. A. Bertran. Chemiluminescent emission data for kinetic modeling of ethanol combustion. *Combustion Science and Technology*, V177(1):1–26, 2005.

[18] M. Bertero and P. Boccacci. *Introduction to Inverse Problems in Imaging*. IOP, 1998.

[19] P. E. Best, P. L. Chien, R. M. Carangelo, P. R. Solomon, M. Danchak, and I. Ilovici. Tomographic reconstruction of FT-IR emission and transmission spectra in a sooting laminar diffusion flame: Species concentrations and temperatures. *Combustion and Flame*, 85(3-4):309–314, IN1, 315–318, June 1991.

[20] R. Bilger. Conditional moment closure for turbulent reacting flow. *Physics of Fluids A: Fluid Dynamics*, 5:436, 1993.

[21] R. W. Bilger. Some aspects of scalar dissipation. *Flow, Turbulence and Combustion*, 72(2):93–114, July 2004.

[22] R. W. Bilger, S. B. Pope, K. N. C. Bray, and J. F. Driscoll. Paradigms in turbulent combustion research. *Proceedings of the Combustion Institute*, 30(1):21–42, Jan. 2005.

[23] B. Böhm, C. Kittler, A. Nauert, and A. Dreizler. Diagnostics at high repetition rates: new insights into transient combustion phenomena. In *Proceedings of the European Combustion Meeting*, pages 5–1, Chania, Crete, Greece, 2007.

[24] J.-Y. Bouguet. Camera calibration toolbox for Matlab. http://www.vision.caltech.edu/bouguetj/calib_doc/, Mar 2009.

[25] P. M. Brisley, G. Lu, Y. Yan, and S. Cornwell. Three-dimensional temperature measurement of combustion flames using a single monochromatic CCD camera. *Instrumentation and Measurement, IEEE Transactions on*, 54(4):1417–1421, 2005.

[26] P. P. Bruyant. Analytic and iterative reconstruction algorithms in SPECT. *Journal of Nuclear Medicine*, 43(10):1343–1358, Oct. 2002.

[27] K. A. Buch and W. J. A. Dahm. Experimental study of the fine-scale structure of conserved scalar mixing in turbulent shear flows. Part 2. Sc [ap] 1. *Journal of Fluid Mechanics*, 364:1–29, 1998.

[28] T. F. Budinger and G. T. Gullberg. Three-dimensional reconstruction in nuclear medicine emission imaging. *Nuclear Science, IEEE Transactions on*, NS-21(2):2–, 1974.

[29] A. Burcat and B. Ruscic. Third millennium ideal gas and condensed phase thermochemical database for combustion with updates from active thermochemical tables. Technical Report ANL-05/20 TAE 960, Argonne National Laboratory, 2005.

[30] A. Caldeira-Pires, D. P. Correia, P. Maia, P. Lacava, and M. V. Heitor. Influence of burner-port geometry in hydrocarbon oxidation and NOx formation mechanisms in methane/air flames. *Fuel*, 81(6):771–783, Apr. 2002.

[31] S. J. Carey, H. McCann, D. E. Winterbone, and E. Clough. Near infra-red absorption tomography for measurement of chemical species distribution. In *1st World Congress on Industrial Process Tomography*, Buxton, Greater Manchester, 1999.

[32] D. Coffin. dcraw (a program for decoding raw digital photos). http://www.cybercom.net/ dcoffin/dcraw/, Mar 2009.

[33] D. P. Correia, P. Ferrao, and A. Caldeira-Pires. Flame three-dimensional tomography sensor for in-furnace diagnostics. *Proceedings of the Combustion Institute*, 28(1):431–438, 2000.

[34] D. P. Correia, P. Ferrao, and A. Caldeira-Pires. Advanced 3D emission tomography flame temperature sensor. *Combustion Science and Technology*, 163:1–24, 2001.

[35] M. De Leo, A. Saveliev, L. A. Kennedy, and S. A. Zelepouga. OH and CH luminescence in opposed flow methane oxy-flames. *Combustion and Flame*, 149(4):435–447, June 2007.

[36] C. Deutsch. A FORTRAN 77 subroutine for determining the fractional area of rectangular grid blocks within a polygon. *Computers & Geosciences*, 16(3):379–384, 1990.

[37] N. Docquier and S. Candel. Combustion control and sensors: a review. *Progress in Energy and Combustion Science*, 28(2):107–150, 2002.

[38] N. Docquier, F. Lacas, and S. Candel. Closed-loop equivalence ratio control of premixed combustors using spectrally resolved chemiluminescence measurements. *Proceedings of the Combustion Institute*, 29(1):139–145, 2002.

[39] J. L. Edwards. *Proper Orthogonal Decomposition and Tomographic Analysis of Combustion Systems for Control Applications*. PhD thesis, Cornell University, USA, 2007.

[40] J. L. Edwards, F. C. Gouldin, and S. Becz. Gas turbine engine combustor control using emission tomography. In *44th AIAA Aerospace Sciences Meeting and Exhibit*, pages 1–33, Reno, Nevada, Jan. 2006.

[41] J. L. Edwards, F. C. Gouldin, and M. A. MacDonald. High speed absorption tomography with advanced reconstruction algorithms. In *41st Aerospace Sciences Meeting and Exhibit*, page 1013, Reno, Nevada, Jan 2003.

[42] G. E. Elsinga, F. Scarano, B. Wieneke, and B. W. van Oudheusden. Tomographic particle image velocimetry. In *6th International Symposium on Particle Image Velocimetry*, Pasadena, California, USA, September 2005.

[43] P. J. Emmerman, R. Goulard, R. J. Santoro, and H. G. Semerjian. Multiangular absorption diagnostics of a turbulent argon-methane jet. *Journal of Energy*, 4(2):70–7, 1980.

[44] J. Feng, K. Okamoto, D. Tsuru, H. Madarame, and M. Fumizawa. Visualization of 3D gas density distribution using optical tomography. *Chemical Engineering Journal*, 86(3):243–250, Apr. 2002.

[45] H. Forkel. *LES of Diffusion Flames*. PhD thesis, TU Darmstadt, 1999.

[46] H. Forkel and J. Janicka. Large-eddy simulation of a turbulent hydrogen diffusion flame. *Flow, Turbulence and Combustion*, 65(2):163–175, 2000.

[47] J. H. Frank, P. A. M. Kalt, and R. W. Bilger. Measurements of conditional velocities in turbulent premixed flames by simultaneous OH PLIF and PIV. *Combustion and Flame*, 116(1-2):220–232, Jan. 1999.

[48] G. Frieder and G. T. Herman. Resolution in reconstructing objects from electron micrographs. *Journal of Theoretical Biology*, 33(1):189–211, 1971.

[49] A. G. Gaydon. *The Spectroscopy of Flames*. Chapman and Hall, 1974.

[50] M. Germano, U. Piomelli, P. Moin, and W. H. Cabot. A dynamic subgrid-scale eddy viscosity model. *Physics of Fluids A: Fluid Dynamics*, 3:1760, 1991.

[51] D. Geyer, A. Kempf, A. Dreizler, and J. Janicka. Scalar dissipation rates in isothermal and reactive turbulent opposed-jets: 1-D-Raman/Rayleigh experiments supported by LES. *Proceedings of the Combustion Institute*, 30:681–689, Jan. 2005.

[52] D. Geyer, A. Kempf, A. Dreizler, and J. Janicka. Turbulent opposed-jet flames: A critical benchmark experiment for combustion LES. *Combustion and Flame*, 143(4):524–548, Dec. 2005.

[53] J. Ghandhi. Spatial resolution and noise considerations in determining scalar dissipation rate from passive scalar image data. *Experiments in Fluids*, 40(4):577–588, Apr. 2006.

[54] G. Gilabert, G. Lu, and Y. Yan. Three-dimensional tomographic reconstruction of the luminosity distribution of a combustion flame. *Instrumentation and Measurement, IEEE Transactions on*, 56(4):1300–1306, 2007.

[55] B. Gillet, Y. Hardalupas, , and A. M. K. P. Taylor. Influence of projection width and spatial gradient of concentration on algebraic reconstruction tomography (ART) for measurements of hydrocarbon fuel/air mixtures. In *42nd AIAA Aerospace Science Meeting and Exhibit*, pages 1–11, Reno, Nevada, January 5-8 2004. AIAA-2004-0648.

[56] B. Gillet, Y. Hardalupas, C. Kavounides, and A. M. K. P. Taylor. Infrared absorption for measurement of hydrocarbon concentration in fuel/air mixtures (MAST-B-LIQUID). *Applied Thermal Engineering*, 24(11-12):1633–1653, Aug 2004.

[57] K. Gkagkas. Laminar Opposed Jet simulation. Personnel Communication, Jan 2007.

[58] R. Gordon. A tutorial on ART (algebraic reconstruction techniques). *Nuclear Science, IEEE Transactions on*, ns-21(3):78–93, 1974.

[59] R. Gordon, R. Bender, and G. T. Herman. Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and x-ray photography. *Journal of Theoretical Biology*, 29(3):471–81, 1970.

[60] H. Guan and R. Gordon. A projection access order for speedy convergence of ART (algebraic reconstruction technique): a multilevel scheme for computed tomography. *Physics in Medicine and Biology*, 39(11):2005–2022, 1994.

[61] R. Hall and P. Bonczyk. Sooting flame thermometry using emission/absorption tomography. *Applied Optics*, 29(31):4590–4598, 1990.

[62] K. M. Hanson and G. W. Wecksung. Local basis-function approach to computed tomography. *Applied Optics*, 24(23):4028–4039, 1985.

[63] Y. Hardalupas and M. Orain. Local measurements of the time-dependent heat release rate and equivalence ratio using chemiluminescent emission from a flame. *Combustion and Flame*, 139(3):188–207, Nov. 2004.

[64] R. Hartley and A. Zisserman. *MuLtiple View Geometry In Computer Vision*. Cambridge university press, 2003.

[65] R. He, C. G. Xie, R. C. Waterfall, M. S. Beck, and C. M. Beck. Engine flame imaging using electrical capacitance tomography. *Electronics Letters*, 30(7):559–560, 1994.

[66] E. Hecht. *Optics*. Addison-Wesley, 2002.

[67] G. T. Herman. *Image Reconstruction From Projections: The Fundamentals Of Computerized Tomography*. Academic Press, 1980.

[68] G. T. Herman and L. B. Meyer. Algebraic reconstruction techniques can be made computationally efficient [positron emission tomography application]. *Medical Imaging, IEEE Transactions on*, 12(3):600–609, 1993.

[69] G. T. Herman and S. Rowland. Resolution in ART. An experimental investigation of the resolving power of an algebraic picture reconstruction technique. *Journal of Theoretical Biology*, 33(1):213–23, 1971.

[70] H. M. Hertz, , and G. W. Faris. Emission tomography of flame radicals. *Optics Letters*, 13:351–, 1988.

[71] L. Hesselink. Optical tomography. In W. J. Yang, editor, *Handbook of Flow Visualization*, chapter 20, pages 307–329. Hemisphere, New York, 1987.

[72] IPCC. Synthesis Report. Summary for Policymakers. In *Fourth Assessment Report*, 2007.

[73] Y. Ishino and N. Ohiwa. Three-dimensional computerized tomographic reconstruction of instantaneous distribution of chemiluminescence of a turbulent premixed flame. *JSME International Journal Series B: Fluids and Thermal Engineering*, 48(1):34–40, 2005.

[74] T. Ito, K. Masuda, and T. Nakazawa. Inner temperature measurement of combustion gas by using infrared emission computed tomography. In *SICE-ICASE, 2006. International Joint Conference*, pages 3342–3345, 2006.

[75] E. T. Jaynes. Information theory and statistical mechanics. II. *Physical Review*, 108(2):171–, Oct. 1957.

[76] S. Kaczmarz. Angenaherte auflosung von systemen linearer gleichungen. *Bull. Acad. Polonaise Sci. et Lettres A*, pages 355–357, 1937.

[77] A. C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. IEEE Press, 1988.

[78] A. Karpetis and R. Barlow. Measurements of flame orientation and scalar dissipation in turbulent partially premixed methane flames. *Proceedings of the Combustion Institute*, 30:665–672, Jan. 2005.

[79] A. N. Karpetis and R. S. Barlow. Measurements of scalar dissipation in a turbulent piloted methane/air jet flame. *Proceedings of the Combustion Institute*, 29:1929–1936, 2002.

[80] A. Kempf. LES validation from experiments. *Flow, Turbulence and Combustion*, 80(3):351–373, Apr. 2008.

[81] A. Kempf, F. Flemming, and J. Janicka. Investigation of lengthscales, scalar dissipation, and flame orientation in a piloted diffusion flame by LES. *Proceedings of the Combustion Institute*, 30:557–565, 2005.

[82] A. Kempf and J. Janicka. A visualisation technique for flame dynamics : Turning combustion-LES into a tool for engineering. *VDI-Berichte*, 1888(1888):463–469, 2005.

[83] A. Kempf, M. Klein, and J. Janicka. Efficient generation of initial- and inflow-conditions for transient turbulent flows in arbitrary geometries. *Flow, Turbulence and Combustion*, 74(1):67–84, Jan. 2005.

[84] A. Kempf, R. P. Lindstedt, and J. Janicka. Large-eddy simulation of a bluff-body stabilized nonpremixed flame. *Combustion and Flame*, 144(1-2):170–189, 2006.

[85] A. M. Kempf. *Large Eddy Simulation of Non-Premixed Turbulent Flames*. PhD thesis, Darmstadt EKT, 2003.

[86] A. Y. Klimenko. Multicomponent diffusion of various admixtures in turbulent flow. *Fluid Dynamics*, 25(3):327–334, 1990.

[87] D. A. Knaus, F. C. Gouldin, and D. C. Bingham. Assessment of crossed-plane tomography for flamelet surface normal measurements. *Combustion Science and Technology*, 174(1):101–134, 2002.

[88] H. S. Ko and K. D. Kihm. An extended algebraic reconstruction technique (ART) for density-gradient projections: laser speckle photographic tomography. *Experiments in Fluids*, V27(6):542–550, Nov. 1999.

[89] J. Kojima, Y. Ikeda, and T. Nakajima. Basic aspects of OH(A), CH(A), and C2(d) chemiluminescence in the reaction zone of laminar methane-air premixed flames. *Combustion and Flame*, 140(1-2):34–45, Jan. 2005.

[90] H. Kolla, J. W. Rogerson, N. Chakraborty, and N. Swaminathan. Scalar dissipation rate modeling and its validation. *Combustion Science and Technology*, 181(3):518–535, 2009.

[91] K. Lange and R. Carson. EM reconstruction algorithms for emission and transmission tomography. *Journal of computer assisted tomography*, 8(2):306–16, 1984.

[92] E. A. Lavinskaya, S. Martemianov, J. Saulnier, and N. A. Fomin. Limited projection laser tomography of combined gasdynamic flows. *Journal of Engineering Physics and Thermophysics*, 77(5):979–989, Sept. 2004.

[93] A. Leipertz, R. Obertacke, and F. Wintrich. Industrial combustion control using UV emission tomography. *Proceedings of the Combustion Institute*, 26:2869–2875, 1996.

[94] J. Lim, Y. Sivathanu, and D. Feikema. Fan beam emission tomography for estimating scalar properties in laminar flames. In *Third Joint Meeting of the U.S. Sections of the Combustion Institute*, pages 1–6, Chicago, Illinois, March 2003. NASA. NASA/TM2003-212302.

[95] R. P. Lindstedt, S. A. Louloudi, and E. M. Vaos. Joint scalar probability density function modeling of pollutant formation in piloted turbulent jet diffusion flames with comprehensive chemistry. *Symposium (International) on Combustion*, 28(1):149–156, 2000.

[96] R. P. Lindstedt, D. S. Luff, and J. H. Whitelaw. Velocity and strain-rate characteristics of opposed isothermal flows. *Flow, Turbulence and Combustion*, 74:169–194, 2005.

[97] R. P. Lindstedt, D. S. Luff, and J. H. Whitelaw. Velocity fields of fuel lean premixed turbulent opposed jet flames. *Proceedings of the Combustion Institute*, 31(1):1459–1466, Jan. 2007.

[98] R. P. Lindstedt and M. P. Meyer. A dimensionally reduced reaction mechanism for methanol oxidation. *Proceedings of the Combustion Institute*, 29(1):1395–1402, 2002.

[99] D. Luff, E. Korusoy, P. Lindstedt, and J. H. Whitelaw. Counterflow flames of air and methane, propane and ethylene, with and without periodic forcing. *Experiments in Fluids*, 35(6):618–626, Dec. 2003.

[100] J. Luque, J. B. Jeffries, G. P. Smith, D. R. Crosley, K. T. Walsh, M. B. Long, and M. D. Smooke. CH(A-X) and OH(A-X) optical emission in an axisymmetric laminar diffusion flame. *Combustion and Flame*, 122(1-2):172–175, July 2000.

[101] C. N. Markides, G. De Paola, and E. Mastorakos. Measurements and simulations of mixing and autoignition of an n-heptane plume in a turbulent flow of heated air. *Experimental Thermal and Fluid Science*, 31(5):393–401, 2007.

[102] E. Mastorakos, A. M. K. P. Taylor, and J. H. Whitelaw. Extinction and temperature characteristics of turbulent counterflow diffusion flames with partial premixing. *Combustion and Flame*, 91(1):40–54, Oct. 1992.

[103] E. Mastorakos, A. M. K. P. Taylor, and J. H. Whitelaw. Extinction of turbulent counterflow flames with reactants diluted by hot products. *Combustion and Flame*, 102(1-2):101–114, July 1995.

[104] S. Matej and R. M. Lewitt. Efficient 3D grids for image reconstruction using spherically-symmetric volume elements. *Nuclear Science, IEEE Transactions on*, 42(4):1361–1370, Aug 1995.

[105] S. Matej and R. M. Lewitt. Practical considerations for 3-D image reconstruction using spherically symmetric volume elements. *Medical Imaging, IEEE Transactions on*, 15(1):68–78, 1996.

[106] K. L. McNesby, R. G. Daniel, J. B. Morris, and A. W. Miziolek. Tomographic analysis of CO absorption in a low-pressure flame. *Applied Optics*, 34(18):3318–3324, June 1995.

[107] E. M. Mikhail, J. S. Bethel, and J. C. McGlone. *Introduction to Modern Photogrammetry*. Wiley, 2001.

[108] D. Mishra, J. P. Longtin, R. P. Singh, and V. Prasad. Performance evaluation of iterative tomography algorithms for incomplete projection data. *Applied Optics*, 43(7):1522–1532, 2004.

[109] E. J. Mohamad, R. A. Rahim, S. Ibrahim, S. Sulaiman, and M. S. Manaf. Flame imaging using laser-based transmission tomography. *Sensors and Actuators A: Physical*, 127(2):332–339, Mar. 2006.

[110] K. Mueller, R. Yagel, and J. F. Cornhill. The weighted-distance scheme: a globally optimizing projection ordering method for ART. *Medical Imaging, IEEE Transactions on*, 16(2):223–230, 1997.

[111] K. Mueller, R. Yagel, and J. J. Wheller. Anti-aliased three-dimensional cone-beam reconstruction of low-contrast objects with algebraic methods. *Medical Imaging, IEEE Transactions on*, 18(6):519–537, 1999.

[112] F. Natterer and F. Wübbeling. *Mathematical Methods in Image Reconstruction*. Society for Industrial Mathematics, 2001.

[113] S. Navarro-Martinez, A. Kronenburg, and F. Mare. Conditional moment closure for large eddy simulations. *Flow, Turbulence and Combustion*, 75(1):245–274, Dec. 2005.

[114] W. B. Ng and Y. Zhang. Stereoscopic imaging and reconstruction of the 3D geometry of flame surfaces. *Experiments in Fluids*, 34(4):484–493, Apr. 2003.

[115] Q. V. Nguyen, R. W. Dibble, D. Hofmann, and S. Kampmann. Tomographic measurements of carbon monoxide temperature and concentration in a bunsen flame using diode laser absorption. *Berichte der Bunsen-Gesellschaft*, 97(12):1634–1642, 1993.

[116] V. N. Nori. *Modeling and Analysis of Chemiluminescence Sensing for Syngas, Methane and Jet-A Combustion*. PhD thesis, Georgia Institute of Technology, 2008.

[117] V. N. Nori and J. M. Seitzman. CH* chemiluminescence modeling for combustion diagnostics. *Proceedings of the Combustion Institute*, 32(1):895–903, 2009.

[118] J. Nygren, J. Hult, M. Richter, M. Alden, M. Christensen, A. Hultqvist, and B. Johansson. Three-dimensional laser induced fluorescence of fuel distributions in an HCCI engine. *Proceedings of the Combustion Institute*, 29:679–685, 2002.

[119] R. Obertacke, H. Wintrich, F. Wintrich, and A. Leipertz. A new sensor system for industrial combustion monitoring and control using UV emission spectroscopy and tomography. *Combustion Science and Technology*, 121(1):133–151, 1996.

[120] C. S. Panoutsos, Y. Hardalupas, and A. M. K. P. Taylor. Numerical evaluation of equivalence ratio measurement using OH* and CH* chemiluminescence in premixed and non-premixed methane-air flames. *Combustion and Flame*, 156:273–291, 2009.

[121] N. Peters. Laminar diffusion flamelet models in non-premixed turbulent combustion. *Progress in Energy and Combustion Science*, 10:319–339, 1984.

[122] N. Peters. *Turbulent Combustion*. Cambridge University Press, 2000.

[123] H. Philipp, A. Plimon, G. Fernitz, A. Hirsch, G. Fraidl, and E. Winklhofer. A tomographic camera system for combustion diagnostics in SI engines. In *International Congress & Exposition*. SAE International, Feb 1995.

[124] H. Pitsch. Large-eddy simulation of turbulent combustion. *Annual Review of Fluid Mechanics*, 38(1):453–482, 2006.

[125] T. Poinsot and D. Veynante. *Theoretical and Numerical Combustion*. RT Edwards, Inc., 2005.

[126] S. B. Pope. A Monte Carlo method for the PDF equations of turbulent reactive flow. *Combustion Science and Technology*, 25:159–174, 1981.

[127] S. B. Pope. Computations of turbulent combustion- progress and challenges. *Proceedings of the Combustion Institute*, 23:591–612, 1991.

[128] S. B. Pope. *Turbulent Flows*. Cambridge University Press, 2000.

[129] M. Ravichandran and F. C. Gouldin. Reconstruction of smooth distributions from a limited number of projections. *Applied Optics*, 27(19):4084–97, 1988.

[130] K. A. Rizos. *Detailed Chemical Kinetic Modelling of Homogeneous systems*. PhD thesis, Imperial College London, UK, 2003.

[131] R. Robinson. Excited species thermodynamic data. Personal Communication, Jan 2007.

[132] P. Sagaut. *Large Eddy Simulation For Incompressible Flows: An Introduction*. Springer, 2002.

[133] R. J. Santoro, H. G. Semerjian, P. J. Emmerman, and R. Goulard. Optical tomography for flow field diagnostics. *International Journal of Heat and Mass Transfer*, 24(7):1139–50, 1981.

[134] C. Schneider, A. Dreizler, J. Janicka, and E. P. Hassel. Flow field measurements of stable and locally extinguishing hydrocarbon-fuelled jet flames. *Combustion and Flame*, 135(1-2):185–190, Oct. 2003.

[135] A. Schwarz. Multi-tomographic flame analysis with a schlieren apparatus. *Measurement Science and Technology*, 7(3):406–413, 1996.

[136] J. M. Seitzman. *Introduction to Spectroscopy and Laser Diagnostics for Gases*, chapter 14, page 229. Stanford University, 2003.

[137] L. A. Shepp and Y. Vardi. Maximum likelihood reconstruction for emission tomography. *Medical Imaging, IEEE Transactions on*, 1(2):113–122, 1982.

[138] Y. Sivathanu, J. Lim, and R. Joseph. Statistical absorption tomography for turbulent flows. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 68(6):611–623, 2001.

[139] Y. R. Sivathanu and J. P. Gore. A tomographic method for the reconstruction of local probability density functions. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 50(5):483–492, 1993.

[140] J. Smagorinsky. General circulation experiments with the primitive equations. *Monthly Weather Review*, 91(3):99–164, 1963.

[141] R. Snyder and L. Hesselink. Measurement of mixing fluid flows with optical tomography. *Optics Letters*, 13:87–, 1988.

[142] O. Stein. *Large Eddy Simulation of Combustion in Swirling and Opposed Jet Flows*. PhD thesis, Imperial College London, 2009.

[143] O. Stein, A. M. Kempf, and J. Janicka. LES of the Sydney-swirl flame series: an initial investigation of the fluid dynamics. *Combustion Science and Technology*, 179:173–189, 2007.

[144] L. K. Su and N. T. Clemens. The structure of fine-scale scalar mixing in gas-phase planar turbulent jets. *Journal of Fluid Mechanics*, 488:1–29, 2003.

[145] J. Sutton and J. Driscoll. A method to simultaneously image two-dimensional mixture fraction, scalar dissipation rate, temperature and fuel consumption rate fields in a turbulent non-premixed jet flame. *Experiments in Fluids*, 41(4):603–627, Oct. 2006.

[146] T. Svoboda, D. Martinec, and T. Pajdla. A convenient multi-camera self-calibration for virtual environments. *PRESENCE: Teleoperators and Virtual Environments*, 14(4):407–422, August 2005.

[147] K. Tanabe. Projection method for solving a singular system of linear equations and its applications. *Numerische Mathematik*, 17(3):203–214, June 1971.

[148] E. D. Torniainen, A. K. Hinz, and F. C. Gouldin. Tomographic analysis of unsteady, reacting flows: numerical investigation. *AIAA journal*, 36(7):1270–8, 1998.

[149] S. D. Tse, R. A. Anthenien, A. C. Fernandez-Pello, and K. Miyasaka. An application of ultrasonic tomographic imaging to study smoldering combustion. *Combustion and Flame*, 116(1):120–135, 1999.

[150] H. Uchiyama, M. Nakajima, and S. Yuta. Measurement of flame temperature distribution by IR emission computed tomography. *Applied Optics*, 24:4111–4115, 1985.

[151] P. Vaishnavi, A. Kronenburg, and C. Pantano. On the spatial length scales of scalar dissipation in turbulent jet flames. *Journal of Fluid Mechanics*, 596:103–132, 2008.

[152] D. Verhoeven. Limited-data computed tomography algorithms for the physical sciences. *Applied Optics*, 32(20):3736–54, 1993.

[153] D. Verhoeven. Multiplicative algebraic computed tomographic algorithms for the reconstruction of multidirectional interferometric data. *Optical Engineering*, 32(2):410–419, Feb. 1993.

[154] V. Vezhnevets. GML MatLab camera calibration toolbox. http://graphics.cs.msu.ru/en/research/calibration/index.html, Apr 2009.

[155] K. Wakai, K. Kamiya, S. Sakai, and S. Shimizu. Instantaneous measurement of two-dimensional temperature and density distributions of flames by a two-band-emission-CT pyrometer. In *Proceedings of the SPIE-The International Society for Optical Engineering*, volume 1762, pages 564–575, San Diego, CA, USA, Jan. 1993. SPIE.

[156] K. T. Walsh, J. Fielding, and M. B. Long. Effect of light-collection geometry on reconstruction errors in Abel inversions. *Optics Letters*, 25(7):457–459, Apr. 2000.

[157] K. T. Walsh, M. B. Long, M. A. Tanoff, and M. D. Smooke. Experimental and computational study of CH, CH*, and OH* in an axisymmetric laminar diffusion flame. *Proceedings of the Combustion Institute*, 27:615–623, 1998.

[158] X. Wan, Y. Gao, Q. Wang, S. Le, and S. Yu. Limited-angle optical computed tomography algorithms. *Optical Engineering*, 42(9):2659–2669, Sept. 2003.

[159] X. Wan, S. Yu, G. Cai, Y. Gao, and J. Yi. Three-dimensional plasma field reconstruction with multiobjective optimization emission spectral tomography. *Journal of the Optical Society of America A*, 21(7):1161–1171, July 2004.

[160] G. Wang and R. Barlow. Spatial resolution effects on the measurement of scalar variance and scalar gradient in turbulent nonpremixed jet flames. *Experiments in Fluids*, 44(4):633–645, Apr. 2008.

[161] G. Wang, A. N. Karpetis, and R. S. Barlow. Dissipation length scales in turbulent nonpremixed jet flames. *Combustion and Flame*, 148(1-2):62–75, Jan. 2007.

[162] G. H. Wang, N. T. Clemens, R. S. Barlow, and P. L. Varghese. A system model for assessing scalar dissipation measurement accuracy in turbulent flows. *Measurement Science and Technology*, 18(5):1287–1303, 2007.

[163] J. Warnatz, U. Maas, and R. Dibble. *Combustion: Physical and Chemical Fundamentals, Modeling and Simulation, Experiments, Pollutant Formation*. Springer Verlag, 2006.

[164] B. Wieneke. Stereo-PIV using self-calibration on particle images. *Experiments in Fluids*, 39(2):267–280, Aug. 2005.

[165] B. Wieneke. Volume self-calibration for 3D particle image velocimetry. *Experiments in Fluids*, 45:549 – 556, 2008.

[166] C. Willert. Assessment of camera models for use in planar velocimetry calibration. *Experiments in Fluids*, 41(1):135–143, July 2006.

[167] A. Williams, M. Wang, C. Ellul, and M. Pourkashanian. Electrical capacitance tomography for high temperature flames. In *Europe Conference on Industrial Furnaces and Boilers*, 2008.

[168] F. A. Williams. Recent advances in theoretical descriptions of turbulent diffusion flames. In S. N. B. Murthy, editor, *Turbulent mixing in nonreactive and reactive flows*, pages 202–208, 1975.

[169] F. A. Williams. *Combustion Theory*. Benjamin/Cummings, 1985.

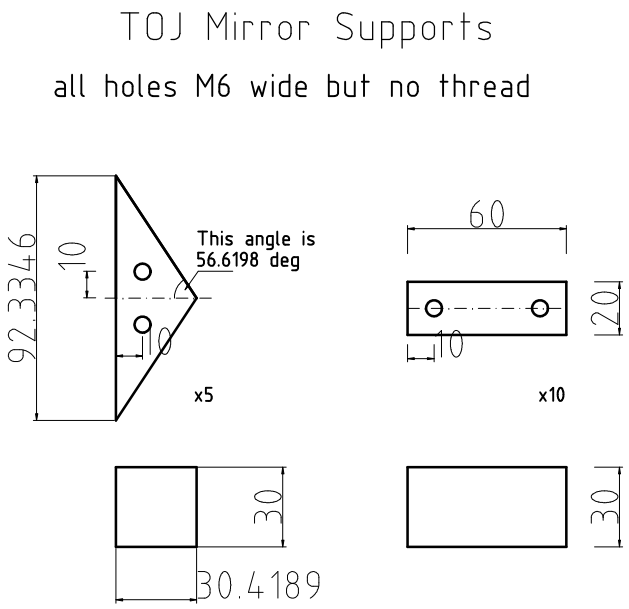
[170] P. Wright, C. Garcia-Stewart, S. Carey, F. Hindle, S. Pegrum, S. Colbourne, P. Turner, W. Hurr, T. Litt, S. Murray, S. Crossley, K. Ozanyan, and H. McCann. Toward in-cylinder absorption tomography in a production engine. *Applied Optics*, 44(31):6578–6592, 2005.

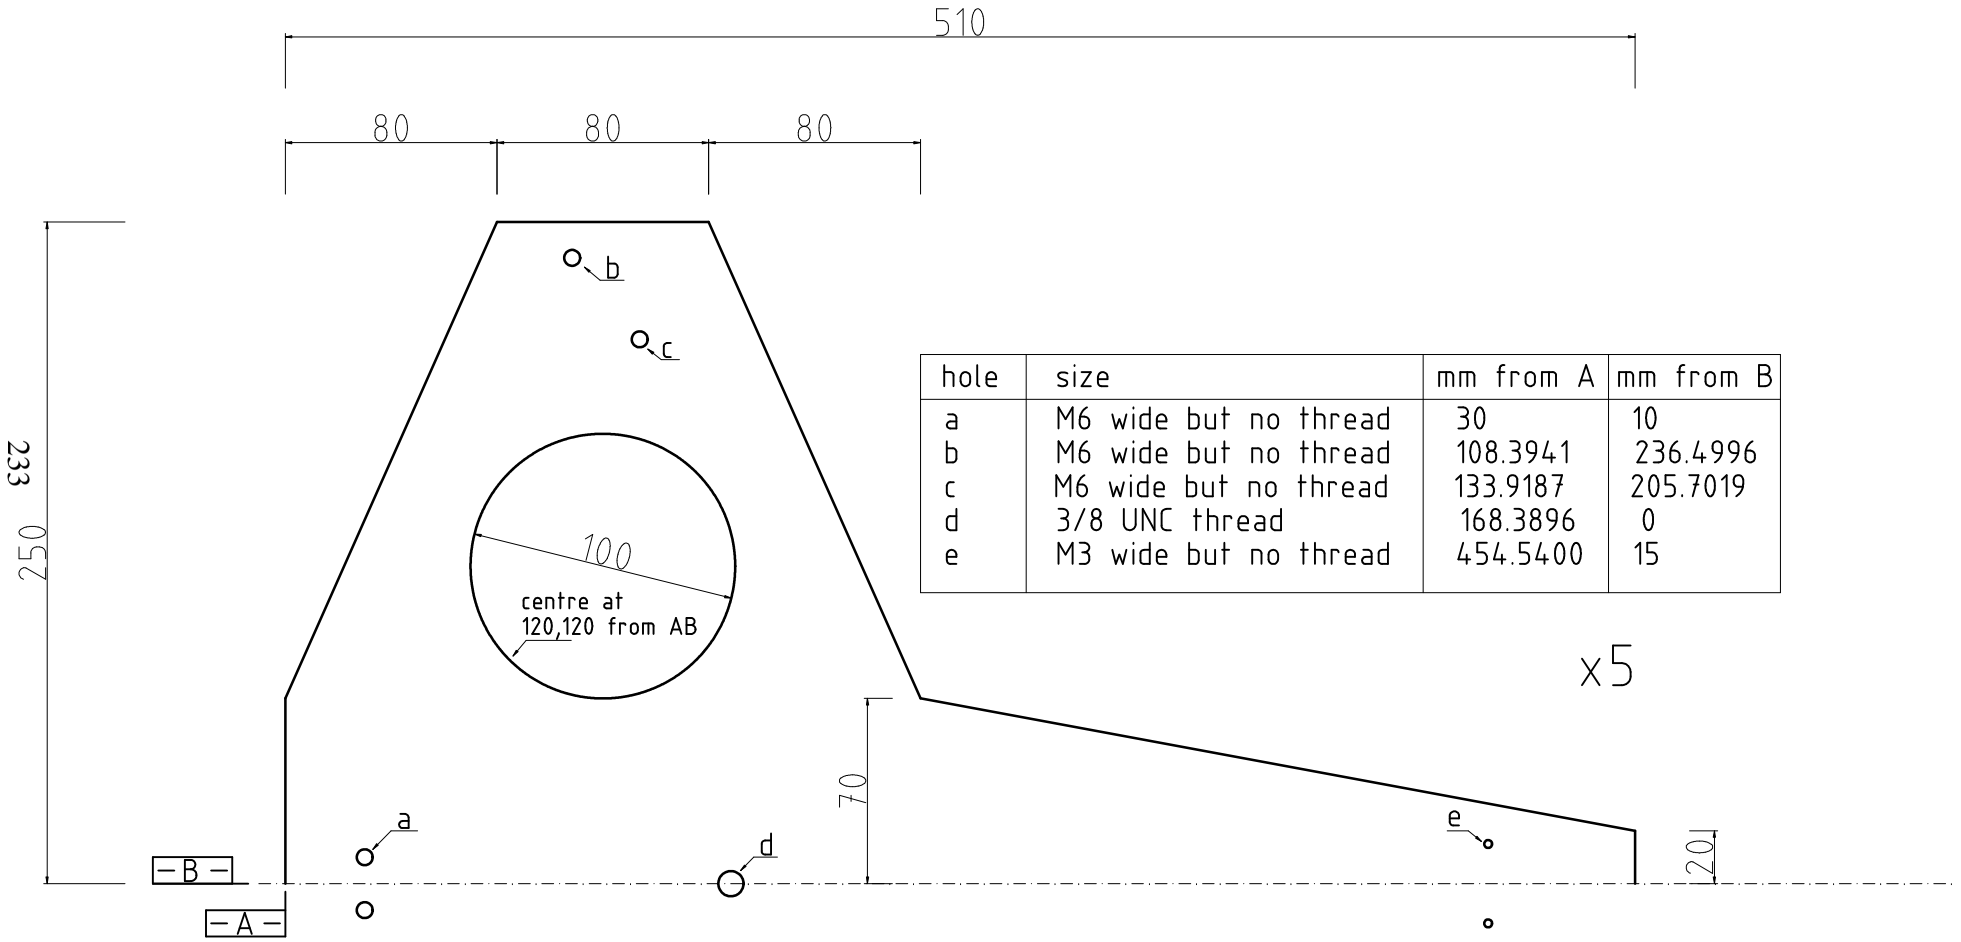
[171] A. Yoshida, M. Narisawa, and H. Tisuji. Structure of highly turbulent premixed flames. *Twenty-Fourth Symposium (International) on Combustion*, pages 519–525, 1992.

[172] Z. Zhang. A flexible new technique for camera calibration. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(11):1330–1334, 2000.

# Appendix A

# Technical Drawings






| hole | size | mm from A | mm from B |
|---|---|---|---|
| a | M6 wide but no thread | 30 | 10 |
| b | M6 wide but no thread | 108.3941 | 236.4996 |
| c | M6 wide but no thread | 133.9187 | 205.7019 |
| d | 3/8 UNC thread | 168.3896 | 0 |
| e | M3 wide but no thread | 454.5400 | 15 |

x5

TOJ Mirror Mount