

**Three Dimensional Flame Reconstruction
towards the Study of Fire-Induced
Transmission Line Flashovers**

by

Colin Crompton

Submitted in fulfilment of the academic requirements for the degree of Master of
Science in Engineering in the School of Electrical, Electronic and Computer
Engineering at the University of KwaZulu-Natal, Durban, South Africa

March 2007

Preface

The research work presented in this thesis was performed by Colin Crompton, under the supervision of Dr. Derek Hoch, at the University of KwaZulu-Natal's School of Electrical, Electronic and Computer Engineering. The financial assistance of the Department of Labour (DoL) towards this research is hereby acknowledged.

Parts of this work were presented by the author at PRASA 2004.

The entire dissertation, unless otherwise indicated as a reference, is the author's own original work and has not been submitted, in part or in whole, to any other university for degree purposes.

Signed: 

Name: COLIN CROMPTON

Date: 28/03/2007

Abstract

The work presented in this thesis focuses on the problem of reconstructing three-dimensional models of fire from real images. The intended application of the reconstructions is for use in research into the phenomenon of fire-induced high voltage flashover, which, while a common problem, is not fully understood. As such the reconstruction must estimate not only the geometry of the flame but also the internal density structure, using only a set of a few synchronised images. Current flame reconstruction techniques are investigated, revealing that relatively little work has been done on the subject, and that most techniques follow either an exclusively geometric or tomographic direction. A novel method, termed the 3D Fuzzy Hull method, is proposed, incorporating aspects of tomography, statistical image segmentation and traditional object reconstruction techniques. By using physically based principles the flame images are related to the relative flame density, allowing the problem to be tackled from a tomographic perspective. A variation of algebraic tomography is then used to estimate the internal density field of the flame. This is done within a geometric framework by integrating the fuzzy c-means image segmentation technique and the visual hull concept into the process. Results are presented using synthetic and real flame image sets.

Acknowledgements

I would like to thank my supervisor, Dr. Derek Hoch, for his academic and financial support and for introducing the project. Thanks must also go to Geoffrey Bennet for his help in the high voltage lab.

On a more personal note I must also express my gratitude to my parents for their support and encouragement during this time. I would also like to thank my girlfriend, Monica, for her encouragement and help with illustration software.

Contents

Title Page	i
Preface.....	i
Abstract.....	ii
Acknowledgements.....	iii
Contents	iv
List of Figures	vii
1. Introduction.....	1
1.1 Fire-Induced High Voltage Flashover Theories.....	5
1.1.1 Reduced Air Density Theory	5
1.1.2 Particle Initiated Flashover	6
1.1.3 Flame Conductivity Theory	6
1.1.4 Summary	7
1.2 Flame Reconstruction as a Research Tool	8
2. Fire Physics and Measurement	10
2.1 Flame Physics	10
2.1.1 The Candle	10
2.1.2 The Radiative Nature of Fire.....	14
2.1.3 Flame Refraction.....	16
2.1.4 The Gaseous Nature of Fire	16
2.2 Image Formation Model.....	17
2.2.1 Soot Volume Model	17
2.2.2 Simplified Radiative Transfer Model.....	18
2.3 Flame Measurement and Diagnostic Tools.....	20
2.3.1 The Thermocouple	20
2.3.2 Two-Wavelength Pyrometry	20
2.3.3 Schlieren Photography	21
2.3.4 Laser Methods.....	22

3.	Computerised Tomography	24
3.1	Projections.....	24
3.2	Fourier Slice Theorem and Filtered Back Projection.....	26
3.3	Algebraic Tomography	30
3.4	Tomographic Segmentation	35
3.4.1	Fuzzy C-Means Image Segmentation	36
3.4.2	The Level Set Technique	39
3.4.3	Level Set Applications.....	44
4.	Flame Reconstruction	46
4.1	Object Reconstruction - the Visual and Photo Hull.....	46
4.2	Flame Modelling.....	50
4.3	Geometric Flame Reconstruction.....	50
4.3.1	Contour-based Reconstruction.....	50
4.3.2	Stereoscopic Reconstruction	52
4.4	Tomographic Flame Reconstruction.....	53
4.4.1	Flame sheets and blobs	54
4.4.2	Visual Hull Limited Tomographic Inversion.....	58
4.5	Summary - The Hybrid Approach	61
5.	The Fuzzy Hull Method	63
5.1	Overview.....	64
5.2	Experimental Set-up.....	64
5.3	Camera Modelling and Calibration.....	66
5.3.1	Pinhole Camera Model.....	67
5.3.2	Camera Calibration	68
5.4	Image Processing	70
5.5	The Reconstruction Methods	72
5.5.1	Simple – SIRT and Fuzzy	73
5.5.2	The Fuzzy Hull Method	74
5.5.3	3D Fuzzy Hull.....	84
5.6	Level Set Reconstruction	91
5.7	Flame reconstruction analysis.....	92
5.7.1	Geometric Analysis.....	92

5.8	Summary	94
6.	Experimental Results	96
6.1	The Setup	96
6.2	Experimental Issues	99
6.2.1	System Configuration	99
6.2.2	EMI	99
6.2.3	Real World Factors	100
6.3	Results	100
6.3.1	3D Renderings.....	101
6.3.2	2D Cross-Sectional Analysis	106
6.3.3	3D Reconstruction Analysis.....	108
6.3.4	Summary	109
7.	Conclusion	111
7.1	Future Work	112
	Bibliography.....	114

List of Figures

Figure 1.1: The 2D geometry used for the electric field simulations in [16].	8
Figure 2.1: The humble candle.	11
Figure 2.2: The elements of a candle flame.	12
Figure 2.3: A Bunsen burner produces a premixed flame.	14
Figure 2.4: Absorption model for Beer-Lambert Law derivation.	15
Figure 2.5: The creation of a Schlieren Image. (based on [18].)	22
Figure 3.1 : The projection of an object at angle θ .	25
Figure 3.2: The Fourier Slice Theorem.	27
Figure 3.3: The Filtered Backprojection process.	28
Figure 3.4: FBP reconstructions of a head phantom.	30
Figure 3.5: Setting up the reconstruction problem algebraically on a regular grid. (Based on [31].)	31
Figure 3.6: SIRT reconstructions for head phantom and flame phantom.	34
Figure 3.7: (b) A 5 view SIRT reconstruction of (a). (c) FCM result after 4 iterations. (d) Modified FCM performed on (c). (e) The SIRT image (b) is masked with (d) to give a segmented tomographic image.	39
Figure 3.8: Level Set evolution of 2 points with constant speed subject to cube boundary.	40
Figure 3.9: The narrow band Fast Marching method. The narrow band of TRIAL points marches through the grid turning DEAD points into ALIVE points.	43
Figure 4.1: The visual hull (shown in grey) of a triangle formed by 3 cameras.	47
Figure 4.2: The photo hull (shown in grey) of the same object as in Figure 4.1. Note how a closer approximation to the original shape is obtained.	49
Figure 4.3: Flame contour arrangement.	51
Figure 4.4: Example of β -spline mesh flame reconstruction, from [5].	52
Figure 4.5: Example of stereoscopic flame reconstruction from [6].	53
Figure 4.6: Demonstration of tomographic ambiguity.	53

Figure 4.7: (a) A flame sheet solution for images I_1 and I_2 . (b) The multiplicative solution.....	55
Figure 4.8: Examples of flame reconstruction from [18] using (a) sheets and (b) blobs.....	58
Figure 4.9: Trilinear basis function covering 8 voxels, with transparent isosurfaces indicating the function values (from [9]).	60
Figure 4.10: Example of flame reconstruction based on tomographic inversion from [9].	61
Figure 5.1: Process overview.....	64
Figure 5.2: Camera configuration. The cameras should be roughly equally spaced – e.g. 60° apart for 3 cameras.....	65
Figure 5.3: The pinhole camera model.	67
Figure 5.4: Calibration object geometry.	70
Figure 5.5: (b) is the (cropped) histogram of image (a).	71
Figure 5.6: Reconstruction of test object from 5 views using SIRT and FCM.....	73
Figure 5.7: Flowchart outlining the Fuzzy Hull algorithm with orphan estimation. .	75
Figure 5.8: Reconstructions of test object #1.....	78
Figure 5.9: Reconstructions of test object #2.....	79
Figure 5.10: (a) Synthetic flame-like test object. (b) Example of generated image.	80
Figure 5.11: 5-view reconstruction using (a) method A, (b) the fuzzy hull method..	81
Figure 5.12: Reconstructions superimposed on original object. On the left is method A and on the right is the fuzzy hull method. (c) and (d) have been rendered with partial transparency.	82
Figure 5.13: Graphical representation of 3D reconstruction accuracy analysis.....	83
Figure 5.14: A set of 4 projections subject to the hull optimisation algorithm.....	86
Figure 5.15: (a) shows the discrete definition of the rays using pixels, while (b) is the same but using oversampling. (c) and (d) are corresponding density renderings of a reconstruction of a candle flame differing only in the ray definition used.	87
Figure 5.16: Cross-sectional reconstruction slice showing the effect of using (b) oversampling and (c) the basis function, compared with a slice reconstructed using (a) neither.	88
Figure 5.17: Surface renderings of 2 reconstructed frames rotated through 180°.	89

Figure 5.18: Density renderings of 2 reconstructed frames rotated through 180°.....	89
Figure 5.19: An example showing a cube and the polygons defined by its corner points.....	92
Figure 5.20: 2D example of an ambiguous case.....	93
Figure 5.21: Rule set for voxel-based surface area estimation.....	94
Figure 6.1: High-voltage arrangement.....	97
Figure 6.2: (a) The five cameras positioned around the grid. (b) The copper rod is connected to the high voltage source in the lab below and held over the grid. (c) The burner.....	98
Figure 6.3: Density and surface renderings of various frames of a reconstructed flame sequence. 160° is an original viewing angle while the rest are novel views...	105
Figure 6.4: (a) shows a raw, unprocessed captured image of a flame, and (b) shows the density rendering of the reconstruction of the corresponding frame. Note that part (b) is based on a relative density reconstruction and will thus not necessarily have the same absolute intensity as the original photograph.	106
Figure 6.5: Comparison of reconstructed cross-section.....	107
Figure 6.6: Several cross-sections taken from the reconstruction of a single flame.	107

Chapter 1.

Introduction

It is known that fires burning underneath high voltage transmission lines can cause flashover – a short circuit from the transmission line to earth. Such transmission line faults are not only inconvenient but can be extremely costly to industry.

An analysis of the Eskom Main Transmission System line faults (in South Africa), for the period 1993 to 1999, was done by Britten [1],[2]. His research showed that 19.3% of all faults were caused by fire, 15.6% by grass fires and 3.7% by sugar cane fires. The only factors more significant than fire were storms (28.8%) and “unknown” (38.2%).

One can thus see that fire-induced flashover is a serious problem, further compounded by the fact that such faults are difficult to prevent. Sugar cane fires are intentional, and therefore controllable, yet the problem remains.

More recently a satellite-based fire detection system has been tested. The CSIR and Eskom developed the Advanced Fire Information System, for the purpose of prediction, detection and assessment of fires in South Africa [3]. Ideally such a system would be able to detect small fires within a large region in near real time. However this ideal satellite does not exist and instead two available satellites, one

geostationary and one polar orbiting, were used. This means that the one satellite can provide detailed high resolution images, but only infrequently, while the other provides low resolution images, but can do so every 15 minutes. It is estimated that the low-res satellite (MSG) can detect fires of a minimum size of 500m^2 , while the high-res satellite (MODIS) has a minimum detection size of 100m^2 . Combining the two satellites, results were obtained showing the detection of 60% of flashover causing fire, with up to 80% of those fires detected before flashover occurred. The system shows great promise, although of course detecting the fire is only the first step – effective communications and systems must be in place in order to quickly react and actually prevent flashover.

However, the actual process of fire-induced flashover is complex and dynamic, and is not completely understood. The three main theories as to the mechanism involved are: the reduced air density theory, particle initiated flashover and the flame conductivity theory. These theories are reviewed in §1.1. A better understanding of the phenomenon could allow for further preventative measures to be developed.

The goal of this project is to create a system that can create a three-dimensional reconstruction of a flame, using photographic images. The primary purpose of the reconstruction would be its use for further research regarding the flame conductivity theory. Three-dimensional electric field models could be estimated and the internal conductivity of the flame analysed. As such, the focus of this project is not on the actual use of the reconstruction for research, but rather on the development of a tool to enable the creation of 3D reconstructions of real world flames.

Keeping in mind the intended application the following specifications were decided upon to guide the development of the system:

- The reconstruction must be volumetric. By this it is meant that the reconstruction space is formed by elementary, regular cells (voxels) which can be definitively classified as being part of the object or not.

- The internal structure of the flame should be represented by an estimation of the density field.
- The method should be general, able to handle any type of visible flame, not require specific laboratory conditions, and should not be too limiting regarding the size of the flame or the reconstruction.
- In order to maintain the generality and portability of the system, camera calibration should be kept to a minimum.
- The system must be able to handle video sequences, not just a single frame, and must therefore be able to create the reconstruction in a reasonable amount of time.
- The video sequences should be instantaneously captured in real time – in other words no axis-symmetric flame assumptions or time-averaged images.
- No expensive equipment should be needed other than several cameras and a PC.

This means that flame analysis techniques using lasers or other specialised equipment are not directly relevant. A brief review of such flame measurement and analysis techniques is presented in §2.3.

In terms of flame reconstruction from images there are two different approaches one can take. The first approach is that of geometric reconstruction. As the name implies these methods are concerned with the physical shape of the fire, creating a reconstruction representing the surface or volume of the flame [4],[5],[6]. The simplest geometric method is the *visual hull*, formed by the intersection of the object's silhouettes.

The second approach involves the use of tomography. Tomography makes use of several side-views, or projections of the object, representing the integral of the internal mass density of the object from the particular view. This density information is then used to estimate the structure of the internal density field of the object. Probably the most common application of tomography can be found in medicine. CAT (Computerised Axial Tomography) uses a series of X-Ray images to construct cross-sectional slices of the human body. The field of tomography is discussed in more detail in chapter 3. Tomography-based flame reconstruction methods have been developed [7],[8],[9], although these reconstruction methods are typically based on only a few images (as opposed to over 100 images in medical applications), and as such require different techniques to those used in medical imaging. Previous work on flame reconstruction from photographic images is reviewed in chapter 4.

This thesis presents a hybrid flame reconstruction technique, termed the *Fuzzy Hull* method. By using elements from both geometric and tomographic approaches this method is able to create a flame reconstruction, from only a few images, that estimates both the geometry and the internal density structure of the flame.

The simplified radiative transfer image formation model (§2.2.2) provides the foundation for a tomographic reconstruction by converting the flame images into an approximation of the flame's ray density integrals (the relative density along the ray forming each image pixel). Algebraic tomography is then used to create an estimate of the flame's internal density field, based on the converted images. Fuzzy c-means images segmentation then classifies the reconstruction voxels as either object or background, based on statistical features, giving a volumetric nature to the reconstruction. The visual hull is used as a boundary for these operations. The evolution and testing of the fuzzy hull method is presented in chapter 5.

Finally, in chapter 6, results are presented from full-scale tests performed with a hot air balloon burner, in an environment simulating a fire burning under a high voltage transmission line.

1.1 Fire-Induced High Voltage Flashover Theories

There would appear to be three main theories relating to the phenomenon of fire-induced flashover. These are outlined below.

1.1.1 Reduced Air Density Theory

An air gap is essentially an insulator, and as such the effectiveness of this insulation is determined by its physical properties. Under normal conditions one can relate the flashover voltage of an air gap to the density of the air by [10]

$$V_t = V_s \frac{D}{H} \quad , \quad (1.1)$$

where D is the relative air density, H is a humidity correction factor and V_s is the flashover voltage under standard temperature and pressure. Expanding the relative density term to

$$D = \frac{p}{p_0} \frac{T_0}{T} \quad (1.2)$$

(where p is pressure, T is temperature, $p_0 = 101.3$ kPa and $T_0 = 293$ °K) and assuming a negligible humidity correction factor gives

$$V_t = V_s \frac{2.892 p}{T} \quad . \quad (1.3)$$

What this implies is that the flashover voltage will decrease, in a linear fashion as the air density is decreased. This density reduction is caused by the heat of the fire, therefore the hotter the flame, the lower the voltage required for flashover (as shown by the inverse temperature relationship of equation (1.3)).

Experimentally it has been shown that while this relationship holds at lower temperatures, it does not apply to the higher temperatures typically found in fire [12] (sugar cane can have an adiabatic flame temperature of over 2000°K [11]). This

indicates that the reduced air density theory is not sufficient to account for the reduction in flashover voltage experienced at such high temperatures.

The effects at these high temperatures are not well understood and further research is required [12],[13].

1.1.2 Particle Initiated Flashover


It has been shown that introducing particles into the flame reduces the air breakdown voltage [10],[12],[14]. A typical source of such particles would be ash from the burning of a solid fuel (e.g. a sugar cane fire). The effect of the particles on the dielectric strength is strongly affected by the concentration and size of the particles, and can be more significant in smaller gaps [14]. However, other factors such as flame turbulence, wind, and the large air gaps under high voltage transmission lines must be considered.

One must conclude that although the presence of particles will reduce the breakdown voltage, this is unlikely to be the dominant breakdown mechanism.

1.1.3 Flame Conductivity Theory

The principle behind the flame conductivity theory is that, because of the conductive nature of fire, the flame acts a conductor, thereby effectively reducing the air gap.

The typical reaction occurring in hydrocarbon-based flame is



However, the mechanism by which this takes place involves a complex series of elementary steps involving highly reactive molecular fragments such as H·, ·OH and ·CH₃ [15]. These atoms and free radicals only have a transient existence, but a steady concentration is maintained as they are continuously formed in chain reactions.

In addition, if temperatures are high enough then the gas atoms or molecules may have enough energy to ionise other atoms and molecules upon collision – thermal ionisation. Thermal ionisation creates a state of matter called plasma, and the degree of ionisation is described by the Saha ionisation equation.

The ionisation of the air above the flame also contributes to lowering the electric field gradient required for breakdown. These high temperature ionisation effects require further research in order to be fully quantified [12][13].

Both of the above-mentioned factors, chemical oxidation and thermal ionisation, are therefore responsible for the concentration of ions and electrons found in a flame. The actual conductivity of the fire can be estimated by the gas conductivity equation

$$\sigma = \rho\mu \quad , \quad (1.5)$$

where ρ is the charge density and μ is the ion motility.

Sukhnandan has carried out simulations to investigate the effect of flame conductivity on the electric field [16]. Using finite element analysis software the 2D electric field was modelled for the case of a fire under a conductor with an air gap. Different flame conductivities were used and the electric field at the conductor surface and flame tip was analysed. His results show that the conductivity of a flame is indeed an important factor in the breakdown process.

1.1.4 Summary

It would appear that all three theories play a part in explaining the fire-induced flashover phenomenon. While the reduced air density theory is an established effect, it does not explain the drastic reduction of air insulation strength at higher temperatures [10]. The introduction of particles alters the electric field and can be viewed as effectively reducing the air gap distance. The degree of breakdown voltage reduction caused by particles can be significant, and varies according to

particle size [10]. Flashover can still be induced however using clean burning LPG flames, and thus the particle effect cannot be seen as the sole mechanism, but rather as an aid to the process. The effect of the flame conductivity in the breakdown process would appear to be that of increasing the electric field gradient in the air gap between the flame and the conductor, thereby making it easier to attain the gradient needed to initiate flashover. Thermal ionisation, both within the flame and in the air, is also a contributing factor, although the precise effect has yet to be fully quantified.

1.2 Flame Reconstruction as a Research Tool

The main purpose of this project is to create a tool able to create a three dimensional reconstruction of fire. This reconstruction can then be used in further research regarding the effect of flame conductivity on the fire-induced flashover phenomenon. As mentioned, work has been done using 2D simulations and a simplified representation of fire, shown below [16].




Figure 1.1: The 2D geometry used for the electric field simulations in [16].

However, such investigations could be more accurate and thorough if a 3D simulation was used, incorporating a 3D reconstruction of the actual flame coupled with actual physical measurements. The electric field could be more completely

modelled, since the flame geometry would be more accurate, with features such as sharp points influencing the electric field formation.

In addition, the approximation of the internal flame density could be used for further analysis. The geometry and varying density of the flame mean that the resistance of the flame is not uniform, and therefore any conductivities measured are only representative of the flame as whole, and might not accurately reflect individual regions.

Chapter 2.

Fire Physics and Measurement

2.1 Flame Physics

2.1.1 The Candle

According to Michael Faraday:

“There is no more open door by which you can enter into the study of natural philosophy than by considering the physical phenomena of a candle. There is not a law under which any part of this universe is governed which does not come into play, and is not touched upon, in these phenomena.” [17]

Indeed, the common candle provides a good example with which to introduce the physical and chemical processes associated with fire.



Figure 2.1: The humble candle.

Consider a steady-state candle flame. Part of the heat generated by the flame is radiated downwards, melting the candle wax to form a small pool of wax around the wick. The liquid wax is drawn up the wick by capillary action and vapourised by the heat. The wax vapour then diffuses across the dark zone, mixing with oxygen. The ensuing combustion in the blue zone generates the most heat in the flame. The combustion process continues up through the yellow region until completion. This is where the main visible light is generated, as the sooty carbon particles are burnt.





Figure 2.2: The elements of a candle flame.


The shape of the flame is caused by convection – the heated gasses expand and ascend. The self-sustaining reaction is able to continue, however, due to the downward and inward flow of radiated heat. Just enough heat reaches the wax to melt the wax at the same rate it is being burned, and the radiation from the blue combustion region ensures the liquid wax is vapourised.

The combustion process is exothermic, meaning energy is released, and typically the process proceeds as shown below: [17]




This is of course only a generalised representation, with many intermediate species forming and reactions occurring along the way, both exothermic and endothermic (requiring energy) in nature. The wax molecules are subject to a thermal “cracking”, being fragmented into various molecules, free radicals and ions. In the dark region

of the flame, where oxygen is scarce, reactions occur without oxidation, creating carbon rich molecules that accumulate to form soot particles: [17]



As the soot particles rise upwards, oxidation occurs on the surface of the particles, creating carbon dioxide:



These incandescent soot particles are responsible for most of the candlelight.

The many complex steps in the chemical process involving highly reactive molecular fragments create a steady concentration of various transient atoms and free radicals. The exact details are beyond the scope of this project, but what should be noted is that these chemical reactions are likely to contribute to the conductivity of the flame.

Although only candles have been looked at here, the physics and chemistry involved are generally valid when looking at organic (carbon based) fuel combustion. The manner in which oxygen is introduced to the reaction can vary, however. Diffusion flames rely on oxygen to diffuse from the air, and premixed flames use a fuel already mixed with oxygen, for example a Bunsen burner.




Figure 2.3: A Bunsen burner produces a premixed flame.

2.1.2 The Radiative Nature of Fire

When one looks at a flame what one is actually seeing is the radiation of visible light, emitted by the reactions occurring within the flame. This means that a flame can be characterised, in terms of temperature and chemical composition, by analysing the spectrum of emission from the flame.

In terms of radiation a *blackbody* is an object that is a perfect emitter of radiation. The spectral radiance, L , of an object at a certain temperature (T) and wavelength (λ) is given by Planck's equation:

$$L(T, \lambda) = \frac{2\xi_{\lambda}\pi c^2}{\lambda^5 (e^{hc/\lambda k_B T} - 1)} \quad , \quad (2.1)$$

where k_B is Boltzmann's constant and h is Planck's constant. The term ξ_{λ} (where $0 \leq \xi_{\lambda} \leq 1$) is the spectral emissivity, which indicates how well the object represents an ideal blackbody. If $\xi_{\lambda} = 1$ then the object is acting like a blackbody and emitting maximum radiance.

Integrating Planck's equation over all wavelengths gives

$$L(T) = \left(\frac{2\pi^4 k_B^4}{15c^2 h^3} \right) T^4 \quad (2.2)$$

which is the Stefan-Boltzman equation. This shows how the total radiance of a blackbody increases exponentially with temperature.

The absorptivity (α_λ) of an object can be defined analogously to emissivity. That is, the absorptivity of an object is the fraction of presented absorbance relative to the absorbance of a blackbody (at the same temperature). This comes about since the absorptivity of a blackbody is equal to its emissivity (according to Kirchoff's law) and a blackbody is therefore also an ideal absorber of radiation.

The reduction in intensity as a beam of light passes through a medium is described by the Beer-Lambert equation. The derivation shown here is based on [18].




Figure 2.4: Absorption model for Beer-Lambert Law derivation.

Consider an infinitesimally thin slab of material of unit area, through which light passes. The absorption process is modelled by using small, completely opaque regions that block the light by an amount proportional to their area. The area of these particles is defined as σ_i , known as the extinction cross-section. Using the density of medium (ρ) one can then represent the total area of extinction as $\sigma_i \rho(s) ds$, giving the total reduction in radiance I_λ as

$$dI_\lambda = -\sigma_t \rho(s) I_\lambda(s) ds \quad . \quad (2.3)$$

Integrating along path L gives

$$\ln I_\lambda(L) - \ln I_\lambda(0) = -\sigma_t \int_0^L \rho(s) ds \quad . \quad (2.4)$$

Calculating absorptivity (assuming blackbody extinction) gives

$$\alpha_\lambda = \frac{I_\lambda(0) - I_\lambda(L)}{I_\lambda(0)} = 1 - \exp\left(-\sigma_t \int_0^L \rho(s) ds\right) \quad . \quad (2.5)$$

Assuming constant density gives the Beer-Lambert Law:

$$\alpha_\lambda = 1 - \exp(-\sigma_t \rho L) \quad . \quad (2.6)$$

This expression for the absorptivity (and consequently the emissivity) of an object allows the formulation of the Soot Volume model of flame image formation (§2.2.1).

2.1.3 Flame Refraction

Viewing the background through a flame will produce a distorted image, due to the refraction of light in the flames. Although the deflection is usually less than 0.1° [19], this is enough to cause noticeable distortions. The refractive index is related to the flame density and a turbulent flame will thus cause a dynamic image warping. The most common solution is to use a uniformly black background. This phenomenon can be exploited though, through Schlieren photography [19],[20], discussed in §2.3.3.

2.1.4 The Gaseous Nature of Fire

Because of the high temperature of fire, the reaction products are mainly in the gaseous phase. As such fire is subject to the laws of fluid mechanics. Although fluid mechanics form the basis of fire modelling and simulation [27], there is not any work

that uses such models for the purpose of physical flame measurement or reconstruction [18].

2.2 Image Formation Model

In order to justify the reconstruction of fire from photographic images one must look at how the images are formed and how this process can be modelled and therefore form the basis of a reconstruction system. The imaging models presented here are based on the work presented in [18] and [8].

2.2.1 Soot Volume Model

The typical luminous yellow colour of a diffusion flame is caused by the glowing carbon-based particles. These soot particles are approximately 10 to 100 nm in diameter and approximate minute blackbodies [15]. If these particles are smaller than the emitted radiation wavelength (visible light is 400-700 nm) then the soot volume model says that the emissivity should be proportional to the soot density [15].

Using Kirchoff's Law one can equate emissivity and absorptivity, allowing one to calculate a flame's emission according to the light absorption of an equivalent soot density. Therefore the Beer-Lambert Law can be re-stated to give the flame emissivity:

$$\xi_{\lambda} = 1 - \exp\left(-\sigma_{\lambda} \int_0^L \rho(s) ds\right) . \quad (2.7)$$

The total emitted radiance can then be calculated, using the emissivity and temperature. However, this means that in order to create an image formation model based on the soot volume model one needs to explicitly account for the temperature and the variations thereof within the flame. The simplest solution is to assume constant temperature, although this is of course not ideal.

2.2.2 Simplified Radiative Transfer Model

The simplified radiative transfer model of fire uses two simplifications:

1. Negligible scattering: For relatively smokeless fires this is a good approximation since the radiance is dominated by the self-emission from glowing soot particles. The total emission thus consists only of self-emission.
2. Constant self-emission: By modelling the brightness of the fire as being dependent only on the density field, one may assume the self-emission to be constant, per unit mass, denoted by Q_0 . This does mean though that temperature is not properly considered and as a result density fields reconstructed using this model will be a function of density and temperature, instead of true mass density. However, the effect of this is dependent on temperature variations, not absolute temperature, and so the smaller the temperature variations the more accurate a density reconstruction should be.

These two simplifications allow the expression of the intensity of a given ray as:

$$I = (1 - \tau)Q_0. \quad (2.8)$$

The total transparency, τ , along a given ray is given by the radiative transfer model as the integral of the density field, $\rho(x)$, along the ray:

$$\tau = \exp\left(-\sigma_t \int_l \rho(x).dx\right), \quad (2.9)$$

where σ_t is a medium dependent constant relating density and transparency, known as the extinction cross-section.

To account for the emission constant, Q_0 , it is noted how equation (2.1) shows that, as the transparency tends to zero, the intensity approaches Q_0 , thus giving a maximum intensity of $I_\infty = Q_0$. Given a digital imaging system with a saturation of I_{max} , we can then say that $I_\infty = I_{max} + 1$. Thus, eliminating Q_0 gives:

$$I = (1 - \tau)I_{\infty} . \quad (2.10)$$

Rearranging this equation in terms of τ :

$$\tau = 1 - \frac{I}{I_{\infty}} , \quad (2.11)$$

and finally, combining equations (2.9) and (2.11) and manipulating gives a transformed intensity:

$$\begin{aligned} I' &= \sigma_t \int \rho(x) dx \\ &= -\ln \left(1 - \frac{I}{I_{\infty}} \right) . \end{aligned} \quad (2.12)$$

This transformed intensity, I' , corresponds to the integral of the density field along the ray, thus showing that computerised tomography solutions can be applied to the problem of flame reconstruction.

It should be mentioned that equation (2.12) assumes a black background. To account for an arbitrary background, equation (2.10) is modified to include a background intensity term:

$$I = \tau I_{bg} + (1 - \tau)I_{\infty} . \quad (2.13)$$

The transformed image intensity then becomes:

$$I' = -\ln \left(\frac{I - I_{\infty}}{I_{bg} - I_{\infty}} \right) . \quad (2.14)$$

This allows one to relate the measured flame intensity (an image pixel) with an integral of the flame density along the ray represented by that pixel. It is this density integral which is required for a tomographic reconstruction of the density field.

2.3 Flame Measurement and Diagnostic Tools

There exists much work regarding the measurement and analysis of flames, for example in the field of combustion fuels. Much of this research is based on two dimensional analyses of flame, focused on anything from temperature and chemical species distribution to flame-front imaging and tracking [21],[22],[23],[24],[25],[26]. If needed, several sets of the 2D data can then be used to create a 3D distribution, or reconstruction, of the flame. The reconstruction of an object from density-based projections is termed *tomography*. The topic of tomography is looked at in detail in chapter 3, while the application of tomography to flame and combustion research is discussed in §4.4.

The fact that normal cameras are being used in this project, to obtain 2D data, makes much of the published research regarding non-photographic flame measurement techniques irrelevant. With this in mind, this section looks briefly at some of the more common or partially relevant flame measurement techniques.

2.3.1 The Thermocouple

The simplest method for measuring the temperature of a flame is to place a thermocouple within the flame. Consisting of two different metals that expand at differing rates, the voltage across a thermocouple will vary according to its temperature. Obviously this invasive method lacks the spatial and temporal accuracy needed for measuring the temperature distribution of a flame, but nonetheless remains a simple and effective technique for elementary temperature determination.

2.3.2 Two-Wavelength Pyrometry

Pyrometry, literally meaning “fire measuring”, is widely used for temperature measurement, due to its accuracy and non-invasive nature [18],[28]. This optical technique makes use of Planck’s equation, which defines the relationship between spectral radiance, L , and the emitted wavelength, λ , of an object. Measuring the

spectral radiance at two different wavelengths, $L(T, \lambda_1)$ and $L(T, \lambda_2)$, and manipulating Planck's equation allows one to obtain the temperature:

$$T = \frac{c_2 \left(\frac{1}{\lambda_1} - \frac{1}{\lambda_2} \right)}{\ln \left(\frac{L(T, \lambda_2)}{L(T, \lambda_1)} \right) + 5 \ln \left(\frac{\lambda_2}{\lambda_1} \right) + \ln \left(\frac{\xi_{\lambda_1}}{\xi_{\lambda_2}} \right)} \quad (2.15)$$

where $c_2 = \frac{hc}{k_B}$. The emissivities, $0 \leq \xi \leq 1$, which indicate how well the object represents an ideal blackbody, are also needed.

Other optical techniques used for temperature measurement include: coherent anti-Stokes Raman scattering, light induced fluorescence, Rayleigh scattering, interferometric and speckle methods [28].

2.3.3 Schlieren Photography

A technique that has some relevance in terms of 3D reconstruction of fire is Schlieren Photography [19],[20]. This method takes advantage of the fact that fire is a transparent medium, which deflects light transmitted through it. Spatial variations of the gas density alter the refractive index of the flame, and the resultant deflection of light is made visible using a Schlieren system. The image is formed by modifying the brightness according to the degree of deflection.




Figure 2.5: The creation of a Schlieren Image. (based on [18].)

A light source is passed through Lens 1, sending parallel light beams through the object. Lens 2 directs the light to the Schlieren filter, before arriving at the camera. The Schlieren filter, or knife-edge filter, stops half the light intensity, so as to remove certain harmonics.

Schwarz presented a complex method [19] where he was able to measure the extinction coefficient, spectral radiance, and refractive index using a Schlieren device and perform a tomographic reconstruction of the temperature distribution of the flame. The intricate apparatus he created consisted of a light source, 120 mirrors, 4 lenses, a rotating mirror, a Schlieren filter and a CCD camera, measuring 1.2m in height, and 2.6m in diameter. The circular apparatus meant that non-axially symmetric flames could be studied, but the ability to only take measurements from one direction at a time means that the flame needs to be relatively stationary. It is this need for a complex system and a lack of temporal resolution that makes the Schlieren method unsuitable for our task.

2.3.4 Laser Methods

Lasers are a common tool in flame diagnostics, being non-invasive, versatile, and can be made to handle the high temperatures and pressures one might find, for example, in a combustion chamber. Measurements that lasers have been used to obtain include: [29]

- chemical species identification and concentration
- temperature,
- pollutants,
- droplet/particle size,
- the velocity and flow of air and fuel.

Planer Laser Induced Fluorescence (PLIF or LIF) is a technique that uses suitable optics to create a laser sheet illuminating a slice of the flame [22],[23]. Seeding of the flame, manipulation of wavelength and spectral analysis form the basis of LIF related techniques.

While lasers are undoubtedly a useful tool for flame analysis, they are not suitable for this project, since a laser-based 3D reconstruction system would be expensive, overly complex and have limited reconstruction capabilities.


Chapter 3.

Computerised Tomography.

Tomography can be defined simply as the reconstruction of an object or image from projections. Originally introduced by Radon in 1917, the most important application of tomography is undoubtedly X-ray computerised tomography, for which Hounsfield received a Nobel prize in 1972 [30]. X-rays are transmitted through the body and a receiver measures their attenuation, giving a proportional mass along the corresponding rays. This data is then used to recreate cross-sections of the human body. There are of course many other wide-ranging applications for tomography, such as non-destructive testing in manufacture, 3D imaging using electron microscopy [31], and the reconstruction of supernova remnants [30], to name a few.

3.1 Projections

The projections used in tomographic reconstruction are ideally formed by the line integral of a certain parameter of the object [31]. Figure 3.1 below illustrates how such projections are formed and defined. To use a typical example, assume that the object is a two-dimensional slice of human tissue through which x-rays are being transmitted in a straight line. The attenuation suffered by the x-rays, forming the projection, is represented by the corresponding line integral through the object.

Figure 3.1 : The projection of an object at angle θ .

The object is described by the function $f(x,y)$, and the one-dimensional projection, P_θ , is parameterised by t . Each ray passing through the object (such as line AB) has the equation

$$x \cos \theta + y \sin \theta = t, \quad (3.1)$$

and the line integral $P_\theta(t)$ can be expressed as

$$P_\theta(t) = \int f(x,y) ds. \quad (3.2)$$

Rewriting using the dirac delta function, $\delta(t)$, gives

$$P_\theta(t) = \iint f(x,y) \delta(x \cos \theta + y \sin \theta - t) dx dy, \quad (3.3)$$

which is known as the Radon transform. This linear expression provides a mathematical description defining the collection of rays forming a parallel projection.

Not all projection geometries are parallel of course, and indeed the standard configuration for a medical CT is a fan-beam spiral/helical scanning system [32]. While a complete volumetric reconstruction is traditionally obtained by stacking a series of 2D slices, this is not always the case. With a cone-beam geometry there is a single point source and a detector plane behind the object. Obviously the different imaging geometries need different projection formation models, but the fundamental principles remain the same.

3.2 Fourier Slice Theorem and Filtered Back Projection

The Fourier slice theorem is a fundamental theorem in computerised tomography, relating the 2D Fourier transform of an object to the 1D Fourier transform of a projection of the object. The 2D Fourier transform of the object function is defined as

$$F(u, v) = \iint f(x, y) e^{-j2\pi(ux+vy)} dx dy, \quad (3.4)$$

and likewise the Fourier transform of a projection of the object at angle θ as

$$S_{\theta}(\omega) = \int P_{\theta}(t) e^{-j2\pi\omega t} dt. \quad (3.5)$$

If one now writes the 2D Fourier transform of the object along the line $v=0$ in the frequency domain, one obtains (from equation (3.4))

$$F(u, 0) = \iint [f(x, y) dy] e^{-j2\pi ux} dx, \quad (3.6)$$

where the integral has been split into two parts. If one looks at the section in the square brackets one can see that it in fact defines a projection (equation (3.2))

parallel to the y axis, at an angle of $\theta=0^\circ$. Rewriting equation (3.6) to show this gives

$$F(u, 0) = \int P_{\theta=0}(x) e^{-j2\pi ux} dx, \quad (3.7)$$

and finally, incorporating equation (3.5) gives

$$F(u, 0) = S_{\theta=0}(u). \quad (3.8)$$

In essence, what this says is that the Fourier transform of a parallel projection of an object gives the slice of the 2D Fourier transform of the object at 90° to the direction of the projection. This is illustrated in Figure 3.2. A more comprehensive proof can be done which shows the Fourier Slice Theorem is valid irrespective of projection angle and coordinate system [31].




Figure 3.2: The Fourier Slice Theorem.

Theoretically, forming the 2D Fourier transform of the object from all the Fourier transforms of the projections (by the Fourier Slice Theorem) and then performing an inverse Fourier transform should yield a perfect reconstruction. Practically, of

course, this is not possible since the image space is discrete, not continuous, there are only a finite number of projections available, and the method is sensitive to noise.

The dominant CT reconstruction algorithm for the last 30 years has been Filtered Backprojection (FBP) [18], based on the Fourier Slice Theorem. The first part of this method, filtering, is used to help reduce noise and to apply the proper weighting to the projections in the frequency domain. The second stage, backprojection, is where the 2D inverse Fourier transform of the filtered projections are summed over the image to give the final reconstruction. The reasons for first inverting each projection individually and then adding them, are that the inversions can be performed quickly since the values are only along a single line, and the process can begin as soon as a single projection is acquired. Also, when summing the contributions of the projections, it is more accurate to perform interpolations spatially than in the frequency domain [31]. While FBP can produce very good reconstructions, many views are required (typically more than a hundred [7]), and as such the method is not used for applications where only a few views are available.

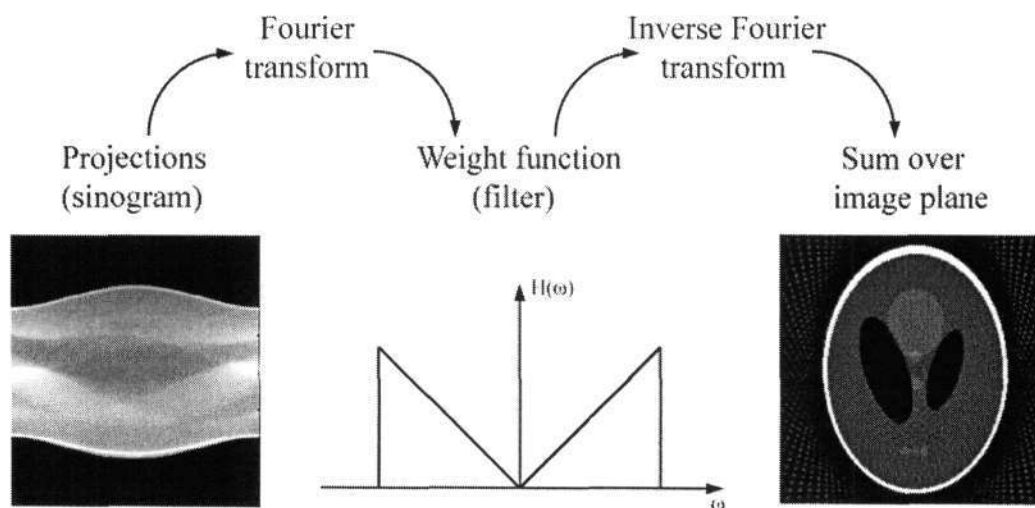


Figure 3.3: The Filtered Backprojection process.

Figure 3.3 shows the filtered backprojection process. A sinogram is simply a collection of all the projections of the object formed into a single image, where each column represents a projection.

A special case is where the object in question is axially symmetric. This is a reasonable assumption to make in flame research when the image of the flame is time averaged [18],[7],[33]. Since all projections are identical, only one is required, and the Radon transform can be modified to give the Abel transform [18],

$$P(t) = 2 \int_0^{\infty} \frac{rf(r)}{\sqrt{r^2 - t^2}} dr, \quad (3.9)$$

where $f(x,y)$ is replaced by $f(r)$ since the density of the object is dependant only on the radial distance r . The reconstruction is obtained using the Inverse Abel Integral Equation [7]:

$$F(r) = -\frac{1}{\pi} \int_r^{\infty} \frac{\partial P(t)}{\partial t} \frac{1}{\sqrt{t^2 - r^2}} dt. \quad (3.10)$$

Figure 3.1 demonstrates Filtered Backprojection using a Shepp Logan head phantom. This is commonly used to evaluate reconstruction techniques, and is meant to represent a cross-sectional slice of the human head. In this case the phantom was generated by MATLAB[®]. Note how only the 150 view reconstruction is precise enough for medical applications.

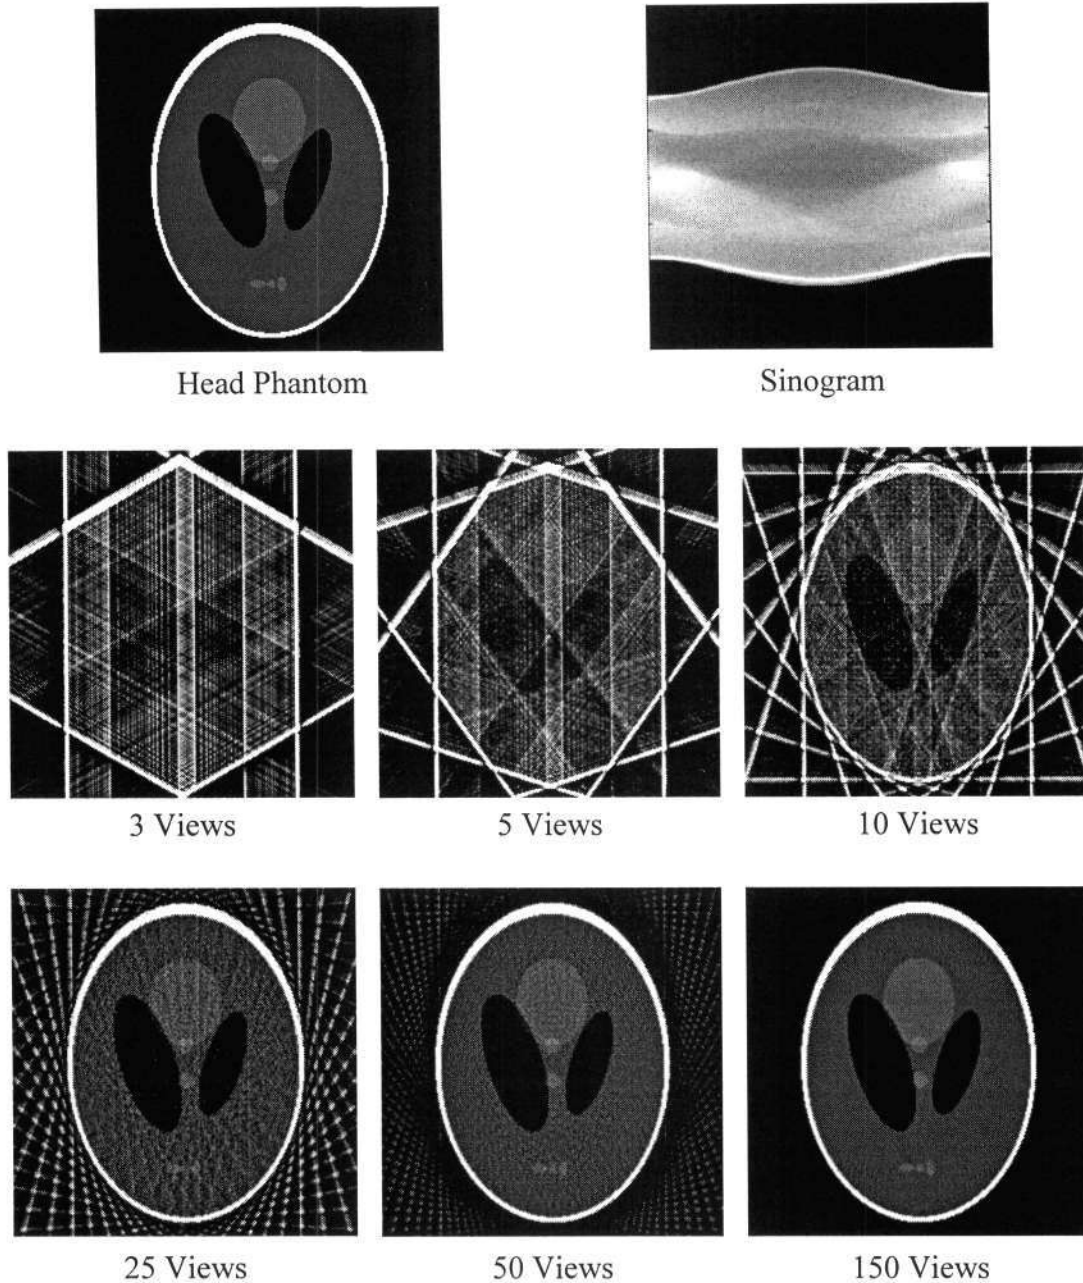


Figure 3.4: FBP reconstructions of a head phantom.

3.3 Algebraic Tomography

Algebraic techniques are useful when only a limited number of views are available (typically less than ten) giving better reconstructions than FBP for few view problems [7], and handling noisy data better [18].

To look at the reconstruction problem from an algebraic perspective, one needs to assume that the reconstructed image consists of a number of unknown values, and then set up equations for these unknowns in terms of the image projections. The task is then reduced to finding values for the image points that best satisfy the equations.

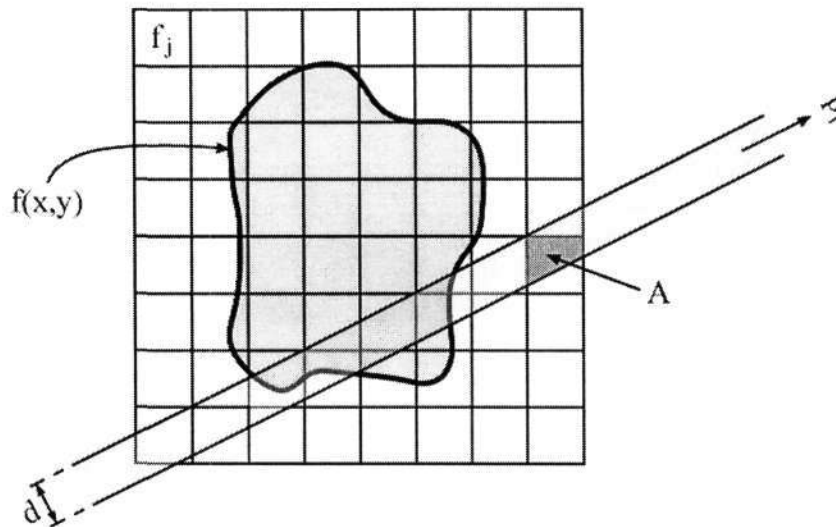


Figure 3.5: Setting up the reconstruction problem algebraically on a regular grid. (Based on [31].)

The image space of the reconstruction is defined by a regular grid, as shown in Figure 3.5. The unknown image is denoted by $f(x,y)$, and f_j indicates the value of image cell j , where there are N image cells in total. Each projection is defined as being composed of a series of rays, denoted p_i , having a finite width d . The rays should be similar in width to the image cells. Each ray has a raysum, equivalent to the line integrals used earlier to model the formation of projections. As such, each ray can be expressed as

$$p_i = \sum_{j=1}^N w_{ij} f_j, \quad (3.11)$$

in a manner analogous to the Radon transform (equation (3.3)). The term w_{ij} is the fraction of cell j falling within ray i , shown as area A in Figure 3.5. The contribution of each cell towards a specific raysum is thus proportional to the area of that cell

intersecting the ray. To simplify calculations it is common to replace w_{ij} with a binary operator [31], which is 1 if the centre of the cell is within the ray (ie. if $w_{ij} > 0.5$).

Because the number of unknowns is large and usually outnumbers the known values, the reconstruction cannot be done analytically and hence an iterative method is the best option. The simplest method of reconstruction, Algebraic Reconstruction Technique (ART), is based on updating each image cell according to the difference between the measured projection and the reconstructed raysum. The correction to cell j in ray i is calculated by

$$\Delta f_j^i = \frac{p_i - q_i}{N_i}, \quad (3.12)$$

where p_i is the measured projection value of the ray, q_i is the reconstructed raysum of the ray, and N_i is the number of cells in the ray. Effectively what is happening is the difference between the actual and reconstructed raysums is smeared back along the ray. An improvement can be made by modifying equation (3.12) to give

$$\Delta f_j^i = \frac{p_i}{L_i} - \frac{q_i}{N_i} \quad (3.13)$$

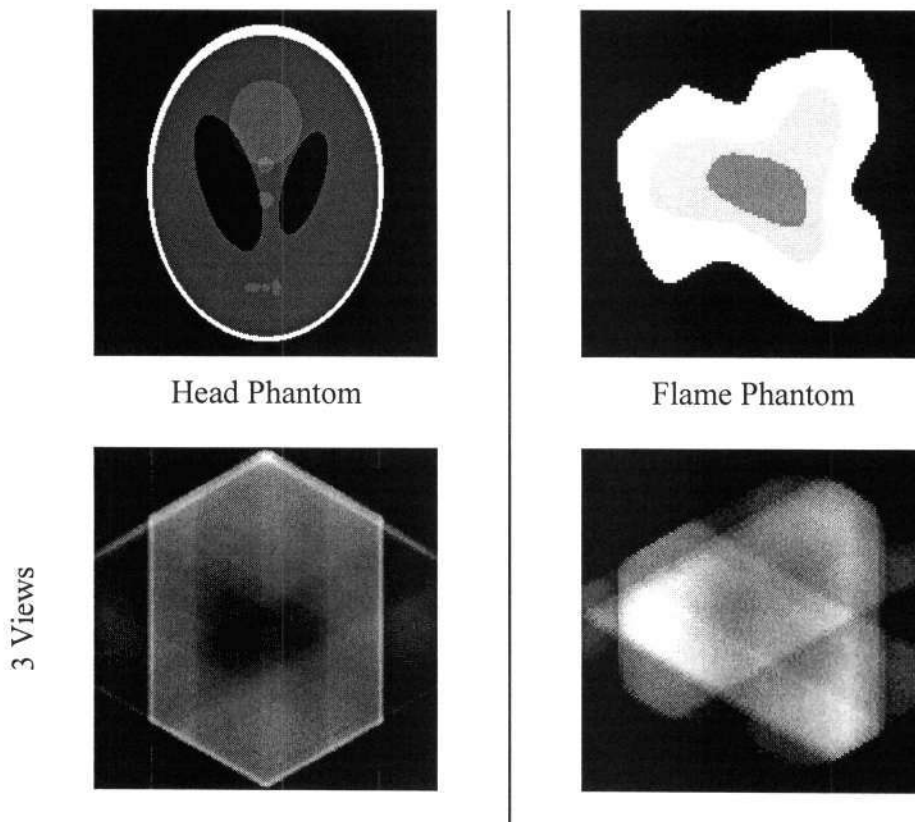
where L_i is the length of the ray in the image.

Reconstructions made using ART have a tendency to suffer from salt and pepper type noise [31]. Improved images can be created by using Simultaneous Iterative Reconstruction Technique (SIRT). Equation (3.13) is still used, but the image cell is only updated after considering the contribution by the all the rays that intersect the cell. Rewriting equation (3.13) to reflect this gives

$$\Delta f_j = \frac{1}{M} \sum_{i=0}^{i=M} \left(\frac{p_i}{L_i} - \frac{q_i}{N_i} \right). \quad (3.14)$$

Both ART and SIRT are iterative techniques, and the image cells of the reconstruction are updated until convergence is reached. The criteria for convergence is usually based on the difference between the measured and reconstructed raysums, or on the degree of correction performed on the image cells in the iteration. An advantage of algebraic techniques is that they can easily be modified, or adapted for a certain application. One possible modification is to add a relaxation parameter to equation (3.14), gradually decreasing the fraction of the calculated Δf that is actually applied to the image cells [31].

SIRT reconstructions for different numbers of projections are presented in Figure 3.6 below, for both the head phantom used to demonstrate FBP, and a more flame-like phantom.






Figure 3.6: SIRT reconstructions for head phantom and flame phantom.

Obviously, more views produce a better reconstruction, but what is more relevant in this case is the performance with only a few views. The SIRT reconstructions tend to produce a smoother image, with fewer artefacts. This smoothness means a loss of detail, but this is irrelevant unless one has many views, in which case SIRT would not be used anyway.

The reconstructions with only a few views may not appear to be very good, but this is because of the nature of the problem, not the method of reconstruction. A few views provide only sparse information with which to recreate an entire cross-section, and as

such the goal is merely to obtain a reasonable estimate. The only way to improve this estimate is if one has prior knowledge. By taking a probabilistic approach one can force the reconstruction to favour certain properties which the mass distribution of the real object is thought to have, such as smoothness, constant density or a certain shape. Such *maximum a posteriori* (MAP) estimates are usually defined in a Bayesian framework [34],[35],[36]. For the purpose of this project the only reasonable prior knowledge that could be incorporated is the tendency of fire to have a degree of local smoothness in terms of its internal density distribution.

While FBP and algebraic techniques form the basis of tomography, there are many possible reconstruction techniques, often incorporating some form of MAP criteria, and obviously these cannot all be discussed here. These include: view interpolation [37], the Expectation Maximisation (EM) algorithm [18], simulated annealing [38][39], wavelets [40], neural networks [41], least squares techniques [42], parametric or deformable models [43] and level set techniques [44].

3.4 Tomographic Segmentation

After a tomographic reconstruction has been created there is sometimes the need to differentiate between various elements within the image. This could mean merely distinguishing the object from the image background, or identifying different regions, for example, tissue classification in a brain scan. Furthermore, applications may call for a 2D or 3D segmentation. Both methods presented here are extendable to any dimension.

Looking at Figure 3.6 one can see that when using only a few views the reconstruction is not well defined and has a generally fuzzy appearance. Therefore any methods based on line/edge detection, region growing and thresholding are probably not a very good option. Statistical models, using a prior model and parameters, generally require supervision and are not really suited to this project. Neural networks are potentially useful, and have been used for the segmentation of

medical images [45][46]. However, they generally require *a priori* information and of course training. The two methods of segmentation discussed here are Fuzzy c-means and the Level Set technique.

3.4.1 Fuzzy C-Means Image Segmentation

Since it is a tomographic reconstruction that needs to be segmented, it makes sense to look at methods used in other applications in which tomographic segmentation is performed. One of the largest places to look is in medical imaging. With X-ray tomography (CT), magnetic resonance imaging (MRI), single photon emission tomography (SPECT), and positron emission tomography (PET), there is a whole range of tomographic images needing to be segmented [47]. A popular solution is the use of soft, or fuzzy, segmentation techniques, where image cells are classified into various sets according to a membership value, indicating the strength of the classification [47],[48].

The fuzzy c-means algorithm (FCM), introduced by J Bezdek [49], is an iterative method that seeks to classify the image cells into a number of classes. Each class has a centroid, which defines the central feature values for that class. Membership values are assigned to each image cell based on the distance of their feature values to that of the class centroid. The typical features used are the statistical mean and standard deviation [50], although any feature and any number of features can be used.

There are several advantages to the FCM algorithm:

- It is an unsupervised method, although this does of course mean one has little control over the final segmentation.
- One can have an image of any dimension and segment into any number of classes using any number of features.
- The membership values are normally distributed.

Consider an image with n pixels or cells, to be segmented into c classes. Let $\{x_1, x_2, \dots, x_n\}$ form set X , where each element x_k is a vector containing the image

features of cell k . Each class, or fuzzy subset, of X , is defined by a centroid v_i , giving the central features of the class. X is partitioned into these subsets by assigning a fuzzy membership value u_{ik} to each image cell, indicating the similarity of cell k to each class i , where the following conditions must be satisfied [51]:

$$0 \leq u_{ik} \leq 1 \quad \forall i, k \quad (3.15)$$

$$\sum_{i=1}^c u_{ik} = 1 \quad \forall k \quad (3.16)$$

$$0 < \sum_{k=1}^n u_{ik} < n \quad \forall i. \quad (3.17)$$

The important point here is equation (3.16), which says that the sum of all the membership values of a particular cell, for all classes, must be equal to one. In other words the membership value does not tell one which class an image cell belongs to, but rather the degree to which it belongs to all classes.

The FCM algorithm attempts to cluster feature vectors by minimizing the energy function

$$J_m(U, v; X) = \sum_{i=1}^c \sum_{k=1}^n (u_{ik})^m d_{ik}, \quad (3.18)$$

by iteratively updating the membership and centroid values (U and V) using

$$u_{ik} = \left(\sum_{j=1}^c \left(\frac{d_{ik}}{d_{jk}} \right)^{\frac{2}{m-1}} \right)^{-1} \quad \forall i, k \quad (3.19)$$

$$v_i = \frac{\sum_{k=1}^n u_{ik}^m x_k}{\sum_{k=1}^n u_{ik}^m} \quad \forall i. \quad (3.20)$$

The term d_{ik} is a similarity measure between centroid i and feature vector k , typically the Euclidean distance, and the factor m is weighting exponent that controls the

crispness of classification. If m is set to 1 then one obtains a crisp classification, with $u_{ik} \in [0,1]$. A larger m gives a greater fuzziness, meaning an increased distribution of membership values. A typical value would be $m = 2$. [51]. The iteration process stops when usually $\sum \|\Delta u_{ik}\|$ or Δv_i is sufficiently small.

Mohamed et al [51] suggest a modification to improve the performance of FCM with a noisy image. The calculation of d_{ik} is modified to include the effect of neighbouring pixels, using

$$d_{ik} = d_{ik} \left(1 - \alpha \frac{\sum_{j \in \text{neighbours}} u_{ij} * p_{kj}}{\sum_j p_{kj}} \right). \quad (3.21)$$

The u_{ik} of neighbouring pixels is considered, weighted relative to their proximity to the central pixel k . The term p_{kj} is a measure of this proximity, where $p_{kj} = \|k - j\|$. The factor α ($0 \leq \alpha \leq 1$) is a constant controlling the strength of the modification, with a suggested value of 0.5. This new method of calculating d_{ik} works as an adaptive filter, improving clustering while maintaining edges.

Once FCM is complete one must defuzzify the results if requiring a hard classification. Obviously the simplest way to do this is to assign the image cell to the class with which it has highest membership. Figure 3.7 below shows the result of an FCM classification on a tomographic image from Figure 3.6, with the aim of separating the object from the background. The first few iterations were performed with normal FCM, followed by several iterations using the modified FCM.




Figure 3.7: (b) A 5 view SIRT reconstruction of (a). (c) FCM result after 4 iterations. (d) Modified FCM performed on (c). (e) The SIRT image (b) is masked with (d) to give a segmented tomographic image.

3.4.2 The Level Set Technique

The Level Set technique was introduced in 1988 by Osher and Sethian [52] for the purpose of tracking a moving interface. Consider a curve, in 2D, or a surface, in 3D, separating two regions, where the interface can propagate in a direction normal to itself with a known speed. Using level sets one can track the evolution of this front. The motion of the interface, or zero level set, is determined using partial differential equations in a dimension one greater than that of the curve/surface being tracked. This increased dimensionality allows cusps and sharp corners to form, and changes in topology are implicitly handled [53]. In contrast, the more traditional active

contour or snake methods use a parameterised interface defined by a set of marker points, making topological changes very difficult to handle [54]. Although the added dimensionality that level sets brings to the problem is very powerful and convenient, it does of course have the disadvantage of increasing the computational complexity and hence processing time. There are other advantages to using a Level Set formulation: [55]

- The techniques used for the numerical solutions of hyperbolic conservation laws can be exploited to create accurate computational schemes.
- Geometric properties such as curvature or the normal vector are easy to obtain.

For a simple illustration, consider an invisible bounding cube and two points inside the cube. Allowing these points to propagate normally at constant speed produces the evolution shown in Figure 3.8, where each stage shown is the zero level set at certain time intervals. Note the topological change when the two regions merge, and how easily this is handled by defining the front as the zero level set of a higher dimension.




Figure 3.8: Level Set evolution of 2 points with constant speed subject to cube boundary.

To find an equation of motion for the evolving front, the initial position of the front is first defined as the zero level set of a higher dimensional function ϕ . It is a requirement that the level set of the evolving function must match the propagating front, written as

$$\phi(x(t), t) = 0. \quad (3.22)$$

Expanding with the chain rule gives

$$\phi_t + \nabla \phi(x(t), t) \cdot x'(t) = 0. \quad (3.23)$$

The speed of propagation, F , is defined in a direction normal to the surface, written as

$$x'(t) \cdot n = F, \quad (3.24)$$

where

$$n = \frac{\nabla \phi}{|\nabla \phi|}. \quad (3.25)$$

This gives the level set equation, defining the evolution of ϕ :

$$\phi_t + F|\nabla \phi| = 0, \quad (3.26)$$

$$\text{given } \phi(x, t = 0). \quad (3.27)$$

This is the initial value level set equation [52].

Alternatively, one could pose the problem as a front advancing towards a boundary, leading to a boundary value formulation. If the restriction is imposed that the speed can be only positive, $F > 0$ (or only negative), then the position of the advancing front can be defined as the arrival time $T(x, y)$ at which the front crosses each point (x, y) . The calculation of the arrival time of this monotonically advancing front is then relatively simple. Using the fact that *distance = speed * time*, one has

$$dx = F(dT) \quad \text{and hence} \quad 1 = F \frac{dT}{dx}. \quad (3.28)$$

In higher dimensions the spatial derivative of T becomes the gradient, giving rise to [56]

$$\begin{aligned} |\nabla T| F &= 1, \\ \text{where } T &= 0 \text{ on } \Gamma, \end{aligned} \quad (3.29)$$

where F is the initial position of the front. This boundary value level set problem is solved using the Fast Marching Algorithm [57].

3.4.2.1 Fast Marching Implementation

The implementation presented here is based on that in [57]. In order to solve equation (3.29) the gradient is approximated by

$$\left[\max(D_{ij}^{-x}T, -D_{ij}^{+x}, 0)^2 + \max(D_{ij}^{-y}T, -D_{ij}^{+y}, 0)^2 \right] = \frac{1}{F_{ij}^2} \quad (3.30)$$

where D^- and D^+ are backward and forward difference operators for point (i,j) . This approximation is easily extended to three dimensions. Since the equation forms a quadratic, one can solve for the value of T at each grid point, taking the largest answer as the solution.

The Fast Marching algorithm is based on the fact that equation (3.30) has an upwind difference structure, meaning that the values of T propagate from smaller values to larger ones. The front is therefore swept forwards in an upwind manner, considering only those points adjacent to the front in a narrow band. This narrow band is marched forward, locking confirmed values of T and bringing new points into the narrow band. The Fast Marching Algorithm is outlined below.

Consider a 2D grid through which a front is propagating. Each grid point is assigned a state value, which can be ALIVE, TRIAL or DEAD. All points through which the front has passed are designated ALIVE, all points adjacent to ALIVE points are marked TRIAL, forming the narrow band, and all other points are said to be DEAD.

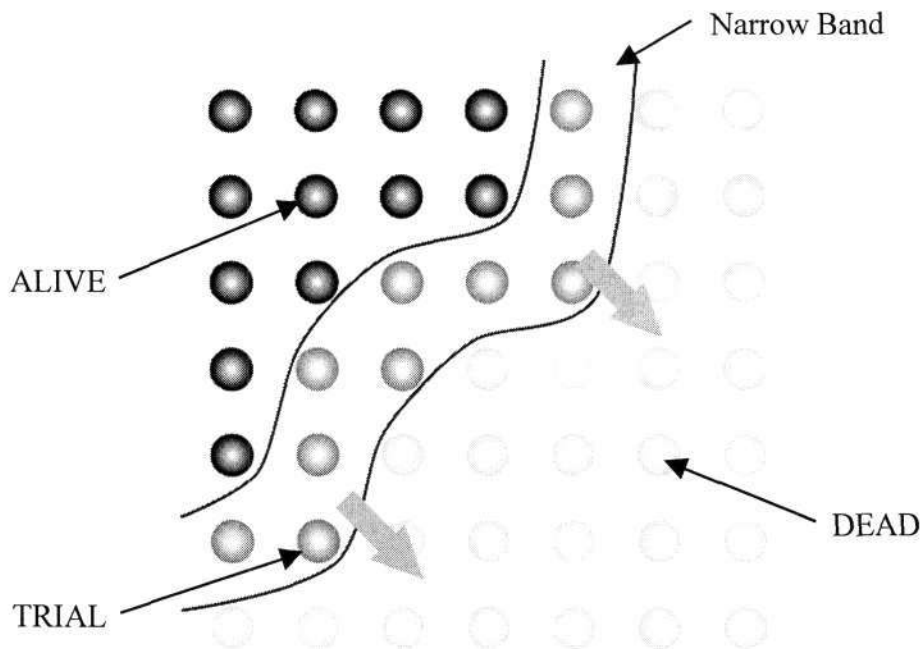


Figure 3.9: The narrow band Fast Marching method. The narrow band of TRIAL points marches through the grid turning DEAD points into ALIVE points.

1. Initialisation

- a. Define initial front at $T=0$.
- b. Mark all grid points as ALIVE, TRIAL or DEAD.
- c. Calculate T for all trial points using equation(3.30).

2. Marching

- a. Take the smallest TRIAL point ($T_{ij}(min)$) and change its status to ALIVE. The T value for that point is now locked.
- b. Check all neighbours of $T_{ij}(min)$ – if any are DEAD change them to TRIAL, forming part of the narrowband.
- c. Recalculate T for all TRIAL values neighbouring $T_{ij}(min)$ according to equation (3.30).
- d. Return to step 2a.

The reason why this algorithm works is that one is always selecting the smallest value in the trial band, which must be correct since the other values in the trial band

are larger (or further away) and thus cannot influence its arrival time. The recalculation of the neighbouring TRIAL points cannot give a smaller value than any of the accepted ALIVE points since the largest solution to the quadratic T calculation is always taken. The front thus marches outwards, continually selecting the minimum trial value and updating its neighbours.

The key to an efficient implementation of the fast marching technique lies in finding a method to quickly determine the lowest valued point in the trial band. In [58] Malladi and Sethain suggest using a variation on a heap algorithm with back-pointers to store the T values, resulting in an overall processing time for the entire algorithm of $O(N \log N)$, where N is the number of points processed in the structure.

The definition of the speed function F depends on the application. For example, in an image segmentation problem, using Fast Marching, F might be based on the image gradient. For problems involving curvature dependant propagation speed one would use a level set approach (initial value) since F can be both positive and negative.

3.4.3 Level Set Applications

The potential applications of the Level Set technique (including Fast Marching) are numerous and varied, and hence only a few will be listed here.

- Image processing [59],
- Shape modelling and reconstruction [60],[61],
- Flame propagation and modelling [27],
- Medical Imaging [62],[63]
- Grid Generation [64],
- Optimal path planning [56],
- Seismic calculations [56],
- Two phase flow (two fluid interface problems) [56],
- Etching and deposition in semi-conductor manufacture [65].

Elangovan [44] uses Level Sets to directly segment tomographic reconstructions. Instead of first creating the reconstruction and then segmenting, he uses Level Sets and works directly with the sinograms (tomographic projections) to find an optimal interface between object and background that best fits the projection data. An iterative process updates the surface model, followed by the object and background density estimations. This technique is useful where normal reconstructions are not suitable for normal segmentation, for example as a result of incomplete sinograms, limited angle tomography, measurement noise or view misregistration. The limitation of this method, however, is the assumption of constant object density.

For a more in-depth look at Level Sets and their applications, the reader is referred to the work by Sethian [55].

Chapter 4.

Flame Reconstruction

This chapter looks at previous work that has been done on the topic of fire reconstruction from normal photographic images. Methods of reconstruction using other techniques such as lasers or Schlieren photography were covered in §2.3. Before focusing on the reconstruction of solely fire however, one should investigate the reconstruction of objects in general.

4.1 Object Reconstruction - the Visual and Photo Hull

The simplest method to reconstruct an object is to use the visual hull concept. Introduced by Laurentini [4], the visual hull of an object is the convex volume created by the intersection of the silhouettes of the object (Figure 4.1). Strictly speaking, the visual hull of an object is the shape that gives the silhouette of the object when viewed from any direction. To create this one would need an infinite number of photos. Therefore the term *visual hull* is more commonly taken to mean the shape formed using N silhouettes.

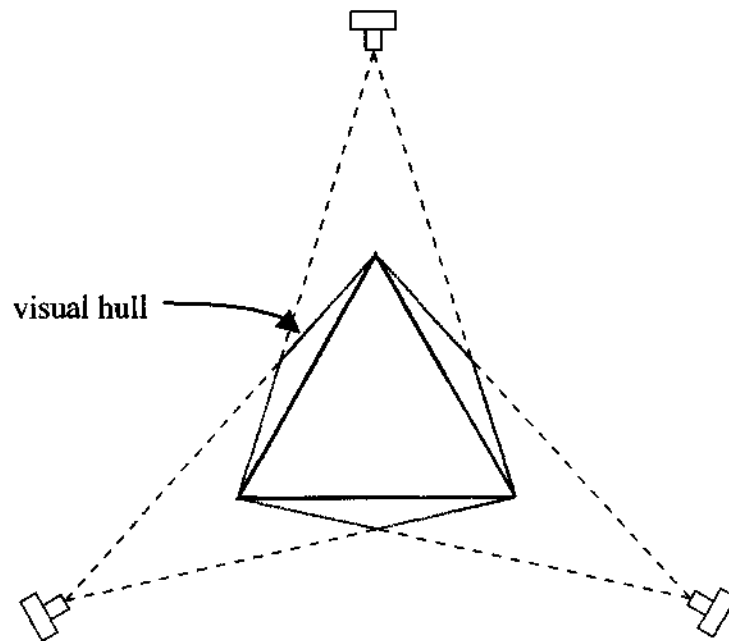


Figure 4.1: The visual hull (shown in grey) of a triangle formed by 3 cameras.

If one compares the ideal reconstruction of an object with its visual hull (using any number of silhouettes), one notices two interesting properties of the hull. Because of the way the visual hull is constructed, it is guaranteed to contain the object. The visual hull can be said to be a maximal estimation of the object, meaning it is the best possible reconstruction, using only silhouette information, which is known to contain the object. Secondly, since the visual hull is formed by the intersection of the projections of the silhouettes this means that the reconstruction is convex. Because concave regions of the object are not apparent in its silhouettes, the visual hull cannot reconstruct such concave areas.

The calculation of the visual hull of an object is relatively fast and simple and can be represented in two different ways:

- **Voxels:** A voxel (volume element) can be thought of as a 3D pixel. A voxel-based reconstruction is easy to calculate and is useful when further processing is to be done on the hull.

- Polyhedral: A mesh is used to define the surface of the hull. Using a polyhedral mesh to represent the hull is more complex than voxels as a mesh generation algorithm needs to be used.

Polyhedral meshes create a surface reconstruction, and can be difficult to manipulate, while voxels provide a volumetric reconstruction that is easy to work with. A voxel representation is also more suitable when considering tomography, and therefore the reconstruction methods used in this thesis are voxel-based.

Typically a visual hull is not very geometrically accurate, since the visual hull is always larger than the object. The accuracy obtainable depends on factors such as the number of silhouettes, the viewing directions, and of course the nature of the shape of the object. Forbes et al [66] effectively increases the number of available silhouettes by using several sets of silhouettes of the same rigid object. He uses a silhouette consistency constraint to calculate the relative poses between the sets, which can then be combined to create a visual hull that is more accurate than using only a single set. However this is not feasible with a dynamic object such as fire.

To enhance the accuracy of the visual hull reconstruction, more information than just the silhouette is needed. The most logical choice is to use colour (or luminance for grey-scale images). To this end Kutulakos and Seitz [67] introduced the concept of the *photo hull*. Whereas the visual hull is consistent only with the silhouettes of the object, the photo hull must also be photo-consistent. By photo-consistent it is meant that the colour of a point in the direction of the camera is the same as that observed when the point is projected to the photograph. This is illustrated in Figure 4.2. The photo hull is thus unique and is found by identifying the spatially largest set of voxels that are photo-consistent with all the photographs taken of the scene [68]. In addition, it is the tightest possible bound on the shape of the true scene that can be inferred from N photographs, in the absence of *apriori* geometric or point correspondence information [67].




Figure 4.2: The photo hull (shown in grey) of the same object as in Figure 4.1. Note how a closer approximation to the original shape is obtained.

The computation of the photo hull can be done using space carving [67],[68],[69]. The general approach used is to remove (carve) voxels that are not photo-consistent with the reference photographs. This is done iteratively (since as a voxel is removed the visibility of other voxels changes) until the entire reconstruction is photo-consistent. Voxel colouring, or generalized voxel colouring is probably the predominant space carving technique [70][71]. For a more detailed look at voxel colouring and object reconstruction in general, the reader is referred to Slabaugh [68].

The problem with any colour-based space carving method is that the colour or radiance needs to be locally computable – the radiance of any point must be independent of the radiance of any other point [67]. A common special case is the Lambertian scene, where a voxel will project a similar colour when viewed from any of the reference views. A Non-Lambertian model can be used, but may require the calibration of light sources and the calculation of surface normals [72]. In addition, objects that are homogeneously textured will not work well with space carving. This

means that transparent objects cannot be reconstructed using such methods. Attempts have been made to handle transparent or semi-transparent objects [73],[74], but they are at best limited and are generally aimed more at handling transparency within the reconstruction of a scene, rather than a transparent object on its own.

It is thus apparent that, given the semi-transparent nature of fire and its lack of texture, the use of a photo hull is not suitable for the reconstruction of fire. Therefore some other method is needed to improve upon the visual hull when considering fire.

4.2 Flame Modelling

It is not unusual in computer vision to fit a model to real world data to create a reconstruction, for example full body motion capture. Within the scope of this project this is obviously not practical, given the dynamic and complex nature of fire. Nguyen et al [27] present a physically based method for the modelling and animation of fire. Using the level set method to track the reaction zone, and the incompressible Navier-Stokes equations to model vapourised fuel and hot gaseous products, they create very realistic fire. However, to try and fit their modelling technique to real world fire would require the measurement of extensive parameters and is not currently feasible.

4.3 Geometric Flame Reconstruction

4.3.1 Contour-based Reconstruction

Yan et al [5],[75] developed a fast and simple method of flame reconstruction. Using three cameras placed equidistantly around the flame, the contours of the flame from the photographs are extracted. The contours are then arranged three dimensionally and β -splines are used to interpolate between them to create a surface mesh. Figure 4.3 below illustrates this contour arrangement. This produces a

smooth, natural looking reconstruction which suits the fluid nature of fire, and is thus probably more accurate than using the angular visual hull. While appealing, such a contour-based approach still produces an essentially convex reconstruction with no internal flame data, and is not likely to handle more complex flame geometries very well.




Figure 4.3: Flame contour arrangement.

This technique was developed for the purpose of monitoring and characterising fossil-fuel combustion systems and is well suited to real-time monitoring applications because of its speed. In addition the need for only three cameras keeps costs low. In terms of characterisation of the flame, several parameters are quantified, namely: volume, surface area, orientation, length, circularity and the uncertainty of these parameters.




Figure 4.4: Example of β -spline mesh flame reconstruction, from [5].

4.3.2 Stereoscopic Reconstruction

Another common tool in computer vision is stereo vision, which emulates the way in which the human vision system works. If one takes two images of the same scene, from a slightly different viewpoint, one can calculate the depth of a point or feature by identifying it in both images (given a set of calibrated cameras). It is this problem of correspondence, finding the same point in both images that provides the challenge in stereoscopic systems. Obstacles to overcome include occlusion, homogeneously coloured regions, multiple matches and non-Lambertian reflectance.

Given these possible problems it is fair to assume that fire is not an ideal candidate for stereoscopic reconstruction. Ng and Zhang [6] have however attempted to reconstruct a flame surface using stereo techniques. The first problem with their work though, it would appear, is the use of only turbulent impinging diffusion flames. This is done to create the necessary detail for stereo matching, but is of course very limiting. Secondly, only two images were used and occlusion was not accounted for, which means the result is only a partial surface reconstruction and not a full three dimensional reconstruction. Therefore one can conclude that stereoscopic reconstruction is not suitable for general 3D flame reconstruction.




Figure 4.5: Example of stereoscopic flame reconstruction from [6].

4.4 Tomographic Flame Reconstruction

The problem with tomography using only a few views is that the task is ill posed – there are many more variables than projection equations. The fewer the number of views the more ambiguous the process of reconstruction becomes. This is illustrated in Figure 4.6. Using a simple data set one can see how any of the three solutions shown are viable and yet completely different. Of course as the number of views increases the degree of ambiguity lessens. What one should understand from this is that any sparse view tomographic technique is at best an estimation, and that a particular method should be chosen because it is more suited to the application, and not only because it performed well in a different application.




Figure 4.6: Demonstration of tomographic ambiguity.

Recently, the use of tomography in combustion research has received a fair amount of attention [7],[33],[76],[77],[78],[79],[80]. However, these techniques often suffer from several problems:

- Poor spatial resolution: Typically because of sensor limitations due to physical constraints [76]. For example, reconstructions limited to cross-sectional slices [77],[81].
- Poor temporal resolution: Flame images may need to be captured consecutively, not synchronously [19]. Alternatively several frames are averaged together to form a time-averaged image [33].
- Axially symmetric flame assumptions: A 3D reconstruction can be inferred from a single image if the flame is assumed to be symmetrical about its axis (for example a non-turbulent Bunsen burner flame) [7], [79],[80].
- Expensive or complex apparatus, such as lasers, fibre optic sensors and the Schlieren apparatus [19],[77],[81].

Typically these techniques are designed for a particular industrial application, and therefore certain limitations may be acceptable in the given case. These limitations usually arise because of physical conditions (e.g. within a combustion chamber), or because of the sensors used, such as photodiodes, fibre optics and capacitive sensors. In addition there is often the need for special combustion chambers and controlled conditions.

The following research is more applicable to this project, using tomographic techniques based on normal photographic images, to produce visually realistic, density-based 3D reconstructions of fire.

4.4.1 Flame sheets and blobs

Hasinoff and Kutulakos [8],[18] show that by modelling fire as a semi-transparent 3D density field one can approach the problem from a tomographic perspective. This was demonstrated in §2.2. In [8] they solve the under-constrained problem by creating a set of density field solutions for every pair of photographs in the form of monotonic curves or sheets. The final reconstruction is formed by a convex combination of the density sheets derived from each pair of input views. As shown

in Figure 4.6 the solution to a two-view tomographic problem may be ambiguous, and the flame sheet method is merely a particular form of solution for this case, that results in a sheet-like, minimal size reconstruction. For specific details on its implementation the reader is referred to [8].

I_1
↓

0.1				0.1
0.5			0.25	0.25
0.25	0.05	0.15	0.05	
0.15	0.15			
	0.2	0.15	0.3	0.35

← I_2

(a)

I_1
↓

0.1	0.02	0.015	0.03	0.035
0.5	0.1	0.075	0.15	0.175
0.25	0.05	0.0375	0.075	0.0875
0.15	0.03	0.0225	0.045	0.0525
	0.2	0.15	0.3	0.35

← I_2

(b)

Figure 4.7: (a) A flame sheet solution for images I_1 and I_2 . (b) The multiplicative solution.

An example of a simple flame sheet is shown in Figure 4.7 above. Part (a) shows two input images, I_1 and I_2 , and the corresponding flame sheet solution. In essence, the sheet is created by starting at the bottom left and filling the image cell with the largest value possible, before moving to the next image cell. Part (b) shows a more conventional solution, formed by the multiplication of corresponding row and column input values.

It should be noted that their goal was to create photo-consistent 3D reconstructions of fire, and visually their results look good. However, from a geometric aspect the reconstructions are not accurate since they are formed from a set of flat 2D solutions. In the reconstruction of a geometrically complex flame one can see how the flame consists of 2D sheets. When considering a problem with only 2 input views, however, the decomposed flame sheet method is likely to be a good choice, preferable to a straightforward multiplicative solution.

In his thesis [18] Hasinoff proposes using blobs to create the reconstruction. By representing each cross-sectional slice as a density field parameterised by a superposition of Gaussian blobs, and by fixing the positions and variances of the blobs, he shows how the reconstruction problem is reduced to a constrained linear least-squares problem.

A normalised, symmetric Gaussian distribution, centred at (μ_x, μ_y) , is described by the equation:

$$G(x, y) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{(x - \mu_x)^2 + (y - \mu_y)^2}{2\sigma^2}\right) . \quad (4.1)$$

This allows a density field to be represented as a weighted sum of r blobs:

$$\rho(x, y) = \sum_{i=1}^r w_i G_i(x, y) . \quad (4.2)$$

Fixing the positions and variances of the blobs allows a forward projection model to be defined that is dependant only on the blob weights. The line integral of blob G over line ℓ can be expressed as

$$\Gamma(G, \ell) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(\frac{-d^2}{2\sigma^2}\right) , \quad (4.3)$$

where d is the distance from the line to the blob centre. The formation of pixel $I(\ell)$ can therefore be described as

$$I(\ell) = \sum_{i=1}^r w_i \Gamma(G_i, \ell) . \quad (4.4)$$

By limiting the number of blobs to less than the total number of 1D input pixels, the system is reduced to an overdetermined set of equations, allowing standard least squares techniques to be used.

The uniform distribution of the blob basis functions is not ideal, meaning many of the blobs, already limited in number, may be being used for empty areas. Therefore, the next stage of Hasinoff's method involves using the initial least squares reconstruction as a guide for repositioning the blobs. This allows the limited number of blobs to be more effectively distributed, thereby enabling regions of greater detail to be more accurately reconstructed. The blob variances can then too be altered, using similar considerations, or by constructing a Delaunay triangulation and using the distance to the furthest neighbour to determine the variances.

Of course there are many possible variations of this method. Not only could different types of basis functions be used, but the manner in which they are distributed or redistributed can also make a significant difference to the quality of the reconstruction.

Although this method is also intended to create graphically realistic renderings of the fire from novel views, the reconstructions created are also more volumetrically

accurate than in [8], being composed of spherical blobs rather than sheets. The success of this type of approach lies in the method of parameterisation, and further research and testing is required to find an optimal solution. A similar least squares approach is discussed in §4.4.2 below.




Figure 4.8: Examples of flame reconstruction from [18] using (a) sheets and (b) blobs.

4.4.2 Visual Hull Limited Tomographic Inversion

Ihrke and Magnor [9] have recently developed a tomographic method for the volumetric reconstruction of fire using at most 8 cameras. Although they state that they are concerned with generating visually accurate animations, not measuring physical properties, they achieve this by creating a reconstruction with an accurate density distribution. The approach they take is to use sparse matrix inversion techniques to find a least squares reconstruction of each slice.

Recall that one can represent the transformed pixel intensity as

$$I_p = \int_c \phi \cdot ds \quad (4.5)$$

where I_p is modelled as the integral of the density field, ϕ , along ray c . In order to solve equation (4.5) one needs to impose some form of structure to the density field. This can be done by defining ϕ as being composed of a linear combination of basis functions:

$$I_p = \int_x \left(\sum_i a_i \phi_i \right) ds, \quad (4.6)$$

where a_i are the coefficients for the density basis functions ϕ_i . Rearranging the equation gives

$$I_p = \sum_i a_i \left(\int_x \phi_i ds \right), \quad (4.7)$$

which describes a linear system and can be written in matrix form as

$$\mathbf{p} = \mathbf{S}\mathbf{a}. \quad (4.8)$$

The matrix \mathbf{S} defines the relationship between each image pixel's projected ray and the basis functions. By definition every pixel is influenced by every basis function, creating a full matrix, but by structuring the basis function such that it only affects a small region, one may eliminate many entries in \mathbf{S} , creating a sparse matrix. In addition, the visual hull is used to restrict the basis functions to those falling within the flame silhouettes, thus further reducing \mathbf{S} .

The simplest basis function is the box function – essentially each voxel is a basis function. This is assumed in traditional algebraic (ART, SIRT) techniques. Ihrke and Magnor [9] use a trilinear basis function, covering 8 voxels as shown in Figure 4.9. The function value at the centre of the cube is 1, decreasing linearly to 0 at the edges. Representing the function value in each voxel with a cubic polynomial allows the analytical integration of the intersection of the image rays with the basis functions.




Figure 4.9: Trilinear basis function covering 8 voxels, with transparent isosurfaces indicating the function values (from [9]).

The linear system (4.8) is then solved for \mathbf{a} using the conjugate gradient method. This iterative method allows a least squares solution to be found. For more information refer to [9].

Although setting up the \mathbf{S} matrix is computationally intensive, it need only be done once, since additional frames require only modification to reflect which entries are contained in the visual hull. Therefore the main computational expense is for solving the linear system in each frame of a video sequence.

While their results look realistic and are reported to be fairly accurate, the problem is that working with such large matrices requires a great deal of computer memory. Although a least squares solution is appealing for sparse view tomography, it is not often used because of the computational resources required. Indeed in [9] it stated that their images needed to be resized to fit the matrix in 2GB of memory. A volume size of $96 \times 96 \times 96$ voxels was the largest mentioned (although it is not clear if this was the largest possible), which is reasonable, although considerably smaller than the 640×480 images being used.




Figure 4.10: Example of flame reconstruction based on tomographic inversion from [9].

4.5 Summary - The Hybrid Approach

As one can see the problem of flame reconstruction is generally tackled from either a tomographic or a geometric perspective. It is the intended application that determines the general direction to be taken. When one wants to generate realistic renderings from novel viewpoints (e.g. for computer graphics) one needs to use tomography. Because of the semi-transparent nature of fire, the perceived images of a flame are dependant on its internal density, and not merely the surface radiance as with opaque objects. This necessitates a tomographic reconstruction of the flame's internal density field. This also means that one would not generally be concerned with defining the actual flame surface, or front, which a tomographic reconstruction would not create. Geometric techniques, using the visual hull or contours, generate an explicit flame surface, and are more suited to monitoring or geometric analysis applications. A definitive flame volume and surface allows for analysis of the geometric properties of the flame.

As computers become faster, with more memory, matrix-based tomographic techniques ([9],[18]) should become more common, and the algorithms used more efficient. For the purposes of this project, it was decided not to pursue a direct matrix solution approach, mainly because of the limitations that it would impose on the reconstruction volume size. In addition the processing time required could be prohibitive when considering a video sequence composed of many frames.

In the case of this thesis, the reconstruction needs to be geometric, having a quantifiable volume and surface, but with an estimate of the internal density structure as well. Hence a hybrid approach is needed, incorporating aspects of both geometric and tomographic flame reconstruction. In chapter 5 a hybrid method is presented using the visual hull, algebraic tomography and fuzzy image segmentation.

Chapter 5.

The Fuzzy Hull Method

The method presented in this thesis, for the reconstruction of fire, shall be referred to as the *Fuzzy Hull* method. As discussed in chapter 4, the research that has been done on flame reconstruction from images is usually for one of two purposes: the generation of photorealistic renderings of fire for computer graphics [8],[9], or flame monitoring and geometric analysis applications [5]. With that in mind the required properties of the method presented here are briefly reviewed.

- **Volumetric:** By volumetric it is meant that each voxel in the reconstruction space must be classified as being a part of the flame volume or not.
- **Density field:** The reconstruction must contain an estimate of the flame density field.
- **Generality:** The reconstruction technique must be able to handle any type of visible flame, should not require specific laboratory conditions, and should not be too limiting regarding the size of the flame or the reconstruction.
- **Non-calibrated:** One should still be able to create a reconstruction even with a non-calibrated (or minimally calibrated) camera set.
- **Speed:** The process should not be too computationally expensive since video sequences need to be processed, and not just a single frame.
- **Cost:** No specialised equipment should be needed, only suitable cameras.

It was thus concluded that the general approach to follow should consist of a tomographic reconstruction that is subject to some form of image segmentation in order to classify the voxels. It was decided that fuzzy c-means would be a suitable segmentation method to use, given that it is used in other tomographic segmentation applications (see §3.4.1). The use of algebraic tomography and fuzzy image segmentation form the basis of the flame reconstruction method presented here. The visual hull is also considered. The implementation thereof is developed throughout the remainder of this chapter.

5.1 Overview

The reconstruction process can be divided into four stages, shown in Figure 5.1 below.




Figure 5.1: Process overview.

The first step is to capture the video sequences. Note that the video is treated as a series of independent images. Before the 3D reconstruction can take place, the images must be suitably processed. The completed reconstruction can then be rendered or analysed from any viewpoint.

5.2 Experimental Set-up

The cameras used are monochrome with a resolution of 640 x 480 pixels. Colour cameras are not necessary since the reconstruction need not be photorealistic. If one were creating a colour reconstruction one would need to reconstruct each channel

(RGB) separately before combining them, as in [9], tripling the computational effort. The bus-powered cameras operate off an IEEE 1394 (Firewire) bus, allowing the simultaneous capture of up to 7 synchronised cameras at 15 frames per second using a single PC. The cameras are able to capture images synchronised to within $125\mu\text{s}$ of each other [82]. In terms of spectral response one is looking at light having wavelengths in the region of 570 – 620 nm (yellow – orange). This gives a fairly linear relative spectral response ranging from 0.92 (at 570 nm) to 0.80 (at 620 nm) [82]. Although this 13% difference is of course a potential source of error, individual flames generally tend to occupy a fairly narrow spectral band, their colour depending on factors such as fuel and temperature, so the spectral response should not present a significant problem.




Figure 5.2: Camera configuration. The cameras should be roughly equally spaced – e.g. 60° apart for 3 cameras.

The cameras are optimally positioned equidistantly within a 180° arc around the flame, as shown in Figure 5.2 above. Although one might think they should be placed around 360° , this is not true since two cameras opposite each other would produce a mirror image of each other, due to the semi-transparent nature of fire. This configuration is optimal (when the cameras are standing on the ground plane, as is the case here) since the unknown space between images is at a minimum.

The cameras are equipped with zoom lenses having a focal length of 5 to 40 mm. When setting up the cameras to record images of fire, one ideally wants the shutter speed to be as quick as possible, since fire is not static and therefore requires a short exposure to obtain a sharp image. In addition one does not want to use the gain function on the cameras, since this introduces and amplifies noise within the image. Therefore the shutter duration and gain are set to their minimum values and the variable aperture on the lenses is adjusted to provide the correct image intensity.

5.3 Camera Modelling and Calibration

The purpose of camera calibration is to provide a link between the 2D images generated by a camera and the real world. In order to relate a point in 3D space and its projection on the image plane, one needs to define a set of extrinsic and intrinsic parameters, calculated by a calibration process. Intrinsic refers to the geometric, optical and digital properties of the camera, while the extrinsic parameters define the location and orientation of the camera. With a calibrated set of cameras one is thus able to calculate the position of a point in 3D space, if that point can be identified in 2 (or more) images. Camera calibration is therefore a critical aspect of many computer vision applications, such as motion capture and object reconstruction, where typically multiple cameras are used.

5.3.1 Pinhole Camera Model

In order to calibrate a camera one needs to model its geometry, defining how an image is formed. The most common way to do this is to use the pinhole camera model, shown in Figure 2.5. This perspective projection model consists of a camera reference frame with origin \mathbf{O} , an image plane π . The optical or z-axis is perpendicular to the image plane intersecting it at the principal point \mathbf{o} , with the focal length, f , being the distance between \mathbf{O} and the image plane. A real world point, $\mathbf{P}(X,Y,Z)$, is shown, projecting onto the image plane as $p(x,y,z)$.




Figure 5.3: The pinhole camera model.

The real 3D world can then be mapped onto the 2D image plane using a similar triangle relation:

$$\begin{cases} x = X \frac{f}{Z} \\ y = Y \frac{f}{Z} \end{cases} \quad (5.1)$$

Since Z is variable the transformation is non-linear. By assigning Z a fixed value, say Z_0 , the transformation becomes linear, having a constant scaling factor of f/Z_0 .

Called weak perspective projection, or scaled orthographic projection, this approximation can be justified if the distance from the camera to the object is considerably greater than the thickness of the object.

For this project it was decided to use the weak perspective projection model, since this simplifies the implementation of the tomographic reconstruction process. This choice can be justified because owing to the dangers of both fire and high voltages it is necessary to place the cameras at some distance from the flame itself, meaning that the distance between camera and flame is much greater than the thickness of the flame. In addition, the use of full perspective projection would be unlikely to produce any significantly better results. One must remember that tomographic reconstruction from only a few views is an approximation, not a precision reconstruction.

5.3.2 Camera Calibration

Many computer vision applications require a precise correlation between a 3D point and its 2D projections, making camera calibration critical. Ideally though, one would like to create systems that do not require explicit calibration. Calibration can be restrictive, requiring specific calibration objects and processes, and should a camera be bumped or altered even slightly, then recalibration is required. An error in the calibration process could potentially render the captured data useless. Therefore it is best if the degree of calibration is kept to a minimum – or rather a compromise needs to be reached between calibration complexity and the accuracy of the end results.

As mentioned previously, a specification of this project is to keep calibration requirements to a minimum, primarily to make the system easier to use. Although a lower calibration accuracy may result in less reconstruction accuracy, this is unlikely to be significant given the reconstruction method and the fact that the reconstruction is only an estimate.

The difference between this application and your typical 3D computer vision system is that one is not dealing specifically with 3D points from the real world, but rather the projection of 2D silhouettes. The most important issue is that the silhouettes are correctly projected to give an optimum and accurate intersection. Considering the experimental camera configuration and the use of weak perspective projection it becomes apparent that the main calibration parameter that needs to be determined is the angle between adjacent cameras (given that the cameras are all positioned vertically and in a planar configuration). From the previous work done on 3D flame reconstruction from images (chapter 4), the only case where proper calibration was done was in the work by Ihrke and Magnor [9]. In the other cases the only information available was usually the angle between the cameras, confirming that this is sufficient information to create a reasonable reconstruction.

An advantage of minimal calibration requirements is adaptability and generality – images can be used from various sources and be subject to the same reconstruction methods, as done by Hasinoff [18]. In addition, the experiments are to be done mostly outdoors, subject to wind, bumping and adjustments to the aperture and zoom of the lenses, all requiring recalibration.

If the angle between the cameras is not known, then a simple calibration object is used with the sole purpose of determining the angles. Using a square object (illustrated in Figure 5.4 below) only a single image from each camera is needed to determine their relative angles. Should it be required one could also estimate the camera rotation and a relationship between real world and reconstruction units.




Figure 5.4: Calibration object geometry.

Simple trigonometry gives

$$\theta = \arctan\left(\frac{x_2}{x_1}\right). \quad (5.2)$$

One can relate the size of the reconstruction to the real world by using the relationship

$$scale = \frac{x_1}{L \cos \theta} \quad \text{pixels per real world unit.} \quad (5.3)$$

Other information, needed to correctly align the silhouette projections, is obtained from the flame images themselves.

5.4 Image Processing

The raw flame images need to be processed before being fed into the reconstruction algorithm. The first step, as discussed in §2.2, is to transform the image intensities in order to give a representation of the flame density along the ray projected by each

pixel. This is achieved using equation (2.12) if the background is dark, or equation (2.14) if there are background intensities.

The transformed images must then be segmented, isolating the flame. Edge or line-based segmentation methods can be problematic when dealing with potentially soft-edged flames. Since the background will already be fairly dark the process is kept simple, using threshold segmentation. The actual threshold value to use is determined automatically by analysing the histogram of the image.

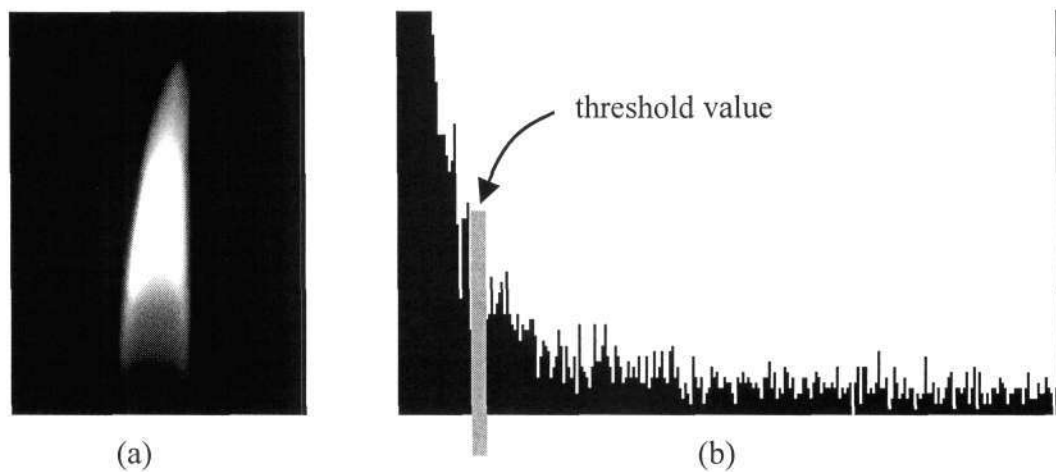


Figure 5.5: (b) is the (cropped) histogram of image (a).

Given a typical flame one would expect the histogram to have certain characteristics. The large, dark background area produces large values on the low end of the histogram, rapidly tailing off, with the rest of the histogram dependant on the flame itself. This indicates that an optimum threshold value would be after the background spikes, where the histogram starts levelling out, with mainly background to the left and flame to the right.

The histogram is smoothed and a derivative based metric is used to find this point. The background spikes produce a large negative gradient as one travels along the histogram. The threshold point is therefore found by determining when the magnitude of the smoothed gradient becomes sufficiently low, after the background spike. Various degrees of smoothing and derivative-based definitions of the

optimum threshold were explored, but it was found that similar results were obtained, as long as the basic principle was followed. Also, an exact comparison between variations proved difficult since fuzzy and indistinct regions of the flame made an ideal threshold value difficult to quantify. Further study is needed to formalise and evaluate this technique.

Once segmented, the images are scaled, ensuring all images are the correct size. The tip of the flame in each image is located, providing a reference point to align the projection of the flame silhouettes. This is simply a search for the upper most (central) point of the segmented flame in each image. It is acknowledged that this is not an ideal method since it relies on the flame having a single, distinct tip, and may fail if, for example, a more complex flame image contains two or more tips in the same row of pixels. The advantage, though, is the provision of a simple self-calibration technique, based on general flame characteristics. In addition, because it need only be performed once, for a single frame in a sequence, one does not need every frame to have a distinct tip.

Finally the images are normalised so that each image has the same sum. This is necessary because the sum of the density seen from any direction must be equal, since one is looking at the same object. In other words, given that the transformed images represent the relative density of the object, all views of the object must give the same total object mass – the same image pixel sum.

5.5 The Reconstruction Methods

Here the evolution of the reconstruction method is presented, from a simple 2D to a full 3D reconstruction.

5.5.1 Simple – SIRT and Fuzzy

By treating the 3D reconstruction as a series of stacked 2D slices, the problem is reduced to two dimensions. Algebraic tomography (SIRT) is used to estimate the density field, after which fuzzy c-means (FCM) image segmentation is used to classify the density field based on each pixel's statistical mean and standard deviation in its local neighbourhood. A non-negativity constraint is imposed during the tomographic stage, since it is not physically possible to have a negative density. Figure 5.6 below shows the result of this technique on a synthetic 2D test object. As one would expect, the results are not spectacular, but they do improve as the number of views is increased.




Figure 5.6: Reconstruction of test object from 5 views using SIRT and FCM.

The problem with this method is that not all the available information is being used. The image pixel intensities guide the tomographic process, but the actual silhouettes of the flame (or test object) are being ignored. By including the silhouette

information one can restrict the space in which the reconstruction is created. Using the silhouettes to define the visual hull of the flame gives a shell outside of which no part of the reconstruction may occupy.

Therefore the concept of the visual hull was incorporated into the process, named the Fuzzy Hull method.

5.5.2 The Fuzzy Hull Method

Using the visual hull provides two main benefits:

- **Accuracy:** Instead of being smeared across the entire reconstruction space, the flame image pixel intensities need only be spread along the portion of the relevant ray intersecting the visual hull. This means that SIRT will be more accurate, the convergence faster, and the fuzzy classification will therefore also be more accurate.
- **Speed:** As well as faster convergence, the fact that only the space occupied by the visual hull needs to be processed means the overall processing time is greatly reduced, since the majority of the reconstruction space can be ignored.

The visual hull provides a reasonable initial reconstruction. Algebraic tomography then provides an estimate of the density field within the hull and fuzzy segmentation is used to refine the volumetric reconstruction. While the tomographic reconstruction stage maintains the integrity of the hull, the fuzzy segmentation does not. This means that the silhouettes of the reconstruction will usually not be the same as the original flame image silhouettes, and are invariably smaller.

Ideally the original and reconstructed silhouettes should be identical, and therefore a silhouette consistency constraint is introduced in an effort to improve the projected reconstruction silhouette and thus the volumetric reconstruction accuracy. First the

orphan rays are identified. The term *orphan rays* describes those rays (pixel projections) that are a part of the original image silhouette, but do not project to the reconstruction. The voxels forming the intersection of the orphan rays and the visual hull are called *orphans*. Because the orphan rays form part of the visual hull it is known that at least one voxel along each orphaned ray should be reclassified as part of the reconstruction. There is however no way to determine exactly which, or how many, voxels need reclassification, so an estimate is made using fuzzy segmentation to decide which orphans are most likely to be a part of the reconstruction. The integration of these reclassified orphaned voxels into the reconstruction is illustrated, along with the entire reconstruction process, in the flowchart below.

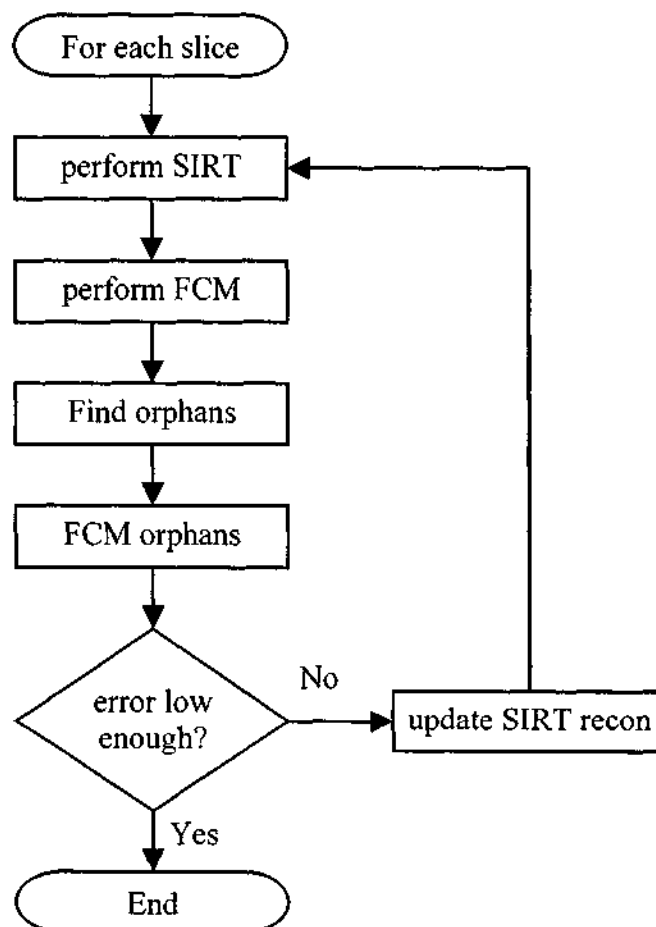


Figure 5.7: Flowchart outlining the Fuzzy Hull algorithm with orphan estimation.

5.5.2.1 Synthetic Data Testing

Synthetic test objects were used to test the developed reconstruction methods. It must be noted that using synthetic objects will of course produce better results than real data, but they are nonetheless useful. Using test objects allows one to determine that a reconstruction method is valid and has been implemented correctly. In addition, synthetic data provides a ground truth against which a reconstruction can be compared, enabling a quantifiable analysis to be done. This means that different algorithms and parameters can be tested and evaluated comparatively. Using real data makes this more difficult, since one cannot do a ground truth comparison and can only really perform an evaluation based on 2D projections and the perceived visual accuracy (i.e. does it look right?).

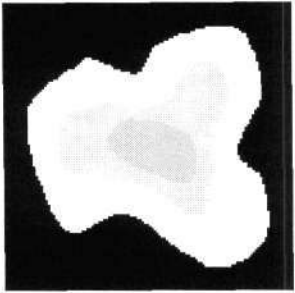
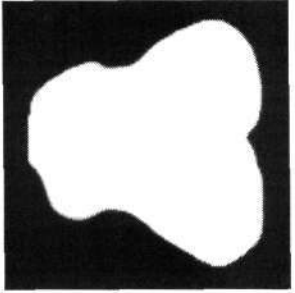
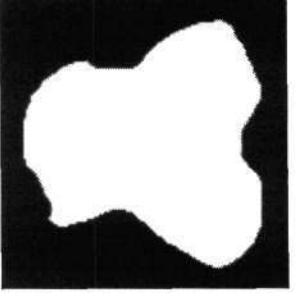
The testing presented in this section focuses on the methods discussed in §5.5 and §5.5.1. The method using the straightforward application of SIRT and FCM is referred to as *method A*. The analysis done is based on the volumetric accuracy, rather than density field accuracy, since the density field estimation is not affected by the fuzzy classification - its accuracy is largely dependant on the number of views, and a better density field reconstruction encourages a better volumetric reconstruction. The results shown here are therefore to test the accuracy with which the reconstructed density field was classified by the fuzzy segmentation.

2D Testing

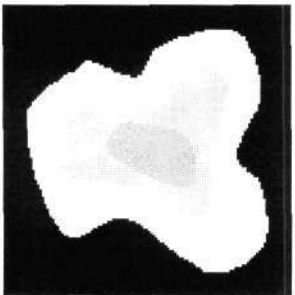
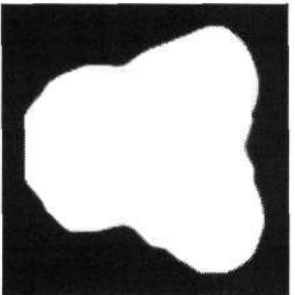
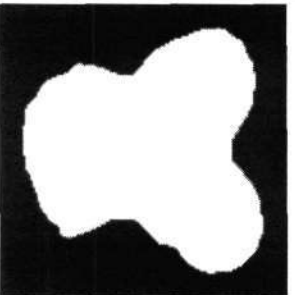
Method A and the Fuzzy Hull method were tested using two different 2D objects, for 3, 5 and 7 views. (The same generated views were used for each method.)

The results were compared using three parameters:

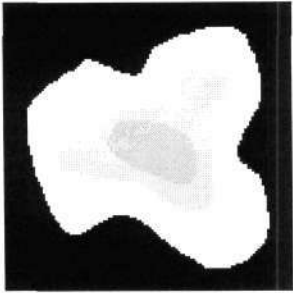
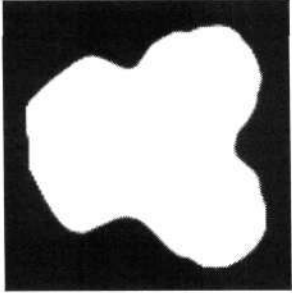
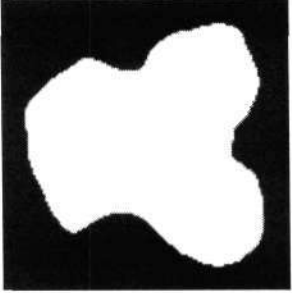
- i. False positive percentage (f_p %): The number of pixels incorrectly classified as object, relative to the total amount of object pixels.
- ii. False negative percentage (f_n %): The number of pixels incorrectly classified as background, relative to the total amount of object pixels.
- iii. Total Error percentage: The total number of incorrectly classified pixels, ($f_p + f_n$).

	Method A	Fuzzy Hull
Test Object #1		
		
f_p %	3.3	1.9
f_n %	3.1	3.3
Total error %	6.4	5.2

(a) 3 views

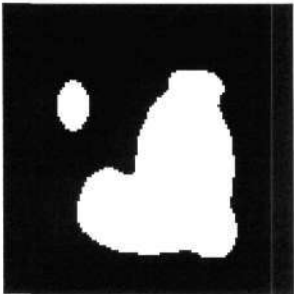
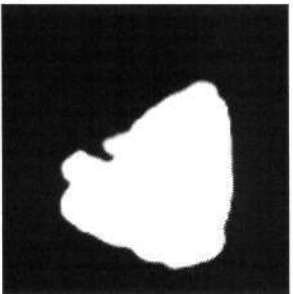

	Method A	Fuzzy Hull
Test Object #1		
		
f_p %	4.8	0.7
f_n %	1.5	2.3
Total error %	6.3	3.0

(b) 5 views

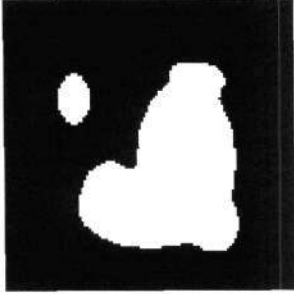
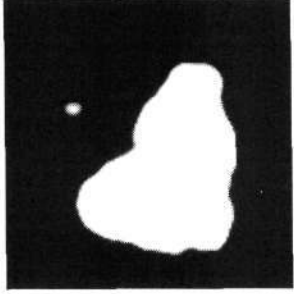
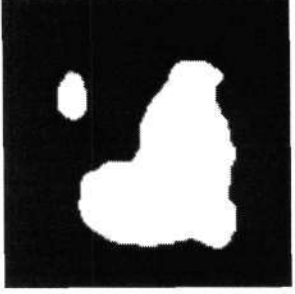
	Method A	Fuzzy Hull
Test Object #1		
		
f_p %	2.5	0.4
f_n %	0.9	1.2
Total error %	3.4	1.6

(c) 7 views

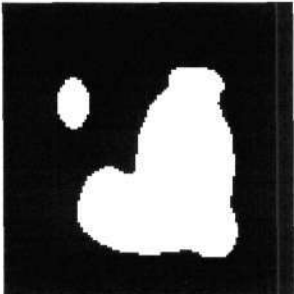
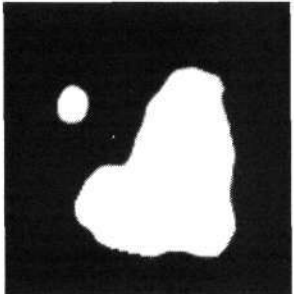
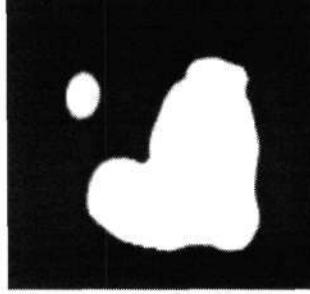
Figure 5.8: Reconstructions of test object #1.

	Method A	Fuzzy Hull
Test Object #2		
		
f_p %	18.5	4.5
f_n %	13.1	9.5
Total error %	31.6	15.0

(a) 3 views

	Method A	Fuzzy Hull
Test Object #2		
		
f_p %	5.0	0.9
f_n %	7.9	3.3
Total error %	12.9	4.2

(b) 5 views

	Method A	Fuzzy Hull
Test Object #2		
		
f_p %	4.1	0.5
f_n %	3.1	3.0
Total error %	7.2	3.5

(c) 7 views

Figure 5.9: Reconstructions of test object #2.

From the above results it is evident that the fuzzy hull method consistently produces more accurate reconstructions. This validates the use of the visual hull and the

silhouette consistency constraint. Both methods give better results as the number of views is increased, as one would expect. Looking at the fuzzy hull (method B) reconstructions one notices that the number of false negatives is always larger than the number of false positives. This indicates that the fuzzy segmentation process is being too strict in its decision-making. Adjusting the fuzzy classification parameter to make $f_p \approx f_n$, although not necessarily increasing overall accuracy, would produce a more balanced result and is thus a good guide to determining an optimal fuzzy classification parameter.

3D Testing

In order to test the reconstruction algorithms with 3D objects a relatively complex flame-like object was created using MATLAB[®]. Figure 5.10 below shows the object and an example of a generated view of the object used as an input to the reconstruction algorithms. Rendering of the reconstructions of this object was also done in MATLAB[®].




Figure 5.10: (a) Synthetic flame-like test object. (b) Example of generated image.

Figure 5.11 below shows reconstructions of the object from 5 views using method A and the fuzzy hull method. Just from a visual inspection one can see that the fuzzy hull algorithm produces a superior result.

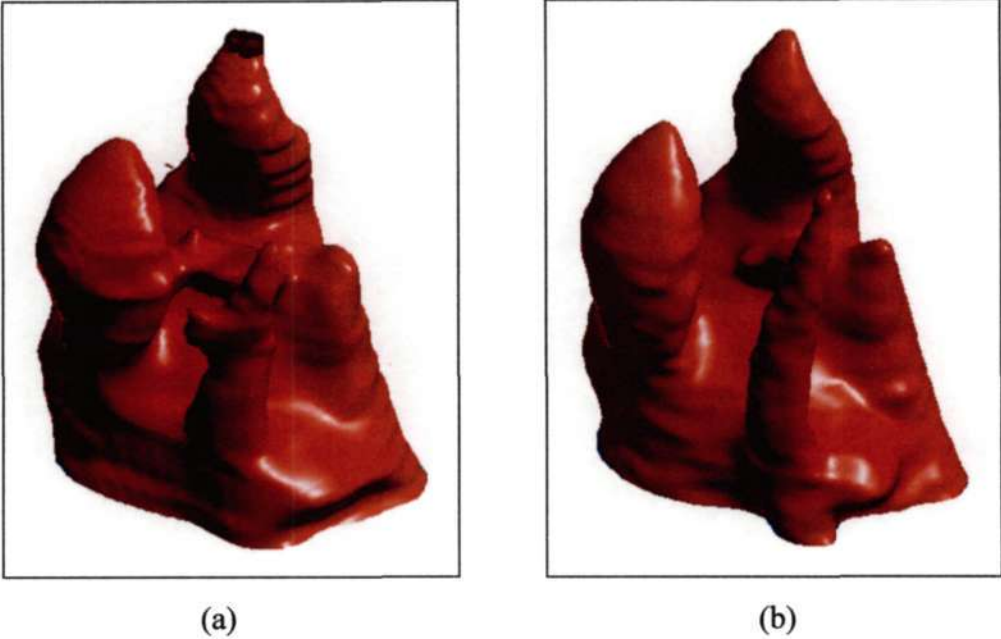



Figure 5.11: 5-view reconstruction using (a) method A, (b) the fuzzy hull method.

In order to obtain a direct visual analysis of the reconstructions they were superimposed on the original, as shown below.






Figure 5.12: Reconstructions superimposed on original object. On the left is method A and on the right is the fuzzy hull method. (c) and (d) have been rendered with partial transparency.

Figure 5.12 shows how the fuzzy hull reconstruction is much closer to the shape of the original object. A more quantifiable analysis is presented in Table 5.1 below and the graphs that follow.

Table 5.1: Analysis of 3D reconstructions.

Views	Method	Time(s)	f_p %	f_n %	% error	% SA error	% Vol error
3	A	56	16.7	4.0	20.7	2.1	12.7
	Fuzzy	34	8.3	4.4	12.7	1.6	3.8
5	A	72	9.4	2.5	11.9	2.7	6.8
	Fuzzy	36	2.9	1.8	4.7	2.0	1.1
7	A	88	8.1	2.5	10.6	3.9	5.7
	Fuzzy	42	1.7	1.7	3.4	2.0	0.1
9	A	102	7.5	2.6	10.1	6.3	4.9
	Fuzzy	47	2.1	1.2	3.3	1.6	0.9

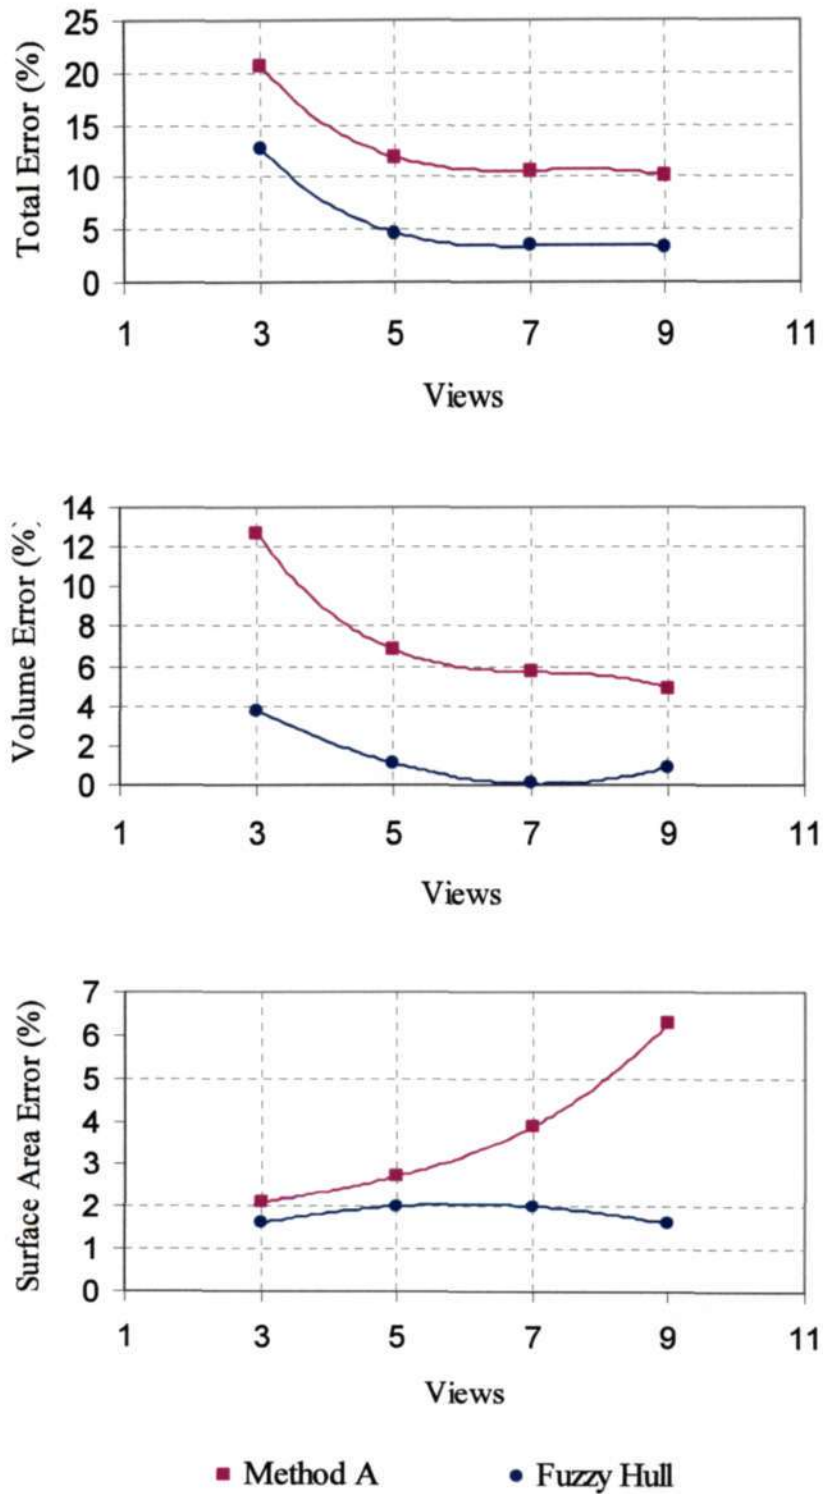


Figure 5.13: Graphical representation of 3D reconstruction accuracy analysis.

From the first graph one can see that the fuzzy hull algorithm consistently achieves a significantly more accurate reconstruction. When looking at the volume and surface area graphs one must note that these do not represent the geometric accuracy of the reconstruction, but rather a measured parameter of the reconstruction. This means that a low surface area error, for example, could be the result of a balance between positive and negative errors. What one should look for in these two graphs is a relatively flat line, indicating consistency. Note how the surface area error of method A actually increases notably with the number of views (as the total error decreases), contrary to what one would expect. This is an indication that these apparently good surface area measurements with 3 or 5 views are due to factors other than reconstruction accuracy.

These graphs were also used to determine the number of cameras to be used in the system. It was noted how the amount of error tends to level off after 5 views, suggesting that 5 cameras would be a good compromise between accuracy, equipment complexity and cost. Obviously these measurements are dependant on other factors such as object geometry and algorithmic details, but they nonetheless remain a reasonable guide for camera quantity selection, providing a rough idea of the accuracy limitations and trends using specific reconstruction principles. It should be remembered that sparse view tomography is in general very dependant on the number of available views.

5.5.3 3D Fuzzy Hull

The final version of the fuzzy hull algorithm was enhanced to better deal with the three dimensional nature of the problem. Instead of viewing the reconstruction as a set of stacked, independent 2D slices, the entire 3D reconstruction space was considered.

In the fuzzy segmentation stage the slices directly above and below are taken into account when calculating the image statistics. This increases the spatial coherence of

the reconstruction, relating adjacent slices and therefore providing a smoother reconstruction.

When testing with real flame images a problem was encountered in that the flame projections did not intersect perfectly, causing sharp edges and density field errors. This incorrect alignment was caused by factors such as calibration inaccuracy and inconsistent image segmentation of the flame images. While these errors were relatively minor on their own, their affect on the reconstruction could be significant. For example, consider a ray that should project to 5 voxels in the visual hull. A slight misalignment causes the ray to only project to 1 voxel instead. This means the density of this voxel will be estimated at 5 times the value it should be, affecting all other rays projecting to that voxel.

One solution is to make the camera calibration as complete and precise as possible, although it would still be difficult to ensure consistent image segmentation in all images. In this case a simpler method was developed to help reduce the effects of misalignment. The algorithm works by comparing the original flame silhouette and the visual hull silhouette. This data is analysed and the image alignment is adjusted, one pixel at a time, based on where the projection errors are occurring. This process effectively improves the alignment of the projections with minimal effort. Figure 5.14 below shows a set of four flame silhouettes, with the projection error indicated by the white regions. After the visual hull optimisation algorithm is used it can be seen how the projection error has been greatly reduced.

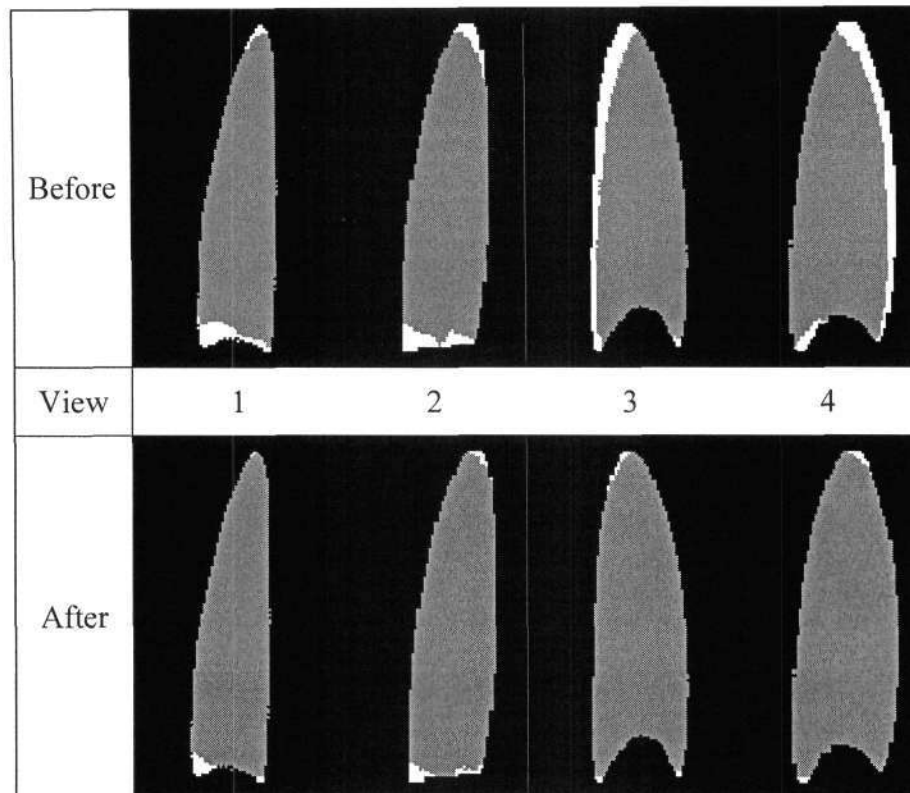


Figure 5.14: A set of 4 projections subject to the hull optimisation algorithm.

The silhouette consistency constraint was abandoned, since the additional processing required to process the orphans was prohibitive, given a relatively small gain in accuracy. Ideally a better way should be found to incorporate a silhouette constraint.

As mentioned previously, the algebraic reconstruction of the density field uses a discrete approximation of the ray projection – each ray is defined as being composed of a set of pixels. However, the actual images (while themselves discrete) are formed by the continuous density field of the real world flame. This discrepancy can result in errors and artefacts in the reconstruction, particularly when the ray width is too small. The wider the ray is, relative to the pixel size, the more accurate the ray definition becomes, and hence also the reconstruction.

An initial solution was to oversample the reconstruction space, replacing each pixel with 4 pixels. This provided a significant improvement in the smoothness and accuracy of the density field estimation, (demonstrated in Figure 5.15 below) but at

the expense of greatly increased processing time and memory requirements (approximately 4 times greater).

It must be noted that if one measures the density field estimation accuracy it is with respect to the discrete ray definition, and is not therefore a good indication of the actual reconstruction accuracy. Hence a visual, qualitative inspection is sometimes necessary.

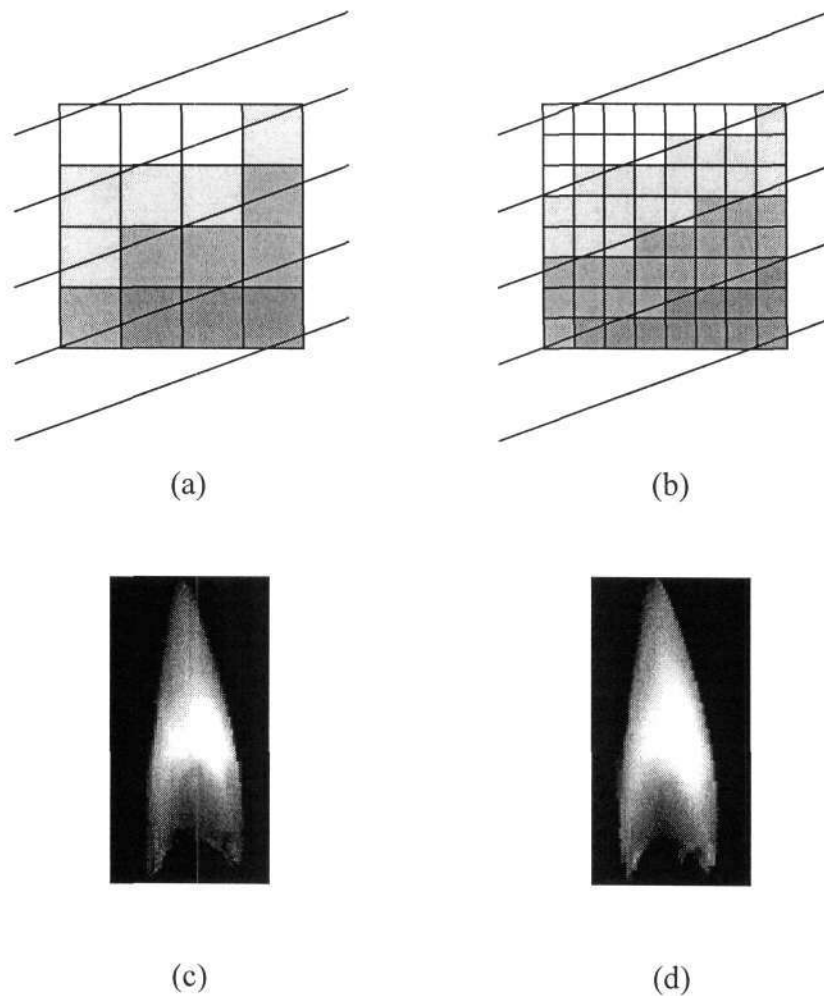


Figure 5.15: (a) shows the discrete definition of the rays using pixels, while (b) is the same but using oversampling. (c) and (d) are corresponding density renderings of a reconstruction of a candle flame differing only in the ray definition used.

Generally, fire tends to have a high degree of local smoothness. What is meant by this is that the density changes gradually, or smoothly. This can be exploited to

improve the reconstruction. As done in [9] a basis function was implemented – a cube, consisting of 8 voxels. These basis functions were arranged so as to overlap once in each dimension. The SIRT algorithm was modified so as to estimate a value for each basis function, instead of each individual voxel. This resulted in a smoother internal density gradient, also helping to counteract the effect of discrete ray definitions. The entire reconstruction process is therefore three dimensional, taking neighbouring voxels in all directions into account in both the tomographic and fuzzy segmentation phases, giving a smoother, more accurate reconstruction in terms of both the internal density field and the flame surface.




Figure 5.16: Cross-sectional reconstruction slice showing the effect of using (b) oversampling and (c) the basis function, compared with a slice reconstructed using (a) neither.

The 3D fuzzy hull algorithm was tested on a real data set, consisting of a video sequence of a candle flame captured with 4 cameras. The reconstructions of two frames are shown below. The surface rendering shows the geometric aspect of the reconstruction, while the density rendering uses the reconstructed density field to create a photo-realistic view of the reconstructed flame. The surface renderings were created with the aid of the VOLPACK C++ library [83].




Figure 5.17: Surface renderings of 2 reconstructed frames rotated through 180°.




Figure 5.18: Density renderings of 2 reconstructed frames rotated through 180°.

Because the reconstructions are modelled using voxels, and not a polygonal mesh, the surface renderings can have a slightly blocky appearance (as above). This can be

improved by performing a three-dimensional smoothing operation on the volume before rendering.

Figure 5.18 shows that the reconstruction algorithm is good enough to create realistic renderings of the flame from any novel viewpoint. Although that is not the objective of this project, it does provide a good qualitative way to assess the algorithm.

Upon further testing it was decided that oversampling the reconstruction space was too resource intensive. A better approach was to change the way in which the relationship between the rays and the reconstruction is defined. Instead of making the simplification of defining the w_{ij} term in equation (3.11) as a binary function, the rays are defined by calculating their relative weight factors (w_{ij}). This means that each voxel in the reconstruction space is defined not by which ray it projects to, but rather by the fraction of the voxel falling within each ray. Although this is more complex and time consuming than the binary simplification, it need only be done once and is faster and more accurate than the oversampling method.

The SIRT algorithm needs to be changed to reflect the new ray definitions, and the update equation for cell j from M projections is amended to

$$\Delta f_j = \frac{1}{M} \sum_{k=0}^M \sum_{i=r_0}^{r_2} (p_i - q_i) \frac{w_{ij}}{W_i} , \quad (5.4)$$

where the term W_i is the weight of ray i (in the visual hull, where a cell (voxel) is of unit weight) and $r_0 .. r_2$ is a simplification to reflect the fact that a maximum of 3 rays can project to any single cell (where ray width = cell width).

An interesting modification can be made to equation (5.4) that forces the reconstruction to be biased towards longer rays – the longer the ray the more it is favoured in the calculation of the update value.

$$\Delta f_j = \frac{1}{\sum_{k=0}^M W_{kr_1}} \sum_{k=0}^M W_{kr_1} \sum_{i=r_0}^{r_2} (p_i - q_i) \frac{w_y}{W_i} , \quad (5.5)$$

This can be useful to help alleviate the problem where misalignment of the projections can cause short rays to dominate a particular voxel, creating artefacts. Note that the terms *long* and *short* refer in this case to the length of the ray within the visual hull.

5.6 Level Set Reconstruction

The use of the level set technique in the reconstruction process was investigated. The advantage of using level sets is that they produce a smooth, continuous surface, or rather, a continuous set of isosurfaces, even in a discrete domain. In addition the formation of the level sets can be controlled by any user-defined speed function, allowing parameters such as curvature and distance to influence the shaping of the level set. Although this means that level sets can be adapted to suit a particular problem, this also means that the process can be very difficult to tune and can require a great deal of trial and error to find optimal parameters.

The primary reason that the level set technique was not adopted, though, is the processing time required. Level sets can be very computationally expensive and in the case where one has a video sequence consisting of many frames this can be a problem. This becomes feasible if one is using high-speed cameras, where each frame is relatively similar to the previous frame. Where consecutive frames are similar it means that the level set does not need to be recreated for each frame, and can instead be used to track the evolution of the flame front from frame to frame. Using level sets in this manner would be a very useful tool in terms of analysing flame surface dynamics.

5.7 Flame reconstruction analysis

5.7.1 Geometric Analysis

The volume of the reconstruction is easily determined by counting the number of voxels forming the reconstruction. Calculating the surface area however, is not so straightforward. If one was using a polygonal mesh then one could sum the areas of the 2D polygons to give the surface, but with voxels one needs to first generate a geometric surface approximation. The most common tool for such a problem is probably the Marching Cubes Algorithm.

Introduced by William E. Lorensen and Harvey E. Kline [84], the Marching Cubes algorithm is able to extract surface information from a three dimensional data grid. The first step is to divide the reconstruction space into a set of cubes. The corners of each cube are then tested and the cube replaced by a polygon, as defined by the corner values. There are 15 possible corner combinations, creating a set of predefined polygons to use.




Figure 5.19: An example showing a cube and the polygons defined by its corner points.

The drawbacks of marching cubes though are a potentially complex model, and the fact that it is possible to have holes in the surface – certain situations create an ambiguity, having more than one possible solution, as illustrated in Figure 5.20 below.




Figure 5.20: 2D example of an ambiguous case.

An alternative to marching cubes is Adaptive Skeleton Climbing [85], which is a multi-resolution isosurface generation method, claiming to reduce the number of triangles generated by marching cubes by a factor of 4 to 25. This method can still suffer from ambiguities.

For this project though it is only required to obtain an estimation of the surface area, so a complete marching cubes surface generation would consume unnecessary processing time. Hence, a simpler, quicker algorithm was developed to give an estimation of surface area.

Of course one could get an estimate by summing the area of the exposed voxel faces, but this would obviously create an over-inflated value (since fire and most objects are not actually composed of little cubes). The method that was developed works by examining the voxel structure and determining the voxels most likely to cause such over-inflation. These voxels are then reshaped to give a more likely contribution towards the surface. The voxels themselves are not altered, merely the contribution that exposed voxel faces make towards the surface area estimate. The voxel contribution is determined according to the number and position of adjacent voxels, as defined in Figure 5.21.

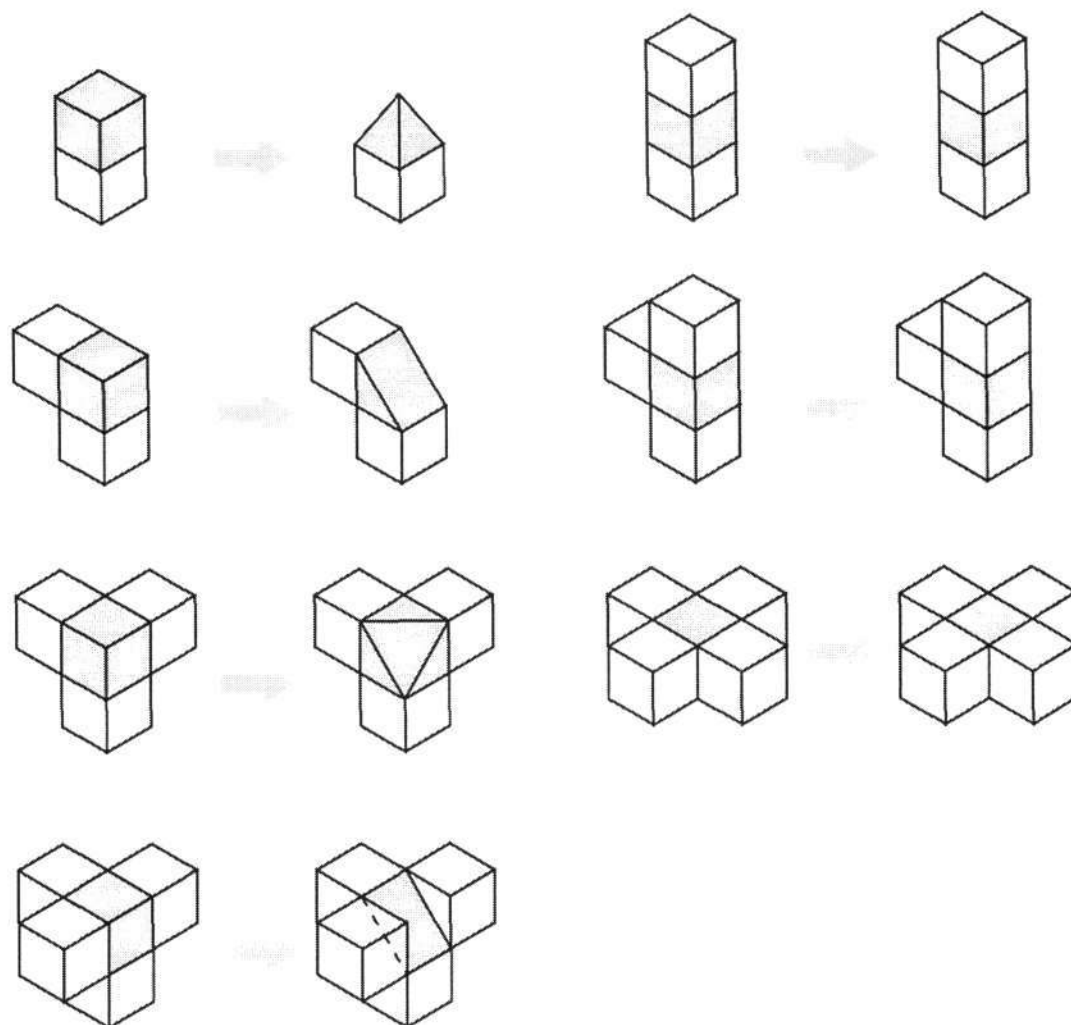


Figure 5.21: Rule set for voxel-based surface area estimation.

5.8 Summary

The 3D Fuzzy Hull method uses algebraic tomography, bound by the visual hull, to create a density field estimation. The use of an area-weighted ray definition and a basis function improves the local smoothness and visual accuracy of the reconstruction. Fuzzy c-means image segmentation then defines the flame volume by segmenting the reconstruction space according to statistical image features. The 3D implementation of the algorithm takes adjacent slices into account at all stages of the process, improving the spatial coherence of the reconstruction.

As has been seen, this method produces reasonable reconstructions, using both a synthetic dataset and real flame images. Although some inaccuracy can be accounted for by the simple calibration and of course the fact that there are only a few cameras, the primary source of error is due to the algebraic tomography itself. Algebraic tomography is not a precision technique, and does not create an optimal reconstruction. It is rather an iterative technique that smears the projections across the reconstruction space in a way that reduces error, suited to cases where only a few views are available.

To create an optimal reconstruction one needs to invert the projection matrix, (as in [9],[18]), using a least-squares sparse matrix inversion technique, thereby creating a solution with minimum error. A state-of-the-art reconstruction method would appear to consist of a sparse matrix inversion density field estimation, using level sets to model and track the flame front. Although this has been deemed unsuitable for this project, it is worth looking at for future work in terms of optimal implementation techniques in order to make such a method feasible.

Chapter 6.

Experimental Results

The original purpose of this project was to develop a tool able to create 3D reconstruction of fire within a high voltage environment, in order to aid in research into the phenomenon of fire-induced flashover. Therefore, full-scale experiments were carried out on the roof of a high voltage lab, using a hot-air balloon burner. This enabled a power line over a fire to be simulated, and although not strictly necessary in terms of the flame reconstruction itself, this was needed to confirm that the flame video sequences could be successfully captured and reconstructed under flashover conditions.

6.1 The Setup

Figure 6.1 below shows the experimental configuration. A hot-air balloon burner, able to be operated remotely, was used to provide the fire. Above the burner was a metal grid, connected to earth. The flames from the burner extended above the grid towards a metal rod, sitting at the potential determined in the lab below. A water resistor was used to limit the current during flashover between the rod and the earthed grid. The voltage was varied via the control panel.

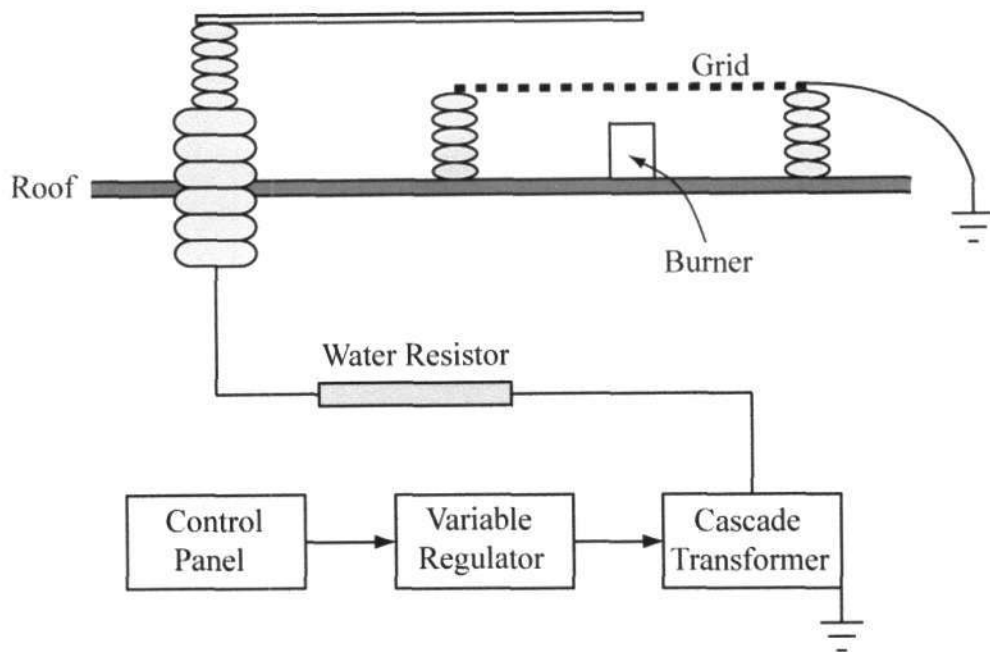



Figure 6.1: High-voltage arrangement.


The configuration of the cameras was as per Figure 5.2. Five cameras were used, placed approximately 36° apart at regular intervals around the flame. This arrangement gives optimum coverage of the flame. The photos below show the setup on the roof more clearly.



(a)



(b)



(c)

Figure 6.2: (a) The five cameras positioned around the grid. (b) The copper rod is connected to the high voltage source in the lab below and held over the grid. (c) The burner.

6.2 Experimental Issues

This section serves to address various issues encountered during testing, both in terms of experimental procedure and setup, and issues affecting the resulting reconstructions. Given the intended application of this project it was felt to be important to identify aspects that could be improved upon in order to allow for more effective future experiments.

6.2.1 System Configuration

This issue, separate from camera calibration, was not given enough attention, making the process more time-consuming and difficult than it could have been. Procedures and tools need to be created for the purposes of camera positioning and lens adjustment (zoom, focus, aperture size). For example, the lens' apertures should ideally be set so that the flame images use as large a dynamic range as possible without saturating. Systematically configuring all the cameras and lenses to be as similar as possible, and to meet certain specifications, will help to reduce error in the final reconstruction, particularly (but not exclusively) in a minimally calibrated system.

6.2.2 EMI

Because a digital (IEEE 1394) camera system was used, there was no visible noise or interference induced in the images themselves at any stage – indeed this was one of the reasons for choosing a digital camera system in the first place. However, the electromagnetic interference produced during flashover was greater than anticipated and occasionally affected the cameras and (or) PC, resulting in a loss of the video capture. The simplest solution is to move the cameras and PC further away from the fire, and if not possible then some further form of shielding may be required.

6.2.3 Real World Factors

Environmental factors, namely wind, proved problematic – although not much can be done about this, it should be noted that fire is strongly affected by wind and must therefore be considered when performing such experiments..

Two other issues arose later, during reconstruction, as a result of real world factors. Firstly, flame occlusion (for example by the earthed grid) was not sufficiently accounted for. Secondly, when processing the captured flame images it was noted that the reflection of the light from the flame by surrounding objects created the possibility of erroneous image segmentation. Methods should be formalised to attempt to handle these potential problems.

Although it does not fit with the original project specifications, the possibility of using an indoor, more permanent experimental setup should be considered. This would allow for controlled conditions, more complete camera calibration, repeatability, and ultimately better reconstruction accuracy.

6.3 Results

All results and reconstructions presented in this chapter are based on the captured video sequences of the flames produced by the hot-air balloon burner.

It is rather difficult to present video reconstruction sequences on paper. Even a single reconstruction frame is best viewed by rotating it in 3D space. For this reason several video clips of the reconstructions can be found at <http://www1.webng.com/ccfire/>.


Nonetheless some results are shown here, using density rendering (photorealistic) and surface rendering to display select reconstruction frames. A 2D comparison of cross-sectional slices is done, as well as an overall 3D reconstruction analysis.

6.3.1 3D Renderings

The density renderings were obtained by integrating the reconstruction along the rays defined by each pixel (ray-tracing) and applying the image formation model principles defined in §2.2.2. The 3D surface renderings were obtained using the VOLPACK [83], which is a fast, portable C++ library for performing volume rendering on a regular three-dimensional grid.

Each of the figures below shows the reconstruction of a different frame. In each case the top row is the density renderings while the bottom row is the surface renderings. The images from the viewing points of 0° , 40° , 80° and 120° are all novel viewpoints, while the final image, at 160° , is an original viewing angle (i.e. viewed from the same direction as one of the original flame images used to create the reconstruction.)

As mentioned previously the flame images were captured at a resolution of 640×480 . An effort was made to ensure the captured flames occupied a large portion of the available resolution - the largest reconstructed flame had a height of around 300 pixels (according to the reconstruction in frame A below). All of the images in Figure 6.3 below use the same scaling factor – they have been scaled to 50% of their original size in order to fit on the page. The reason the frames presented seem to differ in size is due to the dynamic nature of the captured flame – frames A, B and C are in fact consecutive frames from a captured video sequence. A full size reconstruction is shown in Figure 6.4, but because of printing resolutions and other factors it is best to view the videos to get a better sense of the resolution and quality of the reconstructions.




0°

40°


80°

120°

160°



Frame A




0°

40°



80°

120°



160°





Frame B





Frame C




Frame D



Frame E



Frame F



Frame G

Figure 6.3: Density and surface renderings of various frames of a reconstructed flame sequence. 160° is an original viewing angle while the rest are novel views.

Note how the density renderings appear less detailed when viewed from a novel viewpoint. One must remember, however, that the problem of tomographic reconstruction from few views is ill posed, and the algebraic technique employed works on the principle of smearing the projection data back over the reconstruction space. One would therefore expect the result to closely approximate the original image when viewed from the same angle (as in Figure 6.4 below), becoming more vague and less detailed as one moves further away from the original viewpoint.

Of course using more views to create the reconstruction means the solution would be more detailed and accurate, as the “unknown” space between original projection angles is effectively reduced.




Figure 6.4: (a) shows a raw, unprocessed captured image of a flame, and (b) shows the density rendering of the reconstruction of the corresponding frame. Note that part (b) is based on a relative density reconstruction and will thus not necessarily have the same absolute intensity as the original photograph.

6.3.2 2D Cross-Sectional Analysis

Figure 6.5 below shows the same cross-sectional slice taken from different reconstructions of the same flame. They differ in that figure (b) was created using an 8-voxel cubic basis function, promoting better local smoothness and spatial coherence.




Figure 6.5: Comparison of reconstructed cross-section.

The linear artefacts, most noticeable in (a), are as a result of the SIRT algebraic process, reflecting how the ray intensity is smeared back across the reconstruction space. Although improved, this linear effect is still noticeable in (b). Essentially the solution is suffering from overfitting. With only a few input views there simply is not enough information to reconstruct the complexity of the density field at the given resolution. This is evident in Figure 6.6, which shows how the degree of overfitting increases as the relative cross-section size increases.




Figure 6.6: Several cross-sections taken from the reconstruction of a single flame.

Smother solutions can be produced by using a smaller resolution, or larger and fewer basis functions. This of course means a trade-off between smoothness and accuracy, and an optimal solution would be dependant on the specific flame geometry, density distribution and input data resolution. This is not a trivial problem and as such is a good topic for further research.

6.3.3 3D Reconstruction Analysis

In order to perform a quantitative analysis, reconstructions were created using only four views. The result could then be viewed from the direction of the fifth image, thereby allowing one to compare a novel view with a ground truth (the fifth, unused image).

Two metrics were used to quantify the accuracy of the novel view:

- **Silhouette error:** The percentage error of the reconstruction silhouette relative to the ground truth silhouette. It must be noted that this measure is not an indication of absolute accuracy, since image processing is obviously required in order to extract the reference silhouette. The error term was calculated by expressing the sum of false positive and false negative projection pixel classifications as a percentage of the reference silhouette area.
- **Mean pixel error:** The average absolute error per pixel comparing the novel view and the ground truth image (using normal 8-bit greyscale). In order to prevent results being skewed by silhouette errors, only pixels corresponding to both the reconstruction silhouette and the reference silhouette were considered.

A video sequence consisting of 32 consecutive frames was analysed. The flames, created by the hot-air balloon burner, varied significantly from frame to frame in terms of size and geometry. A summary of the results is tabulated below.

Table 6.1: Summary of novel view analysis.

	min.	max.	mean.
Silhouette error %	5.2	25.3	12.6
Mean pixel error	10.1	24.6	15.7

The silhouette error primarily indicates the accuracy of the projection alignment, including the 2D image processing steps of rescaling and segmentation. The relatively high errors reported indicate that calibration and alignment procedure is inadequate for these more complex flames. The average silhouette error for the reconstruction silhouettes compared with the silhouettes of the original 4 input views, over all 32 frames, was 8.7%. This figure gives a rough indication of how much of the mean novel view silhouette error (12.6%) is due to projection misalignment (since the input view silhouettes are known), and hence the degree of improvement potentially possible as a result of better calibration.

The mean pixel error gives an indication of the accuracy of the tomographic density field estimation. Given that only 4 views are used to create the reconstruction one would expect a fair degree of error in this regard. This error is also affected by projection errors, so improved calibration would also improve the mean pixel error.

6.3.4 Summary

The results presented above have shown that the 3D Fuzzy Hull method is able to produce reasonably convincing reconstructions of fire, given only a few views. As far as accuracy is concerned the use of algebraic tomography remains a limiting factor – although this limitation is fundamentally imposed by the use of only a few views (five in this case).

However, the results indicate that there are several causes of inaccuracy that can be worked on to improve the accuracy of the reconstruction. The use of limited camera

calibration gave the system greater portability and generality, and proved adequate for synthetic data and simple flames such as a candle. However, the more complex and dynamic flames produced by the hot air balloon burner made it more difficult to extract calibration information from the flame images, causing greater error in terms of projection alignment. Another area for improvement would be a more systematic method of setting up the cameras and lenses in order to be as similar as possible (as well as preventing pixel saturation). Coupled with calibration errors this caused inaccuracies in the 2D image processing regarding image segmentation and scaling. Another effect not fully accounted for was the partial occlusion of the flame by the earthed grid.

The weakness of algebraic tomography was exposed when reconstructing these larger, more complex flames. The larger flame cross-sections simply had too much detail present along the rays to be reliably reconstructed with only a few views, resulting in overfitting and linear artefacts.

Even with these problems the reconstructions were relatively successful, and the above-mentioned causes of inaccuracy show that there is plenty of scope for improvement within the reconstruction process.

Chapter 7.

Conclusion

This thesis has investigated the three-dimensional reconstruction of fire, from a set of synchronised images. The intended application of this work is for use in research into the phenomenon of flashover induced by fires burning under high voltage transmission lines.

The simplified radiative transfer model was used to model the formation of images of fire, thereby enabling the reconstruction problem to be viewed from a tomographic perspective. In addition, the incorporation of fuzzy c-means image segmentation and the visual hull concept allowed the reconstruction to estimate the geometric structure of the flame, as well as the internal density field.

In order to test the process an experimental setup was created, with a hot air balloon burner positioned under a high voltage conductor. Video sequences of the flame were captured as flashover was induced. The reconstruction of the fire, using the 3D Fuzzy Hull method developed in this thesis, was presented in chapter 6. The resulting reconstructions appeared reasonably convincing, although a complete analysis is difficult since one does not have a 3D ground truth of the flame.

The reconstructions exposed the fact that the extent of camera calibration of the system is inadequate for more complex flames and needs improvement. Also, the fundamental limitation of tomographic reconstruction from only a few views was apparent when viewing the reconstruction from novel viewpoints. The use of algebraic tomography, while adequate for the smaller and less detailed flames used in initial testing, did not perform so well when reconstructing the complex flames produced by the hot air balloon burner.

7.1 Future Work

As previously mentioned, the problem of tomography with few views is severely underdetermined. No matter what technique is employed, the solution remains an estimate. The need therefore arises to find a method of parameterisation of the reconstruction space so as to suit the general structure of the flame. There is much potential work that can be done in this direction, since an optimal parameterisation would most likely be particular to a certain type of fire or fuel. These inferred structural elements or qualities would be used to guide the reconstruction process. Other physically based assumptions or characteristics of fire could also be incorporated into the process. Mathematically based simulations and models of fire (such as [27]) could be useful in this regard.

In addition, the use of variable basis functions would allow a more effective compromise to be found when looking at the smoothness and accuracy of the reconstruction, using larger basis functions to represent areas of relatively constant density and smaller ones to recreate fine geometric detail. Of course the problem then becomes exactly how to distribute the functions and how to determine their relevant sizes. Controlling the number of basis functions can also help to reduce overfitting, forcing the system to be overdetermined.

Sparse matrix inversion techniques, using a least squares approach, would appear to be the best way to actually perform the tomographic reconstruction, and would allow

one to develop parameterisation methods independently. However, in order to make a matrix based approach feasible one could only perform reconstructions on one cross-sectional slice at a time, due to the amount of computer resources required. Reducing the number of basis function would also help in this regard, and should thus be considered in conjunction with the density field parameterisation, as discussed above.

Finally, the use of level sets could be used to represent the flame surface. Combined with high-speed cameras this would allow one to analyse the flame front from both a temporal and geometric perspective.

Bibliography

- [1] F.F. Bologna, A.C. Britten and H.F. Vosloo. "Current Research into the Reduction of the Number of Transmission Line Faults on the Eskom MTS." *2nd South African Electric Power Research Conference*, 2001.
- [2] A.C. Britten. " "Unknown" Category of Eskom Line Faults: Investigation into Light Pollution on Glass Disc Insulators as an Underlying Cause of Unexplained Flashovers." Eskom Internal Report, TRR/CONS6/98, 1998.
- [3] P. Frost and H. Vosloo. "Satellite-Based Early Warnings of Fires, Reducing Fire Flashovers on Transmission Lines." *ESI Africa*, Issue 2, 2006.
- [4] A. Laurentini, "The visual hull concept for silhouette-based image understanding," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 16, no. 2, pp. 150-162, Feb 1994.
- [5] H.C. Bheemul, G. Lu and Y. Yan. "Three-dimensional visualization and quantitative characterization of gaseous flames." *Measurement Science and Technology*, vol.13, no.10, pp.1643-1650, October 2002.
- [6] W.B. Ng and Y. Zhang. "Stereoscopic imaging and reconstruction of the 3D geometry of flame surfaces." *Experiments in Fluids*, vol. 34, no. 4, pp. 484-493, 2003.
- [7] D.P. Correia, P. Ferrão, and A. Caldeira-Pires. "Advanced 3D Emission Tomography Flame Temperature Sensor." *Combustion Science and Technology*, vol. 163, pp. 1 24, 2001.
- [8] S.W. Hasinoff and K.N. Kutulakos. "Photo-Consistent 3D Fire by Flame-Sheet Decomposition." In *Proceedings of 9th IEEE International Conference on Computer Vision*, pp. 1184-1191, 2003.
- [9] I. Ihrke and M. Magnor. "Image-Based Tomographic Reconstruction of Flames." In *Proceedings ACM Siggraph / Eurographics Symposium Proceedings, Symposium on Computer Animation*, pp. 367-375, 2004.
- [10] A. Sukhnandan and D.A. Hoch. "An Investigation of Fire Induced Flashovers of Transmission Lines." Internal report, University of Natal, 2002.
- [11] J.R. Fonseca, A.L. Tan, V. Monassi, W.S. Junqueira, R.P. Silva, L.A.R. Assuncao and M.O.C. Melo. "Effects of agricultural fires on the performance of overhead transmission lines." *IEEE Transactions on Power Delivery*, vol. 5, no. 2, pp. 687-694, 1990.

- [12] W. Liming, G. Zhicheng, H.U. Qihao, L. Xidong, S. Xinfu and N.L. Allen. "Air Breakdown During Fires." *Tsinghua Science and Technology*, vol. 6, no. 5, pp. 497-499, 2001.
- [13] N.L. Allen, D.S.K. Lam and D.A. Greaves. "Tests on the breakdown of air at elevated temperatures in non-uniform electric fields." *IEE Proceedings, Science, Measurement and Technology*, vol. 147, no. 6, pp. 291-295, 2000.
- [14] A. Robledo-Martinez, E. Guzman and J.L. Hernandez. "Dielectric Characteristics of a Model Transmission Line in the Presence of Fire." *IEEE Transactions on Electrical Insulation*, vol. 26, no. 4, pp. 776-782, 1991.
- [15] D. Drysdale. "*An Introduction to Fire Dynamics*." Second Edition, John Wiley & Sons, 1998.
- [16] A. Sukhnandan and D A. Hoch. "Fire Induced Flashovers of Transmission Lines: Theoretical Models." *IEEE Africon*, 2002.
- [17] J.W. Lyons. "*Fire*." Scientific American Books Inc., 1985.
- [18] S. Hasinoff. "*Three-dimensional reconstruction of fire from images*." Master's thesis, University of Toronto, Dept. of Computer Science, 2002.
- [19] A. Schwarz. "Multi-tomographic Flame Analysis with a Schlieren Apparatus." *Measurement Science and Technology*, vol. 7, pp. 406-413, 1996.
- [20] A.F. Ibarreta and C.J. Sung. "Determination of Flame Location Using Rainbow Schlieren Deflectometry." *Third Joint Meeting of the U.S. Section of the Combustion Institute*, Chicago IL, 2003.
- [21] B. Jeanne, E. Samson, B. Renou and A. Boukhalfa. "Bunsen flame analysis using simultaneous tomographic images and PIV in the fresh and burnt gases." *10th International Symposium on Applications of Laser Techniques to Fluid Mechanics*, 2002.
- [22] R. Abu-Gharbieh, T. Gustavsson, C. Kaminski and G. Hamarneh. "Flame Front Matching and Tracking in PLIF Images Using Geodesic Paths and Level Sets." *IEEE International Conference on Computer Vision workshop on Variational and Level Set Methods in Computer Vision*, pp. 112-118, 2001.
- [23] W.G. Bessler, C. Schulz, T. Lee, D.I. Shin, M. Hofmann, J.B. Jeffries, J. Wolfrum, and R.K. Hanson. "Quantitative NO-LIF Imaging in High-Pressure Flames." *Optical and Laser Diagnostics*, vol. 177, pp. 107-114, 2003.
- [24] G.J. Smallwood, O.L. Gulder, D.R. Snelling, B.M. Deschamps and I. Gokalp. "Characterization of flame front surfaces in turbulent premixed methane/air combustion." *Combustion and Flame*, vol. 101, no. 4, pp. 461-470, 1995.
- [25] R.R. Skaggs, R.G. Daniel, A.W. Miziolek, K.L. McNesby, V.I. Babushok, W. Tsang and M.D. Smooke. "Spectroscopic Studies of Inhibited Opposed-Flow

- Propane/Air Flames." In *Proceedings Halon Options Technical Working Conference*, pp. 117-131, 1999.
- [26] J.A. Sanzol. "Video Feedback Control of Pulverized Coal Flow." Master's Thesis, Lulea University of Technology, 2001.
- [27] D.Q. Nguyen, R.P. Fedkiw and H.W. Jensen. "Physically Based Modeling and Animation of Fire." *ACM Transactions on Graphics*, vol. 21, no. 3, pp. 721-728, 2002.
- [28] R.A. Felice. "The Spectropyrometer - a Practical Multi-wavelength Pyrometer." *Eighth International Temperature Symposium*, 2002.
- [29] N. Chigier. "Combustion Measurements." Hemisphere Publishing Corporation, 1991.
- [30] G.T. Herman. "Image Reconstruction from Projections: The Fundamentals of Computerized Tomography." Academic Press, New York, 1980.
- [31] A.C. Kak and M. Slaney. "Principles of Computerized Tomographic Imaging." IEEE Press, 1988.
- [32] G. Wang and M.W. Vannier. "Computerized tomography." *Encyclopedia of Electrical and Electronics Engineering*, vol. 4, pp. 8-24, Webster JG, John Wiley & Sons, 1999.
- [33] D.P. Correia, P. Ferrão, and A. Caldeira-Pires. "Flame 3d tomography sensor for in-furnace diagnostics." *28th International Symposium on Combustion*, University of Edinburgh, Scotland, 2000.
- [34] A. Mohammad-Djafari. "A Bayesian Approach to Shape Reconstruction of a Compact Object from a few Number of Projections." In *Proceedings of the Maximum Entropy Conference*, South Africa, pp. 68-77, 1996.
- [35] A. Mohammad-Djafari. "A full Bayesian approach for the inverse problem." In *Maximum Entropy and Bayesian Methods*, (K. Hanson and R. Silver, ed.), Kluwer Academic Publishers, 1996.
- [36] J. Sachs Jr. and K. Sauer. "Object-oriented methods in Bayesian 3-D tomographic reconstruction from radiographs." *Circuits and Systems, Proceedings of the 36th Midwest Symposium*, vol. 1, pp. 241-244, August 1993.
- [37] C. Hamon, C. Roux, J.L. Coatrieux, and R. Collorec. "An analytical method for 3D tomographic reconstruction from a small number of projections using a simple interpolation scheme." *Acoustics, Speech, and Signal Processing, ICASSP-90*, vol. 4, pp. 1857-1860, 1990.
- [38] A. Mohammad-Djafari. "Shape reconstruction in X-ray tomography from a small number of projections using deformable models." In *SPIE 97 annual meeting*, San Diego, CA, USA, vol. 3170, 1997.

- [39] K. Girodias, H. Barrett, and R. Shoemaker. Parallel simulated annealing for emission tomography. *Physics in Medicine and Biology*, vol. 36, no.7, 921--938, 1991.
- [40] M. Bhatia, W. Karl, and A. Willsky. "A wavelet-based method for multiscale tomographic reconstruction." *IEEE Transactions on Medical Imaging*, vol. 15, no. 1, pp. 92-101, 1996.
- [41] Y.H. Chiu and S.F. Yau. "An artificial neural network for tomographic reconstruction of time varying object." In *Proceedings of Speech, Image Processing and Neural Networks*, ISSIPNN '94, vol. 2, pp. 409-412, 1994.
- [42] J.A. Fessler. "Penalized weighted least-squares image reconstruction for positron emission tomography." *IEEE Transactions on Medical Imaging*, vol. 13, no. 2, pp. 290-300, June 1994.
- [43] L. Staib and J. Duncan. "Model based deformable surface finding for medical images." *IEEE Transactions on Medical Imaging*, vol. 15, pp. 720-731, Oct 1996.
- [44] V. Elangovan and R. Whitaker. "From sinograms to surfaces: A direct approach to the segmentation of tomographic data." In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pp. 213-223, Oct. 2001.
- [45] R. Poli and G. Valli. "Hopfield neural nets for the optimum segmentation of medical images." In *Handbook of Neural Computation*, E. Fiesler and R. Beale, Eds, chapter G.5.5. Oxford University Press, 1996.
- [46] S. Linying, B. Sharp, and D.N. Davis. "Neural Networks for X-Ray Image Segmentation." *ICEIS 1999*, pp. 264-271, 1999.
- [47] D.L. Pham and J.L. Prince. "Adaptive Fuzzy Segmentation of Magnetic Resonance Images." *IEEE Transactions on Medical Imaging*, vol. 18 no. 9, pp. 737-752, 1999.
- [48] M.N. Ahmed, S.M. Yamany, N. Mohamed, A.A. Farag, and T. Moriarty. "A Modified Fuzzy C-Means Algorithm for Bias Field Estimation and Segmentation of MRI Data." *IEEE Transactions on Medical Imaging*, vol. 21, no. 3, pp. 193-199, 2002.
- [49] J.C. Bezdek. "*Pattern Recognition with Fuzzy Objective Function Algorithms*." Plenum Press, New York, 1981.
- [50] S. Chuai-Aree, C. Lursinsap, P. Sophatsathit, and S. Siripant. "Fuzzy C-Mean: A Statistical Feature Classification of Text and Image Segmentation Method." In *Proceedings of International Conference on Intelligent Technology 2000*, Thailand, pp. 279-284, 2000.

- [51] N.A. Mohamed, M.N. Ahmed, and A. Farag. "Modified Fuzzy C-Mean in Medical Image Segmentation." In *Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing*, Piscataway, NJ, USA, vol.6, pp.3429-3432, 1999.
- [52] S. Osher and J. Sethian. "Fronts Propagating with Curvature-Dependent Speed: Algorithms Based on the Hamilton-Jacobi Formulation." *Journal of Computational Physics*, vol. 79, pp.12-49, 1988.
- [53] D. Adalsteinsson and J Sethian. "A Fast Level Set Method for Propagating Interfaces." *Journal of Computational Physics*, vol. 118, pp. 269-277, 1995.
- [54] J. Sethian. "Level Set Methods: An Act of Violence. Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision and Materials Science." Preprint for *American Scientist*, 1997.
- [55] J. Sethian. "*Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science.*" Cambridge University Press, 1999 Cambridge Monograph on Applied and Computational Mathematics (Chapter 1).
- [56] J. Sethian. "Fast Marching Methods and Level Set Methods for Propagating Interfaces." von Karman Institute Lecture Series, *Computational Fluid Mechanics*, 1998.
- [57] J. Sethian. "A Fast Marching Level Set Method for Monotonically Advancing Fronts." In *Proceedings of the National Academy of Sciences*, vol. 93, pp.1591-1595, 1996.
- [58] R. Malladi and J. Sethian. "An $O(N \log N)$ Algorithm for Shape Modeling." In *Proceedings of the National Academy of Sciences*, Vol. 93, pp. 9389-9392, 1996.
- [59] R. Malladi and J. Sethian. "A Unified Approach to Noise Removal, Image Enhancement, and Shape Recovery." *IEEE Transactions on Image Processing*, vol. 5, no. 11, pp.1154-1168, 1996.
- [60] R.T. Whitaker. "A level-set approach to 3D reconstruction from range data." *International Journal of Computer Vision*, vol. 29, pp. 203-231, 1998.
- [61] R. Malladi, J. Sethian, and B.C. Vemuri. "A Fast Level Set based Algorithm for Topology-Independent Shape Modeling." *Math. Imaging and Vision*, vol. 6, pp. 269-290, 1996.
- [62] M. Droske, B. Meyer, M. Rumpf, and C. Schaller. "An adaptive level set method for medical image segmentation." In *Proceedings of the Annual Symposium on Information Processing in Medical Imaging*, (R. Leahy and M. Insana, eds.), Springer, Lecture Notes Computer Science, 2001.

- [63] R. Malladi and J. Sethian. "Level Set Methods for Curvature Flow, Image Enhancement, and Shape Recovery in Medical Images." In *Proceedings of Conference on Visualization and Mathematics*, Berlin, Germany, Springer-Verlag, pp. 329-345, 1997.
- [64] J. Sethian. "Curvature Flow and Entropy Conditions Applied to Grid Generation." *Journal of Computational Physics*, vol. 115, pp. 440-454, 1994.
- [65] J. Sethian and D. Adalsteinsson. "An Overview of Level Set Methods for Etching, Deposition, and Lithography Development." *IEEE Transactions on Semiconductor Devices*, vol. 10, no. 1, pp.167-184, 1997.
- [66] K. Forbes, A. Voigt and N. Bodika. "Using Silhouette Consistency Constraints to Build 3D Models." In *Proceedings of the Fourteenth Annual Symposium of the Pattern Recognition Association of South Africa (PRASA 2003)*, November 2003.
- [67] K. Kutulakos and S. Seitz. "A theory of shape by space carving." *International Journal of Computer Vision*, vol. 38, no. 3, pp.199-218, 2000.
- [68] G.G. Slabaugh. "*Novel Volumetric Scene Reconstruction Methods for New View Synthesis.*" PhD Thesis, Georgia Institute of Technology, School of Electrical and Computer Engineering, 2002.
- [69] K.N. Kutulakos and S.M. Seitz. "What Do N Photographs Tell Us About 3D Shape?" University of Rochester, Rochester, NY, Tech. Rep. 680, Jan. 1998.
- [70] S.M. Seitz and C.R. Dyer. "Photorealistic Scene Reconstruction by Voxel Coloring." In *Proceedings of CVPR '97*, pp. 1067-1073, Puerto Rico, 1997
- [71] W.B. Culbertson, T. Malzbender, and G. Slabaugh. "Generalized voxel coloring." In *Proceedings of the International Workshop on Vision Algorithms*, Lecture Notes in Computer Science, vol. 1883, pp. 100-115, 2000.
- [72] T. Bonfort and P. Sturm. "Voxel carving for specular surfaces." In *Proceedings of the 9th International Conference on Computer Vision*, 2003.
- [73] J.D. Bonet and P. Viola. "Roxels: Responsibly Weighted 3D Volume Reconstruction." In *Proceedings of the IEEE International Conference on Computer Vision 1999*, pp. 418-425, 1999.
- [74] F. Dacheille, K. Mueller and A. Kaufman. "Volumetric Backprojection." In *Proceedings of the IEEE Symposium on Volume Visualization*, pp. 109-117, 2000.
- [75] H.C. Bheemul, G. Lu and Y. Yan. "Digital Imaging Based Three-Dimensional Characterization of Combustion Flames." In *Proceedings of IEEE IMTC (Instrumentation & Measurement Technology Conference)*, USA 2003, vol.1, pp. 420-424, 2003.

- [76] T. Piotrowski, A. Plaskowski, M.S. Beck and M. Byard. "Dynamic Combustion Tomography." *1st World Congress on Industrial Process Tomography*, pp. 167-172, 1999.
- [77] S.J. Carey, H. McCann, D.E. Winterbone and E. Clough. "Near Infra-Red Absorption Tomography for Measurement of Chemical Species Distribution." *1st World Congress on Industrial Process Tomography*, pp. 480-487, 1999.
- [78] P. Vallikul, R. Goulard and C. Mavriplis. "Tomographic Reconstruction of Probability Density Functions in Turbulent Flames." *International Interflam Conference*, pp. 235-243, 1996.
- [79] S.C. Bates, R. Carangelo, K. Knight and M.A. Serio. "Fourier Transform Infrared Hadamard Tomography of Sooting Flames." American Institute of Physics, *Review of Scientific Instrumentation*, vol. 64, no. 5, 1993.
- [80] D.R. Snelling, K.A. Thomson, G.J. Smallwood, O.L. Gülder, E.J. Weckman and R.A. Fraser. "Spectrally Resolved Measurement of Flame Radiation to Determine Soot Temperature and Concentration." *AIAA Journal*, vol. 40, pp. 1789-1795, 2002.
- [81] P. Dunkley, P. Bryanston-Cross, E. Moll and B. Timmerman. "OET a new Flame Front Diagnostic & PIV in Transonic flows." *The 16th Symposium on Measuring Techniques in Transonic and Supersonic Flow in Cascades and Turbomachinery*, 2002.
- [82] The Dragonfly Technical Reference Manual, 2004. <http://www.ptgrey.com>.
- [83] The Volpack Volume Rendering Library, written by Philippe Lacroute, The Leland Stanford Junior University, 1995.
<http://graphics.stanford.edu/software/volpack/>
- [84] W.E. Lorensen and H.E. Cline. "Marching Cubes: A High Resolution 3D Surface Construction Algorithm." In *Proceedings of SIGGRAPH*, vol. 21, no. 4, pp. 163-169, 1987.
- [85] T. Poston, T.T. Wong and P.A. Heng. "Multiresolution Isosurface Extraction with Adaptive Skeleton Climbing." *Computer Graphics Forum*, vol. 17, no. 3, pp. 137-148, September 1998.
-