Image Processing Techniques for Doppler Global Velocimetry

A thesis submitted in partial fulfilment of the requirements of the degree of Doctor of Philosophy at the University of Oxford, Michaelmas Term, 1997.

R.J.Manners
St. Catherine's College.

University of Oxford
Department of Engineering Science
Parks Road
Oxford OX1 3PJ
Tel. 01865 273000

Abstract

A thesis submitted in partial fulfilment of the requirements of the degree of
Doctor of Philosophy at the University of Oxford, Michaelmas Term, 1997.

R.J.Manners

St. Catherine's College

Image Processing Techniques for Doppler Global Velocimetry

There is a demand for a whole field velocimetry technique which offers the capability of rapid characterisation of complex engineering flow fields. This thesis describes a research programme aimed at the development of a reliable Doppler global velocimeter, suited to the measurement of such flows. The programme of work undertaken is discussed with reference both to research undertaken elsewhere and to previous work on the system at Oxford. While much of the underlying technology required for the construction of an accurate and reliable velocimeter has already been studied in Oxford and elsewhere, little attention has been paid by previous workers to the examination of the impact of data processing techniques on attainable flow measurement accuracy. In the present work, a number of image processing methods have been utilised for Doppler global velocimetry data processing. Those methods are described here, together with a theoretical analysis of their expected performance when applied to Doppler global velocimetry data. The expected error resulting from image processing considerations and also from the physical characteristics of the Doppler global velocimetry hardware are quantified in such a way that error estimates may be computed for real measured data frames. The results of the application of the velocimeter to the simple test case of measuring a velocity component of a rotating disc are presented. The velocimeter was subsequently applied to the measurement of a free jet flow and to a transonic flow field in a convergent-divergent nozzle. Correlations with accepted velocity field values were undertaken, and compared to the expected error previously determined. The choice of image processing algorithms was found to be of great importance in terms of Doppler global velocimetry measurement accuracy.

Acknowledgements

The work presented in this thesis would not have been possible without the help and support received from numerous people.

First, I would like to thank my supervisor, Dr. R.W.Ainsworth for his support, guidance and enthusiasm throughout my time of study in Oxford.

I am especially grateful to everyone in the workshop, especially Nigel Brett, for manufacturing and modifying much of the experimental apparatus from my ideas and sketches. The enthusiasm and good humour of the workshop staff has contributed to making my time in the laboratory both productive and enjoyable.

There are many others who have been of invaluable assistance to me over the last three years, not least the other members of the Doppler Global Velocimetry research group: Steve Thorpe and Nathan Quinlan. Special thanks are due to Charles Thomas for his help in producing CFD models for the transonic nozzle flow.

Outside of my immediate colleagues I would like to thank Dr. Fiona Somerset for her dilligent reading of this thesis, and her helpful comments regarding style and English usage. I am also grateful to Macushla McGown for providing the impetus to write this thesis in the first place.

Finally, I would like to acknowledge the Engineering and Physical Sciences Reasearch Council (EPSRC) for the financial support that has made this programme of research possible.

Table of Contents

1 Introduction **1**
- 1.1 Background to This Work . . . 1
- 1.2 Objectives of the Present Research Programme . . . 2
- 1.3 Thesis Overview . . . 4

2 The Doppler Global Velocimetry Technique **7**
- 2.1 Overview . . . 7
- 2.2 The Implementation of the Technique at Oxford . . . 8
- 2.3 Measurement Procedure . . . 12
- 2.4 Doppler Shift . . . 14
- 2.5 The Iodine Vapour Absorption Cell . . . 15
- 2.6 The Argon Ion Laser as a Light Source . . . 19
- 2.7 The CCD Camera . . . 22
- 2.8 Flow Field Seeding . . . 24
- 2.9 Two Component Measurement . . . 27
- 2.10 Summary . . . 33

3 Image Processing Techniques **35**
- 3.1 Overview of Image Processing Techniques . . . 36
- 3.2 Centroid Estimation . . . 46
- 3.3 Image Transformation Techniques . . . 47
- 3.4 Interpolation Methods . . . 55
- 3.5 The Application of Image Processing for Two Component Measurements . . . 63
- 3.6 Summary . . . 64

4 Further Image Processing Techniques **66**
- 4.1 Absolute Velocity Determination and the Calibration of the Iodine Cell . . . 66
- 4.2 Two Component Velocimetry Determination . . . 83

5 Error Sources and Analysis **89**
- 5.1 Overview of Doppler Global Velocimetry Error Sources . . . 89
- 5.2 Camera Pixel Shot Noise . . . 92
- 5.3 Iodine Cell Transmission Calibration . . . 98
- 5.4 Iodine Absorption Filter Instability . . . 101
- 5.5 Adverse Reflections of the Laser Light Sheet . . . 102
- 5.6 Image Misalignment . . . 103
- 5.7 Centroid Estimation . . . 105
- 5.8 Image Transformation Techniques . . . 130
- 5.9 Interpolation Methods . . . 136
- 5.10 The Effect of Viewing Angle on Error . . . 139
- 5.11 Error Considerations of Two Component Measurements . . . 140
- 5.12 Summary and the Development of an Error Modelling Code . . . 142

6 Velocimetry Measurements of a Rotating Disc **152**
- 6.1 Motivation . . . 152
- 6.2 Configuration . . . 152
- 6.3 Flow Measurement and Image Processing Procedure . . . 157
- 6.4 Doppler Global Velocimetry Measurements . . . 161
- 6.5 Error Analysis . . . 168
- 6.6 Two Component Doppler Global Velocimetry Measurement . . . 171
- 6.7 Summary . . . 172

7 Velocimetry Measurements of a Free Jet Flow **174**
7.1 Motivation . . . 174
7.2 Configuration . . . 176
7.3 Description of the Free-Jet Flow Field . . . 178
7.4 Doppler Global Velocimetry Measurements . . . 179
7.5 Correlation of Published Free Jet Flow Field Characteristics with Doppler Global Velocimetry Measurements . . . 193
7.6 Error Analysis . . . 195
7.7 Effect of the Choice of Image Processing Technique . . . 209
7.8 Summary . . . 214

8 Velocimetry Measurements of a Convergent-Divergent Nozzle **216**
8.1 Motivation . . . 216
8.2 Configuration . . . 216
8.3 Doppler Global Velocimetry Measurements . . . 224
8.4 Error Analysis . . . 232
8.5 Summary . . . 234

9 Summary and Conclusions **235**
9.1 Summary . . . 235
9.2 Suggestions for Future Work . . . 238
9.3 Conclusions . . . 239

Nomenclature

A_n	Iodine calibration interpolation polynomial coefficient
D	Free jet exit diameter
$F(s)$	Fourier transform of intensity profile
$G(s)$	Fourier transform of sampled intensity profile
$H(f)$	Fourier transform of interpolation kernel
I_i	Intensity value at position i; Incident intensity
I_t	Transmitted intensity
$I_{\max}$	Maximum intensity count for CCD pixel
K	Number of polynomial warp coefficients; Spreading rate of free jet
L	Mirror separation of laser cavity
M	Number of corresponding points
$\mathbf{M}$	Bilinear warp matrix
N	Degree of polynomial warp; Number of alignment features in x or y direction
N_e	Number of electrons accumulated in CCD pixel
N_{fw}	Maximum number of electrons accumulated in CCD pixel (full well capacity)
$P_i(x, y)$	Polynomial in x and y
R	Intensity ratio
T	Sampling spatial frequency; Temperature
V	Velocity as a scalar
$\vec{V}$	Velocity vector
V_{BG}	Band gap voltage of silicon
W	Alignment feature width
$\mathbf{W}$	Polynomial warp matrix
U	Axial velocity of free jet
U_o	Axial velocity of free jet at exit plane
U, V	Polynomial warp mapping functions
$a(x)$	True image intensity profile
a_i, a_{ij}	Interpolation coefficients
a_i, b_i	Bilinear Warp coefficient i
a_{ij}, b_{ij}	Polynomial warp coefficient (i, j)
$\vec{a}, \vec{b}$	Vectors of bilinear, polynomial warp coefficients
c	Propagation velocity of light
c_k	Interpolation coefficient
$e(x)$	Intensity profile on CCD array
f	Spatial frequency
f_c	Cut-off frequency
f_i	Intensity at position i in discrete domain
f_n	Nyquist frequency
$f(x)$	Intensity at position x in continuous domain; Interpolating polynomial

$g(x)$	Sequence of sampled pixel values
$h(x)$	Interpolation kernel
$h_n(x, y)$	Basis function for orthogonal polynomial warp
$\hat{\imath}$	Illumination direction unit vector
k	Boltzmann constant
k_{conf}	Coefficient of effect of configuration on measurement error
k_i	Coefficient of effect of interpolation method on measurement noise
k_w	Coefficient of effect of warp method on misalignment
$m(i, j)$	Intensity of CCD pixel at (i, j)
$\hat{o}$	Observation direction unit vector
$p(x)$	Pixel sensitivity profile
q	Charge
r	Radial distance from free jet centre-line
$s(x)$	Sampling function
s_c	Spatial cut off frequency
u_i, v_i	Warp control point i (source space)
$\vec{u}, \vec{v}$	Vectors of warp control points (source space)
w	Half of alignment feature width, W
$\bar{x}$	Apparent centroid position
$\bar{x'}$	True centroid position
x_c, y_c	Position of entroid
x_L, x_R	Extent of alignment feature; Extent of bounding box in x direction
x_i, y_i	Warp control point i (destination space)
x_k, y_k	Warp control point k (source space)
$\vec{x}, \vec{y}$	Vectors of warp control points (destination space)
y_L, y_R	Extent of bounding box in y direction
z	Axial distance from free jet exit plane
z_c	Potential core length of free jet

Greek Symbols

$\Delta\nu$	Doppler shift
α_{ij}	Parameters used for construction of orthogonal polynomials
α, β	Bilinear interpolation parameters
δ	Overall feature offset relative to CCD array
δ_k	Error at point k under polynomial warp
δ_L, δ_R	Left and right hand side offset of feature relative to CCD array
$\delta x, \delta y$	Image misalignment
ε	Centroid estimation error
ε_I	Error in intensity value
ε_R	Error in intensity ratio
ε_i	Error source i
$\varepsilon_{\Delta\nu}$	Error in determined Doppler shift
ε_{ν_o}	Error in determination of ν_o
η_i	Noise value at pixel i
θ_i	Observation angle i
ν	Doppler shifted frequency
ν_o	Laser emission frequency
ρ	Density
σ	Standard deviation
ϕ	Transmission fraction of iodine cell

Chapter 1

Introduction

1.1 Background to This Work

The development of measurement technology is driven by the needs of potential users. For many years, designs of aerodynamic and turbomachinery components have been studied experimentally using facilities such as wind tunnels. The use of such an approach has led to the development of suitable instrumentation to allow flow data to be acquired in order to enable the assessment of performance. A number of instruments such as pressure sensing and hot wire probes have been developed to measure flow velocity. While those techniques worked well, the disturbance of flow created by the physical presence of the instrument led to concern over the measurement accuracy. Subsequently, non-intrusive techniques such as Schlieren photography and interferometry methods have provided views of the fluid flow, especially of shock waves. Unfortunately, such methods yield little quantitative data. Due to the non-linear nature of photographic film and of the measurement techniques themselves, it is hard to write suitable software to process the results accurately.

The advent of the laser allowed the development of laser Doppler anemometry (LDA), a technique that measures velocity at a point within the flow to a high accuracy without disturbing the flow. Laser Doppler anemometry techniques commonly employ a heterodyne detection technique whereby shifted scattered light is mixed with unshifted reference light to obtain a difference (Doppler) frequency. As an alternative to heterodyne detection, Fabry-Perot interferometry may be employed to measure the shift in scattered light frequency directly. Such techniques are not applicable to global measurements, since their allowable acceptance angle is too limited. To measure the flow over an entire plane, the laser Doppler anemometry apparatus must be scanned throughout that plane. This may take as long as several hours, during which time the flow must be steady-state so such techniques are of no use for measurement of unsteady flows. A Michelson interferometer may be used to increase the allowable acceptance angle, and this method has been applied to global measurements. The technique produces fringe maps of constant phase shift between direct and time-delayed scattered light. The fringes indicate locations where the Doppler shift of the scattered light is a multiple of the time delay in the interferometer. The difficulty of maintaining alignment of the interferometer and the still relatively narrow acceptance angle of the technique limit its application. The required optical precision of the interferometer

makes subsonic measurements very difficult. Fringe images may be hard to interpret or process.

Subsequently, further experimental work and the requirement for the validation of computational fluid dynamics codes, have provided the impetus for the development of further flow measurement methods. In particular, the ability to perform simultaneous velocity measurements over an extended area was required. In response to these demands, particle image velocimetry (PIV) was developed. Particle image velocimetry uses a photographic or charge-coupled device (CCD) camera to obtain a double-exposure of the particle field within the seeded flow using a double-pulsed laser light sheet. The images are then examined to find corresponding particles, and the distance between these corresponding points is then determined to obtain velocity information. Unfortunately, the autocorrelation computations required for particle image velocimetry image processing can be extremely time consuming. While the technique allows the simultaneous measurement of velocities over the viewing area, it has several drawbacks including a limited viewing area (typically 10 by 15 cm) and a directional ambiguity in the results. In addition, the technique is limited to two velocity components, and is extremely sensitive to seeding particle density.

While these problems may be reduced or even eliminated with further research, Doppler global velocimetry (DGV), explained in detail in Chapter 2, has been proposed as an alternative method. Doppler global velocimetry offers the whole-field measurement characteristics of particle image velocimetry, without many of the drawbacks outlined above. Since measurements may be undertaken of the entire flow field in a very short period of time, Doppler global velocimetry allows the study of unsteady flow fields. Analysis time is short compared with that required for particle image velocimetry. The Doppler global system has a relatively simple optical system by comparison with Michelson interferometer techniques, and Doppler global velocimetry results are intrinsically easier to interpret.

1.2 Objectives of the Present Research Programme

The principles of the Doppler global velocimetry technique were first expounded by Komine, Lee and Cavone [Komine *et al.* 1991]. The focus of that work was the development of a real-time flow measurement intrument. Such an arrangement is intended to allow a representation of the velocity of a flow field to be viewed at video frame rate, at the same time as the measurement is being undertaken. Subsequent workers [Meyers and Komine 1991, Morrison, Gaharan and DeOtte 1994, Ford and Tatum 1994] have also concentrated primarily on this real-time approach to flow visualisation.

Real-time systems, while attractive in principle, suffer from a significant drawback. Due to the constraints of the video frame rate employed, it is not possible to perform any significant image processing operations in real-time. As explained in Chapter 2, the Doppler global system relies on the accurate alignment of a pair of images. In real-time systems, the experimenter must perform this alignment by adjusting certain optical elements manually. It will be seen in Chapter 5 that a high degree of alignment accuracy is required if accurate results are to be obtained. It will further be seen that such accuracy is achievable only with extensive post-processing of the measured data. At present it is impractical to perform such processing at video frame rates due to the limited performance of available computer systems. Additionally, the use of video equipment carries a further penalty — that of increased measurement noise. The effect of this noise will be seen to lead to a significant degradation in achievable measurement accuracy.

The approach taken towards the implementation of a Doppler global velocimetry system at Oxford (for more details see [Ainsworth and Thorpe 1994], [Thorpe and Ainsworth 1994], [Thorpe, Ainsworth and Manners 1995] and [Thorpe 1996]) differs from that outlined above. Rather than concentrating on the development of a system which may allow real-time qualitative visualisation of flow field velocity, research has focussed on the development of a single-shot flow measurement instrument capable of accurate quantitative velocity measurement.

The Oxford system, then, makes no attempt to offer any real-time visualisation of velocity. Rather, it allows the acquisition of data which may be subsequently post-processed to yield velocity information. Such an approach allows a great deal more freedom in terms of the image processing operations which are available to the experimenter. Thus, significantly more accurate image alignment may be achieved, resulting in superior quantitative results. Dispensing with the requirement for real-time performance allows the use of a slow-scan cooled CCD camera, resulting in markedly lower measurement noise. It is the view of the author that the adherence to the development of real-time systems precludes the acquisition of accurate quantitative velocity measurements.

In the longer term, the objective of the current research program is to develop an instrument suited to the measurement of turbomachinery flow fields. The nature of many turbomachinery experiments is such that representative flow conditions are only maintained for a fraction of a second. This being the case, it is arguable that real-time operation of a measurement system is of little or no benefit for such applications.

Due to the focus of previous research on real-time flow visualisation, little attempt has been made in the published literature to estimate and quantify the error performance of such systems. In order, however, to obtain accurate quantitative results it is essential not only to examine suitable image processing algorithms for the alignment of data images closely, but also to study and quantify the error expected in the application of the system.

The current programme of research, then, may be seen as an attempt to develop a single-shot approach to Doppler global velocimetry. Such an approach allows the extensive use of image processing techniques, which have been found to be necessary in order to reliably achieve accurate velocity measurements. The error contributions of all elements of the system are estimated. This error analysis allows the expected velocity error at each point in the measurement frame to be quantified. The expected error may then be compared to the measurement error observed in practice, and the overall system performance assessed.

1.3 Thesis Overview

Doppler global velocimetry is a relatively new measurement technique that has undergone continuous research and development since its inception in 1991. The work conducted at Oxford has occurred in parallel with that of several other groups around the world. The relatively recent development of the technique means that there is only a modest quantity of literature available on specific realisations of the Doppler global velocimetry method. Accordingly, this thesis does not contain a specific chapter devoted to a presentation of an exhaustive literature review. Rather, references to published literature are introduced as and when necessary.

In addition to the publications relating to previous Doppler global velocimetry work undertaken at Oxford [Ainsworth and Thorpe 1994] [Thorpe and Ainsworth 1994], there are a number of publications that have been derived from the work associated with this thesis. [Thorpe, Ainsworth and Manners 1995] and [Thorpe, Ainsworth and Manners 1996] describe the physical realisation of the velocimeter. [Manners, Thorpe and Ainsworth 1996] concentrates on image processing concerns. [Ainsworth, Thorpe and Manners 1996] and [Ainsworth, Thorpe and Manners 1997] present a detailed overview of the state of Doppler global velocimetry research at Oxford and elsewhere. A subsequent application of the Doppler global velocimeter described in this thesis to the time-averaged measurement of a transient supersonic flow is detailed by Quinlan *et al.* [Quinlan *et al.* 1997].

The contents of this thesis may be summarised as follows: Chapter 2 introduces the fundamentals of the underlying technology required in order to implement the Doppler global

system. This material is included in the interest of completeness and in order to enable the reader to readily grasp the significance of the work presented in this thesis. My contribution to the development of data processing techniques for Doppler global velocimetry is detailed in the subsequent chapters. The image processing techniques which may be applied so as to effect accurate image alignment for the case of a non-real-time velocimeter are explained in Chapter 3. Chapter 4 discusses additional post-processing steps which may be taken in order to obtain quantitative velocity measurements, and also to allow the simultaneous measurement of two velocity components. In both Chapter 3 and Chapter 4, a theoretical approach to the data processing operations performed is presented, together with an explanation of how these techniques have been implemented at Oxford.

A detailed theoretical analysis of the sources of error which are present in the velocimeter is developed in Chapter 5. This analysis discusses the fundamental error sources due to the properties of the Doppler global velocimetry hardware. In addition, a detailed evaluation of the expected performance of the image processing techniques utilised is presented, together with estimates of their effect on the overall accuracy of the computed velocity. A numerical model of expected error is developed which allows the velocity error at each point in the image to be estimated.

Chapters 6, 7 and 8 describe a set of experiments that have been undertaken in order to evaluate the performance of the velocimeter in practice. Measured results from the application of the velocimeter to the case of a rotating disc are presented in Chapter 6. These results are used to illustrate the post-processing operations which must be undertaken in order to obtain quantitative velocity measurements from the raw images obtained. Chapter 7 presents further results, this time from the measurement of a free jet flow. In that chapter, the error analysis introduced in Chapter 5 is considered with respect to measured data, and a practical assessment of image processing performance undertaken. The final set of results, from the measurement of transonic flow in a convergent-divergent nozzle, are presented in Chapter 8. Finally, Chapter 9 presents conclusions and summarises the key points of the preceeding chapters.

The principal focus of the work described in this thesis has been the development of a suitable image processing strategy for application to the Doppler global velocimetry technique. This strategy has been analysed numerically in order to determine the effect of the choice of image processing algorithm. A number of experiments have been conducted, ranging from velocimeter measurements of a low velocity rotating disc, through to transonic nozzle measurements. In each case, the data obtained has been correlated with accepted models, or

with computational predictions. This correlation has enabled a comparison between estimated measurement error and predicted error to be made. In this way the performance of the Doppler global velocimeter may be assessed. It is the opinion of the author that such error prediction and analysis is an essential element in developing suitable image processing strategies for velocimetry applications.

Chapter 2

The Doppler Global Velocimetry Technique

In this chapter, the principles of the Doppler global velocimetry flow measurement technique will be outlined. It is not intended that this chapter should present an exhaustive review of Doppler global velocimetry technology. Rather, this material is briefly described in order that the reader may be sufficiently acquainted with the underlying technology to be able to understand the points made in subsequent chapters with regard to image processing techniques and error estimation fully. For a more comprehensive discussion of Doppler global velocimetry technology, the reader is referred to Thorpe [Thorpe 1996].

2.1 Overview

The basic principles of the Doppler global technique of whole field flow velocity determination are straightforward. The flow field to be measured is seeded with suitable particles and illuminated by means of a laser expanded into a planar light sheet spanning the region of interest. It is the seeding particles that are predominantly responsible for scattering the laser light. (It is the velocity of the seeding particles that is measured, rather than that of the flow itself, so it is important that the seeding material follows the flow faithfully). The scattered laser light is subject to a certain Doppler shift: the frequency of the light scattered by the seeding particles differs from that emitted by the laser by a certain value. The magnitude of this shift in frequency is straightforwardly related to the velocity of the scattering body.

The planar region of interest is imaged onto a video or photographic camera by means of an optical system. In order that the Doppler shift may be determined, an iodine vapour cell is used as a frequency-to-intensity converter. Iodine vapour possesses a number of absorption lines, some at appropriate frequencies for use in conjuction with commonly available lasers. Such absorption lines have finite bandwidth and exhibit an approximately linear variation in absorption with respect to frequency. Thus by setting a suitable laser emission frequency, the degree of transmission of the scattered light through an iodine vapour cell may be used to determine the Doppler shift present. An iodine vapour cell may be inserted into the Doppler global velocimetry optical system in such a way that it acts as a frequency-to-intensity discriminator for the entire frame.

In general, the intensity of scattered light will be not be uniform across the field of interest. This may be due to non-uniformities in the laser light sheet intensity or in the density of

seeding particles throughout the flow. Irregularities in the transmission of the optical system across the entire frame may also be present. In order to distinguish intensity variations resulting from the effect of Doppler shift on the transmission of the vapour cell, it is necessary to acquire a pair of images rather than just one. In principal, dividing the image resulting from the transmission of scattered light through the iodine cell (known as the *discriminated image*) by the image obtained without the cell (the *non-discriminated image*) will yield a representation of the cell transmission alone. The cell transmission value at each pixel location may be related to Doppler shift values and hence to velocity. It is essential that the pair of images be accurately aligned prior to division.

Although the principle of the Doppler global velocimeter is simple, development of an accurate instrument requires an understanding of laser physics (to ensure a stable and controllable emission frequency), absorption spectroscopy (for the construction of suitable iodine vapour cells), image acquisition and processing (to allow accurate alignment of images) and fluid dynamics (for successful flow-field seeding).

2.2 The Implementation of the Technique at Oxford

As explained above, it is necessary to acquire a pair of images simultaneously. The first, the discriminated image, is an image of the flow field taken through the vapour cell. The second, the non-discriminated image, is an image of the same scene, but without the vapour cell in the optical path. The optical arrangement suggested by Komine [Komine and Brosnan 1991, Komine *et al.* 1991], and used in most subsequent systems (such as those described in [Ford and Tatum 1994, Morrison, Gaharan and DeOtte 1994]) is shown in Fig. 2.1. This arrangement uses a primary lens to focus the image of the flow field. A transfer lens, sharing a common back-focal plane with the primary lens, is employed to collimate the beam. This makes the overall system less sensitive to path differences between the reference and signal optical paths, and allows the use of a rectangular aperture (or *field stop*) to limit the extent of the observed area. Subsequently, the collimated beam is split into two paths, one passing directly into a camera, the second passing first through a vapour cell and then into a second camera. Monochrome CCD video cameras are commonly employed for Doppler global applications.

It is necessary that the two images be accurately aligned with one another. Komine and others initially accomplished the necessary alignment by careful manual adjustment of the optical elements. More sophisticated approaches are possible, however, using image processing techniques, as will be seen in Chapter 3.

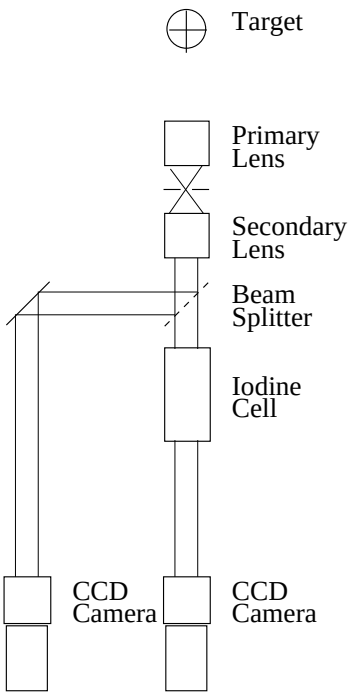


Figure 2.1: Komine's Doppler global velocimetry Apparatus.

In the Oxford system, the optical arrangement of Fig. 2.2 is employed. The flow, illuminated by a planar sheet of laser light, is focused by means of a primary lens. A secondary, or transfer, lens is employed to collimate the image into a parallel beam. A field stop is fitted at the focal plane between two lenses to limit the physical extent of the image. The collimated beam is separated by means of a non-polarisation-sensitive beam splitter, into two paths, one of which passes through an iodine absorption cell. The two beams are then imaged onto separate halves of a single CCD camera using a second non-polarisation-sensitive beam splitter.

The Oxford system is novel in its use of one camera, rather than a pair of cameras as proposed by other workers. The single camera approach avoids the necessity of synchronising of a pair of separate frame grabbers, or the requirement, as proposed by Meyers [Meyers and Komine 1991], for a special dual frame grabber. A similar image recording system employing a single CCD camera to acquire the discriminated and non-discriminated images simultaneously has subsequently been described by Chan *et al.* [Chan *et al.* 1995]. Another approach to the design of Doppler global velocimetry optical systems has been discussed by Meyers [Meyers 1995], who suggests the adoption of a bore-scope optical arrangement which does not incorporate collimating optics.

The emphasis on off-line image processing rather than real-time operation enables the use of rather more sophisticated image processing schemes than those proposed by other workers.

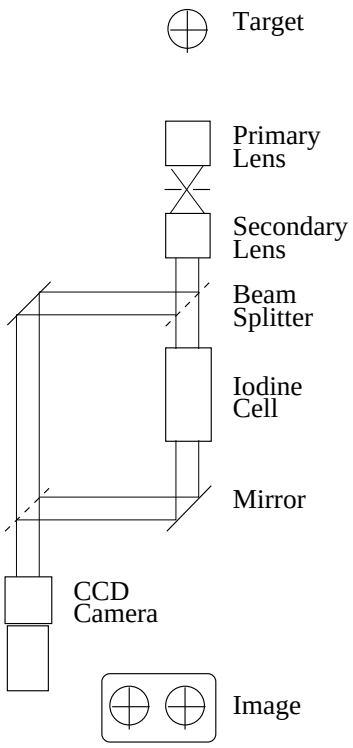


Figure 2.2: The Oxford Doppler global velocimetry optical system.

Image processing methods have been employed to ensure alignment between the discriminated and non-discriminated images. Based on the experimental results obtained for a free-expansion nozzle (Chapter 7), it is a central tenet of the current work that not only is it desirable to use sophisticated image processing in order to reduce error to the lowest possible value, but satisfactory results are not to be expected without the use of such a method, given the requirement for accurate alignment to sub-pixel accuracy for typical flow fields. This is borne out by the realisation of both Meyers [Meyers, Lee and Cavone 1991, Meyers 1992] and Usry [Usry, Meyers and Miller 1992 **(1)**, Usry, Meyers and Miller 1992 **(2)**] that software alignment is necessary. The fact that flow field measurements presented by workers using real-time systems appear to exhibit certain anomalies (for example, the free jet results presented by Komine [Komine and Brosnan 1991] show evidence of extreme misalignment) also supports this view.

The primary lens has a constant back-focal length independent of the object distance, and so forms a real image in the plane of the field stop. The use of a field stop is of particular importance for a single camera system as it limits the extent of the image formed such that the discriminated and non-discriminated images do not overlap on the camera CCD array. The transfer lens is mounted such that its back focal plane is coincident with that of the primary lens. The focus setting of the transfer lens is set to infinity, and hence a pseudo-collimated beam is produced: the light rays associated with a particular point in the object

plane are parallel after passing through the transfer lens, even though different positions in the object plane will produce rays with different angular deviations from the optical axis.

The choice of optical components was found to be significant both in terms of obtaining satisfactory image quality across the frame and also in terms of the ability to acquire images of sufficient intensity in the allowable exposure time. The primary and transfer lenses are derived from a medium format camera, whereas the CCD camera lens is from a 35 mm camera. The use of a photographic lens for the transfer lens, rather than a singlet as used in some previous systems, allows the physical size of the apparatus to be reduced. The nominal diameter of the light beam emerging from the transfer lens may be as large as 25 mm and therefore beam splitters and mirrors of 50 mm diameter are used, in order to minimise the loss of light in the system. The iodine cell windows are also of diameter 50 mm, as described in Section 2.5.

The arrangement of optical components in the Oxford Doppler global velocimeter is such that collimated light is obtained through the iodine cell. There are, however, certain drawbacks to this arrangement — in particular a variation in light intensity reaching different parts of the CCD frame is observed. This phenomenon is known as *vignetting*, and may be seen to be caused by a spillage of light within the optical system. Image processing techniques are employed (Chapter 3) in order to eliminate the effect of vignetting from the measured data, but the loss of intensity may still lead to a degradation of measurement accuracy in the darker areas. The focal lengths of the lenses employed have been chosen to minimise the effect of vignetting.

The primary lens has 60 mm focal length, and a maximum aperture of f/3.5. The transfer lens, a 150 mm f/4.0, is mounted on a sliding platform to allow for flexible focussing. In practice it has been found that the primary lens is best operated at an aperture of around f/5.6 to increase depth of field, but that the transfer lens must be fully open to avoid extensive vignetting. The CCD camera is fitted with a 300 mm f/4.5 lens, focussed at infinity, to increase the capture of the collimated rays into the camera. Non-polarisation sensitive beam splitters, specifically designed for use with the 514.5 nm argon ion laser have been fitted, together with matched coated mirrors (having in excess of 99% reflection at 514.5 nm): these are essential in the optical system since the polarisation state of the particle-scattered radiation is unknown and in general depends on the orientation of the illuminating radiation and the observed direction. All the components are mounted on a heavy optical rack to allow accurate and consistent alignment.

At present an argon ion laser is employed. This laser features an intra-cavity etalon whose temperature may be precisely controlled to ensure frequency stability of the laser. In practice, the laser frequency has been seen to exhibit a significant level of instability. A strategy has been devised to take account of this variation in emission frequency, as explained in Section 2.5.2. When using an iodine cell, accurate temperature control is vital to the accuracy of the technique, since small deviations in temperature markedly affect the transfer properties of the vapour. The iodine cell is also temperature controlled using a Peltier device.

In sum, the Oxford system differs from previous Doppler global velocimetry implementations in several key areas. The use of a single camera leads to a significant simplification in the electronics, and assures simultaneous capture of the image pair. A cooled CCD camera with 15 bit quantisation is employed. Although the choice of this camera precludes real-time operation, it may prove impossible to obtain accurate real-time results without sufficient regard for image processing. Sophisticated image processing techniques have been investigated and implemented, leading to a significant improvement in the quality of results obtained both for test images and for real flows.

The complete experimental apparatus used for Doppler global measurements at Oxford is shown in Fig. 2.3. In addition to the illuminating laser and the receiving optical system discussed above, a second iodine absorption filter, the *reference cell*, is also employed in order to measure the laser emission frequency, ν_o. Three computers are used as part of the experimental configuration: a personal computer to acquire images from the CCD camera, another personal computer for monitoring and recording laser emission frequency, and a UNIX workstation for image post-processing and data visualisation.

2.3 Measurement Procedure

Using the Oxford Doppler global velocimetry system, a single frame containing two images side by side is obtained. The so-called non-discriminated image represents that signal which has not passed through the iodine cell, whereas the discriminated image differs only in that the rays have passed through the iodine cell in order to convert light frequency into an intensity variation. These two parts must be *warped* into the same coordinate space in order to align them spatially so that they may then be divided to yield a representation of the cell transmission.

Optical systems are known to suffer from various types of aberration, as expounded for example by Jenkins and White [Jenkins and White 1976]. In particular it was found that a certain amount of rotation and skew of the images was present in the Oxford system.

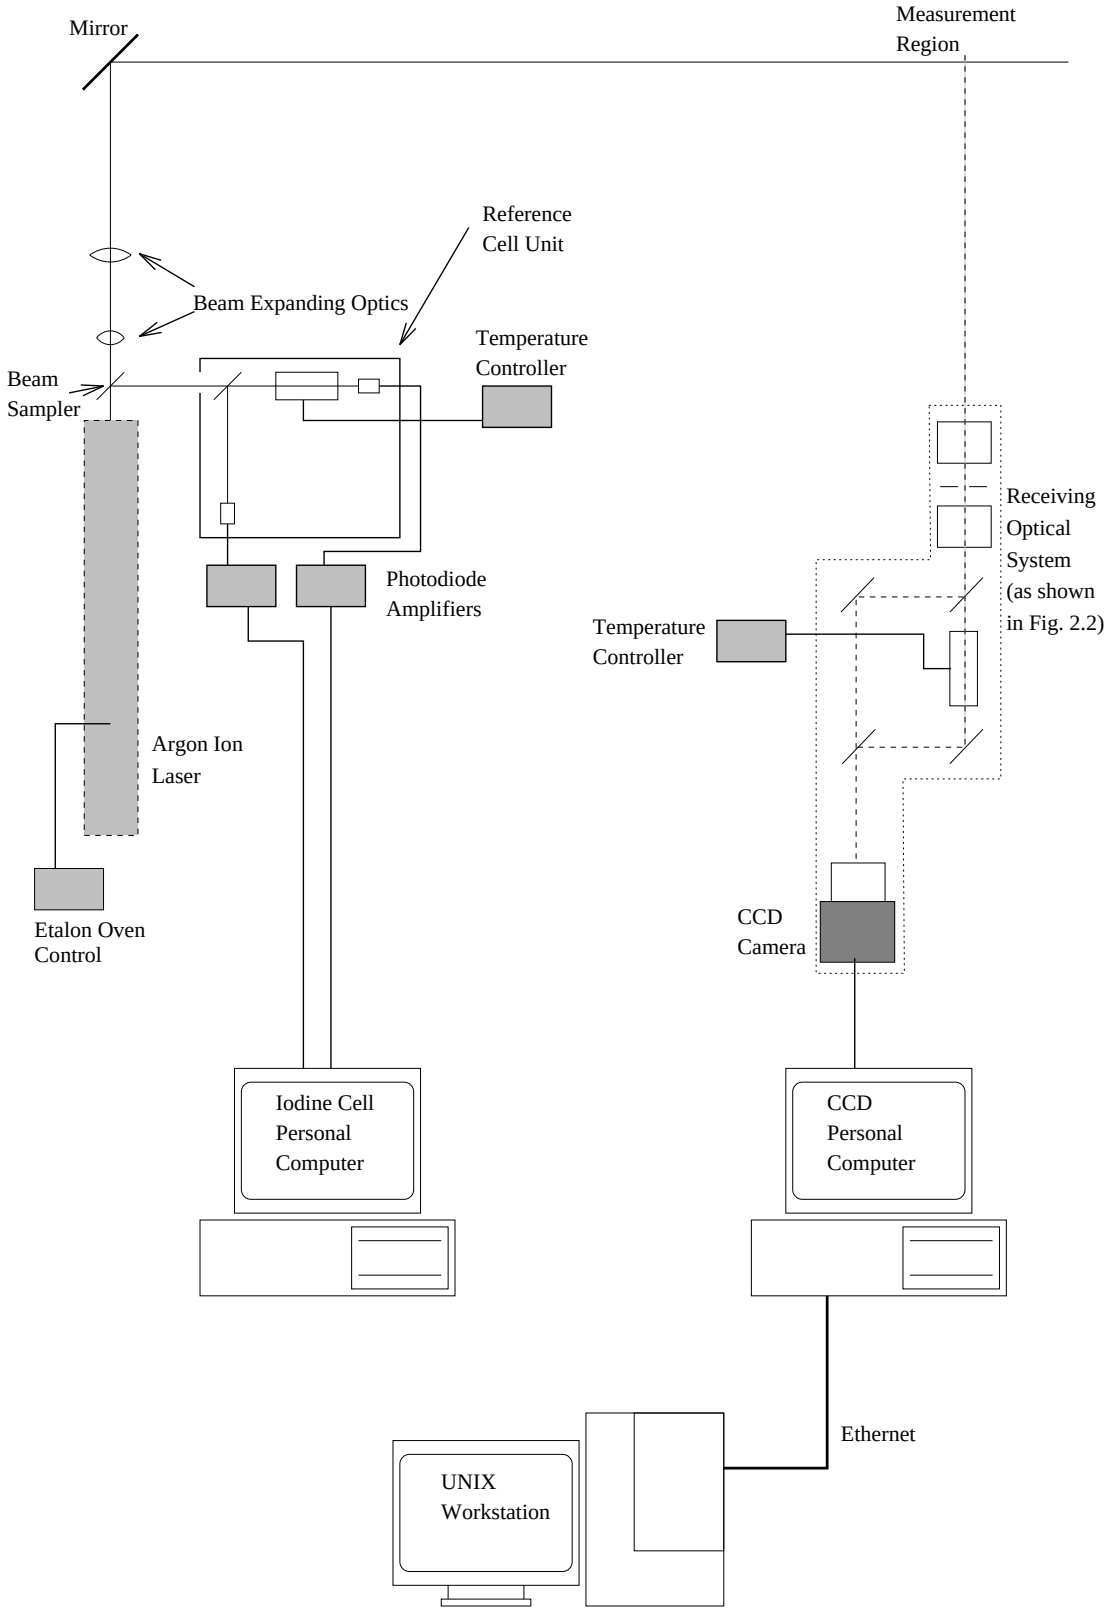


Figure 2.3: Overview of complete experimental apparatus.

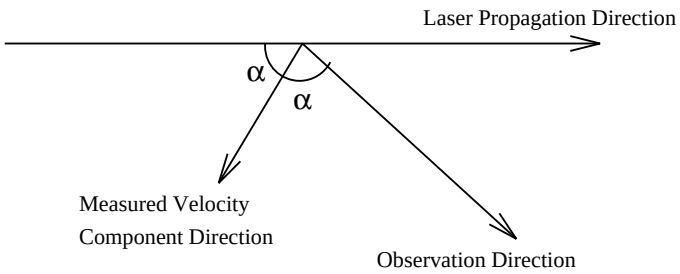


Figure 2.4: The measured component of velocity.

The average relative distortion was found to be approximately one pixel. This represents a significant source of misalignment which prevents the two images from being compared directly. Meyers [Meyers 1992], working with a video rate dual camera system, also concluded that software alignment was necessary, given the degree of physical misalignment present.

Consequently, image processing strategies designed to correct this relative distortion are a prerequisite to obtaining accurate measurements. The procedure for acquiring images and subsequently post-processing the data in order to effect accurate alignment is described in detail in Chapter 3.

Once the cell transmission has been determined across the entire CCD frame, additional processing is required in order to yield quantitative velocity values. The method of performing the required computation is discussed in Section 4.1.

In the following sections, each of the fundamental elements of the Doppler global system is briefly examined in turn. An understanding of the principles of each part of the system is necessary in order for the rationale for the design of the entire velocimeter to be apparent.

2.4 Doppler Shift

It is well known that light scattered from an illuminated moving object will exhibit a Doppler shift. The magnitude of this shift is given by the light propagation direction, $\hat{\imath}$, the observation direction, $\hat{o}$, the incident light frequency, ν_o, and the object velocity vector, $\vec{V}$ by eqn. 2.1 [Yea and Cummins 1964]

$$\Delta\nu = \nu_o \left(\frac{\hat{o} - \hat{\imath}}{c} \right) . \vec{V} \tag{2.1}$$

It may be seen from eqn. 2.1 that the Doppler shift is sensitive to one particular component of velocity, in the direction $\hat{o} - \hat{\imath}$, as shown in Fig. 2.4.

2.5 The Iodine Vapour Absorption Cell

Atomic and molecular vapour cells have been in use for a number of years in order to set frequency standards [Chartier, Robertsson and Fredin-Picard 1991] and also for laser emission frequency stabilisation. Rubidium, caesium, sodium and the other halogens have been widely reported in such applications. Development of similar cells for Doppler global velocimetry applications builds on this experience. The cells take the form of sealed glass chambers incorporating optical windows and containing a small quantity of the element with impurities such as air removed. The absorption properties of such a cell are primarily a function of the number density of the vapour phase absorbing species. For a cell containing solid and vapour phase in equilibrium (a *saturated* cell) the number density may be controlled through the vapour pressure associated with a particular operating temperature. The rapid change in vapour pressure of these elements as a function of temperature makes accurate temperature control essential for the stabilisation of the cell's absorption properties. For a frequency standard application Fredin-Picard [Fredin-Picard 1989] reported the use of a *cold finger* to stabilise a localised area of the cell at a temperature slightly below that of the rest of the cell. Other workers have considered the use of water baths [Chan *et al.* 1995] and heater tape [Morrison, Gaharan and DeOtte 1994] for iodine cell temperature control.

A *starved* iodine cell may be employed in place of the saturated cell used in the present work. Such a cell may be constructed by introducing less iodine during the filling process, so that there is insufficient iodine to allow the formation of a solid phase within the cell. Starved cells have the advantage that their absorption properties are less sensitive to variations in temperature. Practical difficulties in the construction of such cells led to the adoption of a saturated cell for the Oxford Doppler global velocimeter. Accurate control of the cold finger temperature is employed to ensure stable cell absorption properties.

Komine [Komine and Brosnan 1991] first proposed the use of a vapour cell as a frequency-to-intensity converter. Such a cell allows the Doppler global velocimetry apparatus to measure Doppler shift frequencies over a two-dimensional field. The diatomic iodine molecule, like other halogens, has many absorption lines within the visible spectrum (Gerstenkorn and Luc [Gerstenkorn and Luc 1978] list approximately 47,000 lines in the spectral range 490-670 nm), which are typically a few Giga-Hertz in width at room temperature. Some of these lines coincide with the frequency of emission of common lasers, such as the 514.5 nm emission of the argon-ion laser and the 532 nm emission of the frequency-doubled neodymium yttrium aluminium garnet (Nd:YAG) laser. Indeed, the centre frequency of one transmission system lies only 2 GHz above the centre frequency of the gain profile of the argon ion laser. A

saturated iodine cell is employed as a discriminator in the Oxford velocimeter. A typical absorption profile is illustrated schematically in Fig. 2.5.

The absorption features present in the iodine spectrum are of finite width, being approximately Gaussian in shape. The reason for this line width is broadening due to collisional, Doppler and hyperfine phenomena (for a discussion of this see [Chan *et al.* 1995]). The non-linear nature of the iodine cell absorption profile means that such cells require accurate calibration if precise velocity measurements are to be made. While early Doppler global velocimetry systems (for example that of Komine [Komine and Brosnan 1991]) assumed the side wings of the profile to be approximately linear, a more sophisticated calibration approach has been employed in Oxford, as explained in Section 4.1.

In addition to the requirement for accurate temperature control, there are a number of other hazards associated with the use of iodine cells. Previous workers have observed that the presence of contaminants in the cells may be seen to shift the position of the spectral lines. In addition, *outgassing* from the cell may lead to gradual drifts in the frequency of absorption features of the order of tens of kHz over periods of months to years. The iodine cell stability required for Doppler measurements of fluid velocity, where shifts may be of ther order of tens to hundreds of MHz, makes the problems observed by some of these workers of little direct relevance. Techniques of cell construction derived from the field of frequency standards may still be seen as good practice, however.

In order to construct a suitable iodine absorption cell, a number of constraints must be satisfied. The cell must be securely sealed, since the iodine vapour pressure is several torr at operating temperature and contamination by other gaseous species is to be avoided. The length of the cell must be carefully chosen to ensure adequate absorption. The cell in use at Oxford consists of a glass tube onto each end of which optical windows are fused. Parallel windows are employed so that the transmission of the cell is not polarisation sensitive. A glass side arm is fused to the cell to allow degassing of the cell prior to iodine fill. This side arm may also be used as a cold-finger, that is, a region in contact with the iodine vapour on its internal surface which may be easily temperature controlled. The coldest part of the cell in contact with the vapour phase determines where the solid phase accumulates, and thus sets the vapour pressure in the cell. The vapour pressure (and hence number density) of the gaseous species may thus be set by controlling the temperature of this cold-finger. A Peltier device is used in conjunction with a proportional-integral type controller to set the temperature of the cold finger, maintaining it a few degrees below room temperature. Temperature is monitored by means of a thermistor. The body of the cell is not temperature

controlled and assumes room temperature.

The optical system described in Section 2.2 produces a collimated light beam of approximately 25 mm diameter. Windows of 50 mm diameter were therefore used in order to ensure that a sufficient area of undistorted window was available. The cell length was chosen to be 60 mm in accordance with the published details of Komine [Komine and Brosnan 1991].

A strong temperature dependence is observed in such saturated iodine cells — vapour pressure increases by around 1 Pa/K (or around 20%/K). Such a strong sensitivity to temperature implies that the cell control system must offer excellent stability, providing deviations of less that 0.1 K from the set-point temperature [Ainsworth and Thorpe 1994], for velocity calibrations to be accurate to within 2 ms^{-1}. Assuming the cold finger is maintained 4 K below room temperature, it may be found [Thorpe 1996] that an increase in the vapour temperature of 1 K changes the transmission by 0.12%. The cell transmission is approximately 50 times more sensitive to variations in cold finger temperature than to changes in room temperature. The cold finger temperature controller in use at Oxford was observed to deviate from set-point by approximately ± 0.003 K over a 24 hour period, representing variations in transmission of around $\pm 0.02\%$.

Tuning the laser frequency to a wing of the absorption spectrum allows changes in light frequency caused by Doppler shift to be measured as changes in light intensity by passing the scattered light through an iodine vapour cell. A typical arrangement where the laser frequency is tuned to 50% transmission through the cell is shown in Fig. 2.5. An argon-ion laser operating at 514.5 nm may be tuned to coincide with this 50% point. In this way Doppler shifted light, scattered from sites in the flow field will be transmitted to a varying extent through the cell. In the configuration illustrated, positive Doppler shifts will result in decreased transmission, whereas negative Doppler shifts will result in increased transmission.

2.5.1 The Working Cell

As explained above, an iodine vapour cell is used as the frequency-to-intensity transducer for the Doppler global velocimeter. This cell, known as the *working cell*, is installed in the receiving optical system, shown in Fig. 2.2. In order that the working cell may be used to convert Doppler shift frequencies to changes in intensity in such a way that quantitative velocity data may be obtained, it is necessary to calibrate the absorption curve of the cell. The procedure for cell calibration and a discussion of the supporting software may be found in Section 4.1.

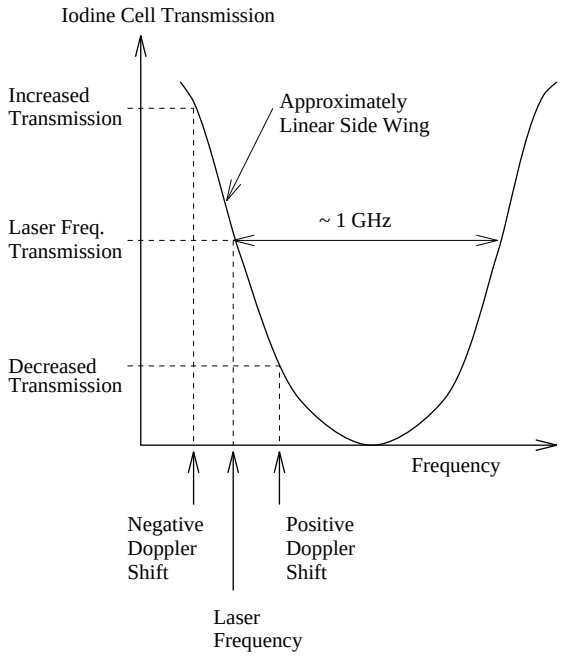


Figure 2.5: The iodine cell transmission profile.

2.5.2 The Reference Cell

In the Oxford Doppler global velocimetry system, a second iodine cell, the *reference cell*, is used. The reference cell is constructed and temperature controlled in the same manner as the main working cell described above. The reference cell arrangement, illustrated in Fig. 2.6, may be used to monitor the laser emission frequency, ν_o, while measurements are being taken. The reference cell is positioned close to the laser itself, and a beam sampler is employed in order to direct approximately ten percent of the laser emission into the reference cell enclosure. Within the enclosure, the beam is further divided using a fifty-percent beam splitter. Each of the resulting beams falls upon a photodiode, one of the beams having first been passed through the reference cell itself. The photodiode outputs are connected to a pair of custom-built amplifiers, and thence to an analogue-digital board installed in a personal computer.

In order to ensure good long term stability of the reference cell, and acceptable immunity to vibration and other disturbances, the reference cell enclosure is itself attached firmly to the laser mounting rail.

It is necessary to calibrate the absorption curve of the reference cell. Once calibrated, the reference cell transmission may be monitored by means of custom-written software, known as `IODINE`, running on the personal computer. Details of the calibration procedure and the means of determining the laser emission frequency, ν_o, may be found in Section 4.1.

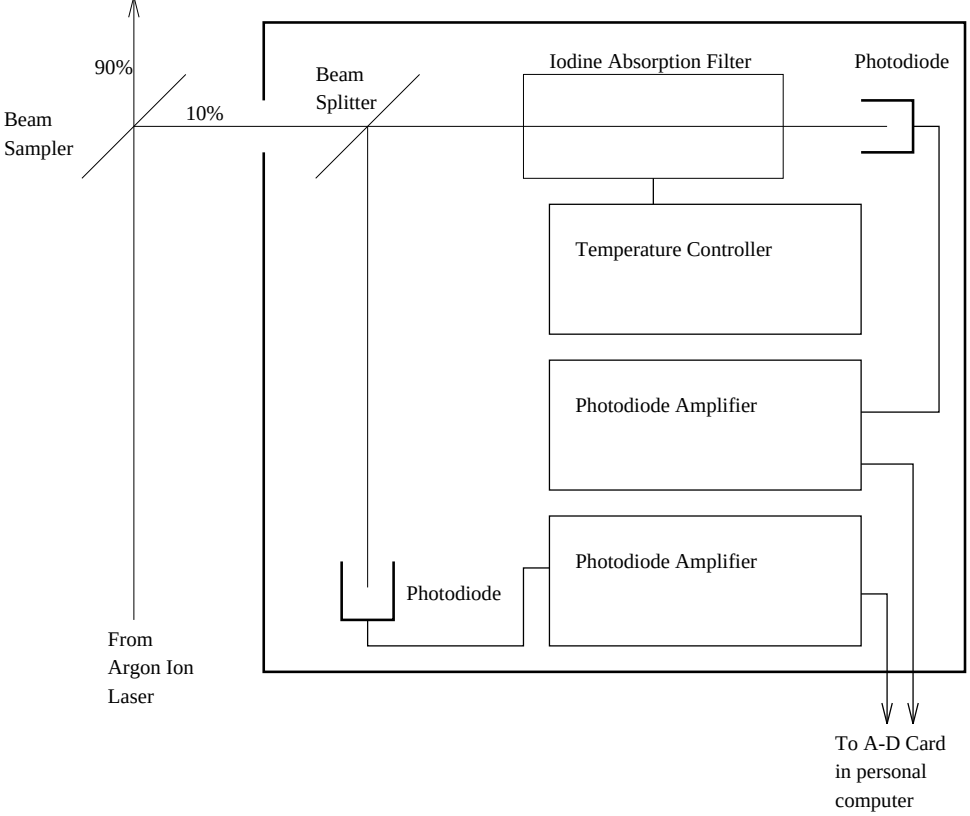


Figure 2.6: The Iodine Reference Cell.

2.6 The Argon Ion Laser as a Light Source

In the published literature, only two types of laser have been considered as suitable monochromatic light sources for Doppler global velocimetry. The use of argon ion lasers was reported in early velocimetry systems [Meyers and Komine 1991] [Komine and Brosnan 1991] [Meyers 1992], and also in some more recent implementations [Ainsworth and Thorpe 1994] [Roehle and Schodl 1994] [Morrison, Gaharan and DeOtte 1994] [Chan *et al.* 1995]. Other research has been based around the use of frequency-doubled neodymium yttrium aluminium garnet (or Nd:YAG) pulse lasers [Komine and Brosnan 1991, Komine *et al.* 1994]. An argon ion laser (Spectra-Physics model 165) is employed for the work discussed in this thesis. A discussion of laser physics is beyond the scope of this thesis. A summary of the principles of laser operation may be found in Thorpe [Thorpe 1996].

In order to illuminate an extended area of the flow field, the highly collimated laser beam is passed through a combination of cylindrical lenses in order to produce a planar sheet of light. Such a technique is commonplace in other flow-visualisation techniques, notably particle image velocimetry. In Doppler global velocimetry it is desirable to ensure that the direction of propagation of the light sheet is consistent across the entire sheet. This may be accomplished by using a combination of positive and negative lenses to expand the laser

beam into a sheet.

The most important characteristic of the argon ion laser in relation to Doppler global velocimetry is the fact that the emitted light frequency may be tuned to an absorption line of molecular iodine. The argon ion laser is a continuous wave laser typically producing several watts of power at predominantly 488 nm and 514.5 nm in its basic configuration. A prism is fitted within the body of the laser in use at Oxford in order that only 514.5 nm radiation is produced. Even so, single line operation of the argon ion laser typically has a line width of 0.01 nm, this being equivalent to 9 GHz bandwidth. The laser source for Doppler global velocimetry needs to have a narrow line width in comparison to the width of the iodine absorption feature, approximately 2 GHz, in order that the iodine cell may be used as a frequency discriminator.

Argon ion lasers may be operated in so-called *single-frequency* mode by the addition of a temperature controlled intra-cavity *etalon*. The etalon is an optical filter (actually a resonant optical cavity) that transmits a narrow range of frequencies. Placing such a device in the laser cavity allows only one longitudinal mode to resonate. The use of an etalon enables commercially available lasers to operate with extremely narrow line-widths of around 10 MHz, while producing power outputs in the region of 0.5 to 5 W. The laser in use in Oxford was set up for single frequency operation by carefully adjusting the optical components of the resonator and the etalon while monitoring the output frequency using a scanning Fabry-Perot interferometer.

The emission from the laser cavity is not continuous in terms of light frequency but made up of distinct *modes*. The mirror separation for the cavity determines a set of resonant light frequencies which are amplified. The separation of these resonant frequencies, or longitudinal modes, is given by Yariv [Yariv 1985]

$$\Delta\nu = \frac{c}{2L} \tag{2.2}$$

The presence of these discrete modes of laser emission at well-defined frequency intervals is exploited in order to perform calibration of the iodine cells (see Section 4.1). For the laser in use at Oxford, $L = 1$ metre, so the mode separation $\Delta\nu = 143$ MHz.

It is necessary to consider the stability of the laser emission frequency. The absolute light frequency is prone to drift due to both thermal instabilities in the laser construction and also due to vibration of the laser cavity mirrors induced by cooling water flows around

the plasma tube. The etalon in use in the Oxford system has a temperature stability of ±0.01 K at normal ambient temperature conditions, giving a frequency stability of ±20 MHz [Spectra-Physics 1987]. A 10 MHz line-width is attainable using the etalon, and, in this configuration, a laser output power of around 400 mW may be obtained.

Long term drifts in the laser emission frequency may be considered in terms of changes in the cavity length, L, due to the effect of thermal expansion. The resonator is constructed using quartz rods of low thermal expansivity but, even so, light frequency drifts of approximately 60 MHz/°C are evident [Spectra-Physics 1987]. Immediately after start-up the temperature of the resonator can change by several degrees, which is sufficient to cause a drift in frequency of several hundred MHz. The expansion of the resonator is made evident by the phenomenon of *mode-hopping*. The transmission characteristics of the etalon are such that only one longitudinal mode is allowed to oscillate, but as the cavity mirror separation changes, the frequency of that mode drifts until the etalon prevents its oscillation. At this point an adjacent mode, differing in frequency by $c/2L$ Hz, becomes active. This abrupt change in frequency is known as a *mode-hop*. After two hours of continuous operation, the resonator temperature has been found to stabilise. Once this steady state is attained, mode hops occur roughly every two hours. Since the longitudinal mode spacing of the laser is 143 MHz, this implies a drift of 1.2 MHz per minute.

Short term disturbances in emission frequency are caused by vibrations induced by the cooling water flows around the plasma tube. Such vibration was found to have a characteristic frequency of 200 Hz and a magnitude of 8 MHz RMS [Thorpe 1996].

Although other workers [Roehle and Schodl 1994] have discussed schemes for the precise stabilisation of laser emission frequency by means of an additional iodine cell (a similar implementation has also been made commercially available [Spectra-Physics 1995]), this is unnecessary for the current application. Rather, the use of a reference cell to monitor laser emission frequency during exposure offers an immunity to error due to laser frequency drift. In this way relatively low cost lasers may be employed for use with the Doppler global system.

The beam, of 1.25 mm diameter, possesses an approximately Gaussian intensity profile [Yariv 1985]. The Gaussian nature of the beam is preserved when it is expanded into a light sheet for Doppler global velocimetry applications. The flow field is thus non-uniformly illuminated, so the measurement images will have regions of high and low intensity irrespective of the seeding particle concentration in the flow.

A final aspect of the argon ion laser which must be considered is that of *speckle*. If a laser is incident on a surface such that light is scattered from that surface towards some observer, it may seem that the intensity of the scattered radiation is not uniform. To the observer it appears that the illuminated surface has a grainy pattern of light and dark spots. This effect is referred to as laser speckle. The cause of this behaviour is the nature of the coherent single frequency radiation produced by the laser. Such light, when scattered, tends to produce interference patterns. Laser speckle is often problematical in certain aspects of Doppler global velocimetry work. Strategies for overcoming the effect of speckle are discussed in Chapters 3, 5 and 6.

2.7 The CCD Camera

Whereas previous workers have concentrated on the real-time applications, using video cameras operating at 25 or 30 frames per second, the present system is designed for off-line processing. This allows the use of a Peltier-cooled, low noise, variable integration time CCD camera. Recording of the images is performed electrically rather than photographically, since the discriminated and non-discriminated image intensities need to be divided and additional image processing is required in order to achieve spatial correspondence between the two images.

As explained in Chapter 5, the noise present in the signal obtained from a CCD camera may be considered to be made up of three components: thermal noise, read-out noise and shot noise. Thermal noise may be markedly diminished by cooling the CCD array itself. Read-out noise is related to the frequency of read-out of the charges on the array — a slow read-out rate may be employed to reduce the level of such noise. Shot noise is statistical in nature, and is fundamental to the principle of operation of the CCD device. Such shot noise may not be reduced in magnitude by any such means.

While a full analysis of the impact of noise introduced by the camera may be found in Chapter 5, it is appropriate to briefly discuss the source of such noise in this section. The pixels of a CCD array store charge in proportion to the number of photons they have received. The CCD chip contains an amplifier that converts this charge into an output voltage. This voltage is then digitised to provide a grey level as the final numerical output from each pixel. However, in the absence of light reaching the CCD, charge still accumulates in the pixels due to the leakage of electrons from the wafer upon which the CCD is constructed. This phenomenon is usually called *dark current*, or thermal noise, and is extremely temperature sensitive [EEV 1987]. For video rate cameras operating at 25 Hz frame rate, the exposure time is of the order of 20 ms, and is usually short enough for the accumulation of dark current

to be disregarded. If longer integration times are to be used, however, the dark current needs to be reduced.

The levels of dark current present in a CCD chip may be minimised by reducing the operating temperature of the CCD chip. Both Peltier and liquid nitrogen cooled cameras are widely available. Cooling the chip allows long integration periods to be used and this makes it possible for images of particle laden flow fields to be acquired. If no cooling were used, the CCD would soon saturate as a result of the accumulation of dark current.

The high read-out rate implicit in the use of video rate cameras has implications for the resolution of the analogue-digital converters used to digitise pixel voltages, and also has a deleterious effect on the noise present in the pixel data. Komine *et al.* [Komine *et al.* 1994] have considered the level of noise present in the pixel values obtained using various video cameras, concluding that levels of noise of around 0.8% are attainable under optimum conditions. Commercially available video rate frame grabbers generally incorporate 8 bit analogue-digital converters. This resolution limits the ultimate accuracy of the video based systems.

A 25 Hz frame rate means that pixel values are produced at approximately 10 MHz, so precludes the use of high resolution frame grabbers. The read-out rate influences a second source of noise in the digitised pixel values: read-out noise. This noise is associated with the conversion of the pixel charge into a voltage by means of a transistor incorporated into the CCD chip [Astromed 1992]. For video rate systems, Komine *et al.* [Komine *et al.* 1994] have found this noise level to be approximately 1%.

The read-out noise may be minimised by employing a slow read-out rate from the camera, allowing a longer time period for the conversion of charge into voltage.

Furthermore, the use of dual video rate cameras for Doppler global applications requires the construction of custom-built frame grabbers and synchronisation circuitry, as reported by Meyers [Meyers and Komine 1991, Meyers, Lee and Cavone 1991]. This is necessitated by the difficulty in synchronising and and grabbing images from two video cameras. A single camera system has been employed in order to avoid this extra complication.

The use of a slow-scan cooled camera (Wright Instruments Ltd. [Wright 1993]) with high resolution analogue-digital conversion offers the benefit that thermal and read-out noise are reduced to such low levels that they may be disregarded in comparison with the fundamental shot noise present. A far better signal to noise ratio than that which may be obtained using conventional video cameras is attained, together with a greater sensitivity to low intensity scattered light. The CCD camera interfaces with its own frame grabber, which is installed

in a PC. The camera controlling software was supplied by the manufacturers of the system. A 512 by 512 pixel frame may be digitised to a depth of 15 bits, rather than the 8 bit quantisation normally employed. The signal to noise ratio of conventional video cameras is insufficient to justify the use of greater than 8 bit resolution. This increase in the number of quantisation levels available with the Peltier cooled camera should allow greater velocity measurement accuracy, especially where image intensity is low. The camera in use in Oxford is cooled by means of a four stage Peltier system to 200 K. The full-well capacity of the pixels is specified as 640,000 electrons. The read-out rate may be set to either 30 or 250 pixels/s, giving an estimated read out noise of 9 electrons or 39 electrons respectively. Wright Instruments claim the dark current specification of the camera to be 0.25 electrons/pixel/s under typical conditions.

An investigation of the CCD properties of the Wright Instruments camera undertaken by Thorpe [Thorpe 1996] showed a measured RMS uncertainty in pixel intensity of 0.2%. The shot noise was estimated to account for 0.15% RMS uncertainty. In general, the measured noise was found to be approximately 25% higher than the estimated shot noise and to follow the pattern theoretically predicted for shot noise alone. In assessing the rôle of camera noise in Doppler global measurement uncertainty, the assumption may be made that that the pixel grey level accuracy is limited by shot noise.

The pixels forming the CCD array are not identical in terms of their photo-response, as a result of manufacturing inperfections. The pixels thus have different sensitivites to incident light. Thorpe found variations in response of up to $\pm 2\%$ along a typical pixel row. Such variations need to be compensated for in the image processing procedure.

2.8 Flow Field Seeding

Optical anemometry techniques use scattered light in various ways in order to determine flow velocity. It is a common requirement of all such techniques that sufficient light be scattered so as to enable reliable measurement by an appropriate detector. Historically, particles have been introduced into the flow field to enable visualisation. Such particles must be sufficiently small to follow the fluid velocity faithfully, and yet large enough to scatter the light power required for detection.

The advent of laser Doppler anemometry in the 1960s gave rise to a closer inspection of the performance of seeding particles in terms of their flow following characteristics. The laser Doppler anemometry technique, in common with the particle image velocimetry technique, measures the velocity of individual seeding particles. Thus a relatively sparse concentration

of particles is required for such methods. Commercially available seeding generators for gas flow are typically developed for laser Doppler anemometry applications. It has been argued that such devices may not be suited to Doppler global velocimetry, which requires significantly greater seeding concentrations in order to register sufficient light in each CCD pixel [Komine *et al.* 1994].

The choice of seeding particles for a particular flow field is influenced both by the capacity of those particles for following the accelerations within the flow and by their light scattering ability. These two considerations can present conflicting particle diameter requirements, and hence a compromise is ultimately required. Seeding of air flows has been reported using a number of different materials, for example: olive oil, water/glycerine mixture, pine pollen, powdered milk, latex and smoke. Liquid materials require atomisation into appropriately sized particles. For example, in the Laskin atomiser proposed by Echols and Young [Echols and Young 1963], high speed air jets are injected into a reservoir of the liquid. The shearing action of the jets creates the particles which float to the surface of the liquid in air bubbles. It is important that the rate of evaporation of the material is sufficiently low that the particles, once formed, do not change size or disappear completely. Solid seeding particles are commercially available in various diameters: for example, latex particles, in suspension in water or isopropanol, are available in the diameter range 0.3 to 2.7 μm. Such particles may be deployed by atomisation of the emulsion containing the particles themselves. Both solid and liquid particles tend to accumulate in low flow velocity regions and can contaminate optical windows, making optical measurements through such windows difficult.

The accuracy of laser anemometry techniques is ultimately limited by the flow-following characteristics of the seeding particles employed. Consequently, the physical properties of such particles require careful consideration. Ideally the seeding particles would have the same density as the fluid but a different refractive index, thereby offering excellent flow-following characteristics and light scattering efficiency. A consideration of the behaviour of such seeding particles may be made according to the analysis presented in Schlichting [Schlichting 1979]. A number of other publications have presented considerations of the flow-following characteristics of seeding particles. For example, Melling and Whitelaw [Melling and Whitelaw 1973] were able to conclude that upper limits on particle size of 2.6 μm and 0.8 μm were found to be able to follow turbulence levels of 1 kHz and 10 kHz respectively to within 1%. For particle image velocimetry applications to supersonic flows, Bryanston-Cross and Epstein [Bryanston-Cross and Epstein] concluded that particles of less than 0.5 μm were necessary in order to follow the sudden fluid accelerations experienced while traversing an oblique shock. For the case of turbulent subsonic flows, Melling [Melling 1986]

suggested that an upper limit on diameter of 1 μm should be adopted. For Doppler global velocimetry measurements, Komine [Komine and Brosnan 1991, Komine *et al.* 1991] reports the use of olive oil seeding, atomised in a Laskin nozzle, to produce seeding of a mean diameter of around 1.5 μm, for application to a near sonic free jet. Usry, Meyers and Miller [Usry, Meyers and Miller 1992 **(1)**, Usry, Meyers and Miller 1992 **(2)**] report the use of an evaporation/condensation seeding generator producing approximately 0.7 μm diameter particles for a Doppler global velocimetry measurement of a 70 ms^{-1} flow over a delta wing. For a point-wise measurement, of a 100 ms^{-1} free jet flow, Roehle and Schodl [Roehle and Schodl 1994] report the use of a Laskin nozzle to produce approximately 1 μm diameter particles.

The other aspect of the choice of seeding particles which must be taken into account is their light-scattering properties. The mathematical theory describing the scattering of radiation by small spherical particles was first described by Mie [Mie 1908], such phenomena subsequently being termed *Mie scattering*. This scattering regime predominates when the particle diameter is greater than the wavelength of the scattered light. For the case of the 514.5 nm emission of the argon ion laser, it may be said that scattering from particles in excess of around 0.5 μm diameter may be described by Mie theory. In brief, the scattered light power from a particular particle is dependent on the particle diameter, the relative refractive index of the particle material compared to the fluid, the wavelength of the illumination, the relative orientations of the incident light ray and the scattered ray, and the properties of the detection optics. Scattering exhibits a complex directional pattern, with peak scattering efficiency in the forwards scatter direction. Relatively high scattering efficiency is observed in the backwards direction and several local minima are found in certain off-axis directions. Calculations have been undertaken by Thorpe [Thorpe 1996], using the commercial Mie scattering code `STREU` [Invent GmbH 1991], for a number of particle sizes. The number of local scattering minima was found to vary with particle size, with just one minimum for the case of 0.5 μm particles, 50^o from the illumination axis. In Doppler global applications, typical observation directions range from 55^0 to 135^o from the illumination direction. Within this range, a 2 μm particle was found to scatter approximately an order of magnitude more than a 1 μm particle. The spatial variation in scattered intensity implies that a careful choice of Doppler global observation direction is essential if adequate intensity is to be obtained.

Thorpe [Thorpe 1996], performing the so-called light balance calculations of Bryanston-Cross and Epstein [Bryanston-Cross and Epstein] was able to conclude that full-well conditions could be obtained for the CCD camera used in the Oxford system in around ten seconds for the case of 1 μm glycerine/water particles observed in the 90^o position for the free jet flow

experiment described in Chapter 7. In practice, somewhat longer exposures were required. It is interesting to consider the case of the video rate systems where the typical integration period is around 20 ms. Such an exposure constraint would require a significantly higher seeding density, a much higher laser emission power, or the use of larger particles. A high density of seeding particles is to be avoided in order to minimise the problem of *multiple scattering*, whereby light scattered from a seeding particle is incident on a second particle and is rescattered with an additional Doppler shift. Significant levels of multiple scatter can lead to velocity determination error.

Commercially-available seeding generators are generally intended for use in laser Doppler anemometry applications. The higher seeding particle concentration required for successful Doppler global velocimetry measurements necessitated the construction of a high output seeding generator, capable of supplying seed particles above atmospheric pressure. One of the simpler designs reported, the Laskin atomiser [Laskin 1948], is reported to produce 1 μm particles from a glycerine/water mixture [Echols and Young 1963]. A comparison of such a generator with a commercially available generator [DANTEC 1984] revealed that the Laskin atomiser was capable of producing 55 times as many particles. A Laskin atomiser for use with glycerine/water mixture was constructed. This atomiser was used for the free jet measurements discussed in Chapter 7. For transonic applications, such as those presented in Chapter 8, the 1 μm glycerine/water particles were found to be unable to follow the flow accurately. Therefore, for this application, 0.5 μm latex spheres were introduced into the test flow by means of a commercially-available seeder.

2.9 Two Component Measurement

It is clear from eqn. 2.1 that the magnitude of the Doppler shift observed by the velocimeter is dependent not only on the flow velocity vector $\vec{V}$, but also on the angle of observation $\hat{o}$, and illumination $\hat{\imath}$. Hence, by choosing different values of $\hat{o}$ or $\hat{\imath}$, the velocimeter may be made sensitive to velocity vectors in different directions. This aspect of the Doppler shift equation may be exploited in order to perform measurements of flow where either two or three velocity components are resolved.

Several possible configurations for two component measurements will be considered. First, take the situation where two Doppler global velocimetry optical systems are used, as shown in Fig. 2.7. In this case both optical systems are mounted in the same plane. There are then two observation vectors $\hat{o}_1$ and $\hat{o}_2$ making angles θ_1 and θ_2 respectively with respect to the illumination vector $\hat{\imath}$. Defining the axes such that

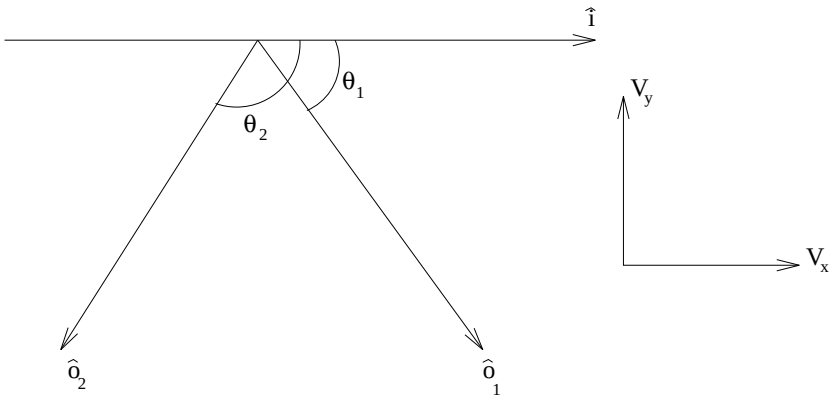


Figure 2.7: Two optical systems used for two component measurement

$$\hat{\imath} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \tag{2.3}$$

the true velocity vector in two components may be expressed as

$$\vec{V} = \begin{pmatrix} v_x \\ v_y \end{pmatrix} \tag{2.4}$$

The observation unit vectors may be written thus

$$\hat{o}_1 = \begin{pmatrix} \cos\theta_1 \\ -\sin\theta_1 \end{pmatrix} \tag{2.5}$$

$$\hat{o}_2 = \begin{pmatrix} \cos\theta_2 \\ -\sin\theta_2 \end{pmatrix} \tag{2.6}$$

Using eqn. 2.1 and denoting the Doppler shift of the first imaging system by $\Delta\nu_1$ and of the second by $\Delta\nu_2$

$$\frac{c\Delta\nu_1}{\nu_o} = \begin{pmatrix} \cos\theta_1 - 1 \\ -\sin\theta_1 \end{pmatrix} . \begin{pmatrix} v_x \\ v_y \end{pmatrix} = (\cos\theta_1 - 1)v_x - (sin\theta_1)v_y \tag{2.7}$$

$$\frac{c\Delta\nu_2}{\nu_o} = \begin{pmatrix} \cos\theta_2 - 1 \\ -\sin\theta_2 \end{pmatrix} . \begin{pmatrix} v_x \\ v_y \end{pmatrix} = (\cos\theta_2 - 1)v_x - (sin\theta_2)v_y \tag{2.8}$$

Solving eqns. 2.7 and 2.8 simultaneously in order to obtain expressions for the two orthogonal velocity components v_x and v_y yields

$$v_x = \frac{c(\Delta\nu_1 \sin\theta_2 - \Delta\nu_2 \sin\theta_1)}{\nu_o(\sin\theta_1 - \sin\theta_2 + \cos\theta_1 \sin\theta_2 - \sin\theta_1 \cos\theta_2)} \tag{2.9}$$

$$v_y = \frac{c(\Delta\nu_1(\cos\theta_2 - 1) - \Delta\nu_2(\cos\theta_1 - 1))}{\nu_o(\sin\theta_1 - \sin\theta_2 + \cos\theta_1\sin\theta_2 - \sin\theta_1\cos\theta_2)} \tag{2.10}$$

It has thus been seen that by mounting two Doppler global velocimetry optical systems such that the illumination vector and the two observation vectors lie in the same plane, and by setting the systems to observe the same region of the flow field from different angles, measurements may be obtained for two velocity components — one velocity component in the plane of the laser light sheet and the other perpendicular to it. The Doppler shift equation, eqn. 2.1, may be used to compute the components as required.

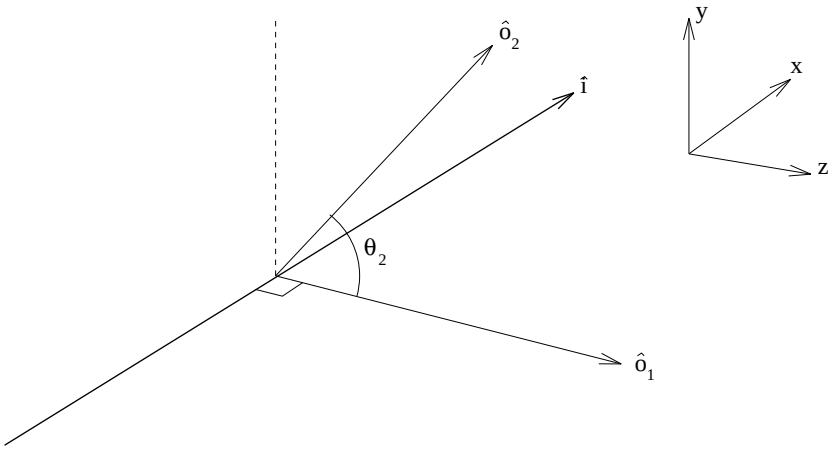


Figure 2.8: Another configuration of two optical systems used for two component measurement

Clearly it is possible to orientate the optical systems differently, with one system mounted out of the plane formed by the illumination vector and the other observation vector. In certain cases, it is possible to resolve two velocity components in the plane of the laser light sheet with such a configuration. This arrangement, where one optical system is positioned so that its optical axis is at 90° to the light sheet while the second is mounted directly above it, observing the same position in the light sheet, with its optical axis inclined by θ_2 is shown in Fig. 2.8.

Now three components must be considered in the vector calculations. Defining the axis such that the illumination vector is aligned with the x direction

$$\hat{\imath} = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \tag{2.11}$$

the velocity vector may be written

$$\vec{V} = \begin{pmatrix} v_x \\ v_y \\ v_z \end{pmatrix} \tag{2.12}$$

The observation directions may be stated thus

$$\hat{o}_1 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \tag{2.13}$$

$$\hat{o}_2 = \begin{pmatrix} 0 \\ \sin\theta_2 \\ \cos\theta_2 \end{pmatrix} \tag{2.14}$$

Applying eqn. 2.1 once again

$$\frac{c\Delta\nu_1}{\nu_o} = \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix} . \begin{pmatrix} v_x \\ v_y \\ v_z \end{pmatrix} = v_x - v_z \tag{2.15}$$

$$\frac{c\Delta\nu_2}{\nu_o} = \begin{pmatrix} 1 \\ -\sin\theta_2 \\ -\cos\theta_2 \end{pmatrix} . \begin{pmatrix} v_x \\ v_y \\ v_z \end{pmatrix} = v_x - v_y \sin\theta_2 - v_z \cos\theta_2 \tag{2.16}$$

There are now two equations for three unknowns (v_x, v_y, v_z) and a solution seems impossible. It is, however, often the case that no out-of-plane velocity vectors are present. This is the case for the spinning disc measurement of Chapter 6 and also for the free jet measurement of Chapter 7, when the laser light sheet is positioned along the axis of the jet. For such cases, it is possible to specify that $v_z = 0$ and hence

$$v_x = \frac{c\Delta\nu_1}{\nu_o} \tag{2.17}$$

$$v_y = \frac{c(\Delta\nu_1 - \Delta\nu_2)}{\nu_o \sin\theta_2} \tag{2.18}$$

Two schemes have been considered — the first allowing the measurement of one component of velocity in the plane of the laser light sheet and one component perpendicular to the light sheet, and the second allowing the measurement of two components within the light sheet. Clearly it is possible to implement many such configurations, determining the desired

components of velocity in the manner outlined above. The choice of observation angles will depend on a number of factors, including the practicality of mounting the optical systems in the desired locations and the accuracy required of each velocity component determined.

Using multiple optical systems with a single light sheet is a convenient means of obtaining two component measurements. In order for such measurements to be made, it is necessary to ensure that the shutters of the two CCD cameras employed are synchronised so that each optical system will begin its exposure at the same instant. In the case of the Oxford system, a connector is available on the CCD camera for this purpose.

In Chapter 3, computational schemes are detailed that enable the discriminated and non-discriminated images to be spatially aligned with one another. When performing two component measurements, it is necessary to ensure that the images obtained in the two CCD cameras of the two optical systems are aligned in a similar manner. As explained in Chapter 3, the algorithm chosen for image transformation can be readily extended for application to two component measurements. Section 4.2 describes the software which has been implemented to perform the velocity component calculations and also to allow visualisation of the two component velocity data so obtained. A discussion of the measurement procedure for acquiring two component measurements may also be found in that section. A consideration of the expected error in two component velocity measurements may be found in Section 5.11.

While the use of multiple optical systems makes it possible to acquire simultaneously velocity data which may be processed in order to measure two velocity components, such complexity and expense may not be justified in certain cases. There are a great many situations in the laboratory where a test facility allows repeatable flow conditions to be obtained. In this case, there is no reason why it is necessary to use two optical systems and acquire the data for two component measurements at once. Rather, it may be possible to set the system up in one configuration, obtain measurements, and then reconfigure the system to a new configuration, obtaining a new set of measured data. In this way, for the case of repeatable flow conditions, two component measurements may be obtained using a single optical system.

Repeatable flows allow a different approach to two component measurements to be pursued. Just as the component of velocity to which the Doppler global velocimetry system is sensitive may be varied by means of changing the observation direction $\hat{o}$, so it may also be changed by varying the direction of illumination $\hat{\imath}$. Thus, provided the laser light sheet remains in the same plane, it may be possible to reorientate the direction of propagation within that light sheet by means of mirrors, in order to obtain two component measurements.

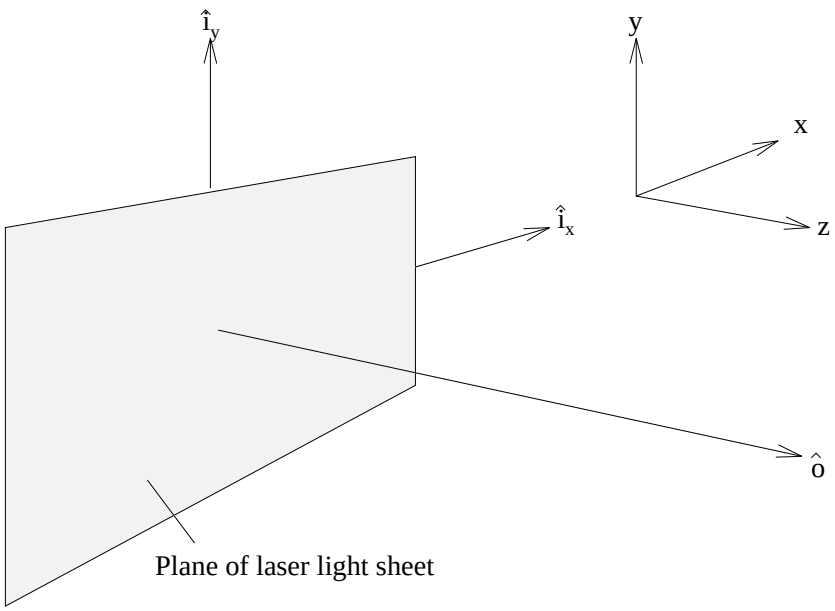


Figure 2.9: Changing the orientation of the light sheet in order to obtain two component measurements

As an example, consider the case shown in Fig. 2.9. Here, one observation direction $\hat{o}$ is used, perpendicular to the plane of the light sheet. The light sheet, while being maintained in the same plane, may be set up so that the direction of propagation of the laser emission is either in the x direction ($\hat{\imath}_1 = \hat{\imath}_x$) or in the y direction ($\hat{\imath}_2 = \hat{\imath}_y$). Now

$$\hat{\imath}_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \tag{2.19}$$

$$\hat{\imath}_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \tag{2.20}$$

$$\hat{o} = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \tag{2.21}$$

So

$$\frac{c\Delta\nu_1}{\nu_o} = \begin{pmatrix} -1 \\ 0 \\ 1 \end{pmatrix} . \begin{pmatrix} v_x \\ v_y \\ v_z \end{pmatrix} = v_3 - v_1 \tag{2.22}$$

$$\frac{c\Delta\nu_2}{\nu_o} = \begin{pmatrix} 0 \\ -1 \\ 1 \end{pmatrix} . \begin{pmatrix} v_x \\ v_y \\ v_z \end{pmatrix} = v_3 - v_2 \tag{2.23}$$

Once again, if there are no out-of-plane velocity values expected, $v_3 = 0$ and,

$$v_1 = -\frac{c\Delta\nu_1}{\nu_o} \tag{2.24}$$

$$v_2 = -\frac{c\Delta\nu_2}{\nu_o} \tag{2.25}$$

If simultaneous measurement of two components is not required, this latter approach of varying $\hat{\imath}$ has a number of attractions. It is easier to obtain a significant angular variation in $\hat{\imath}$ than in $\hat{o}$, and hence less velocity error is expected in the orthogonal components obtained by such a method. In addition, varying the vector $\hat{\imath}$ does not necessitate the alignment of two separate optical systems. Unfortunately, such an approach may not be applied if the flow regime under consideration is not repeatable. The technique of changing the laser sheet propagation direction has been applied to the case of the rotating disc. The results of this experiment may be found in Chapter 6.

The method of using multiple optical systems in order to obtain two component measurements may be readily extended by the addition of a third optical system. In this way both components of velocity within the plane of the laser light sheet and also the out-of-plane component may be simultaneously measured. Such an experiment has not yet been conducted, although it is worthy of note that the image processing techniques described in this thesis are readily applicable to such a case.

2.10 Summary

In this chapter the fundamentals of the technology underlying the Doppler global velocimeter have been introduced and an overview of the physical configuration of the Oxford system has been presented. This material has been introduced so that subsequent chapters in this thesis relating to my work in applying image processing techniques (Chapter 3), post-processing for quantitative velocity determination (Section 4.1) and post-processing for two component measurements (Section 4.2) may be readily understood.

It is beyond the scope of this thesis to include an exhaustive review of published works relating to each aspect of the Doppler global technique. For a more thorough review of Doppler global technology, the reader is referred to Thorpe [Thorpe 1996].

With an understanding of the physical basis for the Doppler global technique, and an appreciation of the extent of previous work which has been undertaken in the development of video frame-rate velocimeters, the importance of accurate alignment of the discriminated

and non-discriminated frames is evident. The next chapter discusses the theory behind such image processing operations, and describes the computer code that has been written to achieve image alignment. The real-time systems implemented by other workers have been unable to offer fully-calibrated quantitative velocity results. This may be achieved by undertaking post-processing to take account of the profile of the iodine absorption curve and hence determine velocity values, as discussed in Chapter 4.

The introduction of effective data processing techniques for Doppler global velocimetry has been assisted by the development of a comprehensive error analysis scheme, as described in Chapter 5. This scheme may be used to determine the expected measurement accuracy and allows assessment of both physical and data processing performance. Accordingly, the presentation of measured results in Chapters 6, 7 and 8 is augmented by a consideration of expected and estimated measurement error. In place of the somewhat *ad hoc* approach of earlier workers, numerical and analytical techniques for error assessment serve as a cornerstone of the work discussed in this thesis.

Chapter 3

Image Processing Techniques

As described in Section 2.3, the Doppler global velocimetry technique of flow measurement requires a pair of images to be acquired. In the case of the Oxford system, the pair of images are captured on a single frame. One of the images, known as the *discriminated image*, is formed as a result of viewing the flow through the iodine absorption cell. The other image, the *non-discriminated image*, is identical except for the fact that the iodine absorption cell is not present in the optical path. In principle, division of the discriminated image by the non-discriminated image should yield a representation of the absorption of the iodine across the entire flow field under consideration.

In order to obtain accurate quantitative results, however, it is necessary to carefully align the discriminated and non-discriminated images prior to division. Image processing techniques have been extensively applied, as part of the current work, in order to effect this alignment. In this chapter the procedure which has been adopted for producing a *results image* representing the transmission of the signal through the iodine absorption cell for all points within the field of interest is described. A number of exposures must be made in order to compute the alignment transformations and also to compensate for non-uniform illumination, seeding and light transmission through the optics. The procedure for perfoming Doppler global velocimetry measurements with the Oxford system is described, and the purpose of each of the individual exposures discussed. The numerical and computational procedures which are subsequently employed in order to process each of these CCD camera exposures and produce a results image are presented. A detailed assessment of the performance of the image processing regime presented in this chapter may be found in Chapter 5. In addition, the experimental results presented in Chapters 6 to 8 are examined with close attention to the choice of image processing routine. In this way the effectiveness of the strategy outlined in this chapter has been assessed both theoretically and practically.

The generation of the results image represents a significant part of the process of determining velocity data from the CCD frames obtained. The process by which the results image may be employed to compute quantitative velocity data is discussed in Section 4.1.

3.1 Overview of Image Processing Techniques

Early publications of Meyers [Meyers, Lee and Cavone 1991] [Meyers 1992] and also of Usry [Usry, Meyers and Miller 1992 **(1)**] [Usry, Meyers and Miller 1992 **(2)**] concluded that it is essential that the discriminated and non-discriminated images are accurately aligned with one another prior to division. Large errors in the measured velocity may result from relatively small errors in the alignment of the two images. Those early studies attempted neither to quantify the magnitude of the expected velocity error, nor to investigate the requirements for image alignment. Image alignment error is discussed in detail in Chapter 5 of this thesis, where a numerical model for expected velocity error in the presence of image misalignment is developed. The error analysis calculations that have been performed, and also practical experience with the Oxford velocimeter, suggest that for typical flow field conditions it is necessary to undertake alignment to *sub-pixel* levels of accuracy. That is, the misalignment throughout the frame should be significantly less than the width of a single pixel in magnitude. Such an alignment criterion suggests that the manual alignment methods employed in early systems are inadequate and that a comprehensive image processing strategy for Doppler global velocimetry is necessary.

Throughout this thesis a distinction has been made between the terms *frame* and *image*. In this context a frame refers to the data resulting from a single exposure of the CCD camera employed in the Doppler global velocimetry system. Each frame contains a pair of images: the discriminated and non-discriminated images.

As a consequence of aberrations in the optical system such as astigmatism and imperfect alignment of optical components, the two images may differ in a number of ways. One image may be slightly rotated relative to the other, may be of a slightly different magnification relative to the other, or may be distorted in a more complex way due to the behaviour of the lenses in the system. Previous workers, such as Komine [Komine and Brosnan 1991] have described a technique whereby the ratio of the two images is continuously computed and displayed while the optical components of the system are adjusted. By imaging a characteristic test pattern and adjusting camera and lens positions until the features of the pattern are seen to be eradicated from the ratio of the images, coarse alignment of the images is possible.

Meyers [Meyers 1992] described a similar optical system to that proposed by Komine and Brosnan [Komine and Brosnan 1991] and discovered that the differences in the two optical paths were such that serious geometrical distortion was present in the two images. He therefore recognised the need to use image processing algorithms, and suggested the use of geomet-

rical image transformations, known as *warps*, in order to rectify the problem. The extensive computational overheads associated with carrying out such numerical procedures meant that the Meyers system was unable to provide accurate real-time results. Similarly, other workers who use real-time video systems, for example Komine *et al.* [Komine and Brosnan 1991, Komine *et al.* 1991] and Ford and Tatum [Ford and Tatum 1994] have not reported the use of any such image processing. Suitable image processing algorithms are described in detail by, for example, Wolberg [Wolberg 1991], and have been principally developed for creating animated sequences for motion pictures.

The detailed error analysis presented in Chapter 5 makes a quantitative examination of all the sources of error expected in the Doppler global velocimetry system. A significant part of this error analysis concerns the problem of image alignment. That analysis not only confirms the findings of earlier workers that alignment is a critical factor in obtaining accurate quantitative velocity measurements, but also provides a framework within which the different image processing algorithms presented in this chapter may be assessed. In addition, the experimental results presented in Chapters 6 to 8 are subjected to rigorous error analysis, allowing the importance of alignment and the choice of image processing technique to be assessed in practice.

It is clear then that the accuracy with which the images must be aligned, in order to yield accurate quantitative velocity measurements, is such that a sophisticated image processing approach is called for. To this end, image processing software has been developed in the present programme of work which exploits image transformation techniques, otherwise known as warps, in order to effect accurate alignment of the image pair. This image processing software, `warp`, has been written in the programming language **C++** [Stroustrup 1991] to run on a UNIX workstation, under the X-window [Nye 1990, Reiss and Radin 1988] system.

There follows a step by step discussion of the image processing strategy implemented in order to effect accurate alignment. The reader may find it helpful to refer to the images presented in Chapter 6 which illustrate the intermediate stages of the image processing procedure, as implemented in the `warp` software, followed when aligning the Doppler global velocimetry system in order to measure the velocity of a rotating disc.

The image processing method may be split into two distinct parts. When first setting up the Doppler global velocimetry system in a new position, it is necessary to compute the warp coefficients. Subsequently, provided that neither the position of the system nor its optical configuration is altered, these warp coefficients may be employed in order to transform the pair of images so as to effect accurate alignment. It is only necessary to compute the

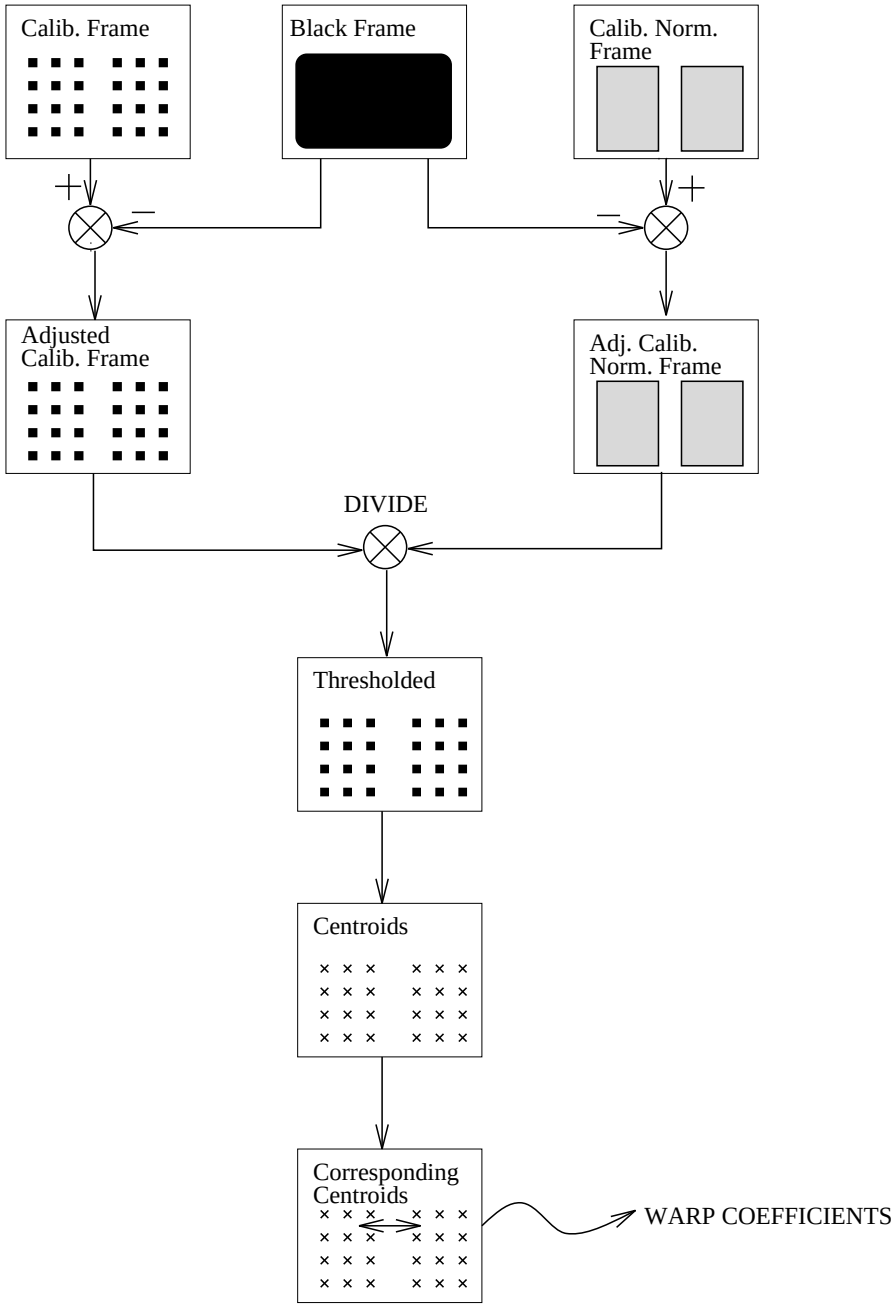


Figure 3.1: Flow diagram showing the procedure for the computation of the warp coefficients.

coefficients once when setting up the system in a new configuration or location.

The computation of the warp coefficients proceeds as shown in Fig. 3.1. An alignment card consisting of a regularly spaced array of white squares on a black background is used. In order to create such an alignment card, a design produced using a graphics software program is printed using a laser printer. The resulting image is then fixed to a metal plate, in order to ensure that it remains flat. The spacing and size of these dots has been found to have a significant effect on the accuracy of the alignment obtained from the image processing technique — this aspect will be discussed in Chapter 5. The alignment card is placed in the intended plane of measurement and the system focussed on it. It has been seen that adjusting the aperture of the lenses in the optical system can affect the degree of distortion of the images. In addition it has been seen that the colour of the illuminating light also has some effect, due to *chromatic abberations* in the lenses employed. For these reasons, the alignment card is imaged at the intended working aperture, under argon-ion laser illumination. The resulting data is known as the *calibration frame*.

The computation of the warp coefficients requires the accurate estimation of the positions of the centroids of the squares in the calibration frame. It is then necessary to find corresponding pairs of centroids in the discriminated and non-discriminated images. These pairs of centroids may then be used in order to compute the warp coefficients by the methods outlined in Section 3.3 below.

In order to locate the centroids accurately, the effects of pixel sensitivity non-uniformity, vignetting (typically a drop in intensity at the edges of the frame), and non-uniform illumination must be taken into account. A *calibration normalisation frame* is thus required. This is simply a frame aquired by means of viewing a white target, positioned in the same plane as the calibration target, and similarly illuminated.

In addition, it has been seen in Section 2.7 that a dark current effect is present in the CCD array. This means that even in the absence of any light reaching the CCD array, a certain signal level will be produced. To compensate for this, and also for the non-zero intensity of the black of the alignment target, an image of a black target, positioned and illuminated as before, is acquired. This frame is known as the *black frame*. The calibration frame, calibration normalisation frame and the black frame should all be acquired with the same exposure and the same CCD camera gain setting.

The black frame is subtracted from both the calibration frame and the calibration normalisation frame. The adjusted calibration frame must be divided by the adjusted calibration

normalisation frame to compensate for intensity non-uniformities.

As suggested by Meyers [Meyers 1992] Sobel edge detection may be employed in order to highlight the dots and cancel any 'DC offset'. In the work undertaken in Oxford, the use of such edge-detection was discontinued since it proved to be computationally costly, unnecessary, and led to emphasis of noise, especially the laser speckle pattern. In place of edge-detection, the dots are highlighted by the use of amplitude threshold segmentation [Pratt 1991] after the edges of the image outside the calibration region have been masked off. Although Meyers found this approach unsatisfactory, it has proved to work well at Oxford when a calibration normalisation image is used in order to remove intensity non-uniformities prior to centroid estimation.

The image, with dots highlighted by amplitude segmentation, is subjected to a novel algorithm which attempts to reject any lines in the horizontal and vertical directions and remove high frequency noise while retaining the dots themselves. Such horizontal and vertical lines may be introduced into the data at the edges of the frame due to the presence of the field stop in the Doppler global velocimetry optical system. High frequency noise components are sometimes present in cases where the threshold for amplitude segmentation is set to an intensity value close to that of the intensity of the dots in the image — in such a case CCD shot noise may manifest itself as spurious features on the thresholded image. The use of an additional algorithm to reject lines and high frequency noise has been seen to significantly enhance the robustness of the image processing scheme.

This algorithm is applied independently to both the discriminated and the non-discriminated images. Initially each column of pixels is considered. Iterating along the columns, pixel values are considered in a binary sense: either they are greater in intensity than a user-specified threshold or less than that threshold. In the context of white squares on a black background, the presence of high intensity values provides evidence for the existence of a feature, while low intensity provides evidence of the absence of a feature. Some *a priori* knowledge of the expected size of features in terms of pixels is extremely useful in this application. This information may be provided by the user of the software. The algorithm considers the number of consecutive pixels of either high or low intensity. A small number of consecutive pixels of high value is typically a noisy dot in the image which should be discarded before any further processing is undertaken. A small number of pixels of low value typically represent a gap in a feature, due once again to measurement error, which should be considered to be 'filled in'. Very large numbers of consecutive high intensity pixels are often the result of a line appearing in the image, perhaps representing the edge of the field stop. These should

once again be disregarded. The algorithm is then applied throughout the rows of the image, iterating across the columns. The intersection of the remaining features in each column and the remaining features in each row is then computed. By this means, only those regions of high pixel intensity which are likely to represent the squares on the alignment target are left. Pixel noise and lines in the image are rejected.

Vertical and horizontal dividing lines are then superimposed on the image, through the spaces in the rectangular array of features, to isolate each feature in a *bounding box*. Each of these bounding boxes may then be subjected to centroid estimation in order to find the coordinates of the centre of the feature so contained. The advantage of this technique is that centroid estimation of the array of dots may be performed entirely automatically, without the need for user intervention. The use of the calibration normalisation frame and the black frame ensures that the regions surrounding the features themselves have a pixel value close to zero, and hence these regions do not unduly distort the positions of the detected centroids. The centroid analysis technique is evaluated in Chapter 5. For segmentation and edge-detection methods, Pratt [Pratt 1991] has been found to be a useful reference.

Unfortunately, the automated centroid location technique becomes more difficult to perform successfully if the images are rotated away from the vertical orientation. For this case, an editing facility has been implemented which allows the user to remove any noise or spurious features, or to add features which the automatic image processing algorithm had failed to recognise. This allows the bounding box technique to be used for most calibration images, even in the case of extreme viewing angles or poorly aligned optical systems.

Having found the centroids of the dots in the calibration frame, the correspondence between points in the discriminated and non-discriminated parts must be established. Again, the intention has been to ensure that this process requires the minimum of user intervention. The method of establishing correspondence depends on the type of warp to be used. Both the *polynomial warp pseudo-inverse solution* and the *polynomial warp with orthogonal polynomials* (Section 3.3.3) may take any number of corresponding point pairs greater than or equal to six. The *bilinear warp* (Section 3.3.2) may only take four corresponding point pairs.

In the interests of ensuring that image processing is automated, the pixel offset between the signal and reference images is determined by an iterative routine. This algorithm supposes an initial offset and subtracts one image from the other, finding the average absolute pixel value difference. This process is repeated with different offset values, effectively moving one image relative to the other in a spiral of increasing size. After a certain limit, the absolute average pixel differences of each candidate offset are compared, and the lowest value selected as the

correct offset. This method may be seen as a means of performing a limited cross-correlation to determine offset.

Given an offset from one image to the other, the software subsequently iterates through the points in the non-discriminated image, finding the closest centroid in the discriminated image when the offset has been taken into account. A maximum error distance may be specified to avoid centroids being wrongly matched in the case that a point is missing from the discriminated image. For the bilinear warp, the user must manually specify four pairs of points; these should be selected so that they form the corners of the region of interest. Once the correspondence has been established, the warp coefficients may then be computed. The three warping techniques considered in the current work are described in Section 3.3 below.

It would be possible to effect a warp so that the discriminated image was transformed to align with the non-discriminated image. The `warp` application, however, does not follow this approach. Rather, both the discriminated and non-discriminated images are warped to the same regular array of control points. This has the effect of removing perspective effects when the optical system is not positioned orthogonal to the plane of measurement. In addition, as described in Section 2.9, when using multiple optical systems to enable measurement of more than a single velocity component it is important that the images from the separate optical systems are aligned to a common set of coordinates. The approach employed here readily lends itself to the measurement of two or more velocity components, as expounded in Section 4.2.

Having computed the warp coefficients as described above, they may subsequently be used to align data images. Unless the system is reconfigured or otherwise disturbed, the same set of warp coefficients may be employed repeatedly. Indeed, the mechanical stability of the optical system is such that the same set of warp coefficients has been successfully employed over a period of several months. Fig. 3.2 shows the procedure for warping the actual data images.

Just as a calibration normalisation frame was found to be necessary in order to compensate for intensity non-uniformities in the calibration frame, a *normalisation frame* is required for the *data frames* too. In the same way, another *black frame* must be acquired in order to combat the effect of the black level which is present. The procedure, however, is slightly different to that described previously. Once again, the normalisation frame, black frame and data frame must be acquired with the same exposure and CCD camera gain setting. In this case, however, it is more important to ensure that the laser frequency is the same for each of the three images. If this is not the case, a severe distortion of the measured velocity

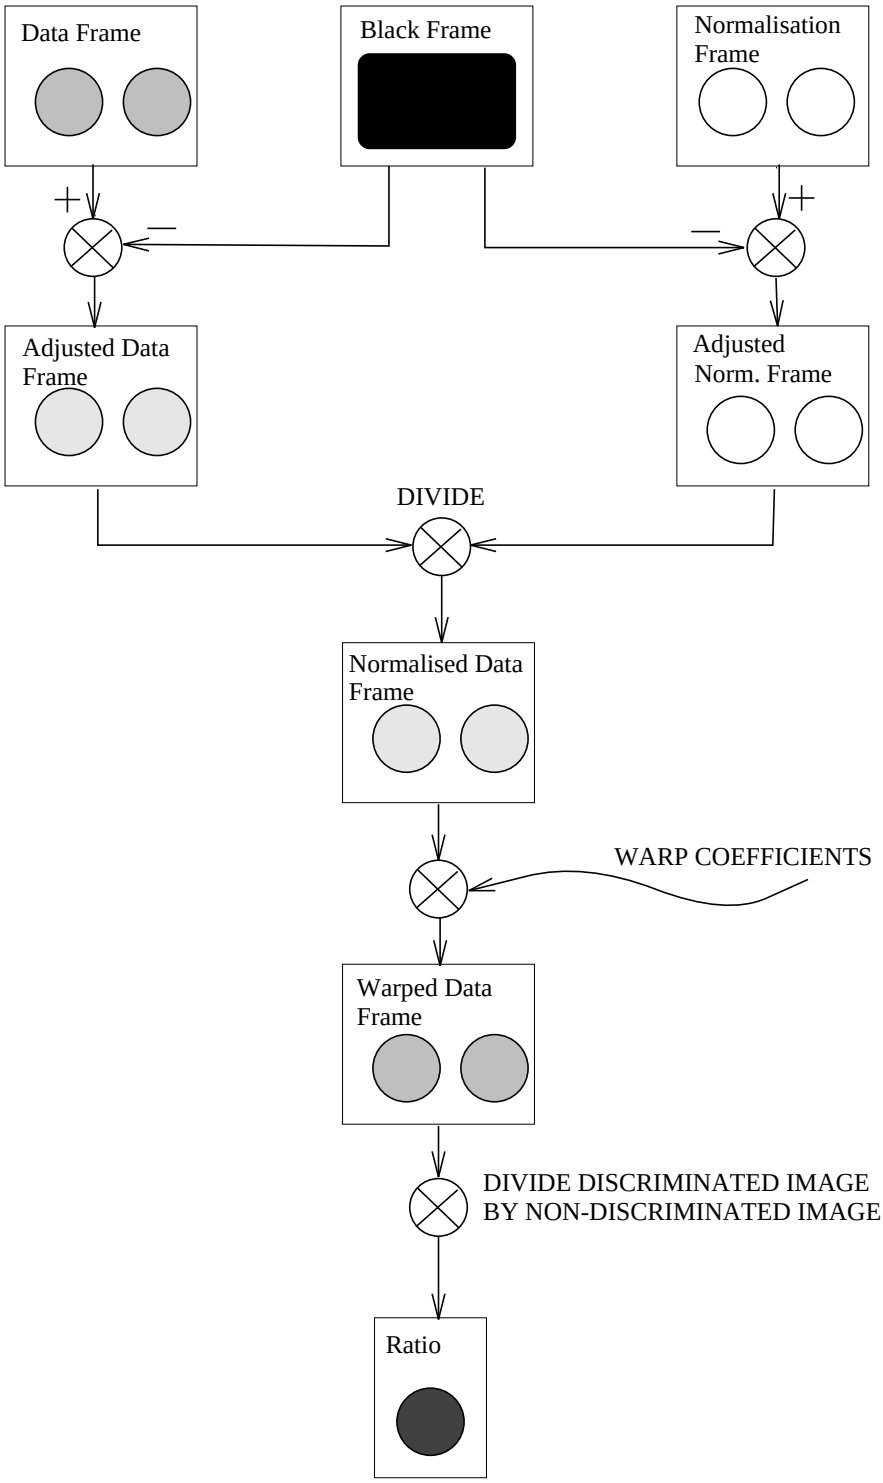


Figure 3.2: Flow diagram showing the procedure for the computation of the measured flow field, the warp coefficients having already been determined.

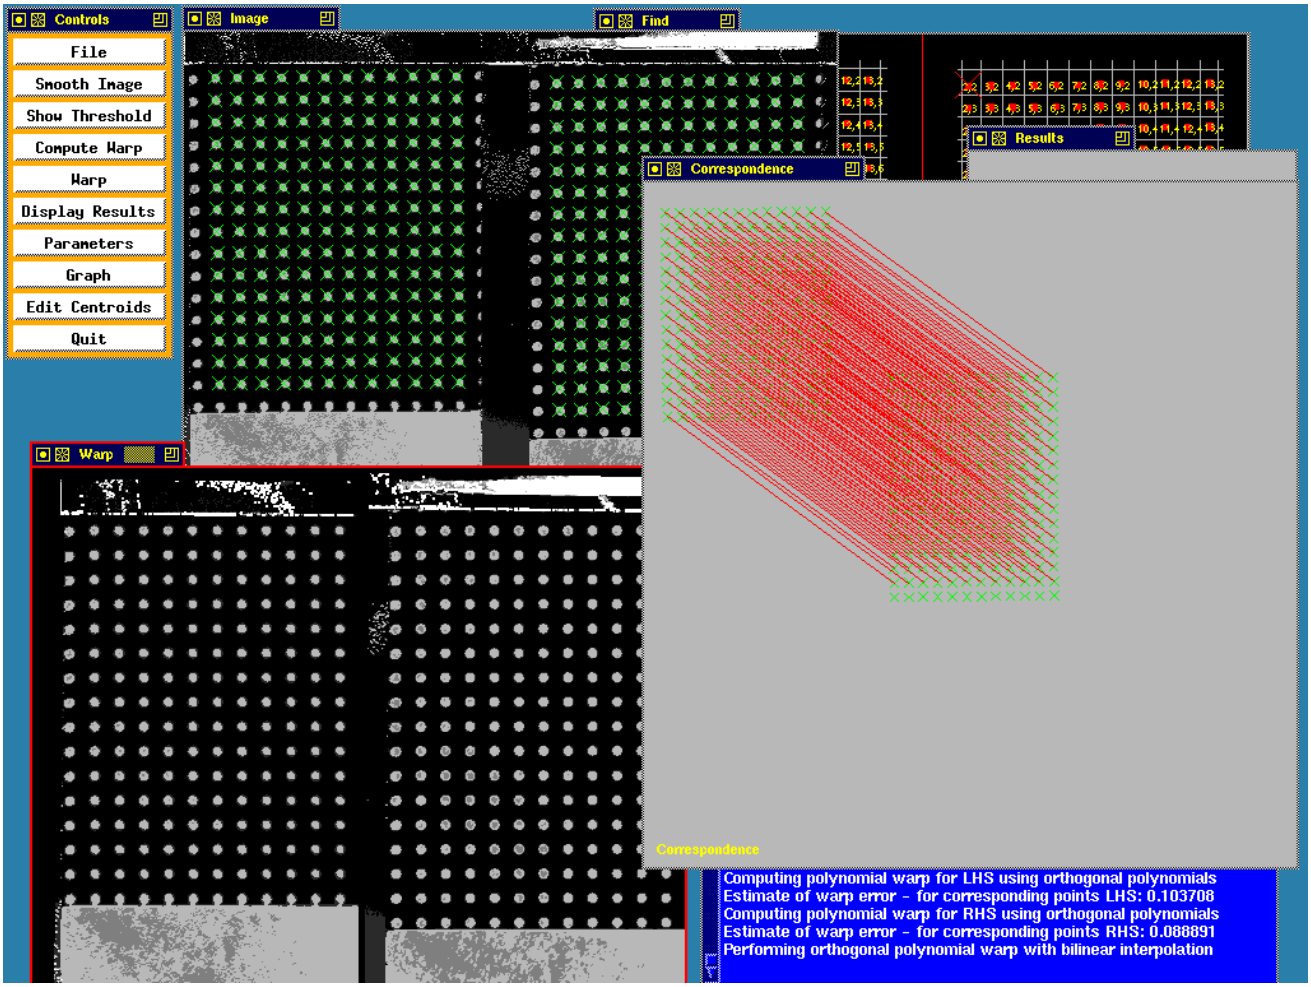


Figure 3.3: Typical `warp` screen display.

is possible. The program `norm_adjust`, described in Section 4.1, may be employed when necessary to adjust the normalisation and black frames in order to remove the effect of a slight drift in laser frequency.

The normalisation frame is generally acquired as the image of a slowly rotating PTFE disc of uniform intensity, placed in the measurement plane. PTFE was found to have suitably uniform scattering characteristics. The rotation of the disc is required to reduce the effect of laser speckle, described in Section 2.6, which may otherwise mar the uniformity of the frame. The rate of rotation should be slow, however, so as not to introduce any significant Doppler shift of its own into the discriminated image. The arrangement for illuminating the normalisation disc is explained for the case of the rotating disc measurement in Chapter 6. The black frame is simply an image of the working section, acquired in the absence of any laser illumination. This black frame is used to compensate for the effect of CCD array dark current effects (Section 2.7) and stray illumination.

The black frame is subtracted from both the data frame and the normalisation frame. The adjusted data frame is then divided by the adjusted normalisation frame. Warping is then performed using the coefficients determined as outlined above — the method of implementing warps is described in Section 3.3 below. The resulting warped image may then be displayed.

As already described, the discriminated and non-discriminated images are warped to a regular coordinate system. The warp produces a mapping of the integer coordinates of the pixels in the warped images to non-integer coordinate values in the original unwarped images. Clearly, an interpolation scheme is necessary in order to determine pixel values at each point in the image. Four such schemes, *nearest neighbour*, *bilinear*, *cubic* and *exponential*, are considered in Section 3.4.

Finally, having performed the warping and the interpolation necessary to align the images, the discriminated image is divided by the non-discriminated image to generate a representation of the absorption of the iodine cell across the frame. To exclude spurious noisy data from the low intensity region surrounding the field of interest, the results image ratio is set to unity for those pixels where the intensity of the non-discriminated image is below a certain user-specified threshold. The techniques of Section 4.1 may then be applied to obtain absolute velocity data.

In Chapter 6, the acquisition of the CCD frames described above is shown for the case of actual Doppler global velocimetry data (the measurement of a single velocity component of a rotating disc). In that chapter, the image processing technique is executed one step at a

time and the resulting display of the `warp` application is illustrated at each stage.

The processing strategy employed has the potential to improve significantly the accuracy with which Doppler global measurements may be obtained. However, the accuracy to which alignment of the discriminated and non-discriminated images may be performed is limited. There are three reasons for this:

- There is a fundamental limit on the accuracy to which the position of the alignment features may be determined.
- The warping technique is unable to model certain geometrical distortions which may be present in the image.
- The interpolation procedure employed may introduce further error into the determined velocity.

The contribution of each of these effects is assessed, and the expected accuracy of alignment extensively analysed in Chapter 5.

3.2 Centroid Estimation

As outlined in Section 3.1, centroid detection is performed for each computed bounding box. The calibration normalisation image is subtracted from the calibration image in order to reduce the value of the area surrounding each dot to a value close to zero.

Due to the presence of CCD camera shot noise, variations in illumination intensity and CCD camera shutter timing the area of the bounding box surrounding the feature will have some small non-zero intensity value. Since the location of the centroid is not known *a priori*, it is not possible to ensure that the bounding box is symmetrical about the centroid of the feature so contained. In order that the non-zero intensity outside of the feature does not introduce undue error into the estimate of centroid position, a threshold may be specified. Any values less than this threshold may then be set to zero. As discussed in Chapter 5, the value of this threshold must be carefully chosen in order to minimise error.

For a bounding box spanning x_L to x_H in the x-direction and y_L to y_H in the y-direction, centroid detection then proceeds in the conventional manner

$$Q = \sum_{i=x_L}^{x_H} \sum_{j=y_L}^{y_H} m(i,j), \tag{3.1}$$

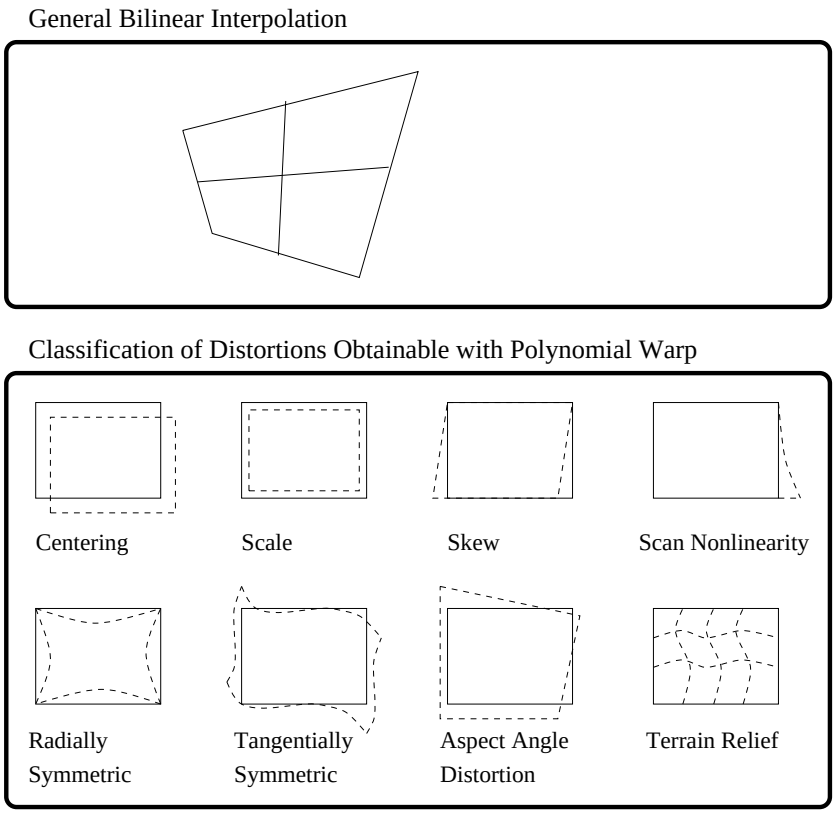


Figure 3.4: Illustration of the capabilities of bilinear and polynomial warps.

$$x_c = \frac{1}{Q} \sum_{j=x_L}^{x_H} \sum_{i=y_L}^{y_H} (m(i,j)x), \tag{3.2}$$

$$y_c = \frac{1}{Q} \sum_{j=x_L}^{x_H} \sum_{i=y_L}^{y_H} (m(i,j)y), \tag{3.3}$$

where $m(i,j)$ is the 'mass' at each point — in other words $m(i,j)$ represents the pixel intensity at position (i,j).

3.3 Image Transformation Techniques

3.3.1 Theoretical Limitations

Two forms of warping have been implemented — *bilinear* and *polynomial*. The polynomial warp used in the present work is of degree two. By reference to Fig. 3.4 it may be seen that the polynomial warp is able to represent a more general class of non-linear image distortions than the simple bilinear warp. If non-linear distortion such as astigmatism is present in the optics, polynomial techniques are therefore expected to yield more accurate results.

Wolberg [Wolberg 1991] provides a detailed discussion of warping algorithms.

3.3.2 Bilinear Warp

Given any four points, (u_0, v_0), (u_1, v_1), (u_2, v_2) and (u_3, v_3), and their respective function values x_0, x_1, x_2 and x_3, any intermediate coordinate $X(u, v)$ may be computed by the expression

$$X(u, v) = a_0 + a_1 u + a_2 v + a_3 uv, \tag{3.4}$$

where the a_i coefficients are computed by solving

$$\begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 1 & u_0 & v_0 & u_0 v_0 \\ 1 & u_1 & v_1 & u_1 v_1 \\ 1 & u_2 & v_2 & u_2 v_2 \\ 1 & u_3 & v_3 & u_3 v_3 \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ a_2 \\ a_3 \end{bmatrix} \tag{3.5}$$

Eqn. 3.5 may be written

$$\vec{x} = \mathbf{M}\vec{a}. \tag{3.6}$$

Since the bilinear mapping may be seen to be a *separable* transformation — the product of two one dimensional mappings operating along orthogonal axes — the one dimensional mapping may easily be extended into two dimensions

$$\vec{y} = \mathbf{M}\vec{b}. \tag{3.7}$$

So, inverting M

$$\vec{a} = \mathbf{M}^{-1}\vec{x}, \tag{3.8}$$

$$\vec{b} = \mathbf{M}^{-1}\vec{y}. \tag{3.9}$$

3.3.3 Polynomial Transformations

The mapping functions U and V may be chosen to be global bivariate *polynomial transformations* of the form

$$u = \sum_{i=0}^{N} \sum_{j=0}^{N-i} a_{ij} x^i y^j, \tag{3.10}$$

$$v = \sum_{i=0}^{N} \sum_{j=0}^{N-i} b_{ij} x^i y^j. \tag{3.11}$$

where a_{ij} and b_{ij} are the constant polynomial coefficients. These polynomial transformations are low order global mapping functions operating on the entire image. They are intended to account for sensor-related spatial distortions such as centering, skew, and pincushion effects. These errors are characterised as low-frequency distortions. The global effects of the polynomial mapping will not account for high-frequency deformations that are local in nature. Since the nature of optical distortions due to imperfections in the manufacture of lenses and mirrors tends to have a global effect across the frame, no such high-frequency distortions are expected. Commonly, these transformations have been used with values of N in the range one through to four. By specifying corresponding positions in the signal and reference images the polynomial coefficients may be inferred.

The simplest solution for the polynomial coefficents is the so-called *pseudoinverse solution*. Suppose there is a known correspondence between M points in the observed and reference images. A spatial transformation of degree N is chosen. For two variables (x and y), such a polynomial has K coefficents, where

$$K = \sum_{i=0}^{N} \sum_{j=0}^{N-i} 1 = \frac{(N+1)(N+2)}{2} \tag{3.12}$$

For example, a second-degree approximation has $N = 2$ and $K = 6$. Where $M \geq 6$ it may be stated that

$$\begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ \dots \\ \dots \\ u_M \end{bmatrix} = \begin{bmatrix} 1 & x_1 & y_1 & x_1 y_1 & x_1^2 & y_1^2 \\ 1 & x_2 & y_2 & x_2 y_2 & x_2^2 & y_2^2 \\ 1 & x_3 & y_3 & x_3 y_3 & x_3^2 & y_3^2 \\ \dots & & & & & \\ \dots & & & & & \\ 1 & x_M & y_M & x_M y_M & x_M^2 & y_M^2 \end{bmatrix} \begin{bmatrix} a_{00} \\ a_{10} \\ a_{01} \\ a_{11} \\ a_{20} \\ a_{02} \end{bmatrix} \tag{3.13}$$

A similar equation holds for v and b_{ij} so

$$\vec{u} = \mathbf{W}\vec{a}, \tag{3.14}$$

$$\vec{v} = \mathbf{W}\vec{b}. \tag{3.15}$$

In order to solve for both $\vec{a}$ and $\vec{b}$ the inverse of $\mathbf{W}$ must be computed. Since $\mathbf{W}$ has dimensions $M \times K$ it is not a square matrix and thus has no inverse. Instead, both sides may be premultiplied by $\mathbf{W}^T$. This serves to cast $\mathbf{W}$ into a $K \times K$ square matrix that may be readily inverted. Thus

$$\mathbf{W}^T\vec{u} = \mathbf{W}^T\mathbf{W}\vec{a}, \tag{3.16}$$

$$\mathbf{W}^T\vec{v} = \mathbf{W}^T\mathbf{W}\vec{b}. \tag{3.17}$$

Expanding the matrix notation, eqn. 3.18 may be obtained. For notational convenience limits of summation and subscripts are omitted. All summations run from $k = 1$ to M, operating on data x_k, y_k and u_k.

$$\begin{bmatrix} \sum u \\ \sum xu \\ \sum yu \\ \sum xyu \\ \sum x^2u \\ \sum y^2u \end{bmatrix} = \begin{bmatrix} M & \sum x & \sum y & \sum xy & \sum x^2 & \sum y^2 \\ \sum x & \sum x^2 & \sum xy & \sum x^2y & \sum x^3 & \sum xy^2 \\ \sum y & \sum xy & \sum y^2 & \sum xy^2 & \sum x^2y & \sum y^3 \\ \sum xy & \sum x^2y & \sum xy^2 & \sum x^2y^2 & \sum x^3y & \sum xy^3 \\ \sum x^2 & \sum x^3 & \sum x^2y & \sum x^3y & \sum x^4 & \sum x^2y^2 \\ \sum y^2 & \sum xy^2 & \sum y^3 & \sum xy^3 & \sum x^2y^2 & \sum y^4 \end{bmatrix} \begin{bmatrix} a_{00} \\ a_{10} \\ a_{01} \\ a_{11} \\ a_{20} \\ a_{02} \end{bmatrix} \tag{3.18}$$

So,

$$\vec{a} = (\mathbf{W}^T\mathbf{W})^{-1}\mathbf{W}^T\vec{u}, \tag{3.19}$$

$$\vec{b} = (\mathbf{W}^T\mathbf{W})^{-1}\mathbf{W}^T\vec{v}. \tag{3.20}$$

Unfortunately this pseudo-inverse solution to the least-squares problem presents another difficulty: multiplying eqns. 3.14 and 3.15 with $\mathbf{W}^T$ squares the condition number, hence reducing the precision of the coefficients by half. This loss of precision proves to be significant for Doppler global velocimetry image processing applications, as demonstrated in Section 5.8.1.1.

The pseudo-inverse solution proves to be indentical to the classic *least squares formulation* with ordinary polynomials. Again, for $N = 2$, the coefficients a_{ij} may be determined by minimising

$$E = \sum_{k=1}^{M} \delta_k^2 = \sum_{k=1}^{M} \left(U(x_k, y_k) - u_k\right)^2 \tag{3.21}$$

$$= \sum_{k=1}^{M} (a_{00} + a_{10}x_k + a_{01}y_k + a_{11}x_k y_k + a_{20}x_k^2 + a_{02}y_k^2 - u_k)^2. \tag{3.22}$$

This may be achieved by determining the partial derivatives of E with respect to coefficients a_{ij}, and equating them to zero, For each a_{ij}

$$\frac{\partial E}{\partial a_{ij}} = 2 \sum_{k=1}^{M} \delta_k \frac{\partial \delta_k}{\partial a_{ij}} = 0 \tag{3.23}$$

By considering the partial derivative of E with respect to all six coefficients, eqn. 3.24 may be obtained. Limits of summation and subscripts have been omitted for clarity.

$$\begin{aligned}
\sum u &= a_{00}M + a_{10}\sum x + a_{01}\sum y + a_{11}\sum xy + a_{20}\sum x^2 + a_{02}\sum y^2 \\
\sum xu &= a_{00}\sum x + a_{10}\sum x^2 + a_{01}\sum xy + a_{11}\sum x^2 y + a_{20}\sum x^3 + a_{02}\sum xy^2 \\
\sum yu &= a_{00}\sum y + a_{10}\sum xy + a_{01}\sum y^2 + a_{11}\sum xy^2 + a_{20}\sum x^2 y + a_{02}\sum y^3 \\
\sum xyu &= a_{00}\sum xy + a_{10}\sum x^2 y + a_{01}\sum xy^2 + a_{11}\sum x^2 y^2 + a_{20}\sum x^3 y + a_{02}\sum xy^3 \\
\sum x^2 u &= a_{00}\sum x^2 + a_{10}\sum x^3 + a_{01}\sum x^2 y + a_{11}\sum x^3 y + a_{20}\sum x^4 + a_{02}\sum x^2 y^2 \\
\sum y^2 u &= a_{00}\sum y^2 + a_{10}\sum xy^2 + a_{01}\sum y^3 + a_{11}\sum xy^3 + a_{20}\sum x^2 y^2 + a_{02}\sum y^4
\end{aligned} \tag{3.24}$$

This is a symmetric 6×6 system of linear equations whose coefficients are all summations from $k = 1$ to M which are evaluated from the original data. By inspection, this may be seen to be equivalent to the system of equations 3.18 derived for the pseudo-inverse formulation. Known as the *normal equations*, the system may be compactly expressed as

$$\sum_{i=0}^{N} \sum_{j=0}^{N-i} a_{ij} \left(\sum_{k=1}^{M} x_k^i y_k^j x_k^l y_k^m \right) = \sum_{k=1}^{M} u_k x_k^l y_k^m \tag{3.25}$$

for each $l = 0, \ldots, N$ and $m = 0, \ldots, (N - l)$, (i, j) are running indices along rows, where each row is associated with a constant (l, m) pair.

The least squares procedure operates on an overdetermined system of linear equations ($M > K$). As a result, an approximating mapping function is computed. Intuitively, the polynomial of order K must yield the best compromise among all M control points. (If $M = K$, then no compromise is necessary and the polynomial mapping function will actually pass through

the given points.) It will be seen in Chapter 5 that the greater the number of control points used to define the polynomial warp, the better the noise performance of that warp.

The normal equations (eqn. 3.25) may be solved if M is greater than K. As K increases, however, the solution becomes increasingly unstable and inaccurate. Numerical accuracy is hampered by linear dependencies among the equations. As a result, direct solution of the normal equations is generally not the best method of obtaining a solutions. Here, *orthogonal polynomials* are introduced.

A set of polynomials $P_1(x, y)$, $P_2(x, y)$, ..., $P_K(x, y)$ is orthogonal over points (x_k, y_k) if

$$\sum_{k=1}^{M} P_i(x_k, y_k) P_j(x_k, y_k) = 0 \qquad i \neq j \tag{3.26}$$

The mapping functions U and V may be rewritten in terms of orthogonal polynomials

$$u = \sum_{i=1}^{K} a_i P_i(x, y), \tag{3.27}$$

$$v = \sum_{i=1}^{K} b_i P_i(x, y), \tag{3.28}$$

where a_i and b_i are unknown coefficients for orthogonal polynomial P_i. Introducing orthogonal polynomials allows the coefficients to be determined without solving a linear system of equations.

Such polynomials may be constructed from a set of linearly-independent basis functions spanning the space of polynomials. These basis functions may be denoted $h_1(x, y)$, $h_2(x, y)$, ..., $h_K(x, y)$. The polynomials may be generated by use of the Gram-Schmidt orthogonalisation process, which generates orthogonal functions by the following incremental procedure

$$\begin{aligned}
P_1(x, y) &= \alpha_{11} h_1(x, y) \\
P_2(x, y) &= \alpha_{21} P_1(x, y) + \alpha_{22} h_2(x, y) \\
P_3(x, y) &= \alpha_{31} P_1(x, y) + \alpha_{32} P_2(x, y) + \alpha_{33} h_3(x, y) \\
&\ldots \\
P_K(x, y) &= \alpha_{K1} P_1(x, y) + \alpha_{K2} P_2(x, y) + \ldots + \alpha_{KK} h_K(x, y)
\end{aligned} \tag{3.29}$$

In order to correspond the the polynomial warp function introduced above, the basis functions $h_i(x, y)$ may be chosen to be:

$$\begin{aligned} h_1(x,y) &= 1 \\ h_2(x,y) &= x \\ h_3(x,y) &= y \\ h_4(x,y) &= x^2 \\ h_5(x,y) &= xy \\ h_6(x,y) &= y^2 \end{aligned} \tag{3.30}$$

The α_{ij} parameters are determined by setting $\alpha_{i1} = 1$ and applying the orthogonality property of eqn. 3.26 to the polynomials. That is, the α_{ij} values are selected such that the product $P_i P_j = 0$ for $i \neq j$. The following is thus obtained

$$\alpha_{ii} = -\frac{\sum_{k=1}^{M} (P_1(x_k, y_k))^2}{\sum_{k=1}^{M} P_1(x_k, y_k) h_i(x_k, y_k)} \qquad i = 2, \ldots, K \tag{3.31}$$

$$a_{ij} = -\alpha_{ii} \frac{\sum_{k=1}^{M} P_j(x_K, y_k) h_i(x_k, y_k)}{\sum_{k=1}^{M} (P_j(x_k, y_k))^2} \qquad \begin{matrix} i = 3, \ldots, K \\ j = 2, \ldots, i-1 \end{matrix} \tag{3.32}$$

The equations above are the result of isolating the α_{ij} coefficients once P_i has been multiplied with P_1 and P_j respectively. This multiplication serves to eliminate all product terms except those associated with the desired coefficient.

Having computed the α_{ij} coefficients, the orthogonal polynomials are determined — they are simply linear combinations of the basis functions. Now it is possible to solve for the a_i coefficients of eqn. 3.27 using the least-squares approach. The error E may be defined

$$E = \sum_{k=1}^{M} \left(\sum_{i=1}^{K} a_i P_i(x_k, y_k) - u_k \right)^2 \tag{3.33}$$

The coefficients may then be determined by taking the partial derivatives of E with respect to the coefficients and setting them equal to zero. This results in the following system of linear equations

$$\sum_{i=1}^{K} a_i \left(\sum_{k=1}^{M} P_i(x_k, y_k) P_j(x_k, y_k) \right) = \sum_{k=1}^{M} u_k P_j(x_k, y_k) \qquad k = 1, \ldots, K \tag{3.34}$$

Applying the orthogonality property of eqn. 3.26

$$a_i = \frac{\sum_{k=1}^{M} u_k P_i(x_k, y_k)}{\sum_{k=1}^{M} (P_i(x_k, y_k))^2}. \tag{3.35}$$

Similarly for b_i

$$b_i = \frac{\sum_{k=1}^{M} v_k P_i(x_k, y_k)}{\sum_{k=1}^{M} (P_i(x_k, y_k))^2}. \tag{3.36}$$

Performing the least squares procedure by means of orthogonal polynomials offers several advantages. First, the determination of coefficients a_i and b_i does not require the solution of a system of equations. Instead, a closed form is obtained. This proves to be a faster and more accurate solution. (The computational cost of this method is $O(MK^3)$.) Second, additional orthogonal polynomial terms may be added to the mapping function to increase the fitting accuracy of the approximation.

The mean-squared error may be defined to be

$$E_{\mathrm{MS}} = \frac{1}{M} \sum_{k=1}^{M} \left(\sum_{i=1}^{K} a_i P_i(x_k, y_k) - u_k \right)^2. \tag{3.37}$$

If E_{MS} exceeds some threshold, the number of orthogonal polynomial terms in the mapping functions may be increased to reduce error, *without* recomputation of all the polynomial coefficients. Thus a means of adaptively determining the degree of polynomial necessary to satisfy some error-bound may be implemented.

In practice numerical accuracy is improved by normalisation of the coordinate values x in the range $-1 < x < +1$. Once the warping calculations have been performed, the normalisation may be inverted. This refinement is implemented in the warping algorithm used in `warp`.

3.3.4 Other Methods

These warping techniques have been implemented *globally* — a single set of warp coefficients is computed for the entire image. An alternative approach would be to implement *local* methods, whereby the region is divided into segments and warp coefficients computed for each segment (see, for example, [Thorpe 1996]). Local warping would appear to be preferable since it allows the modelling of local distortions. However, the distortion expected in the optical chain falls within the class of transformations which may be implemented using the

polynomial warp, as shown in Fig. 3.4. Distortions in the optical path may be due to poor physical alignment of the optical components or due to manufacturing defects in those components. Defects observed in lenses and mirrors are global in nature — having an effect across the entire frame. No high spatial frequency local distortion is expected. Local warping leads to difficulties in evaluation of the warp across segment boundaries. It is not anticipated that local distortions in the optics will be significant; thus, local warping techniques may lead to unnecessary discontinuities at segment boundaries, resulting in significant error in the overall result. Local warping techniques also have the effect of increasing the mean alignment error due to inaccurate centroid location as the segments are made successively smaller. Section 5.8.1.1 presents a consideration of the effect of centroid estimate inaccuracy on overall image misalignment under different warping algorithms.

Wolberg [Wolberg 1991] describes a *weighted least-squares* technique for computation of the polynomial warp. Such a technique takes account of the distance between control points and approximation positions so that the error contributions of distant control points have less influence than those which are closer to the approximating position. Such a scheme could be utilised to reduce warp inaccuracy due to poorly located centroids.

Since error may be evaluated for each pair of control points, it is possible to discard any points which exceed a certain error bound. Since higher order local distortion is not expected in the optical chain, the presence of such pairs indicates that significant error has been encountered in the centroid location stage of the the image processing.

Higher order polynomial methods may be implemented in order to model more complex non-linear distortions. However, the use of such higher-order techniques may lead to over-fitting and increased susceptibility to noise due to errors in the computed centroid locations.

3.4 Interpolation Methods

Interpolation is the process of determining the values of a function at positions lying between its samples. This may be achieved by fitting a continuous function through the discrete input samples, so that input values may be evaluated at arbitrary positions rather than only at the sample points themselves. Whereas sampling generates an infinite bandwidth signal from one that is band-limited, interpolation reduces the bandwidth of the signal by applying a low pass filter to the discrete signal. Once again, Wolberg [Wolberg 1991] discusses a number of different interpolation schemes — although little attempt is made in that work to compare the performance of the techniques with one another.

Interpolation in one dimension for equally spaced sample points, may be formulated as a *convolution* operation

$$f(x) = \sum_{k=0}^{K-1} c_k h(x - x_k), \tag{3.38}$$

where h is the *interpolation kernel* weighted by coefficients c_k and applied to K data samples x_k. The kernel h is nearly always symmetric: $h(-x) = h(x)$. It is also generally the case that the c_k coefficients are the data samples themselves.

The concept of an interpolation kernel as suggested by convolution is useful for the comparison of interpolation algorithms. Analysis of the kernel may be used to determine the numerical accuracy of the interpolated function. Typically the performance of the kernel in the passband and stopband is studied. The ideal kernel will have unity gain in the passband and zero gain in the stopband in order to transmit and suppress the signal's spectrum in these respective frequency ranges. Ideal filters generally have a wide extent in the spatial domain — for instance, the **sinc** function has infinite extent. Such *infinite impulse response filters (IIR)* may only be realised approximately, as *recursive filters* with feedback. Although recursive filters may give superior results to *finite impulse filters (FIR)* for a given number of coefficients, they are more difficult to design and implement. Using an FIR filter, each output value is computed as the weighted sum of a finite number of neighbouring input elements. FIRs find widespread use in image processing.

The ideal interpolation kernel in the frequency domain is the **rect** function, since this has unity gain in the desired passband $-f_n < f < +f_n$, where f_n represents the Nyquist frequency, and zero gain elsewhere. The quality of an interpolation technique may thus be evaluated by inspection of the Fourier transform of its kernel. In the spatial domain, the inverse Fourier transform yields a **sinc** function. Unfortunately, such a function has infinite extent and, thus, may not be evaluated.

Truncating the kernel in the spatial domain may be shown to lead to ringing in the frequency domain. Truncation is equivalent to multiplying by a **rect** function, whose Fourier transform is a **sinc** function. Multiplication in the spatial domain is equivalent to convolution in the frequency domain. Thus, in the frequency domain, the truncated **sinc** function becomes a **rect** function convolved with a **sinc**, leading to ringing.

The ideal interpolation scheme, then, must attempt to evaluate an approximation to the **sinc** function in the spatial domain, without the harsh truncation which may cause ringing.

In the following discussion, analysis is applied to the one dimensional case. Interpolation in two dimensions is a simple extension of the one dimensional results. Data samples are assumed to be equally spaced along each dimension.

3.4.1 Nearest Neighbour

The *nearest neighbour* algorithm is the simplest method from a computational standpoint. By this method each output pixel is assigned the value of the nearest sample point in the input image. Although this is strictly not an interpolation scheme, the method may be examined in the same manner as the more sophisticated approaches described below. The following interpolating polynomial achieves this

$$f(x) = f(x_k) \qquad \frac{x_{k-1} + x_k}{2} < x \leq \frac{x_k + x_{k+1}}{2} \tag{3.39}$$

Nearest neighbour may be performed by convolving the image with a one pixel width rectangle in the spatial domain. The interpolation kernel is defined as

$$h(x) = \begin{bmatrix} 1 & 0 \leq |x| < 0.5 \\ 0 & 0.5 \leq |x| \end{bmatrix} \tag{3.40}$$

This kernel is known as a *box filter* or *sample-and-hold function*. The kernel and its Fourier transform are shown in Fig. 3.5.

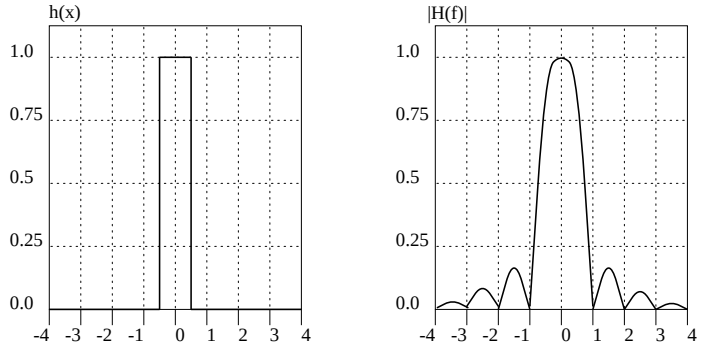


Figure 3.5: Nearest neighbour kernel and Fourier transform.

Convolution in the spatial domain with a rectangle function is equivalent to multiplication with a **sinc** function in the frequency domain. Because of its prominent side lobes and infinite extent, a **sinc** function makes a poor low pass filter. Thus the nearest neighbour algorithm has a poor frequency domain response.

3.4.2 Linear

Linear interpolation is a first degree method that passes a straight line between every two consecutive points of the input signal. For an interval (x_0, x_1) and function values f_0 and f_1 for the endpoints, the interpolating polynomial is

$$f(x) = a_1 x + a_0 \tag{3.41}$$

where a_0 and a_1 are determined by solving

$$\begin{bmatrix} f_0 & f_1 \end{bmatrix} = \begin{bmatrix} a_1 & a_0 \end{bmatrix} \begin{bmatrix} x_0 & x_1 \\ 1 & 1 \end{bmatrix} \tag{3.42}$$

This gives rise to the interpolating polynomial

$$f(x) = f_0 + \left(\frac{x - x_0}{x_1 - x_0} \right) (f_1 - f_0) \tag{3.43}$$

In the spatial domain, linear interpolation is equivalent to convolving the sampled input with the following kernel

$$h(x) = \begin{cases} 1 - |x| & 0 \leq |x| < 1 \\ 0 & 1 \leq |x| \end{cases} \tag{3.44}$$

This kernel is known as a *triangle filter* or *Chateau function*. The kernel and its Fourier transform are shown in Fig. 3.6.

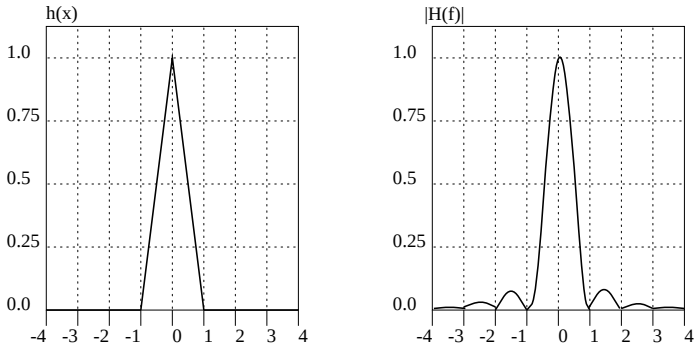


Figure 3.6: Linear interpolation kernel and Fourier transform.

Linear interpolation offers better image quality than nearest neighbour by accomodating first-degree fits. Since it produces reasonably good results at moderate cost, it is the most

widely used interpolation algorithm. As a first order fit is employed here, the order of the expected error due to the use of a linear interpolation scheme is second order.

Although second-degree interpolating polynomials would appear to be the next step, it has been shown that all such filters of even degree are space-variant. This may be attributed to the fact that the number of sampling points to each side of the interpolated point always differs by one.

3.4.3 Cubic

Cubic convolution is a third-degree interpolation scheme originally suggested as an efficient approximation to the theoretically optimum **sinc** interpolation function. The kernel is composed of piecewise cubic polynomials defined on the unit subintervals $(-2,-1)$, $(-1,0)$, $(0,1)$ and $(1,2)$. Outside the interval $(-2,2)$, the interpolation kernel is zero. Each interpolated point is thus a weighted sum of four consecutive input points. A symmetric, space-invariant interpolation kernel of the form

$$h(x) = \begin{cases} a_{30}|x|^3 + a_{20}|x|^2 + a_{10}|x| + a_{00} & 0 \leq |x| < 1 \\ a_{31}|x|^3 + a_{21}|x|^2 + a_{11}|x| + a_{01} & 1 \leq |x| < 2 \\ 0 & 2 \leq |x| \end{cases} \tag{3.45}$$

is required. The values of the coefficients may be determined by application of the following constraints:

1. $h(0) = 1$ and $h(x) = 0$ for $|x| = 1$ and 2.
2. h must be continuous at $|x| = 0$, 1 and 2.
3. h must have a continuous first derivative at $|x| = 0$, 1 and 2.

The first constraint states that when h is centered on an input sample, the interpolation function is independent of neighbouring samples, permitting f to actually pass through the input points. In addition it establishes that the c_k coefficents of eqn. 3.38 are the data samples themselves.

From the first two constraints

$$1 = h(0) = a_{00} \tag{3.46}$$

$$0 = h(1^-) = a_{30} + a_{20} + a_{10} + a_{00} \tag{3.47}$$

$$0 = h(1^+) = a_{31} + a_{21} + a_{11} + a_{01} \tag{3.48}$$

$$0 = h(2^-) = 8a_{31} + 4a_{21} + 2a_{11} + a_{01} \tag{3.49}$$

From the third constraint

$$-a_{10} = h'(0^-) = h'(o^+) = a_{10} \tag{3.50}$$

$$3a_{30} + 2a_{20} + a_{10} = h'(1^+) = 3a_{31} + 2a_{21} + a_{11} \tag{3.51}$$

$$12a_{31} + 4a_{21} + a_{11} = h'(2^-) = h'(2^+) = 0 \tag{3.52}$$

There are now seven equations for the eight unknowns. By allowing $a = a_{31}$ to be a free parameter that may be controlled by the user, the family of solutions may be obtained

$$h(x) = \begin{bmatrix} (a+2)|x|^3 = (a+3)|x|^2 + 1 & 0 \leq |x| < 1 \\ a|x|^3 - 5a|x|^2 + 8a|x| - 4a & 1 \leq |x| < 2 \\ 0 & 2 \leq |x| \end{bmatrix} \tag{3.53}$$

The heuristics applied to derive this kernel are motivated from properties of the ideal reconstruction filter, the **sinc** function. Requiring h to be concave upwards at $|x| = 1$, and concave downwards at $x = 0$, gives that $-3 < a < 0$. The choice of $a = -1$ may be made in order to match the slope of the **sinc** function at $x = 1$. Alternatively, $a = -0.5$ may be chosen by making the Taylor series approximation of the interpolated function agree with the original signal to as many terms as possible. The kernel and its Fourier transform are shown in Fig. 3.7.

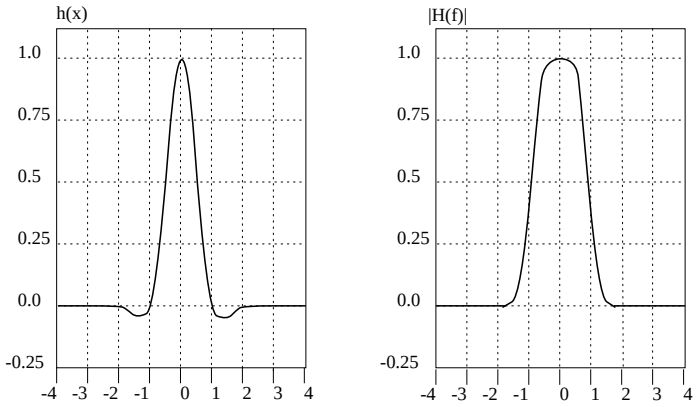


Figure 3.7: Cubic convolution kernel and Fourier transform.

Cubic convolution may give rise to values outside the range of input data. Consequently care must be taken when applying this technique to image processing.

As a third order fit is employed here, the order of the expected error due to the use of a cubic interpolation scheme is forth order. The cubic interpolation scheme is thus expected offer superior performance to the linear method.

Press *et al.* [Press *et al.* 1992] describe useful techniques for implementing a cubic interpolation scheme efficiently.

3.4.4 Exponential

Exponential functions may be used to implement a superior class of interpolation functions. The hyperbolic tangent **tanh** has several desirable properties. First, it converges quickly to ± 1. Second, its transition from -1 to 1 is sharp. This transition may be further sharpened by use of **tanh(**kx) for $k > 1$. The function is infinitely differentiable everywhere.

Using **tanh(**kx), a new function that resembles the ideal low-pass filter, **rect** (a box in the frequency domain), may be defined

$$H_k(f) = \left(\frac{\tanh(k(x + f_c)) + 1}{2}\right)\left(\frac{\tanh(k(-x + f_c)) + 1}{2}\right) \tag{3.54}$$

where f_c is the cut-off frequency, set to 0.5 to conform to the Nyquist rate. Function H_k is treated as the desired spectrum of the reconstruction filter. For low values of k, H_k is smooth and resembles a Gaussian function. As k is made larger, H_k will have increasing sharper corners, tending towards a **rect** function, as shown in Fig. 3.8.

Having established the desired spectrum H_k, the actual kernel is obtained by computing the inverse Fourier transform, giving $h_k(x)$, as shown in Fig. 3.9. The kernel has infinite extent. However, unlike the **sinc** function that decays at a rate of $1/x$, h_k decays at an exponential rate. This means that truncation of the kernel may be performed with a negligible penalty in reconstruction quality. A seven point kernel has been reported to offer excellent results for general image processing applications [Wolberg 1991].

The published literature offers no practical comparison of the relative performance of the different interpolation schemes discussed above. In Chapter 5 each of the schemes is modelled and analysed in order to determine the optimal method for use with Doppler global velocimetry data. The effect of using different interpolation schemes in practice is discussed in Chapter 7 for the case of the measurement of a free jet flow using the Oxford velocimeter.

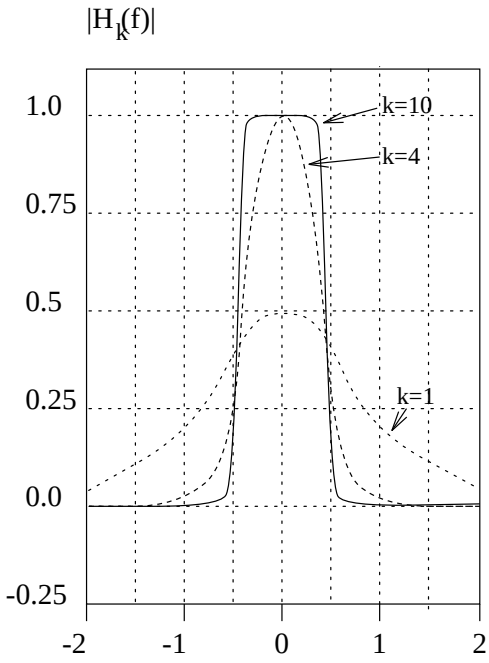


Figure 3.8: Spectrum of $H_k(f)$

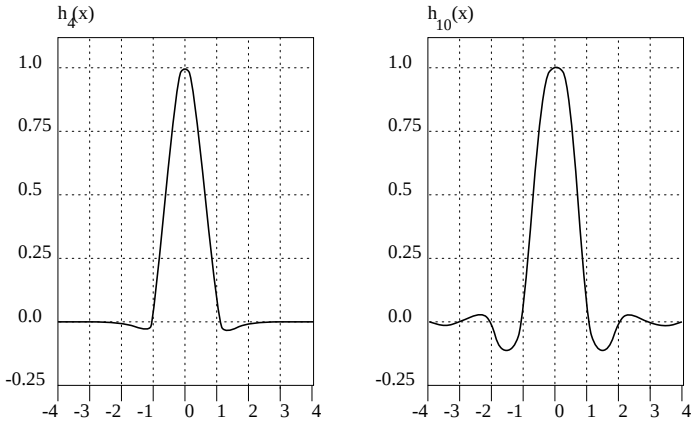


Figure 3.9: Exponential interpolation kernels

The process of finding the position of the alignment card dots may be used to effect alignment of the discriminated and non-discriminated images, as described above. This method has more general application however. Suppose that, as described in Section 2.9, a pair of optical systems is used in order to obtain two component velocity measurements. There are then three alignment tasks to be performed rather than just one. For both optical systems, the discriminated and the non-discriminated images must be aligned. In addition, the images from each system must be aligned with one another in order that two component velocity results may be obtained from corresponding points in the flow field.

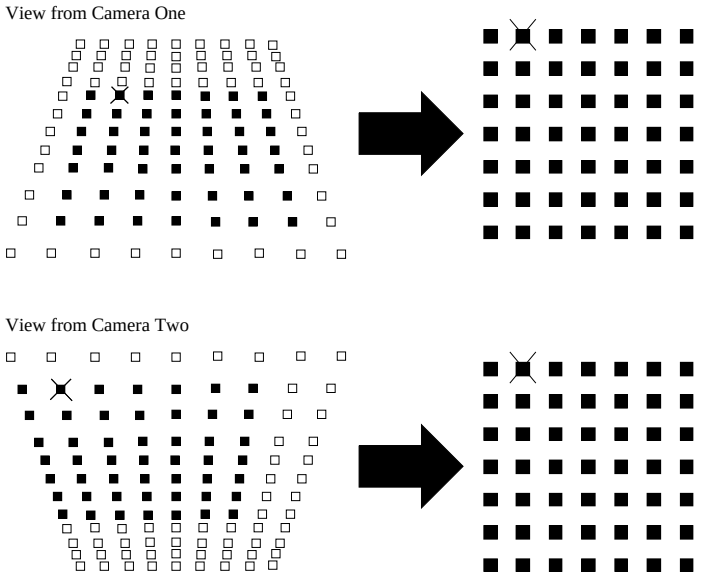


Figure 3.10: Alignment of images from two optical systems

The nature of the image processing strategy described above makes it well suited to this task. Rather than warping the discriminated image into the same space as the non-discriminated image or vice-versa, each image is warped onto a common regular grid. Hence, if the same alignment dots are chosen to form the basis of the alignment for each optical system, not only will the discriminated and non-discriminated images be aligned for each system, but, since they are warped to a common grid, the two optical systems will be aligned to one another at the same time.

Fig. 3.10 illustrates how the alignment may be effected. A representation of the view of an alignment target from two different optical systems is shown. Due to the different observation angles, perspective effects will be present in such a configuration. A single alignment dot should be marked differently to the others to enable the experimenter to observe the correspondence between the two images. In the case of the alignment cards used for the Oxford system, one of the dots is printed a little larger than the others. In Fig. 3.10 one dot

is represented with a cross through it. This dot may be seen in both images and hence the correspondence of the dots should be readily apparent.

When processing the data, the manual editing facility of the `warp` software may be used to limit the alignment features to be used for computing the image transformation to those present in the images obtained from both optical systems. In general, there will be dots in one image which are not visible in the other and *vice versa* due to the different observation angles employed. The same physical dots on the alignment card should be used for computing the warp vectors for each optical system. The warp is then performed to a regular grid in such a way that the two separate optical systems will be aligned with one another.

Such an alignment technique applied to two component measurements may be readily extended to allow the alignment of a third optical system. In this way, three component Doppler global velocimetry systems may be set up using the same image processing methodology. The optical systems should differ in observation angle sufficiently to ensure that accurate velocity components may be obtained (see Section 5.11). Therefore, the region which may be viewed in all the optical systems is significantly smaller than the CCD array. A compromise between image resolution and the accurate determination of multiple components of velocity must be achieved.

3.6 Summary

In this chapter a complete image processing methodology for the alignment of Doppler global velocimetry images has been outlined. The scheme centres on the use of a calibration frame, obtained by imaging a special alignment test card. A calibration normalisation frame is used to overcome the effects of vignetting and non-uniform pixel sensitivity. Additionally, a calibration black frame is employed in order to compensate for non-zero black levels.

The resulting adjusted calibration frame is then subject to a novel automated algorithm, which isolates each calibration feature in its own bounding box. The centroids of the calibration features may then be estimated, and the correspondence between the features in the discriminated and non-discriminated parts of the image may be computed.

Once a set of corresponding feature pairs has been obtained, the image transformation may be effected by means of either the bilinear or polynomial warps. Such transformations result in a mapping from non-integer coordinates to the integer-spaced pixels of the aligned output frame. The use of interpolation is therefore necessary. Four interpolation techniques have been discussed: nearest neighbour, bilinear, cubic, and exponential.

Having computed and saved the warp vectors for a given velocimeter configuration, it is possible to use them to align the images obtained in the acquision of flow data. Just as was the case for the alignment procedure, each velocity measurement requires the acquisition of three frames: a data frame, a normalisation frame and a black frame. If the velocimeter configuration is not changed, it is only necessary to compute the warp vectors once regardless of how many measurements are to be made. The `warp` software is used to compute the transmission ratio of the iodine cell across the entire area of interest. As explained in Section 4.1, this results frame may be further processed in order to obtain quantitative velocity data.

In Chapter 5, the expected error resulting from centroid estimation, image transformation and interpolation will be assessed, and the different warping and interpolation schemes will be compared and evaluated.

An example run of the image processing software, `warp`, showing each of the stages and the intermediate images which are formed, may be found in Chapter 6. An evaluation of the image processing performance achieved in practice is presented in Chapter 7.

Chapter 4

Further Image Processing Techniques

4.1 Absolute Velocity Determination and the Calibration of the Iodine Cell

In Chapter 3, image processing techniques for ensuring accurate alignment of the discriminated and non-discriminated images were presented. The goal of that procedure was to compute a representation of the transmission ratio through the iodine absorption filter for each point in the image. In order, however, to obtain accurate quantitative velocity measurements, it is necessary to consider the shape of the iodine absorption feature itself.

This chapter describes the use of a second iodine cell, the reference cell (discussed in Section 2.5.2), to determine the laser emission frequency ν_o during Doppler global velocimetry measurements. Experimental techniques for the determination of the absorption profile of both the reference cell and the working cell of the Doppler global optical system are presented. Examples of cell calibrations are described, together with the polynomial curve fit approximations used in subsequent processing algorithms to convert intensity to velocity data. The repeatability of the calibrations is also considered.

The principles of operation of the `IODINE` and the `viewresults` software are described. The `IODINE` software makes use of the reference cell apparatus in order to determine the laser emission frequency ν_o. The `viewresults` software, on the other hand, provides a post-processing stage for the ratio images obtained from `warp`, allowing quantitative velocity measurements to be obtained from the Doppler global technique.

4.1.1 Calibration Technique

It is possible to determine the transmission profile of an iodine cell by scanning a tunable monochromatic light source across the bandwidth of the absorption feature, while measuring the cell transmission as a function of the frequency of the incident light. The argon-ion laser is undoubtedly the most convenient suitable light source for this purpose. The principal difficulty in determining the transmission profile is establishing the light frequency corresponding to each of the transmission measurements.

In the context of determining iodine cell transmission profiles and subsequently employing those profiles for velocity determination, the relative light frequency — that is the frequency deviation with respect to some arbitrary datum — may be considered. It is not necessary

to determine the absolute value of laser emission frequency.

A number of other workers have examined the problem of calibrating iodine cells for use with Doppler global velocimetry systems. The use of a Fabry-Perot inteferometer is discussed by Ainsworth and Thorpe [Ainsworth and Thorpe 1994] and Chan *et al.* [Chan *et al.* 1995]. Such an instrument may be employed to determine the relative frequency of the laser emission during the calibration procedure.

Roehle and Schodl [Roehle and Schodl 1994] describe a calibration method which relies on examining the hyperfine structure of the iodine absorption line that is coincident with the 514.5 nm emission of an argon ion laser. The frequency of the features in this structure are dependent purely upon the physics of the iodine molecule. For this reason, these frequencies have been employed in the definition of the meter, and are hence reported in the literature ([Metrologia 1984] for example) to an accuracy of several kHz.

A further procedure for determining the iodine transmission curves is described by Komine *et al.* [Komine *et al.* 1991]. Using this technique, the Doppler global velocimetry system itself is employed to observe the Doppler-shift frequency exhibited by the light scattered from a disc rotating at a known angular velocity. This technique does not rely on tuning the light frequency across an absorption feature, but rather employs a rotating disc to create a range of optical frequencies simultaneously in the vicinity of the absorption line.

Each of the above mentioned calibration techniques suffers drawbacks in practical application. A Fabry-Perot interferometer is insufficiently accurate to allow reliable calibrations to be undertaken, due to the temperature sensitivity of such devices and the non-linearity of the Piezo-electric mirror drives typically employed. The Fabry-Perot interferometer available at Oxford, typical of its type, has a linearity of around 1% and a temperature sensitivity of 200 MHz/°C. The linearity specification implies approximately 20 MHz error in the frequency calibration over the free spectral range of 2 GHz. Improved Fabry-Perot interferometers are commercially available, but the cost of such an instrument has precluded its use.

The detection of the hyperfine structure of the iodine absorption feature in order to establish light frequency during a calibration procedure requires complex apparatus and active control of the laser cavity mirror separation using Piezo-electric mirror mountings and an associated control system, as described by Roehle and Schodle [Roehle and Schodl 1994]. In addition, the components of hyperfine structure are not regularly spaced over the entire bandwidth of the absorption feature. In particular, there are no components coincident with the region best suited for Doppler global measurements, where transmission varies most rapidly.

The use of rotating discs for cell calibration has been rejected for two reasons: such an approach would be inconvenient once the Doppler global velocimetry system had integrated into a flow-measurement experiment and it may prove impossible to obtain the required range of velocities from the disc surface, when a high-velocity flow is to be investigated.

The technique employed in the present work, developed by Thorpe [Thorpe 1996], and known as *mode-hop based calibration*, allows accurate cell calibrations to be undertaken with the Doppler global system *in situ* in a flow measurement experiment. No additional apparatus beyond that which is necessary for Doppler global measurements to be undertaken is required. The features of the argon-ion laser itself have been exploited in order to determine relative frequency — in particular, the longitudinal mode spacing is known to be 143±0.4 MHz [Spectra-Physics 1987].

The reference iodine cell (which is described in Section 2.5.2) is used as a frequency standard for all of the Oxford-based Doppler global experiments. Iodine absorption features have been utilised in various frequency standard applications, such as those reported by Gill and Thompson [Gill and Thomson 1986] and Chartier, Robertsson and Fredin-Picard [Chartier, Robertsson and Fredin-Picard 1991]. The characteristics of the transmission profile of the reference cell itself must first be established. After this calibration procedure, the reference cell allows the relative laser frequency to be determined during the acquisition of Doppler global measurement frames by monitoring the transmission ratio using the photodiode pair discussed previously. Because the iodine cells used in the Oxford system are highly stable the calibration, once performed, remains valid over long time scales — of the order of months.

The transmission level measured for a given iodine cell is the combined transmission of the iodine vapour in the cell and the cell windows. The windows used in the present cell are plane glass, without any coatings which might minimise reflective loss. The reflective loss at each window to gas interface is approximately 4% for light at normal incidence [Jenkins and White 1976]. The windows of the iodine cell therefore exhibit a transmission of $0.96^4 = 0.85$, there being four interfaces for the two windows.

The photodiodes used to monitor the input and output light intensity levels in the reference cell enclosure are sensitive to the spatial alignment of the beam incident upon them. Due to variations in the fabrication of individual components, the response of the pair of photodiodes is not identical. In order to prevent spatial sensitivity influencing measurements of light intensity, and hence ensure that the long term calibration validity of the reference cell is not compromised by mechanical instability in the apparatus, the reference cell apparatus has

been securely mounted adjacent to the laser on optical rails. The working cell installed in the Doppler global velocimetry optical system is also mounted securely. The working cell is calibrated *in situ*, using the CCD camera itself to measure transmission.

As discussed in Section 2.6, the argon-ion laser employed for Doppler global measurements emits radiation at specific frequencies, or modes, of known separation. The use of an intra-cavity etalon in the laser allows the emission to be restricted to just one of these longitudinal modes. The mode frequencies are determined by the laser cavity mirror separation itself. For the laser in use in the current work, the manufacturers [Spectra-Physics 1987] specify a longitudinal mode separation of 143 ± 0.4 MHz. If the laser can be induced to change its emission from one mode to another adjacent mode, the frequency ν_o will be changed by this known separation. Such a mode transition (or mode-hop) may be induced by adjusting the temperature of the intra-cavity etalon. Each time the laser mode-hops to a new frequency, a step change in iodine transmission is observed.

The laser employed at Oxford features no active control of the cavity mirror separation, which is thus subject to a thermally induced variation as the resonator expands and contracts with ambient temperature. This effect is made evident by a drift of approximately 40 MHz/^{o}C in the laser emission frequency. The thermally induced variation in emission frequency, however, typically decays to around 1.2 MHz per minute after two hours continuous steady-state operation.

4.1.2 Calibration of the Reference Cell

In order to undertake calibration of the reference cell, the laser is initially left to operate at constant excitation current and cooling water flow rates for a period of at least several hours, allowing any thermal transients to decay. The single longitudinal mode of the laser is selected by the intra-cavity etalon, which is tuned so that the laser frequency corresponds to the edge of the iodine absorption line.

The etalon operating temperature is then gradually adjusted by manipulating the etalon oven control, so that the laser frequency increases as a series of mode-hops. The response time of the etalon controller is such that it takes around ten minues for the laser frequency to change by the 2 GHz necessary to scan across the bandwidth of the iodine absorption curve. As explained above, each mode-hop represents a change in emission frequency by 143 MHz. The reference cell apparatus, shown in Fig. 2.6, utilises a pair of photodiodes as input and output light intensity transducers. The recording of the photodiode voltages is achieved by means of a 16 bit analogue-digital converter fitted to the reference cell personal computer.

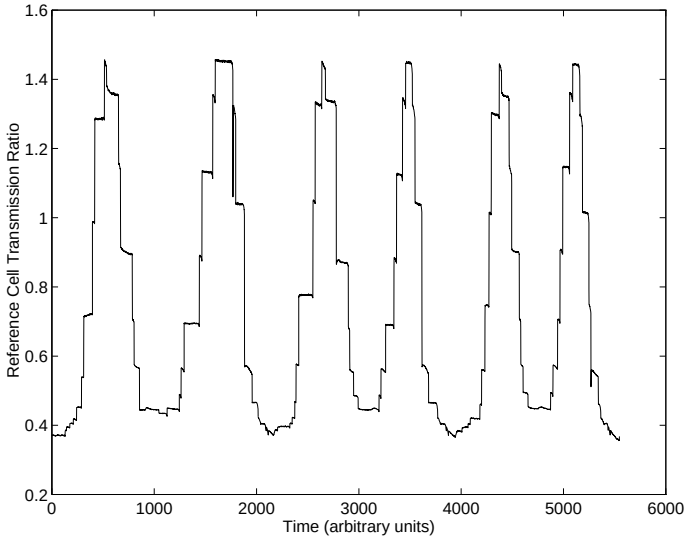


Figure 4.1: Reference cell transmission ratio versus time during the calibration run.

Once the full bandwidth of the absorption curve has been covered, the etalon oven control is gradually adjusted in the opposite direction so that the laser emission frequency mode-hops back to the initial value. This procedure is repeated three times, so that the absorption profile is scanned six times in total. Fig. 4.1 shows the ratio of the photodiode voltages versus time during such a calibration run. This ratio, once analogue-digital converter offset voltages have been taken into account, may be taken as a measure of the transmission of the reference cell. Each of the horizontal regions in the ratio plot represents a mode, with the vertical regions representing a rapid mode-hop.

It may be seen that the ordinate value does not represent transmission *per se*. The reason for this is that the two photodiode amplifiers are not set to equal gains. Since, in general, the transmitted intensity will be lower than the incident intensity, it is appropriate to choose a higher gain for the transmitted photodiode amplifier in order to minimise the effect of electrical noise on the signal and make the best possible use of the resolution of the analogue-digital converter. Taking account of the individual photodiode sensitivities and the characteristics of the beam splitter employed would allow a true transmission ratio to be obtained. This is feasible, but unnecessary in practice, since all that is required of the reference cell is a repeatable calibration of the iodine cell profile. Subsequent use of the reference cell data to find the laser emission frequency ν_o requires only that the reference cell characteristics be consistent with those determined during calibration.

The six forward and backwards scans of the bandwidth of the absorption feature provide some insight into the thermal stability of the laser and indicate whether thermal transients in the resonator are significant. Approximately 17 longitudinal modes of the laser are required to

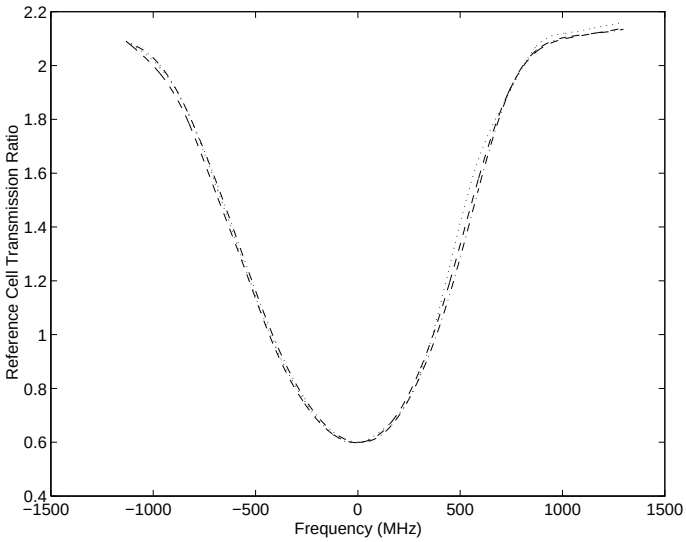


Figure 4.2: Reference cell transmission profiles for the three calibration runs.

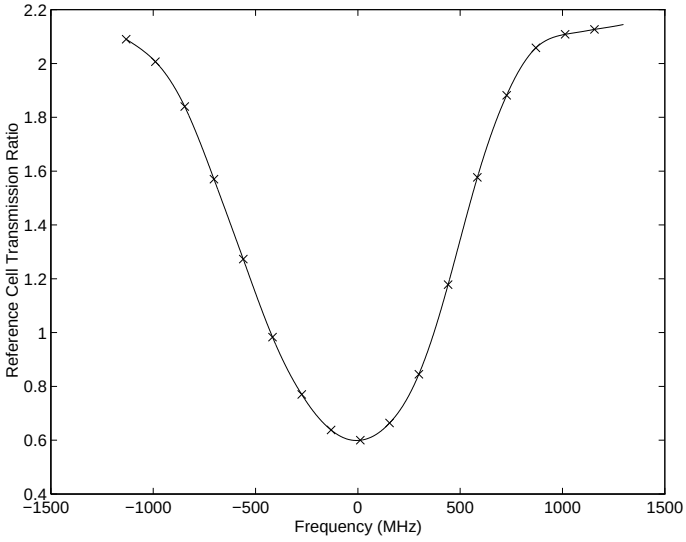


Figure 4.3: Mean of the three measured reference cell transmission profiles.

define the transmission characteristics of the reference cell, this being equivalent to 2.43 GHz.

The reference cell transmission data must be converted to a form such that a measurement of the photodiode voltages may be converted into a relative laser emission frequency value as explained in Section 4.1.5. The cell calibration data is therefore cast into the form of a polynomial curve fit to the data associated with each longitudinal mode of the laser. A least-squares approach is used for establishing the coefficients of the polynomial. A 6th order polynomial has been employed, as expressed by

$$\nu_o = \sum_{n=0}^{5} A_n \left(\frac{I_t}{I_i}\right)^n \tag{4.1}$$

where I_t represents the transmitted intensity, I_i the incident intensity, and A_n the polynomial coefficients.

Making use of the fact that the modes of the laser are known to be separated by 143 MHz, and utilising the polynomial curve fit of eqn. 4.1, the reference cell transmission profile (once again with arbitrary values on the ordinate axis) may be determined. An example of such a transmission profile is shown in Fig. 4.2. It may be seen from this figure that the three absorption profiles determined from the three forwards and backwards traverses of the absorption feature are in close agreement with one another. The mean of the three computed transmission profiles is shown in Fig. 4.3. The use of the curve-fit allows photodiode voltages from the reference cell arrangement to be interpreted as a measure of laser frequency, ν_o.

In converting Doppler-shift frequencies into intensity variations, the Doppler global technique makes use of only one side of the absorption feature at any one time. It is desirable to calibrate both wings of the absorption profile, since there are circumstances where it is more convenient to employ one wing than the other. This is undertaken by performing two separate polynomial curve fits to cover the entire bandwidth.

The accuracy of the curve-fit has been established by calculating the light frequency using the curve-fit and comparing the result to the measured data points. The RMS difference between the data and the curve-fit was found to be 2.9 MHz. For the calibration to be useful in determining laser emission frequency, the laser must be tuned into the region of validity of the curve-fit.

4.1.3 Calibration of the Working Cell

Having mapped out the transmission curve for the reference cell, it is also necessary to establish the corresponding transmission for the working cell in the Doppler global velocimetry system itself. This may be done with the aid of the newly-computed reference cell transmission.

The transmission function of the working cell may be determined with the cell *in situ* in the Doppler global imaging system. This may be achieved by acquiring a series of CCD frames of a white target illuminated by the argon ion laser emission while simultaneously scanning the emission frequency across the bandwidth of the iodine absorption line, as was done for the reference cell calibration. The previously calibrated reference cell is used to determine the emission frequency ν_o for each frame.

The laser is initially tuned to one side of the absorption line using the intra-cavity etalon oven control. The controller is then adjusted so that the laser frequency increases by mode-hopping. As the laser mode-hops across the appropriate frequency range, the reference cell photodiode voltages are monitored and simultaneously a succession of imaging system CCD frames are acquired. Since the reference cell calibration is used to determine the emission frequency corresponding to each measurement, the stability of the frequency of each of the longitudinal modes is unimportant when calibrating the working cell.

The CCD camera features a connector which enables a number of signals to be accessed externally. One of these signals indicates whether the camera shutter is open or closed. This signal is connected to a further input of the personal computer analogue-digital board.

In order to obtain reference cell transmission values for each acquired frame, the personal computer is set up to run the `IODINE` code, which computes and displays the average reference cell transmission ratio during the period of the CCD camera exposure. The 'shutter open' signal is thus used as a trigger for the data acquisition of the transmitted and incident light levels from the photodiodes.

Subsequent data processing of the CCD image may be used to determine the transmission of the working iodine absorption cell. This may be calculated simply by averaging the pixel values in both the discriminated and non-discriminated images and then dividing the former by the latter.

The data determined for the working cell transmission must be expressed in a form which allows the ratio of discriminated to non-discriminated intensity to be converted into a value

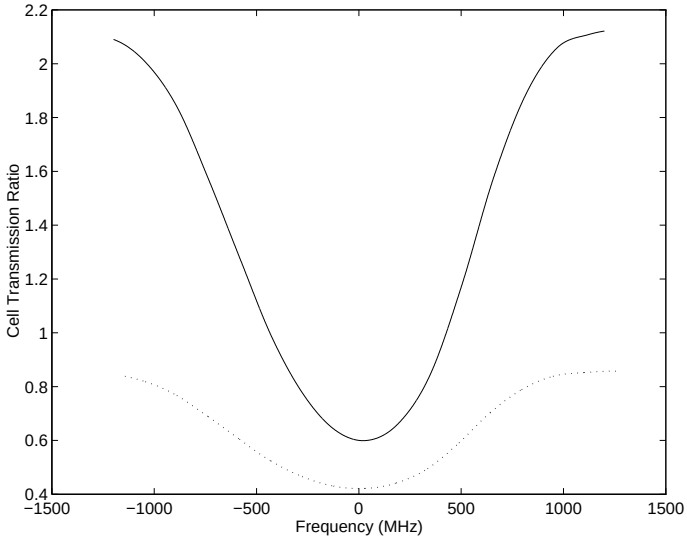


Figure 4.4: Reference cell transmission profile (solid line) versus working cell transmission profile (dotted line).

corresponding to the frequency of the Doppler-shifted light scattered from the seeding particles in the flow under consideration. The same curve fit procedure as that described in section 4.1.2 is thus employed again in order to obtain transmission curves for the working cell. An example of such a transmission curve together with the associated reference cell curve is shown in Fig. 4.4.

It is immediately apparent that the absolute value of the working cell data differs from that of the reference cell. As explained previously, the reference cell photodiode amplifiers may be set to different values of gain in order to minimise the impact of noise in reference cell transmission determination. This has the effect of changing the overall value of the reference cell data so that it may not be regarded as a measure of transmission. The absolute value of the reference cell data is unimportant for the purposes of determining the laser emission frequency ν_o provided that the properties of the reference cell apparatus remain stable. In contrast, both the transmitted and the incident light intensities for the working cell are measured by means of the same CCD array. The working cell data may thus be regarded as a true measure of transmission. Even in the case of the working cell, the configuration of lenses and mirrors in the Doppler global velocimetry system will affect the transmission value determined, so it is important to recalibrate the cell if the optical system is adjusted.

Using the working cell curve of Fig. 4.4, Doppler-shifted frequencies may be established for each pixel in a Doppler global measurement image. The use of the reference cell apparatus and its corresponding calibration enables the laser emission frequency ν_o to be determined. It is thus possible to compute the Doppler-shift, and hence velocity at each pixel in the

acquired measurement images.

4.1.4 Repeatability of the Calibration Data

Thorpe [Thorpe 1996] describes experiments which were conducted to determine the repeatability of the calibration over a 72 hour period. In that work, repeat calibrations were conducted for the temperature stabilised iodine cell, employing the photodiode system for determining transmission. While undertaking the calibrations, and also during the period between them, the cold-finger of the iodine cell was continually temperature stabilised.

The calibration profiles obtained led to the conclusion that excellent repeatability was evident between the two sets of calibration points. The density of the iodine vapour was found to be consistent between the two calibrations — any change in density would manifest itself as inconsistencies in the two sets of calibration curves. The RMS difference between the interpolated transmission profiles was found to be approximately 0.3%.

Further to the work of Thorpe, an investigation has been undertaken of the longer term stability of the iodine cells in use. This investigation has been extended to consider both the reference and working cells. Any change in the physical configuration caused by mechanical disturbance of the system is expected to be evident in such a study.

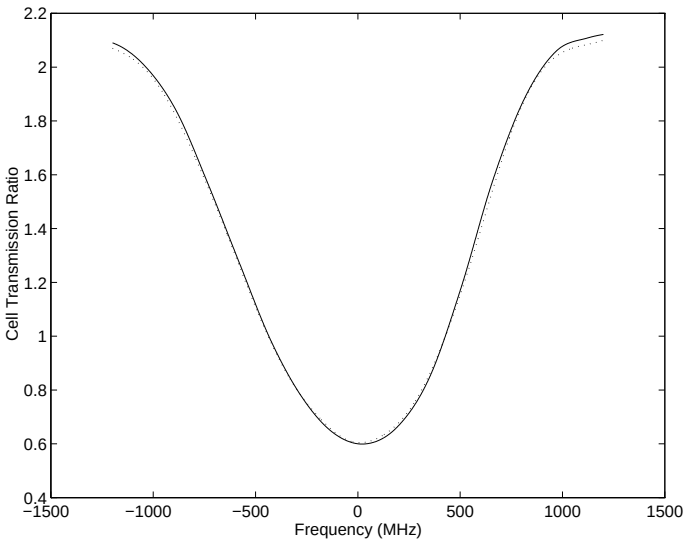


Figure 4.5: Reference cell transmission profiles for the two calibration runs. (Run 1 is represented by a dotted line, run 2 by a solid line.)

A number of calibration runs were performed at an interval of around six months. The results of a second calibration are presented here to allow a comparison to be made between the data, and hence to allow an assessment of the long term validity of the cell calibration to be made. Figs. 4.5 and 4.6 respectively allow direct comparison of the reference and

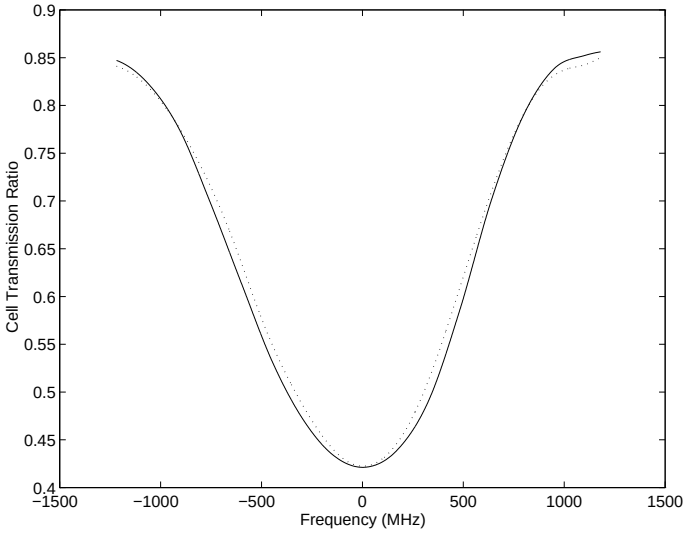


Figure 4.6: Working cell transmission profiles for the two calibration runs. (Run 1 is represented by a dotted line, run 2 by a solid line.)

working cell curves obtained in the two calibration runs.

It may be seen that the computed reference cell data from the two calibration runs represent a good fit. Indeed the RMS difference between the reference cell absorption curve obtained in the first calibration run with that obtained six months later has been found to be about 1.8%. The depth of the absorption curve increased slightly during the six month period, probably due to slight variations in temperature of the cell, resulting from seasonal ambient temperature changes or from a slight drift in the temperature controlling arrangement. Stabilisation of the ambient room temperature may be useful in reducing seasonal changes in cell performance; but these results indicate that the physical stability of the reference iodine cell, and the long term stability of the temperature stabilisation technique are sufficiently good that reference cell calibrations once computed remain valid for a period of at least six months — assuming that the physical arrangement of the illumination optics is not disturbed during that period.

Whereas the reference cell curves determined by the two calibration runs were seen to fit each other closely, the two working cell curves shown in Fig. 4.6 diverge more significantly. The reason for this is that the Doppler global velocimetry receiving optical system was changed in the period of time between the two calibration runs. It was this change in configuration — including the incorporation of three new lenses, different beam splitters and a new optical rail — that necessitated the recalibration. The changes introduced were intended to reduce vignetting, increase the efficiency of light transfer through the system and reduce optical aberrations. As part of the re-assembly procedure, the temperature controller was disconnected from the iodine cell. (In regular use, the reference and working cells are

temperature controlled at all times). As stated above, the transmission ratio determined for the working cell is also dependent on the properties of the optical components in the system. The results of Fig. 4.6 illustrate the fact that cell recalibration should be undertaken whenever the optical configuration is changed. The reference cell apparatus, in contrast, was undisturbed between the two calibration runs. The close fit found for the reference cell data, combined with the previous findings of Thorpe, are sufficient to suggest that the iodine cell calibrations, once performed, may be regarded as valid over long periods of time, provided the physical arrangement of the optical elements is not disturbed. Indeed the transmission profile of the reference cell has been seen to exhibit excellent stability. Comparison of the working cell curves obtained underlines the point that recalibration should be performed if the apparatus is reconfigured.

4.1.5 The Reference Cell Monitoring Software

During calibration, a custom-written program, `IODINE`, was used in order to display reference cell transmission values and to compute mean reference cell transmission values continuously during a CCD camera exposure period. Once calibration is complete, this `IODINE` program is employed to graphically display the laser frequency in real-time.

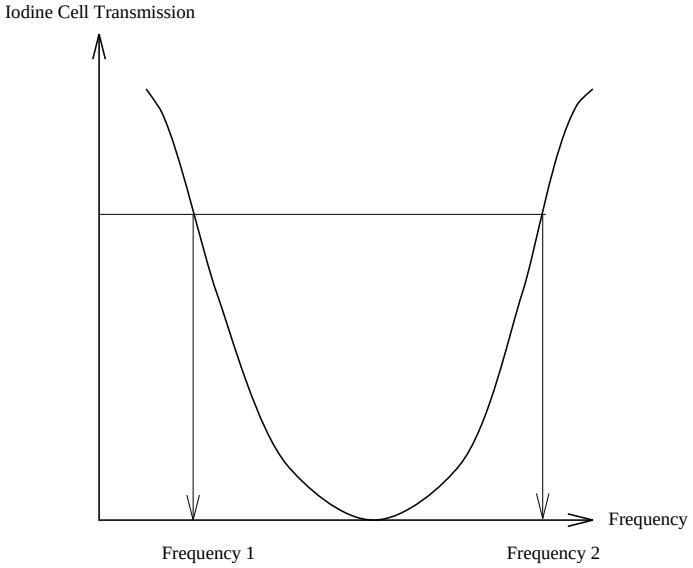


Figure 4.7: Determination of laser frequency from reference cell transmission ratio.

The `IODINE` program is run continuously whenever the Doppler global velocimetry system is in use. There are two modes of operation. In *real-time* mode, the software continuously acquires photodiode voltage data from the analogue-digital board, forms an average of a number of samples, and displays the resulting transmission ratio on the screen. The display shows a graphical representation of the reference cell transmission curve derived from the calibration data obtained as outlined in Section 4.1.1.

The `IODINE` software first loads the interpolated reference cell data from disk. Subsequently, the photodiode voltages are repeatedly sampled and a mean formed of 100 readings, to reduce the effect of measurement noise. Having compensated for known analogue-digital converter offset values, the ratio of these mean photodiode voltage values may be taken as a measure of cell transmission ratio. From this transmission ratio, it is possible to compute the laser frequency as shown in Fig. 4.7. For reasons of efficiency, a binary chop algorithm is used. For each value of transmission ratio, clearly there are a pair of possible laser frequencies. Both these frequency values are shown on the reference cell transmission curve in real-time. It is the responsibility of the experimenter to ascertain which of the frequency values is in fact the correct one. In addition to the display of the ratio and the reference cell transmission curve, three graphs are continuously updated on-screen. One of these shows transmission ratio against time, while the other two show the corresponding pair of possible frequency values.

The experimenter can adjust the laser etalon oven control in order to set the frequency to the desired portion of the iodine absorption profile, using the `IODINE` software in real-time mode to see the effect. Mode-hops and instability are immediately apparent by reference to the graph displays. This real-time indication of laser frequency has been found to be of great value both for setting laser frequency and for ensuring that the laser is operating in a stable single mode. When setting frequency prior to undertaking velocimetry measurements, the `IODINE` software real-time display enables the experimenter to readily observe the laser emission frequency. When the laser is operating in a stable single mode, the frequency display remains static, whereas mode-hopping behavior is evident as an oscillation of the displayed frequency on-screen.

The `IODINE` software also has a second mode of operation, *exposure mode*. In this latter state, the 'shutter open' signal from the CCD camera is employed to trigger the data acquisition process, just as it was for the working cell calibration. A mean transmission ratio for the period of exposure may thus be obtained. This value is subsequently required for the computation of quantitative velocity values when undertaking Doppler global flow measurements.

4.1.6 Absolute Velocity Computation

The `warp` application performs all the image processing necessary to align the discriminated and non-discriminated images accurately and to divide them so as to yield a single results image. The value of each pixel in the results image represents the working cell transmission ratio at that point. It is the task of another X-window application, `viewresults`, to take this

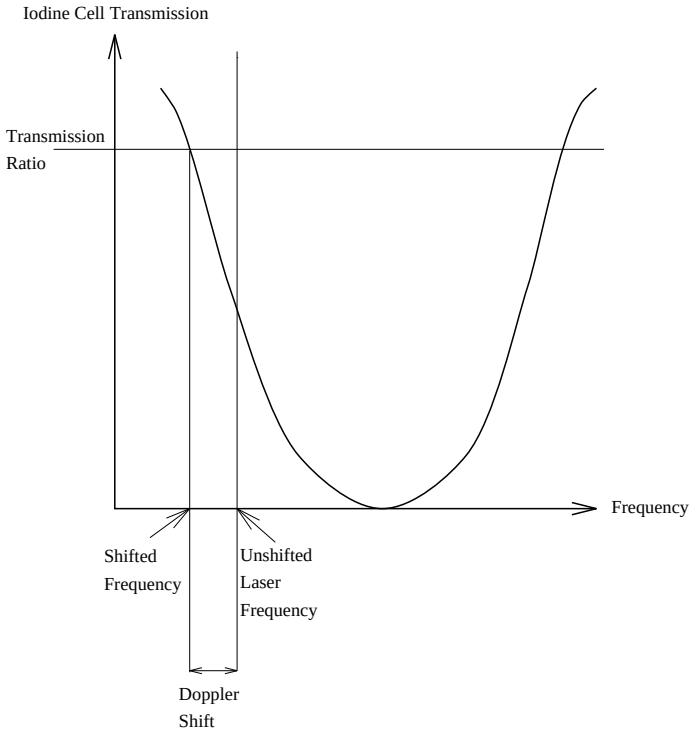


Figure 4.8: Determination of Doppler shift from working cell transmission ratio.

data and produce a true velocity representation of the flow field. In addition, `viewresults` features a number of tools useful for the visualisation of the velocity field, either in a colour-map representation or as a surface plot. In the development of the graphic visualisation tools of `viewresults`, [Foley *et al.* 1990] and [Watt 1989] were found to be useful reference works.

In order to compute absolute velocities, the `viewresults` program must be provided with three items of data. One of these, clearly, must be the results image produced by `warp`. The reference cell transmission ratio during the period of exposure as computed by `IODINE` is also required. Finally, the experimenter must inform the application which of the two possible laser frequency values is appropriate.

Initially, `viewresults` loads from disk the interpolated reference and working cell transmission curves, determined from the calibration procedures of Sections 4.1.2 and 4.1.3. The reference cell transmission ratio for the data frame is input by the user. In the same way as for the `IODINE` program, a binary-chop technique is used to determine the two possible laser frequencies in arbitrary units, using the reference cell transmission curve.

Once the unshifted laser frequency ν_o has been determined, the results image is loaded from disk. It is necessary to follow the binary-chop procedure, this time using working cell data, for each pixel in this image in order to determine the frequency value, in arbitrary units, at

every point.

In determining the Doppler-shifted frequency, it is important to consider the effect of the normalisation image which is used as described in Chapter 3, to compensate for non-uniform intensities. The normalisation image is acquired at the laser no-shift frequency ν_o. The data image (once adjusted for black level offset) is divided by this normalisation image (also similarly adjusted). It is thus the case that at any laser emission frequency, the no-shift ratio in the results image will be equal to unity. The effect the normalisation process has on transmission ratios observed in the results image must be taken into account before velocities may be resolved.

Fortunately, the effect of normalisation on the apparent transmission ratios obtained may be easily taken into account. As explained above, the laser emission frequency ν_o may be obtained from the reference cell transmission value using the `viewresults` program. The working cell curve, as determined by the calibration procedure, may be used to determine the working cell transmission at that frequency. All values in the results image may then be multiplied by this value in order to compensate for the effects of image normalisation. Thus, the effect on the resulting ratio resulting from the use of a normalisation image may be removed from the results image during processing.

Once again, it is ambiguous which wing of the absorption curve should be used. In order to resolve this ambiguity, the assumption is made that the shifted frequency occupies the same wing of the absorption curve as the unshifted frequency ν_o. As a consequence of this, care must be taken when working near the centre of the absorption curve to ensure that the Doppler shift resulting from the flow velocity is not of a value that may shift the frequency from one wing of the profile to the other.

Now that the frequency of the shifted laser emission for each point in the results image has been determined it is trivial to subtract the unshifted frequency from each point in order to obtain Doppler shift data. The fact that the original calibration points determined by the mode-hop behaviour of the argon ion laser were separated in frequency by 143 MHz enables the conversion of these Doppler shift quantities from arbitrary values into the appropriate units.

Eqn. 2.1 may now be employed in order to compute velocity from these Doppler shift values. By assuming the viewing configuration of Fig. 4.9, eqn. 2.1 may be simplified to

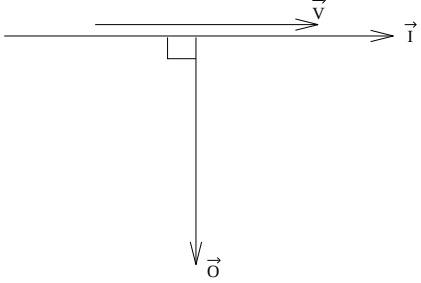


Figure 4.9: Assumed viewing configuration for velocity determination.

$$\Delta\nu = \frac{\nu_o V}{c} \tag{4.2}$$

Hence

$$V = \frac{c\Delta\nu}{\nu_o} \tag{4.3}$$

Taking account of different viewing configurations is straightforward, but is currently left to the user, rather than being automatically incorporated into the calculation.

The velocity for each point in the flow field having been computed, `viewresults` displays the velocity data as a colour-map representation. A surface representation is also shown. This surface plot may be interactively manipulated by the user and hence may be viewed from different 'observation positions'. A typical `viewresults` screen, showing the results of the measurement of the velocity of a rotating disc, is shown in Fig. 4.10.

The final program related to the processing of iodine cell data is `norm_adjust`. Unlike `warp` and `viewresults`, this is a very simple piece of code. Clearly, for satisfactory determination of quantitative velocity it is essential that the normalisation, black, and data frames are acquired at the same laser emission frequency. In some cases this may not prove possible, as the laser may prove to be unstable or may drift in frequency. The purpose of `norm_adjust` is simply to allow the black and normalisation images to be adjusted as if they had been exposed at the same laser frequency.

Suppose that a data frame is acquired at reference cell transmission value r_1 and the corresponding normalisation frame at a reference cell transmission value of r_2. Using the iodine absorption curves determined from the calibration process, it is trivial to compute, in arbitrary units, the laser emission frequency ν_o for each of the frames, and hence the corresponding working cell transmission values w_1 and w_2. The pixel values in the discriminated portion of the normalisation frame may then be multiplied by the value w_1/w_2 in order to

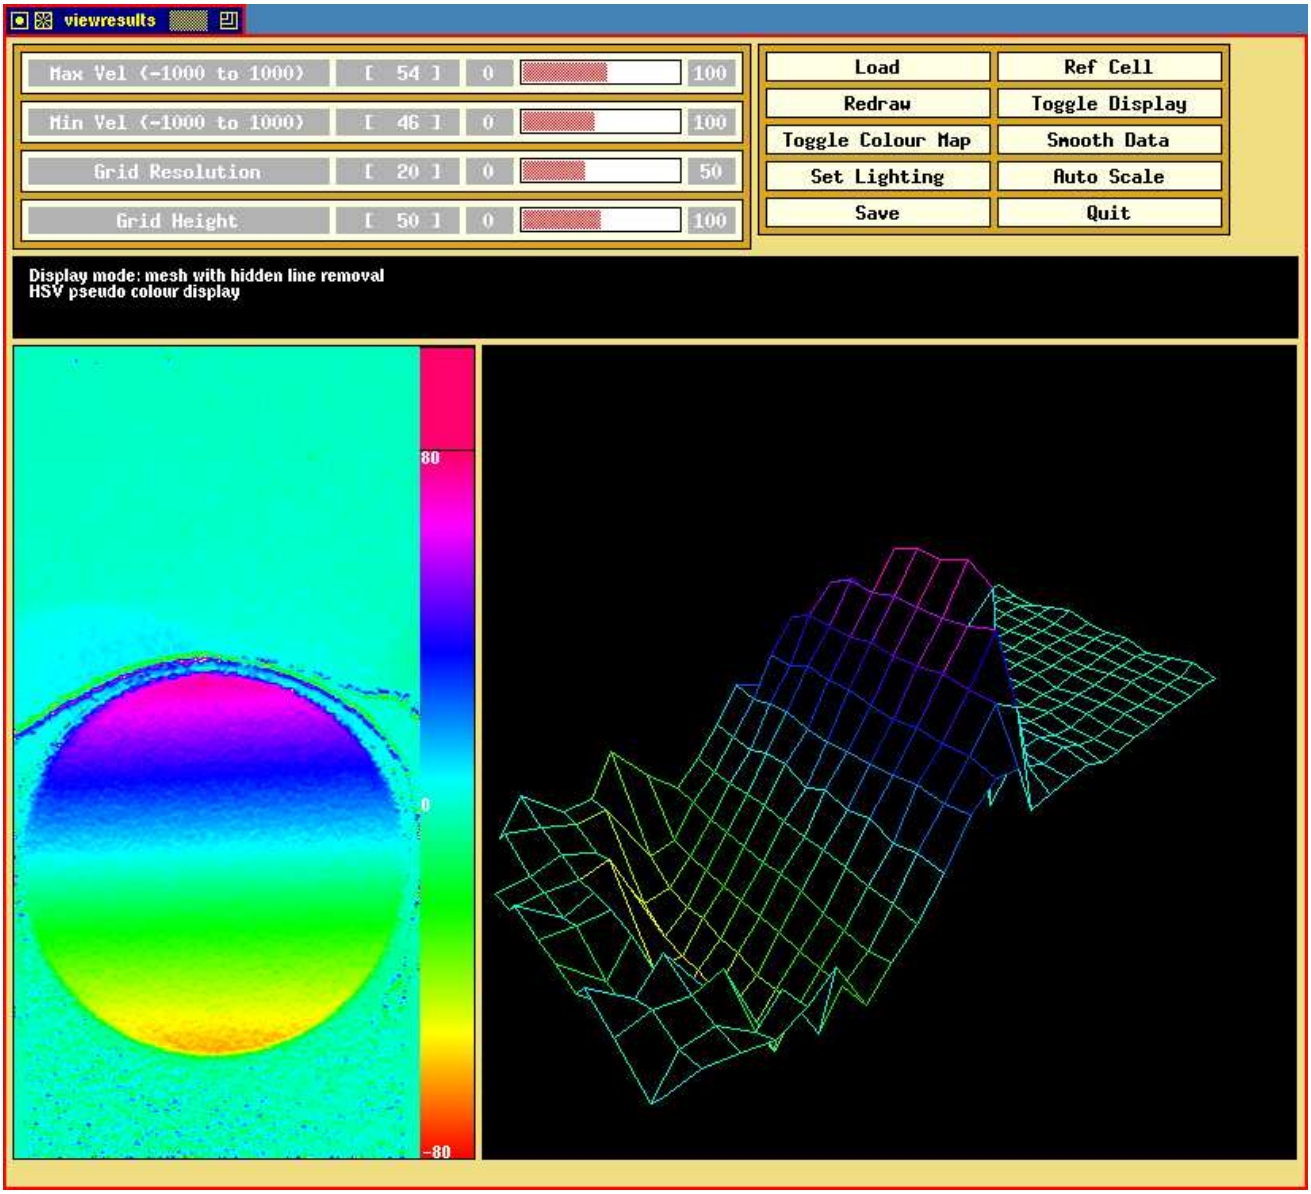


Figure 4.10: The `viewresults` display.

correct for the difference in laser emission frequencies between that frame and the data frame itself. This simple adjustment, using the reference cell transmission values for each frame provided by `IODINE`, has been found in practice to improve quantitative accuracy.

4.1.7 Summary

The acquisition of accurate quantitative velocity measurements using a Doppler global velocimetry system depends upon an accurate knowledge of the transmission of the imaging cell as a function of light frequency. In the case of the Oxford system, the reference cell transmission must also be determined. A simple technique, using the Doppler global apparatus itself and the same argon-ion laser that is used for measurements, enables both of these cells to be accurately calibrated *in situ*. Unlike other calibration techniques, this method requires no additional hardware. The calibrations are cast in the form of polynomial curve fits that may subsequently be used in the `IODINE` and `viewresults` software in order to convert transmission values into relative frequencies, and then into velocities.

4.2 Two Component Velocimetry Determination

The principles underlying two component Doppler global velocimetry methods are outlined in Section 2.9. In this section the steps which must be taken in order to obtain such two component measurements are described, together with the software which has been written to allow such results to be readily interpreted.

4.2.1 Flow Measurement Procedure

Two component measurements may be made by the use of a pair of optical systems arranged so that they are able to observe the same section of the flow field, illuminated by the laser light sheet, from different angles. It may be seen from the Doppler shift equation (eqn. 2.1) that the measured component of velocity is dependent on the direction of illumination $\hat{\imath}$ and the direction of observation $\hat{o}$. In section 2.9, two possible arrangements of optical systems that allow different orthogonal components of velocity to be obtained are considered. A third arrangement suitable for repeatable flow conditions is also discussed — this arrangement utilises a single optical system and requires that the direction of illumination $\hat{o}$ be altered instead.

Considering first the case of two optical systems, the procedure for obtaining two component measurements is as follows. A laser light sheet is arranged to cover the region of interest within the flow field. The two optical systems are then set up so that the region of the light sheet within which two component measurements are required is imaged in both systems.

The choice of observation angles should be made after careful consideration of the required accuracy of the velocity components to be resolved — an analysis of the expected accuracy of such measurements with respect to observation angle may be found in Section 5.11. The CCD cameras of each optical system should be connected in such a way that the shutters are synchronised with one another. Using the Wright Instruments cameras [Wright 1993], this is possible by means of an electrical connection on the camera body.

Just as was the case for single component Doppler global measurements, it is necessary to align the systems in order that the discriminated and non-discriminated images may be corresponded with one another spatially. In addition, for the case of two optical systems, it is necessary to correspond the images obtained using one system with those obtained from the other.

In order to effect this alignment, a test card consisting of white dots on a black background is employed. This test card is placed in the plane of the laser light sheet, and exposures are obtained using each optical system. As for single component measurements, a calibration normalisation and a calibration black frame should be obtained for each system.

The `warp` software is employed in order to compute warp vectors for each of the optical systems. As explained in Section 2.9, if care is taken to ensure that the same physical alignment features are used to compute the warp for each optical system, then not only will the discriminated and non-discriminated images be aligned for each optical system, but the two systems will also be aligned spatially with one another, so that direct comparison is possible.

Once the warp vectors have been computed, flow data may be acquired in the same manner as for single component measurements, ensuring that the shutters of the two CCD cameras are synchronised. Suitable black and normalisation images should be obtained for each optical system.

The data from each optical system may be processed independently, using the `warp` software with the appropriate pre-computed warp vectors. The iodine cell transmission profile may then be applied using the `viewresults` code. For each measurement undertaken, a pair of aligned results images will be obtained, each depicting a certain component of measured velocity. These images may be processed further as described in Section 4.2.2 below.

The alternative arrangement suitable for repeatable flows, using a single optical system but manipulating the laser light sheet propagation direction between exposures, requires a slightly different measurement procedure.

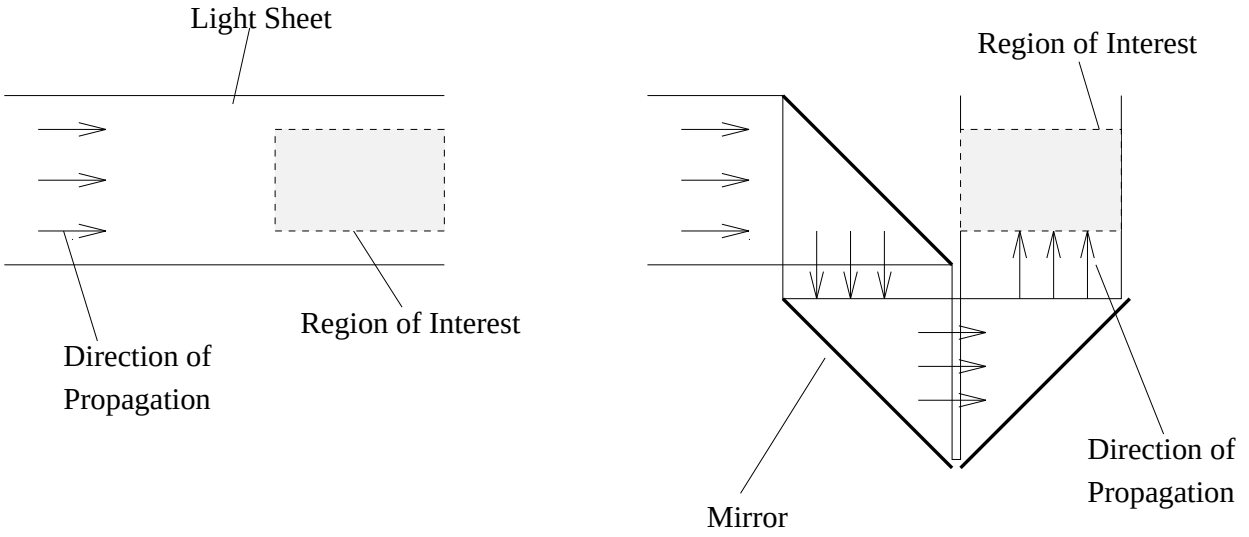


Figure 4.11: The use of mirrors to alter the light sheet propagation direction î

In this case, since there is only one optical system, there is no need to align a pair of optical systems. The acquisition of a calibration image is required only in order to align the discriminated and non-discriminated images. The two exposures required for two component measurements are made without repositioning the optical system. Thus only one set of warp vectors need be computed, as is the case with single component measurements. Image processing then follows the scheme presented in Chapter 3.

The acquisition of the first data exposure is made in the conventional way for single component work (see Chapter 3). After this exposure is complete, however, it is necessary to manipulate the laser light sheet in order that its direction of propagation may be changed. In doing so, it is important to ensure that the plane of the light sheet is not altered, but only the propagation direction. It is generally most convenient to manipulate the propagation direction by introducing mirrors into the illumination system as shown in Fig. 4.11. The second exposure is taken using the altered propagation direction.

As in the case of measurements involving two optical systems, after image processing with `warp` and `viewresults` is complete, two results images are obtained. The same software (Section 4.2.2) may be used to post-process these results in order to obtain the desired components.

In undertaking either of these two component schemes, it is important to measure the angles of observation and of illumination accurately. These angles must be specified by the user in order that the direction of the measured component of velocity is known for each data frame acquired. The angles are specified relative to the coordinate frame shown in Fig. 4.12.

4.2.2 Computation of Two Component Flow Field

Whichever of the two approaches to two component flow measurement is used, whether multiple observation directions or multiple illumination directions, the result is the same: a pair of velocity images each depicting the measured velocity component in a certain direction. Provided that the experimenter has measured the angle(s) of the illumination vector $\hat{\imath}$ and of the observation vector $\hat{o}$, these velocity images may then be combined in order to obtain a measurement of two orthogonal components of velocity. The software `2component` has been written in order to do this.

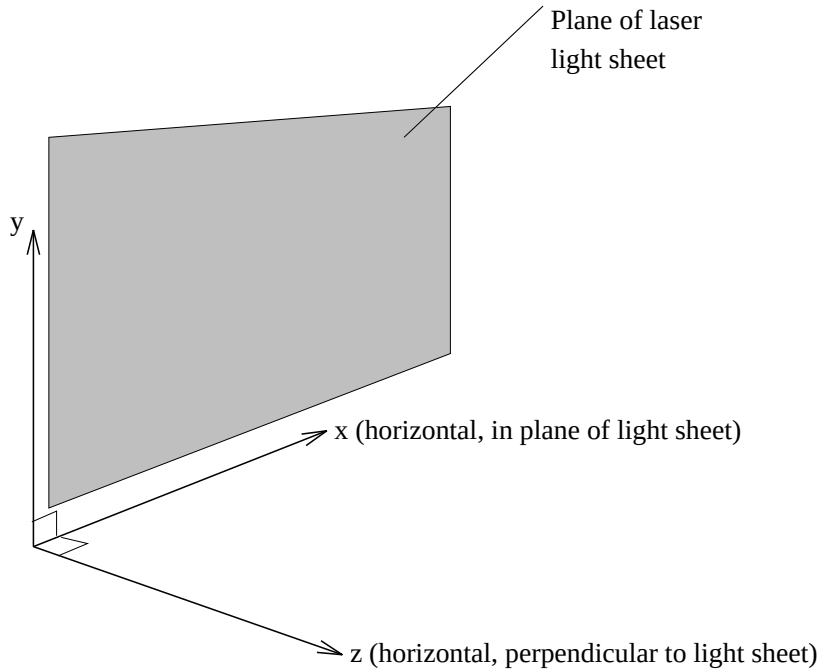


Figure 4.12: Coordinate frame used for specifying illumination and observation angles.

The program `2component` is a straightforward piece of code, written to implement the theory outlined in Section 2.9. The software loads from disk two velocity component images, produced by the `viewresults` application. The user must then specify the viewing angles and the illumination angles for each image, relative to the frame of reference shown in Fig. 4.12. Depending on the configuration used, it may be necessary to prompt the user to verify that the out-of-plane component of velocity is known to be zero. The desired orthogonal components are computed as shown in the analysis of Section 2.9. The resulting two component velocity data is then saved to disk.

An X-window application `2component_vis` may then be used to allow the two component velocity data obtained to be visualised. This code divides the field of view into a number of square cells. The mean velocity in each of the orthogonal directions within each cell may then be computed. The velocity in each cell may then be shown as a line, the angle of which indicates the velocity direction and the length of which indicates the magnitude of the velocity. Such a two component image may be seen in Chapter 6, Fig. 6.27.

4.2.3 Summary

Even the earliest publications (such as Komine and Brosnan [Komine and Brosnan 1991]) relating to the Doppler global technique were quick to recognise that two or even three component measurements were attainable. Indeed, velocity measurements of a free jet flow observed from three different positions [Komine and Brosnan 1991, Komine *et al.* 1991] have been published. No previous workers, however, have presented any information regarding the processing of such data. No attempt has previously been made to take account of the optical system observation direction and to process the resulting measured components such that representations of orthogonal velocity components may be obtained. Without suitable data processing, two or three component Doppler global measurements are of little use.

In this section, two alternative procedures for undertaking such two component measurements have been presented. Building on the theory presented in Section 2.9, the operation of the two additional pieces of software required for processing two component results, `2component` and `2component_vis`, has been explained. The problem of quantifying the error expected in two component Doppler global measurements is examined in Section 5.11. The results of applying the two component Doppler global velocimetry technique to the case of a rotating disc are presented in Chapter 6.

It is perhaps appropriate at this point to review briefly the flow of data through the various software programs which have been developed to support the Doppler global velocimetry technique. Data is acquired in the form of CCD camera frames, measured viewing and observation angles, and iodine transmission profiles. This data is processed using the `warp` application in order to effect spatial alignment of the images; the `viewresults` program for computing quantitative velocity values and allowing the visualisation of single component data; the `IODINE` program for determining laser emission frequency values; the `norm_adjust` program for correcting normalisation images acquired at incorrect emission frequency values; and the `2component` and `2component_vis` programs for computing and visualising two component measurements. The flow of data through this system of applications is shown in Fig. 4.13.

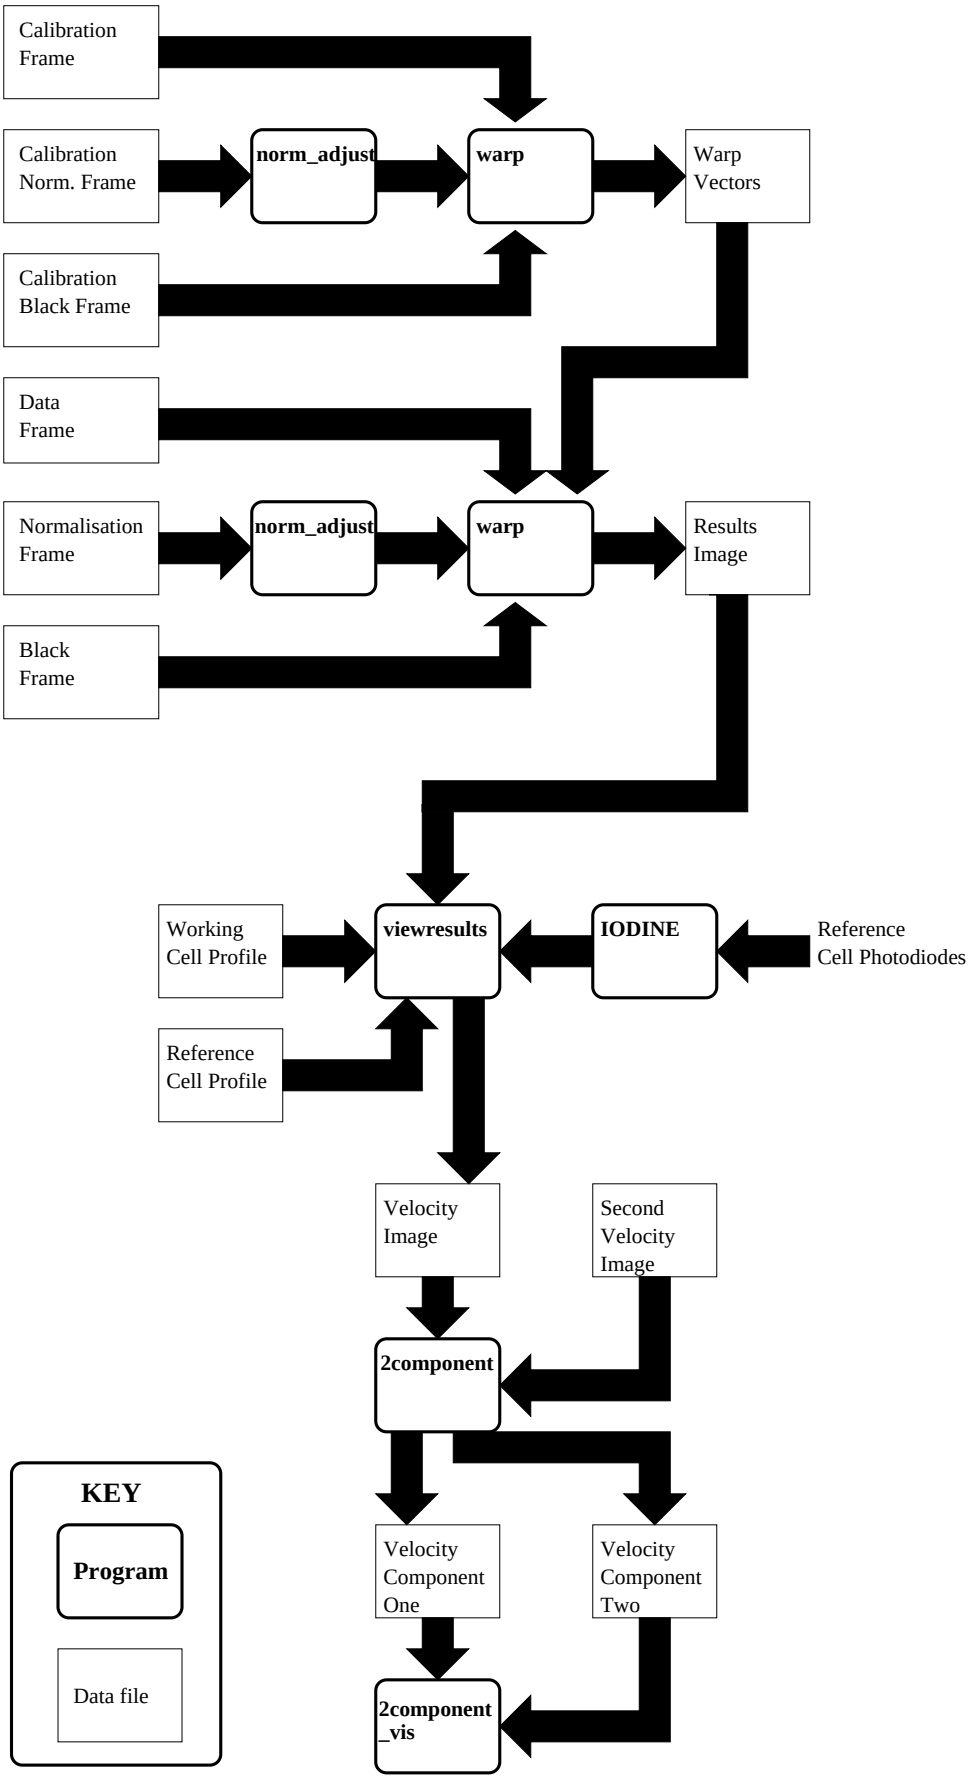


Figure 4.13: Overview of Doppler global velocimetry software and the flow of data.

Chapter 5

Error Sources and Analysis

5.1 Overview of Doppler Global Velocimetry Error Sources

The existing literature, aside from those publications derived from the current work, makes little reference to those factors that influence the accuracy of the quantitative velocity measurements that may be obtained using the Doppler global velocimetry system. No discussion of expected measurement uncertainty is presented for the video camera based systems such as those of Komine *et al.* [Komine and Brosnan 1991, Komine *et al.* 1991], Meyers [Meyers 1992] and Ford and Tatum [Ford and Tatum 1994]. Although Roehle and Schodl [Roehle and Schodl 1994] considered the errors associated with the determination of laser frequency during a measurement and the determination of the iodine cell transmission calibration, they applied these estimates only to a point-wise realisation of the Doppler global technique. Accordingly, the analysis of Roehle and Schodl fails to consider the uncertainties likely to be introduced into global methods due to the use of an imaging system. Rather, they concentrate on using an argon-ion laser to illuminate a seeded round free jet, placing much emphasis on the stabilisation of the laser emission frequency. In their video camera based Doppler global velocimetry systems, Meyers [Meyers 1992] and Komine *et al.* [Komine *et al.* 1991] report a qualitative comparision with laser Doppler anemometry results, but offer no assessment of Doppler global measurement accuracy. Meyers was able to conclude, however, that image processing of the acquired discriminated and non-discriminated images was necessary in order to obtain sufficiently good alignment to ensure satisfactory quantitative measurement accuracy. Subsequent work focussed on the development of real-time video based systems, such as that of Morrison, Gaharan and DeOtte [Morrison, Gaharan and DeOtte 1994] and Ford and Tatum [Ford and Tatum 1994], has not reported the implementation of any such schemes.

A thorough investigation of the primary sources of error in the Oxford system has been conducted. This work has enabled the Doppler global velocimetry measurement error-bounds for each pixel to be quantified. The analysis of error sources has provided a valuable indication of where improvements in Doppler global technology would be most useful in order to improve quantitative accuracy.

The most significant sources of measurement error in the Oxford system have been identified as:

- Camera pixel shot noise.
- Iodine absorption filter transmission calibration.
- Determination of the laser emission frequency ν_o.
- Iodine absorption filter instability.
- Adverse reflections of the laser light sheet.
- Misalignment of discriminated and non-discriminated images.

These individual error sources may be examined briefly in turn before a more detailed analysis is embarked upon later in this chapter.

As stated in Section 2.7, even in the best case, the digitised pixel grey levels in the measurement image have an inherent noise level which is shot noise limited. The use of a cooled slow-scan camera has enabled the error in recorded pixel values to approach the theoretical shot noise limit. The effect of pixel noise on measured velocity values is assessed in Section 5.2 below.

Errors in the calibration of the laser emission frequency are introduced both by temperature fluctuations in the laser resonator, which lead to a drift in output emission frequency of approximately 60 MHz/^{o}C, and by variations in the longitudinal mode spacing of the laser, used to define the light frequency during the calibration procedure as discussed in Section 4.1.1, which is specified by the laser manufacturer as 143 MHz, with an uncertainty of $\pm$0.4 MHz.

During the measurement procedure, the laser emission frequency ν_o must be known relative to the fixed iodine absorption feature, as this represents the no-shift or no-velocity frequency relative to which fluid velocities are determined. Failure to determine this frequency accurately will yield an offset in the Doppler shift frequency determined for each measurement pixel. It was shown in Section 2.5 that since the iodine vapour pressure exhibits a strong temperature dependence, the transmission of the iodine absorption filters themselves is acutely temperature sensitive. Drift in the iodine cell properties will effect the determination of ν_o (reference cell) and also of the individual pixel Doppler shift values $\Delta\nu$ (working cell).

Those aspects relating to the use of an iodine cell as a frequency-to-intensity converter are considered in Sections 5.3 and 5.4.

Eqn. 2.1 is an expression for Doppler shift in terms of the incident light propagation direction, the observation direction, and the velocity vector itself. Implicit in the use of Doppler shift for Doppler global velocimetry measurements is the assumption that the observation and incident light propagation vectors, $\hat{o}$ and $\hat{\imath}$ are known. The situation is altered, however, when reflections of the light sheet occur, changing the vector $\hat{\imath}$. This may be a problem when measuring enclosed flows such as the convergent-divergent nozzle flow of Chapter 8. Although errors due to reflection of the light sheet will be discussed in Section 5.5, it is not convenient to place a quantitative value on such inaccuracy, since it depends on the geometry of the measurement itself and the reflectivity of any boundary which the light sheet strikes.

The process of aligning the discriminated and non-discriminated images can lead to significant velocity error. The magnitude of this error component is extremely sensitive to the presence of large intensity gradients in the images, which may cause image misalignment to be the most significant source of velocity measurement error. The precision of alignment achieved is limited by the accuracy with which the position of alignment features may be determined in the calibration frame. Even in the absence of such alignment feature centroid estimation error, the subsequent resampling of the images onto an aligned pixel grid may give rise to some degree of error in the resulting measured velocity field. In Section 5.6 the effect of misalignment on velocity determination is considered. Sections 5.7, 5.8 and 5.9 discuss the impact on misalignment of centroid estimation, image transformation and interpolation techniques respectively.

Estimates of the magnitude of the errors associated with fundamental measurement uncertainties are presented below. These individual error sources may be combined into an overall assessment of the consequent error in measured Doppler shift frequency at each pixel within the images. The error sources are determined as RMS quantities, and, following the work of Coleman and Steele [Coleman and Steele 1995], may be combined to form the overall estimate of the Doppler shift frequency error

$$\varepsilon_{\Delta\nu} = \sqrt{\sum_{i=1}^{n} \varepsilon_i^2} \tag{5.1}$$

where the ε_i terms on the right hand side represent the error sources discussed above.

The expected error estimates made in this chapter allow an estimated measurement error to be computed for every pixel in the velocity image. An error model has been implemented in software, as explained in Section 5.12. This code may be used in conjunction with acquired Doppler global velocimetry data, allowing the estimated error to be compared with that observed in practice.

5.2 Camera Pixel Shot Noise

The digitised value recorded by a CCD camera is affected by three sources of noise: thermal noise, read-out noise and shot noise.

Thermal noise (also described as 'dark current') is a leakage of electrons into the potential wells that form the pixels of the CCD array. This type of noise is highly sensitive to the temperature of the CCD chip itself. Published results [EEV 1987] give

$$\textbf{dark current} = Ae^{\frac{-V_{BG}}{2kT/q}} \tag{5.2}$$

where T is the operating temperature of the CCD chip.

At room temperature, dark current is typically of the order of 10^5 electrons per pixel per second. Such a value of dark current would saturate the CCD device in a few seconds, full well capacities usually being of the order 10^5 electrons. At a CCD operating temperature of 200 K, as used by the Peltier-cooled camera in use in Oxford, the dark current drops to approximately 1 electron per pixel per second, which may be considered negligible by comparison with the level of shot noise for typical exposure durations. Thus, exposures up to several minutes in length may be made without risking the introduction of significant thermal noise errors.

CCDs suffer from defects in the manufacturing process, seen as variations in the properties of individual pixels, such as the full-well capacity. This may lead to a non-uniform pixel sensitivity across the CCD array. As described in Chapter 3 this effect is compensated for by the use of a normalisation image.

Read-out noise is the error associated with converting the accumulated pixel charge (which may be up to 640,000 electrons for the TK-512 CCD chip) into an equivalent voltage. Subsequently, this voltage must be digitised, leading to further error. Read-out noise is known to increase at high read-out rates [Astromed 1992], such as those found in video-rate frame-grabbers. By reducing the read-out rate of the pixel voltage levels, the error is minimised.

The camera employed at Oxford uses such a technique.

The reduction in read-out noise achieved by the use of slow read-out speeds allows an increased number of bits to be employed in the digitisation process. Video-rate frame-grabbers usually use 8 bit precision, whereas the frame-grabber in use for the Oxford system employs 15 bits per pixel. Read-out noise in the Oxford system is typically 30 electrons RMS [Wright 1993], representing 0.005% of the full-well capacity of the pixels. This may be considered negligible compared with shot noise, which is discussed below. Meyers [Meyers 1992] and Komine *et al.* [Komine *et al.* 1994], who both report the use of video-rate camera technology, suffer measurement noise of 0.5% under optimum measurement conditions. Ford and Tatum [Ford and Tatum 1994] have reported computing the mean of 45 video frames in order to reduce pixel noise for the measurement of the velocity of a rotating disc.

For the camera employed at Oxford, the fundamental limit on pixel measurement accuracy is shot noise, an inherent statistical noise. It is well known [Astromed 1992] that the RMS shot noise associated with the measurement of a number of electrons in a pixel may be expressed

$$\varepsilon_{N_e} = \sqrt{N_e} \tag{5.3}$$

where N_e is the number of electrons accumulated in the pixel.

For cooled, slow-scan cameras, shot noise is the only significant contributor to camera pixel noise. For a digitised intensity I, measured in terms of analogue-digital converter counts, the number of electrons in the pixel may be expressed

$$N_e = \frac{I}{I_{\mathbf{max}}} N_{\mathbf{fw}} \textbf{ electrons} \tag{5.4}$$

where $N_{\mathbf{fw}}$ represents the maximum number of electrons or full-well capacity of the pixel and $I_{\mathbf{max}}$ is the corresponding digitisation level — 32767 for 15 bit analogue-digital conversion. The shot noise associated with measuring I is thus

$$\varepsilon_{I_{\mathbf{shot}}} = \left(\frac{I}{I_{\mathbf{max}}} N_{\mathbf{fw}}\right)^{\frac{1}{2}} \frac{I_{\mathbf{max}}}{N_{\mathbf{fw}}} = \left(\frac{I I_{\mathbf{max}}}{N_{\mathbf{fw}}}\right)^{\frac{1}{2}} \textbf{ grey levels} \tag{5.5}$$

The pixels of the TK-512 CCD chip have a full-well capacity of 640,000 electrons, giving a minimum shot-noise level of 0.125% under full-well conditions.

When considering the effect of shot noise, it is worth recalling from Chapter 3 that interpolation must be performed as a final stage of the image processing procedure. The values at a single pixel in the results image, then, are derived from a number of pixels in the CCD array itself. The two-dimensional interpolation may thus be considered to be a weighted summation of the pixel values surrounding the point where the interpolated value is required. A number of interpolation schemes have been discussed, ranging from a simple bilinear scheme through to more complex cubic and exponential methods.

For the case of the nearest neighbour technique, the output pixel value is always set to the value of the single nearest source image pixel. The RMS value of any CCD measurement noise will be unaffected by this technique, and will be passed through to the interpolated image.

Bilinear interpolation is more interesting in this respect. Each output pixel is a linear combination of four source pixels. Using the interpolation parameters $0 \leq \alpha \leq 1$ and $0 \leq \beta \leq 1$ and taking the input pixels x_1, x_2, x_3, and x_4 to produce an output pixel y

$$y = \alpha(\beta x_1 + (1-\beta)x_2) + (1-\alpha)(\beta x_3 + (1-\beta)x_4) \tag{5.6}$$

For each input pixel, the RMS shot noise is expected to be the same. Thus, $\mathrm{var}(X_i) = \mathrm{var}(X)$ for all i. It is well known that for a constant k, $\mathrm{var}(kX) = k^2\mathrm{var}(X)$. Hence

$$\mathrm{var}(Y) = (\alpha^2\beta^2 + (\alpha - \alpha\beta)^2 + (\beta - \alpha\beta)^2 + (1-\alpha)^2(1-\beta)^2)\mathrm{var}(X) \tag{5.7}$$

The worst case is clearly when the position of the output pixel coincides with the position of one of the input pixels, x_1, x_2, x_3 or x_4. In this case, the bilinear interpolation scheme reduces to nearest neighbour. The best case, however, is when the output pixel lies in the middle of the four input pixels. At that point, $\alpha = \beta = 0.5$. For that case, eqn. 5.7 gives

$$\mathrm{var}(Y) = \frac{\mathrm{var}(X)}{4} \tag{5.8}$$

Thus the standard deviation of the noise present after bilinear interpolation has been performed is, at best, half that of the source pixels, and, at worst, the same as that of the source pixels.

Since the cubic and exponential interpolation schemes represent a weighted mean of a greater number of CCD camera pixels, it may be expected that the effect of CCD camera shot noise may be diminished further by their use. However, as the number of points considered by the interpolation scheme increases, so does the number of calculations required at each point. Numerical error may become a factor here. This is best investigated by means of a numerical simulation employing the interpolation code used in the `warp` application. All the numerical simulations presented in this chapter have been implemented either in Matlab code [Mathworks 1993] or, for more expensive computations, in the **C++** language [Stroustrup 1991].

The simulation technique proceeds as follows. An array of regularly spaced points representing the input pixels is set up. Each of these 'pixels' is set to a value of zero then a value for the expected noise, sampled from a Gaussian distribution representing noise of known RMS value, is added. Interpolation may be performed.

It should be noted that the noise (which has been seen to be dominated by shot noise) should really be sampled from a Maxwell-Boltzmann distribution rather than a Gaussian. The use of a Gaussian distribution is more convenient, however, and may be seen to make little difference to the validity of the simulation when the shot noise amounts to a large number of electrons per pixel, as in the case for our CCD array.

The simulation iterates through a number of values of offset (α and β above) in order to cover all cases. For each (α,β) pair, the bilinear, cubic and exponential techniques are applied in turn. The RMS error of the values of the output pixels computed by the four techniques may then be computed. It is then trivial to compute the RMS error of the output pixels divided by the RMS error of the input pixels. This dimensionless coefficient, k_i, may be used as a measure of noise-immunity for the four schemes.

The code may be set for a number of iterations and the noise values of the input pixels recomputed with each iteration. In this way a large number of iterations may be performed in order to overcome any statistical irregularity in the noise values computed.

After running for 1000 iterations, the coefficients k_i of noise-immunity were obtained as follows.

Bilinear	$k_i = 0.58$
Cubic	$k_i = 0.73$
Exponential	$k_i = 0.79$

For the bilinear method it was analytically determined that the best case value was 0.50, and the worst case 1.00. The value of 0.58 suggested by the simulation lies within these bounds. The most interesting result is that the cubic and exponential schemes, rather than providing improved immunity to CCD measurement noise, are inferior in this regard to the bilinear method. This suggests that numerical errors resulting from limited machine precision are more significant in the case of these complex techniques.

While the shot noise exhibited by the CCD array of the Oxford camera has been assessed, and the effect of interpolation algorithms examined, there is another aspect of the image processing technique which must be considered. Each pixel in the discriminated or non-discriminated images is derived from a number of CCD frames. In Chapter 3 the method of processing data frames, black frames, and normalisation frames was described. The intensity of the data, normalisation, and black frames may be denoted as I_{data}, I_{norm} and I_{black} respectively. The intensity at each pixel in the discriminated and non-discriminated images may then be described by

$$I = \frac{I_{\text{data}} - I_{\text{black}}}{I_{\text{norm}} - I_{\text{black}}} \tag{5.9}$$

It must be borne in mind that each of I_{data}, I_{norm} and I_{black} is subject to shot noise. The resulting shot noise present in the adjusted data frame may be denoted

$$\varepsilon_{I_{\text{adjdata}}} = \left(\varepsilon^2_{I_{\text{data}}} + \varepsilon^2_{I_{\text{black}}}\right)^{\frac{1}{2}} \tag{5.10}$$

Similarly, for the adjusted normalisation frame

$$\varepsilon_{I_{\text{adjnorm}}} = \left(\varepsilon^2_{I_{\text{norm}}} + \varepsilon^2_{I_{\text{black}}}\right)^{\frac{1}{2}} \tag{5.11}$$

It is necessary to consider the effect of dividing the adjusted data frame by the adjusted normalisation frame. Suppose that the two pixels under consideration may be represented as having nominal values of I_1 and I_2, with shot noise $\varepsilon_{I_{s_1}}$ and $\varepsilon_{I_{s_2}}$ respectively. The ratio R may then be expressed

$$R = \frac{I_1 + \varepsilon_{I_{s_1}}}{I_2 + \varepsilon_{I_{s_2}}} \tag{5.12}$$

So, utilising a series expansion,

$$R = \frac{I_1 + \varepsilon_{I_{s_1}}}{I_2}\left(\frac{1}{1+\frac{\varepsilon_{I_{s_2}}}{I_2}}\right) = \frac{I_1 + \varepsilon_{I_{s_1}}}{I_2}\left(1 - \frac{\varepsilon_{I_{s_2}}}{I_2}\right) + O\left(\frac{\varepsilon_{I_{s_2}}^2}{I_2^2}\right) \tag{5.13}$$

Taking the second order terms to be negligible, the expression for ratio can be simplified to the form

$$R \approx \frac{I_1}{I_2} - \frac{I_1}{I_2^2}\varepsilon_{I_{s_2}} + \frac{\varepsilon_{I_{s_1}}}{I_2} \tag{5.14}$$

In this expression, $\frac{I_1}{I_2}$ represents the ratio in the absence of any noise; the remaining terms represent errors associated with the noise itself. The RMS error in the ratio value may thus be expressed

$$\varepsilon_R = \left(\left(\frac{\varepsilon_{I_{s_1}}}{I_2}\right)^2 + \left(\frac{I_1\varepsilon_{I_{s_2}}}{I_2^2}\right)^2\right)^{\frac{1}{2}} \tag{5.15}$$

This expression may be used in conjunction with the intensity values of the adjusted data image, the adjusted normalisation image, and their respective RMS noise estimates, to derive an RMS noise estimate for each pixel in the discriminated and non-discriminated images.

Thus, obtaining the values of $\varepsilon_{I_{adjdata}}$ and $\varepsilon_{I_{adjnorm}}$ from eqns. 5.10 and 5.11 respectively, the following expression for the RMS error expected in the discriminated and non-discriminated images may be obtained

$$\varepsilon_{I_s} = \left(\left(\frac{\varepsilon_{I_{\text{adjdata}}}}{I_{\text{adjnorm}}}\right)^2 + \left(\frac{I_{\text{adjdata}}\varepsilon_{I_{\text{adjnorm}}}}{I_{\text{adjnorm}}^2}\right)^2\right)^{\frac{1}{2}} \tag{5.16}$$

It is now necessary to determine how the error introduced into the discriminated and non-discriminated images by CCD array shot noise will effect the resulting ratio of discriminated to non-discriminated image, and hence the velocity determined by the Doppler global system. The ratio R of discriminated to non-discriminated image values at each point may be expressed

$$R = \frac{I_d + k_i\varepsilon_{I_{s_d}}}{I_n + k_i\varepsilon_{I_{s_n}}} \tag{5.17}$$

where I_d and I_n represent the discriminated and non-discriminated intensities at the point under consideration uncorrupted by noise, while $\varepsilon_{I_{s_d}}$ and $\varepsilon_{I_{s_n}}$ represent the shot noise in the two images at that point determined from eqn. 5.16. The constant k_i represents the effect on the resultant noise of the choice of interpolation scheme, as discussed above.

Eqn. 5.15, which was employed above to determine the effect of shot noise on the discriminated and non-discriminated images, may be used once more to yield an expression for the error in the resulting ratio

$$\varepsilon_R = \left(\left(\frac{k_i \varepsilon_{I_{s_d}}}{I_n} \right)^2 + \left(\frac{I_d k_i \varepsilon_{I_{s_n}}}{I_n^2} \right)^2 \right)^{\frac{1}{2}} \tag{5.18}$$

It may be seen that the shot noise present at each pixel in the CCD frame affects the data frame, the normalisation frame, and the black frame alike. An expression for the resulting noise introduced into the discriminated and non-discriminated image has been obtained in the form of eqns. 5.10, 5.11, and 5.16 above. Having obtained a value for the noise at each pixel in the discriminated and non-discriminated images, eqn. 5.18 may be employed to yield an estimate of the effect of the CCD array shot noise on the final ratio obtained once image processing is complete.

Having derived an expression for the expected shot noise at each pixel in terms of CCD array grey level counts, an expected value of ratio error and hence velocity error may also be deduced. The analysis outlined above has been implemented in such a way as to allow the velocity error expected due to the presence of CCD camera shot noise to be estimated for each pixel of the measured image. In Chapters 6, 7, and 8 these error estimates are presented together with the associated data images.

5.3 Iodine Cell Transmission Calibration

The calibration method described in Section 4.1 is used to determine the transmission fraction of the working and the reference iodine cells as a function of the laser emission frequency. The iodine cells themselves are stabilised by use of a temperature controlled cold finger, as described in Section 2.5. In this section it shall be assumed that the iodine cell transmission function is invariant during the calibration run; the actual stability of the cells is considered in Section 5.4. In order to obtain a complete numerical model of expected error across the CCD frame it is essential to consider all sources of such error. Much of the work on the error expected due to the behaviour of the iodine cell is derived from that presented by Thorpe [Thorpe 1996]. Although a complete error estimate incorporating the iodine cell calibration

error is developed in the course of this chapter, it may be useful as a rule of thumb to bear in mind that a typical value of the frequency-to-velocity conversion factor is 2 MHz/ms^{-1}.

There are three factors which may introduce an error into the calibration of the iodine cell transmission functions:

- Uncertainty in the longitudinal mode spacing of the laser.
- Thermally-induced drift in laser frequency.
- The use of a curve-fit representation in data processing.

The determination of the laser emission frequency during the calibration run suffers due to the first two sources of error listed above. The calibration technique relies on the use of a known longitudinal mode spacing to define frequency. This separation is specified as 142.5 – 143.3 MHz [Spectra-Physics 1987], implying an uncertainty of ± 0.4 MHz. In addition to this uncertainty, ambient temperature variations change the laser cavity length slightly, leading to a possible drift in emission frequency. The quartz rod resonator, which determines mirror separation, minimises such effects by virtue of its low coefficient of thermal expansion. The thermal properties are such that they result in a change of frequency of 60 MHz/°C. After three hours of continuous laser operation, the typical drift in emission frequency of a particular longitudinal mode has been experimentally determined to be 1 MHz/minute.

The curve fit representation of the iodine cell transmission calibration data has been found to have an RMS error of 2.9 MHz with respect to the measured data points. This value is assumed to be typical of the curve-fit error at all points in the calibration data, and is incurred each time the curve-fit is used for the determination of Doppler shift frequencies in the data analysis.

A period of approximately 12 minutes is required in order to perform the mode-hop calibration technique, during which time the laser frequency modulates across the bandwidth of the iodine absorption profile. The thermally induced drift in the laser frequency during this period is estimated to be 12 MHz. Since the mode spacing of the laser is nominally 143 MHz and the iodine absorption profile has a bandwidth of approximately 2100 MHz, 15 or so laser modes are required to calibrate the profile. The aggregate error from mode separation uncertainty is thus $0.4 \times 15 = 6$ MHz. It may be concluded that the overall error in frequency is $6 + 12 = 18$ MHz. This error is incurred across a total calibration bandwidth of 2100 MHz, implying an error of one part in 117.

This same level of uncertainty applies to the determined Doppler shift frequency $\Delta\nu$, such that this error may be expressed as

$$\pm \frac{\Delta\nu}{117} \textbf{ MHz} \tag{5.19}$$

This represents an error of less than 1% in the frequencies measured by the system.

The laser emission frequency is seen to vary typically by 8 MHz RMS over a characteristic period of 5 ms, due to the effect of vibrations induced by the cooling water flow. In principle, such a frequency variation could cause a significant uncertainty in the velocity measurements obtained from the system. In practice, however, the measurements of the reference iodine cell photodiode voltages and the integration time of the CCD camera are sufficiently long, being of the order of seconds, so as to smooth out this jitter. The measurements of light frequency using both the reference iodine cell and the working cell consequently represent the average light frequency over these relatively long time-scales.

In processing Doppler global velocimetry results, the reference iodine cell is used to determine the laser emission frequency ν_o during the measurement. The light frequency $\nu_o+\Delta\nu$ incident at each pixel is then computed from the intensity measured at that pixel, using the working cell calibration data. The effects of thermal drift and longitudinal mode spacing separation uncertainty clearly affect the determination of the Doppler shift frequency. The error relating to the curve-fit accuracy of the iodine cell calibration is incurred each time the calibration equations are utilised. The error in the determined value of ν_o can thus be expressed as $\varepsilon_{\nu_o} = 2.9$ MHz.

The error in determining the Doppler shift frequency is affected by ε_{ν_o}; and is also dependent on the effects of the uncertainty in laser mode separation, the drift in the laser frequency during the calibration run, and the curve-fit error in the image recording system iodine cell. The magnitude of the mode separation and drift errors have been estimated previously to be $\Delta\nu/117$. Thus the error in the light frequency determined for each pixel may be expressed

$$\varepsilon_{\nu_o+\Delta\nu} = \left(2.9^2 + \left(\frac{\Delta\nu}{117}\right)^2\right)^{\frac{1}{2}} \tag{5.20}$$

The Doppler shift frequency associated with each pixel is calculated by subtracting ν_o from the pixel values of $\nu_o + \Delta\nu$. Hence the error in calculated Doppler-shift frequency may be found by combining the errors associated with these values

$$\varepsilon_{\Delta\nu} = \left(2.9^2 + 2.9^2 + \left(\frac{\Delta\nu}{117}\right)^2\right)^{\frac{1}{2}} \tag{5.21}$$

This equation is employed to form an estimate of the measurement error at each pixel caused by the error in the iodine cell calibration. Such estimates are calculated and presented for each pixel of the velocity measurements in Chapters 6, 7, and 8.

5.4 Iodine Absorption Filter Instability

The previous section considered errors encountered in the calibration of iodine cells. The stability of the transmission function of the iodine absorption filters employed in the Oxford Doppler global velocimetry system is of crucial importance in obtaining accurate quantitative measurements of fluid velocity. The filters act as transducers, converting variations in light frequency into variations in intensity which may readily be measured. If the operating conditions of the cells change between the time that their transmission function is determined and the time when a measurement is conducted, the relative light frequencies calculated will be subject to error, leading directly to an error in the determined velocity. The operating condition of the cells is determined by the temperature of the cold-finger, which establishes the cell vapour pressure; and by the ambient temperature, fixing the temperature of the body of the cell and hence the vapour temperature. Both of these factors affect the density of the iodine vapour within the cell and thereby influence its transmission, as discussed in Section 2.5.

The iodine cell's instability has two effects upon the accuracy of Doppler global velocimetry measurements. Instability in the reference cell corrupts the measurement of the laser emission frequency ν_o. Instability in the imaging system iodine cell leads to an incorrect calculation of the light frequencies determined for the data from the pixels within the measurement frames. These two errors have a direct impact on the determination of Doppler shift frequencies, and hence of velocity.

A strategy has been developed whereby the measured iodine cell temperature instabilities may be converted into estimates of the corresponding error in measured Doppler shift. As shown in Section 2.5, the known vapour pressure of iodine as a function of temperature may be used to calculate the density of iodine vapour within the cells. The density is calculated from the ideal gas law.

Having obtained a calibration by the methods of Section 4.1.1, the effect of small changes in both cold-finger and ambient temperatures may be estimated using the vapour pressure

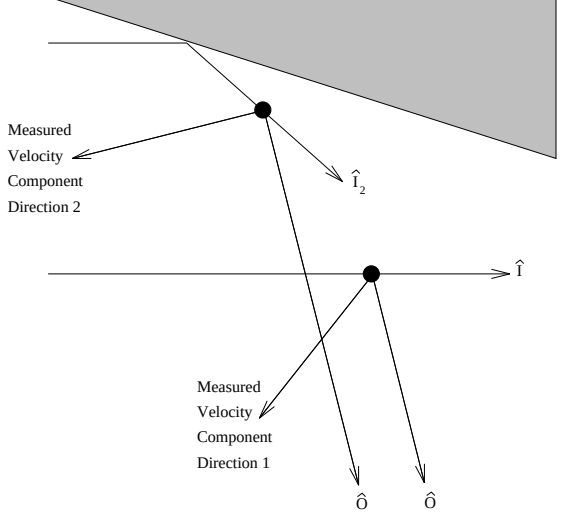


Figure 5.1: The effect of scattering from side walls.

equation and Beers Law. Such calculations reveal that the sensitivity of the cells to cold finger temperature is such that the light frequency deduced from the cell will be in error by approximately 25 MHz/K. Similar calculations based on varying ambient temperature show an error of approximately 1 MHz/K.

The measured variation in cold finger temperature of 0.003 K (Section 2.5) indicates that an error of 0.075 MHz is introduced into the overall measurement accuracy by the cold-finger temperature control performance. A typical ambient temperature variation of ± 1.5 K indicates a further error of ± 1.5 MHz. Thus, the instability in the transmission characteristics of the iodine cells caused by variations in the temperature of the cold-finger is sufficiently low that it may be regarded as negligible in comparison to the effect of ambient temperature.

5.5 Adverse Reflections of the Laser Light Sheet

The use of eqn. 2.1 to determine velocity in terms of the measured Doppler shift frequency assumes that the observation vector $\hat{o}$ and the illumination vector $\hat{i}$ are both known. In the case of the measurement of enclosed flows, or flows which are in close proximity to a physical boundary of some sort, it may be the case that the laser light sheet impinges on a solid boundary. A proportion of the light sheet will be reflected from such a boundary and may then propagate through the region to be measured. The propagation vector $\hat{i}_2$ of such reflected light will be different to the assumed propagation vector $\hat{i}$ which is used in velocity calculation.

The light scattered from a boundary may interact with seeding particles in the flow and be re-scattered, Doppler shifted. The Doppler shift of this re-scattered light will be different to that expected for the supposed illumination direction. Thus, close to any boundaries where

the intensity of scattered light from the walls themselves is highest, the mean Doppler shifted frequency will not be as expected, leading to velocity determination errors in this region.

The situation is illustrated by Fig. 5.1. The incident light sheet propagation direction $\hat{\imath}$ results in velocity component 'one' being measured. Incident light scattered from the side wall, on the other hand, may be diverted so that its direction of propagation is given by $\hat{\imath}_2$. Velocity component 'two' is then measured at that point. Since it is assumed that all the incident illumination propagates in direction $\hat{\imath}$, a velocity determination error occurs here.

Unlike the other error sources discussed in this chapter, the error due to stray reflections is entirely dependent on the geometry of the flow to be measured. As such, it is not possible to quantify accurately the likely significance of such an error source in the general case.

In principle, if the proportion of the light sheet scattered from a given boundary were known, it would be possible to modify the velocity calculation software, `viewresults`, to take account of the stray reflection. Attempting to model the scatter from the side walls and compensate for this in the image processing software is not an attractive option, however. Such an approach would require that the geometry of each flow measurement region be accurately known and that the scattering properties of all the materials in contact with the beam be readily determined. The approach would be cumbersome and error-prone.

In practice, it is more satisfactory to attempt to minimise the intensity of any stray reflection through the use of matt paints, and by restricting the extent of the light sheet in order to avoid it impinging directly on a wall.

An assessment of the significance of stray scattering as a source of error for the case of the convergent-divergent nozzle may be found in Chapter 8.

5.6 Image Misalignment

In order to be able to undertake any predictions of expected velocity measurement error in Global doppler velocimetry results, it is important to have an understanding of the effect of image misalignment on the velocity values obtained.

Suppose that there is a small misalignment between the discriminated and non-discriminated images, of δx in the x direction and δy in the y direction. For such a small misalignment, a good estimate of the resulting velocity error may be obtained by considering the gradients of the iodine cell transmission profile, the velocity field, and the intensity at the point (x, y).

The intensity of the discriminated image at location (x, y) may be denoted by $I_d(x, y)$ and that of the non-discriminated image by $I_n(x, y)$. At that point, the intensity gradients of the non-discriminated image are given by $\frac{\partial I_n}{\partial x}$ and $\frac{\partial I_n}{\partial y}$ in the x and y directions respectively. The velocity at (x, y) is taken to be v, and the velocity gradients are given by $\frac{\partial v}{\partial x}$ and $\frac{\partial v}{\partial y}$.

If $\phi(v)$ represents the transmission fraction of the iodine cell for a certain velocity value v, the ratio R of discriminated to non-discriminated image intensity may be expressed

$$R(x, y) = \frac{\phi\left(v + \frac{\partial v}{\partial x}\delta x + \frac{\partial v}{\partial y}\delta y\right).I_n(x + \delta x, y + \delta y)}{I_n(x, y)} \tag{5.22}$$

Considering the iodine cell transmission profile,

$$\phi\left(v + \frac{\partial v}{\partial x}\delta x + \frac{\partial v}{\partial y}\delta y\right) = \phi(v) + \frac{\partial \phi}{\partial v}\left(\frac{\partial v}{\partial x}\delta x + \frac{\partial v}{\partial y}\delta y\right) \tag{5.23}$$

The normalised image intensity may be expressed thus

$$I_n(x + \delta x, y + \delta y) = I_n(x, y) + \frac{\partial I_n}{\partial x}\delta x + \frac{\partial I_n}{\partial y}\delta y \tag{5.24}$$

So

$$R(x, y) = \frac{\left\{\phi(v) + \frac{\partial \phi}{\partial v}(\frac{\partial v}{\partial x}\delta x + \frac{\partial v}{\partial y}\delta y)\right\}\left\{I_n(x, y) + \frac{\partial I_n}{\partial x}\delta x + \frac{\partial I_n}{\partial y}\delta y\right\}}{I_n(x, y)} \tag{5.25}$$

In the absence of misalignment, it is expected that $R(x, y) = \phi(v)$, so the error in ratio due to misalignment ε_R is

$$\varepsilon_R = \frac{\left(\frac{\partial I_n}{\partial x}\delta x + \frac{\partial I_n}{\partial y}\delta y\right)\phi(v) + \frac{\partial \phi}{\partial v}\left(\frac{\partial v}{\partial x}\delta x + \frac{\partial v}{\partial y}\delta y\right)\left\{I_n(x, y) + \frac{\partial I_n}{\partial x}\delta x + \frac{\partial I_n}{\partial y}\delta y\right\}}{I_n(x, y)} \tag{5.26}$$

Clearly, the error in terms of velocity rather than of the ratio of the discriminated to non-discriminated intensities is

$$\varepsilon_V = \frac{\partial v}{\partial \phi}\varepsilon_R \tag{5.27}$$

In the analysis above, it has been convenient to consider the function $\phi(v)$. This function represents transmission through the iodine cell in terms of the velocity being measured.

Thus $\phi(v)$ may be seen to depend not only on the properties of the iodine cell, but also on the directions of illumination and observation employed for Doppler global velocimetry measurements.

5.7 Centroid Estimation

The accuracy of centroid estimation may be considered both analytically and through numerical simulation. While an analytical approach offers some insight into the likely behavior of the algorithm, a numerical method is necessary to investigate fully the effect of CCD measurement noise on the accuracy of centroid estimation. Both approaches have been pursued here.

Here the focus is on the problem of one-dimensional centroid location: results are readily applicable to the case of two dimensions.

5.7.1 Integer Width Feature

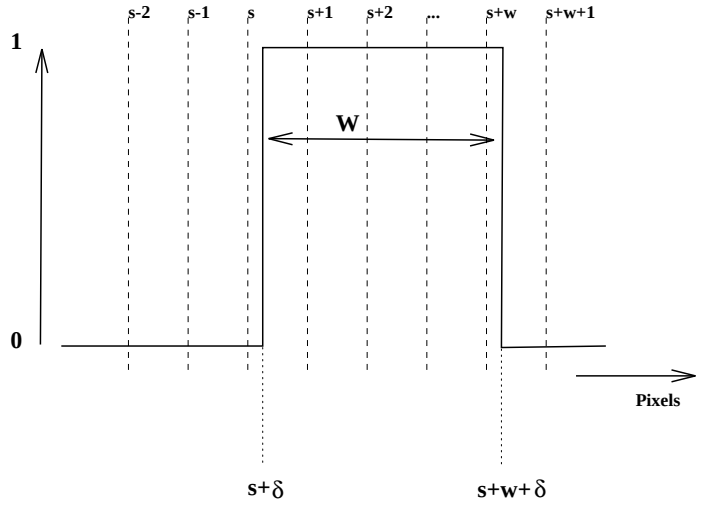


Figure 5.2: Rectangular feature — integer width.

Suppose there is a rectangular feature as shown in Fig. 5.2. In terms of pixels, this feature is of integer width W. The offset of the feature relative to the CCD pixel array is δ pixels, where $0 \leq \delta < 1$.

The centroid is given by

$$\bar{x} = \frac{\sum_{x=-\infty}^{\infty} (f(x).x)}{\sum_{x=-\infty}^{\infty} f(x)} \tag{5.28}$$

Integrating the feature for each pixel bin

- For $x < s$: bin value is 0
- For $s \leq x < s+1$: bin value is $(1-\delta)$
- For $s+1 \leq x < s+W$: bin value is 1
- For $s+W \leq x < s+W+1$: bin value is δ
- For $s+W+1 < x$: bin value is zero

So,

$$\sum_{x=-\infty}^{\infty} f(x) = (1-\delta) + (W-1) + \delta = W \tag{5.29}$$

and, taking the x values at the centre of each bin,

$$\begin{aligned}
\sum_{x=-\infty}^{\infty} (f(x).x) &= (1-\delta)\left(s+\frac{1}{2}\right) + \left(s+1+\frac{1}{2}\right) + \left(s+2+\frac{1}{2}\right) + ... \\
&\quad + \left(s+(W-1)+\frac{1}{2}\right) + \delta\left(s+W+\frac{1}{2}\right) \\
&= s + \frac{1}{2} + \delta W + (W-1)\left(s+\frac{1}{2}\right) + \left(\frac{W-1}{2}\right)(2+(W-2)) \\
&= \delta W + Ws + \frac{W}{2} + \left(\frac{W-1}{2}\right) W
\end{aligned} \tag{5.30}$$

Finally,

$$\begin{aligned}
\bar{x} &= \frac{\delta W + Ws + \frac{W}{2} + \left(\frac{W-1}{2}\right) W}{W} \\
&= s + \frac{W}{2} + \delta
\end{aligned} \tag{5.31}$$

The centroid estimate $\bar{x}$ for a rectangular feature of integer width, in the absence of any noise, is as expected. For such features, there is no theoretical error in centroid estimation due to offset relative to the CCD array.

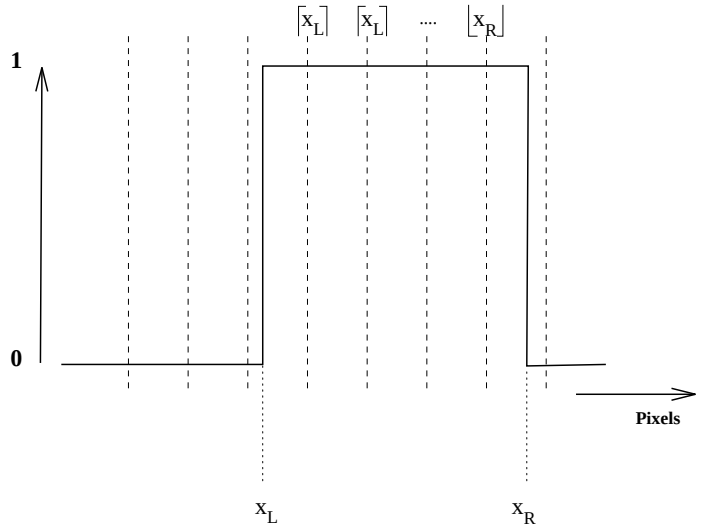


Figure 5.3: Rectangular feature — non-integer width.

5.7.2 Non-Integer Width Feature

The theoretical estimated centroid value of the integer width feature in the absence of noise has been determined above. In general, however, it is not possible to ensure that the alignment features of the calibration card are imaged onto the CCD array in such a way that each feature covers an integer number of pixels. It is thus necessary to relax the assumption of integer width, and deduce the resulting theoretical centroid estimation error.

The value of $\lfloor x \rfloor$ is defined to be the largest integer smaller than or equal to x, whereas $\lceil x \rceil$ is defined to be the smallest integer greater than or equal to x. Clearly $\lceil x \rceil - \lfloor x \rfloor = 1$ for any non-integer x.

So,

$$\begin{aligned} \sum_{x=-\infty}^{\infty} f(x) &= (\lceil x_L \rceil - x_L) + (x_R - \lfloor x_R \rfloor) + (\lfloor x_R \rfloor - \lceil x_L \rceil) \\ &= x_R - x_L \end{aligned} \tag{5.32}$$

and,

$$\sum_{x=-\infty}^{\infty} f(x).x = (\lceil x_L \rceil - x_L)\left(\lfloor x_L \rfloor + \frac{1}{2}\right) + (x_R - \lfloor x_R \rfloor)\left(\lfloor x_R \rfloor + \frac{1}{2}\right) + \sum_{i=\lceil x_L \rceil}^{\lfloor x_R \rfloor - 1}\left(i + \frac{1}{2}\right) \tag{5.33}$$

$$\begin{aligned}
\sum_{i=\lceil x_L \rceil}^{\lfloor x_R \rfloor - 1} \left(i + \frac{1}{2}\right) &= \left(\lceil x_L \rceil + \frac{1}{2}\right) + \left(\lceil x_L \rceil + \frac{1}{2} + 1\right) + \left(\lceil x_L \rceil + \frac{1}{2} + 2\right) + ... \\
&\quad + \left(\lfloor x_R \rfloor + \frac{1}{2} - 2\right) + \left(\lfloor x_R \rfloor + \frac{1}{2} - 1\right) \\
&= \frac{1}{2}(\lfloor x_R \rfloor - \lceil x_L \rceil) + \left(\frac{\lfloor x_R \rfloor - \lceil x_L \rceil}{2}\right)(2\lceil x_L \rceil + (\lfloor x_R \rfloor - \lceil x_L \rceil - 1)) \\
&= \frac{1}{2}(\lfloor x_R \rfloor - \lceil x_L \rceil)(\lceil x_L \rceil + \lfloor x_R \rfloor) \qquad (5.34)
\end{aligned}$$

So,

$$\begin{aligned}
\sum_{x=-\infty}^{\infty} f(x).x = (\lceil x_L \rceil - x_L)\left(\lceil x_L \rceil - \frac{1}{2}\right) + (x_R - \lfloor x_R \rfloor)\left(\lfloor x_R \rfloor + \frac{1}{2}\right) \qquad (5.35) \\
+ \frac{1}{2}(\lfloor x_R \rfloor - \lceil x_L \rceil)(\lceil x_L \rceil + \lfloor x_R \rfloor)
\end{aligned}$$

and,

$$\begin{aligned}
\bar{x} &= \frac{\lceil x_L \rceil^2 - \lceil x_L \rceil - 2x_L\lceil x_L \rceil + x_L + 2x_R\lfloor x_R \rfloor + x_R - \lfloor x_R \rfloor^2 - \lfloor x_R \rfloor}{2(x_R - x_L)} \\
&= \overbrace{\frac{(x_R + x_L)(x_R - x_L)}{2(x_R - x_L)}}^{\bar{x'}} \\
&\quad + \underbrace{\frac{\lceil x_L \rceil^2 - \lceil x_L \rceil - 2x_L\lceil x_L \rceil + x_L + 2x_R\lfloor x_R \rfloor + x_R - \lfloor x_R \rfloor^2 - \lfloor x_R \rfloor - x_R^2 + x_L^2}{2(x_R - x_L)}}_{\epsilon} \qquad (5.36)
\end{aligned}$$

where $\bar{x'}$ is the exact centroid and ϵ is the error term. Let

$$\begin{aligned}
x_L &= \lfloor x_L \rfloor + \delta_L \\
&= \lceil x_L \rceil - 1 + \delta_L \qquad (5.37) \\
x_R &= \lfloor x_R \rfloor + \delta_R \qquad (5.38)
\end{aligned}$$

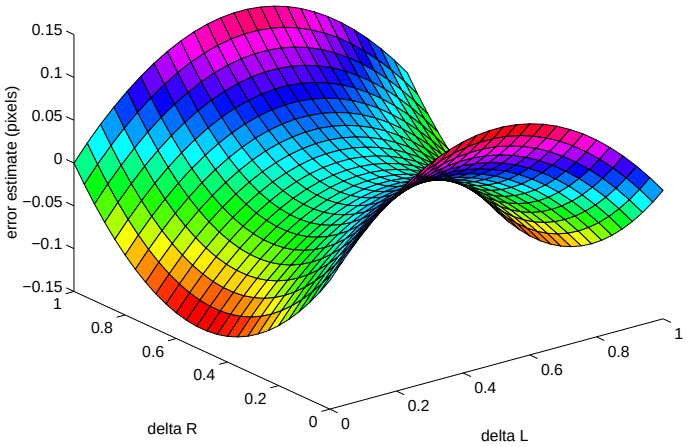


Figure 5.4: Plot of ϵW against δ_L and δ_R

Simplifying ϵ gives

$$\begin{aligned} \epsilon &= \frac{\delta_L - \delta_L^2 + \delta_R^2 - \delta_R}{2(x_L - x_R)} \\ &= \frac{(\delta_R^2 - \delta_L^2) - (\delta_R - \delta_L)}{2(x_L - x_R)} \end{aligned} \tag{5.39}$$

Evaluating this function for different values of δ_L and δ_R gives a maximum value for ϵ of $0.1250/W$. The form of ϵW may be seen in Fig. 5.4. It may be seen that for the case where $\delta_L = \delta_R$ (where the model reduces to the integer feature width model) the error reduces to zero as shown in Section 5.7.1.

Thus, even in the absence of any CCD array shot noise, an intrinsic error source is present in the estimation of centroid location, due to the fact that it is impossible for the experimenter to ensure that each calibration card feature is imaged onto the CCD array in such a way that it covers an integer number of pixel cells. The effects of relaxing the no-noise assumption may now be investigated, and compared to the error source determined in this section.

5.7.3 The Effect of Feature Width in the Presence of Noise

Now, taking noise into account, it is possible to compute the relationship between the feature width in pixels and the error in the estimated centroid position. Once more, it shall be assumed that the feature is an integer width in terms of pixels, while the offset relative to the CCD array is δ.

With a reasonably large feature, it may be assumed that the mean value of the noise across the CCD array is zero, so that the 'mass' of the function may be treated as a constant. For notational convenience, the width will be expressed in terms of w, the 'half-width', where $2w = W$.

Thus the estimated centroid may be expressed

$$\begin{aligned}\bar{x} &= [-w(\delta+\eta_{-w}) - (w-1)(1+\eta_{-(w-1)}) - (w-2)(1+\eta_{-(w-2)}) - ... \\ &\quad +(w-2)(1+\eta_{(w-2)}) + (w-1)(1+\eta_{(w-1)}) + w((1-\delta)+\eta_w)]/(2w) \end{aligned} \quad (5.40)$$

where η_i is the noise value at pixel i. The η_i values may be looked upon as distributions of mean zero and variance σ^2.

Simplifying

$$\bar{x} = \frac{\sum_{i=-(w-1)}^{(w-1)}(i(1+\eta_i)) + w((1-\delta)+\eta_w) - w(\delta+\eta_{-w})}{2w} \quad (5.41)$$

Now subtracting δ (the correct centroid position) to yield the centroid error ϵ

$$\epsilon = \frac{\sum_{i=-(w-1)}^{(w-1)}(i(1+\eta_i)) + w + w(\eta_w - \eta_{-w}) - 4w\delta}{2w} \quad (5.42)$$

It is well known, for variates X, Y, and Z, that if $Z = X \pm Y$, then $\mathrm{var}(Z) = \mathrm{var}(X) + \mathrm{var}(Y)$. Also, if $Z = kX$, then $\mathrm{var}(Z) = k^2\mathrm{var}(X)$. Hence the variance of ϵ may be obtained

$$\begin{aligned}\sigma_\epsilon^2 &= \frac{\left(\mathrm{var}(\sum_{i=-(w-1)}^{(w-1)} i) + \mathrm{var}(\sum_{i=-(w-1)}^{(w-1)} i\eta_i) + \mathrm{var}(w\eta_w - w\eta_{-w} + w) + \mathrm{var}(4w\delta)\right)}{4w^2} \\ &= \frac{\left(\sum_{i=-(w-1)}^{(w-1)}(i^2\sigma^2) + 2w^2\sigma^2\right)}{4w^2} \\ &= \frac{\sigma^2\left(\sum_{i=-(w-1)}^{(w-1)} i^2 + 2w^2\right)}{4w^2} \\ &= \frac{1}{12}\frac{\sigma^2(2w^2+3w+1)}{w} \end{aligned} \quad (5.43)$$

Expressing this in terms of the width W

$$\sigma_\epsilon^2 = \frac{1}{12}\frac{\sigma^2(W^2+3W+2)}{W} \quad (5.44)$$

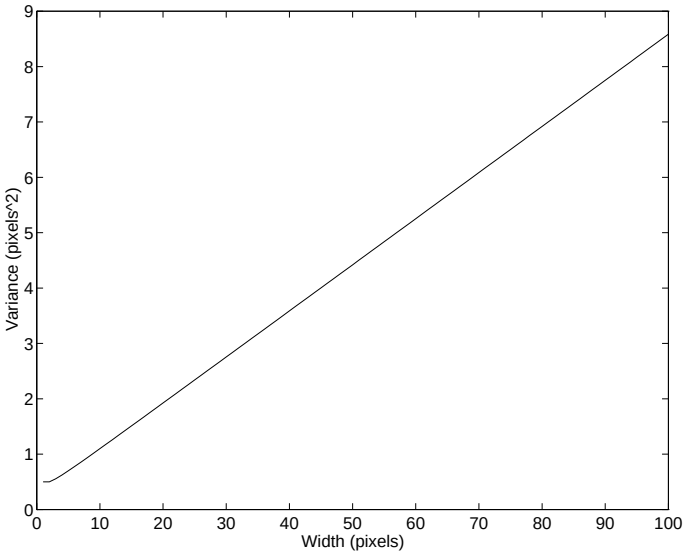


Figure 5.5: Theoretical variance of centroid error versus feature width.

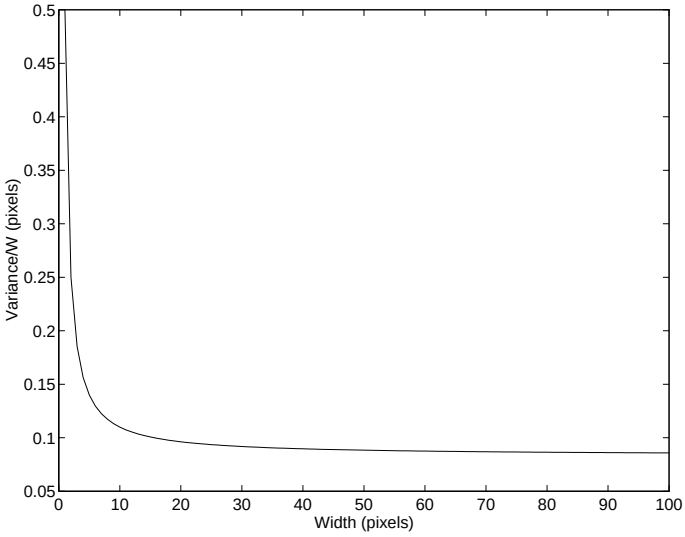


Figure 5.6: Theoretical variance of centroid error divided by feature width, versus feature width. This shows the effect of changing the CCD array resolution.

It may be seen, as shown in Fig. 5.5, that the variance of the centroid error increases with feature width. In general, then, for integer width features, increasing the feature size tends to lead to greater centroid estimation error due to the presence of noise.

Note that increasing the sampling frequency (ie: employing a finer pixel pitch) will lead to an increased centroid estimation error in terms of pixels, but since those pixels themselves will be closer together, the absolute centroid error will improve. This effect may be seen in Fig. 5.6.

5.7.4 Averaging of Frames

It is possible to reduce the effect of noise upon centroid estimation accuracy by grabbing a number of frames of the calibration target, rather than just one, and finding the mean of these frames.

Suppose n frames are averaged. Considering a single pixel, each frame may have intensity I_i with noise η_i, of mean value zero, but variance σ^2. So

$$I' = \frac{\sum_{i=1}^{n} I_i}{n} \tag{5.45}$$

and,

$$\begin{aligned} \sigma'^2 &= \frac{\sum_{i=1}^{n} \sigma^2}{n^2} \\ &= \frac{n\sigma^2}{n^2} \\ &= \frac{\sigma^2}{n} \end{aligned} \tag{5.46}$$

In general, then, the variance of the noise decreases as frames are averaged. In terms of the standard deviation σ, for n frames averaged

$$\sigma' = \frac{\sigma}{\sqrt{n}} \tag{5.47}$$

This reduction in the standard deviation of the individual pixel noise values as increasing numbers of frames are grabbed and averaged may be employed in order to diminish the adverse effect of CCD camera shot noise on the estimates of feature centroid location.

5.7.5 Subtraction of Black Level

It has been seen that a CCD array exhibits dark current even when exposed to no light. In addition, in the determination of centroid location, ambient lighting may be expected to give a non-zero intensity value outside of the feature itself. Centroid location for each feature is accomplished by first finding a bounding box which surrounds a single feature. Since the centroid of the feature is obviously not known *a priori*, it is not possible to guarantee that the feature is centered in the bounding box. Clearly, if the level of the black offset is not compensated for, the centroid of the region of the bounding box may differ from that of the feature itself.

The further the feature is from the centre of the bounding box, and the greater the value of the black level, the greater the extent of potential errors from poor black level compensation. Such compensation may be accomplished by acquiring a second frame of an entirely black target using the same exposure time as for the calibration image, and subtracting this frame from the original calibration image. In the absence of any pixel noise, the black level will be entirely removed, and error due to poorly-aligned bounding boxes eliminated.

However, as already indicated, shot noise is introduced by the CCD array, which ensures that the black level may not be entirely removed by this method. Performing averaging of frames (both for the calibration image and for the black image) will reduce the effective level of shot noise as explained in Section 5.7.4. As explained in Section 3.2, a threshold may be specified by the user of the `warp` software — pixels of value lower than this threshold are then set to zero. This technique appears to be attractive for the removal of shot noise from outside the feature itself, although caution should be exercised, since such an approach may itself introduce error in the centroid estimate of the feature. The reason for this is that if either of the two 'part-full' pixel bins at the feature edges has a value less than the threshold value, it will be set to zero, hence distorting the estimated centroid position.

Thus there is an optimal setting for the black level threshold. Too low a threshold will allow shot noise in the nominally black region of the bounding box to adversely influence the centroid estimate. Too high a level of threshold may cause the values in the 'part-full' bins at each extreme of the feature to be set to zero. The practical influence of this parameter is best investigated numerically — the reader is referred to the analysis of Section 5.7.7.

5.7.6 Fourier Method

Alexander and Ng [Alexander and Ng 1991] discuss the *systematic error* of centroid estimation. There are two sources of error in centroid estimation: systematic error due to the nature of the algorithm and errors arising due to the presence of noise. Systematic error is reproducable given a particular lens line-spread function and a particular offset of the feature relative to the CCD array. Alexander and Ng consider only the systematic error of sampling in a limited number of pixels.

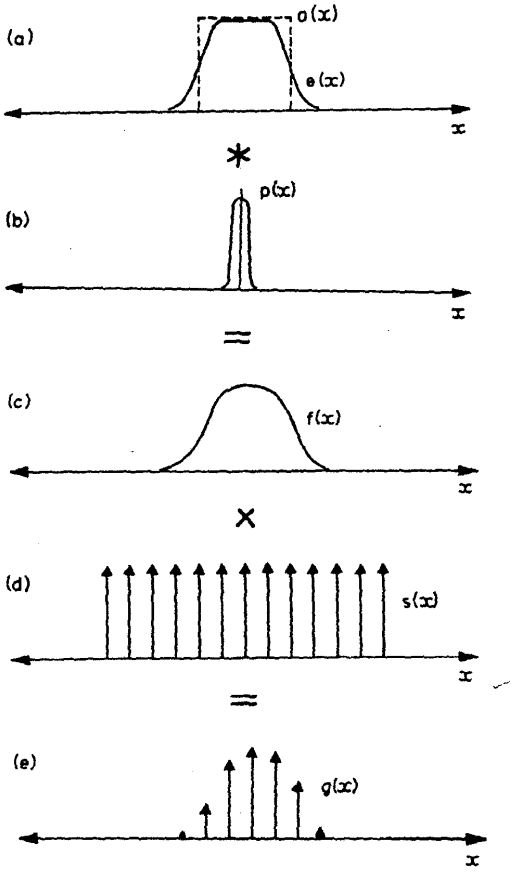


Figure 5.7: The process of sampling an image. a) Intensity profile $e(x)$ of light incident on the CCD surface, with $a(x)$ shown dotted; b) pixel sensitivity profile $p(x)$; c) waveform $f(x)$ resulting from convolution of pixel sensitivity profile with stripe image intensity profile; d) sampling function $s(x)$; e) sample waveform $g(x)$ obtained by multiplying sampling function by $f(x)$ (after Alexander & Ng, [Alexander and Ng 1991])

.

Fig. 5.7 illustrates the process of forming a sampled image of a single feature. The function $e(x)$ is the intensity profile of the image on the CCD array, obtained by convolving $a(x)$, the intensity profile of the feature, with the line-spread functions of the camera lenses. Convolving $e(x)$ with $p(x)$, the pixel sensitivity profile, yields the function $f(x)$. Multiplying $f(x)$ by the sampling function $s(x)$, gives $g(x)$, the sequence of sample values.

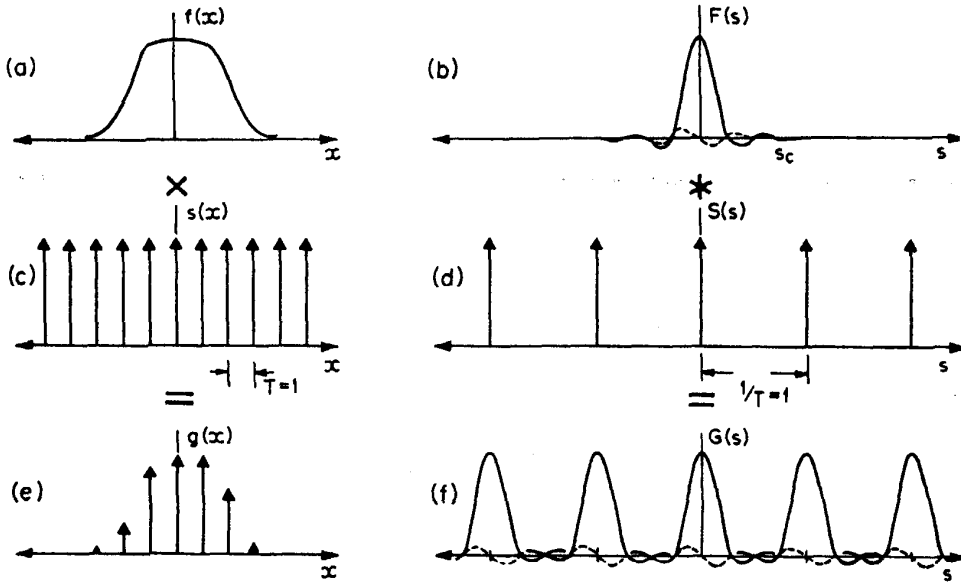


Figure 5.8: Sampling in the spatial domain and the spatial frequency domain. a) Intensity profile $F(x)$; b) Fourier transform $F(s)$ of $f(x)$; c) Sampling function $s(x)$; d) Fourier transform of sampling function $S(s)$; e) Waveform $g(x) = f(x)s(x)$; f) Fourier transform of $g(x)$, $G(s) = F(s) \otimes S(s)$ (after Alexander & Ng, [Alexander and Ng 1991])
.

A systematic error is said to exist if

$$\overline{x_f} \neq \overline{x_g} \tag{5.48}$$

The relationship between the centroid of a continuous function and its sampled equivalent may be understood by considering the Fourier transform of the two waveforms. Fig. 5.8 shows, on the left, waveforms corresponding to Fig. 5.7; and on the right, their Fourier transforms. Whereas multiplication of $f(x)$ by $s(x)$ gives the sampled function $g(x)$ in the spatial domain, convolution of $F(s)$ with $S(s)$ gives the function $G(s)$ in the spatial frequency domain. The function $G(s)$ exhibits the transform of $f(x)$, $F(s)$, replicated at intervals of $1/T$, the sampling frequency.

Now, the centroid of $f(x)$ and $g(x)$ may be related to their Fourier transforms as follows

$$\overline{x_f} = -\frac{F'(0)}{2\pi j F(0)} \tag{5.49}$$

$$\overline{g_f} = -\frac{G'(0)}{2\pi j G(0)} \tag{5.50}$$

The centroid of $f(x)$ is dependent, then, only on the behaviour of $F(s)$ and $F'(s)$ at the origin. Similarly, the centroid of $g(x)$ is dependent only on the behaviour of $G(s)$ and $G'(s)$ at the origin.

So, if

$$G(0) = F(0) \tag{5.51}$$

and,

$$G'(0) = F'(0) \tag{5.52}$$

then $\bar{x_g} = \bar{x_f}$. In other words, the centroid of the sampled function $g(x)$ will then be equal to the centroid of the corresponding continuous function $f(x)$.

Referring to Fig. 5.8(f) once more, it is clear that eqns. 5.51 and 5.52 will be satisified if the replicated transforms of $f(x)$ do not overlap to such a degree that any part of a replicated transform is present at the origin. This will be the case if the cut-off frequency s_c is less than $1/T$, the sampling frequency, as shown. Thus the degree of systematic error is determined by the cut-off frequency of the lenses in the optical chain. Choosing a small aperture (large f number) decreases the cut-off frequency and thus the 'overlap' at the origin of the replicated waveform $G(s)$. Above a certain critical f number, there will be no overlap at the origin of $G(s)$, so systematic error may be eliminated. In other words, if $F(s)$ does not cut off below the sampling frequency, systematic error will, in general, occur.

Alexander and Ng provide a detailed analysis of the nature of systematic error when the conditions of eqns. 5.51 and 5.52 are not met. The paper concludes that the nature of the systematic error in estimated centroid location may be expressed as

$$\epsilon \approx \frac{F_e'(1/T)\sin(2\pi d/T)}{\pi F_e(0)} \tag{5.53}$$

where $f(x)$ is $f_e(x)$ shifted by the sampling offset d from the origin. Further, $f(x) = f_e(x-d)$, and f_e is even, with Fourier transform $F_e(s)$.

The function ϵ has the form of $\sin\theta$, with θ ranging from $-\pi$ to $+\pi$ as the offset d is changed from -0.5 to 0.5 (for $T = 1$). The magnitude of ϵ is directly proportional to the slope of $F_e(s)$ at the sampling frequency.

The work of Alexander and Ng [Alexander and Ng 1991] shows that the line-spread function of the optical system is significant to the accuracy of estimated centroids. A large value of f number should be chosen in order to minimise this particular error term, and high quality optics are clearly preferred. A systematic error of around 0.1 pixel, in the worst case, is expected. This value was found to be largely insensitive to feature width.

5.7.7 Numerical Simulation of Centroid Error

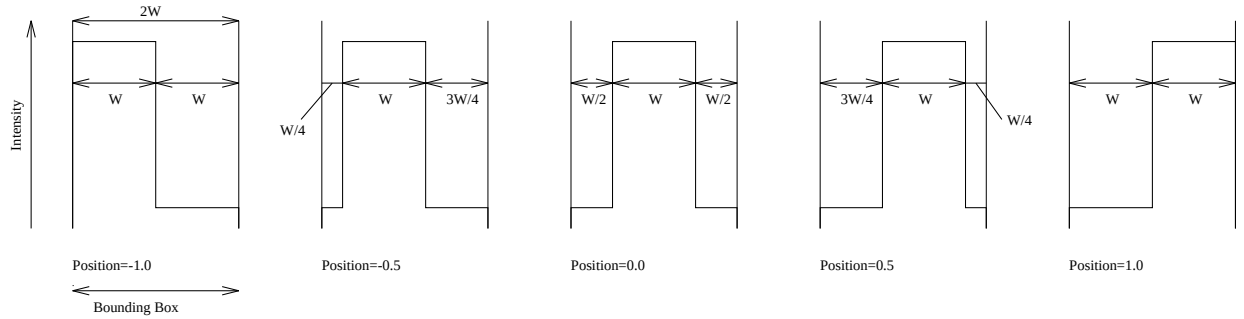


Figure 5.9: Diagram illustrating the meaning of the parameters employed in the numerical simulation of centroid estimation

In order to investigate further the influence of sampling on attainable centroid estimation accuracy, a numerical simulation was devised. The simulation enables the computation of the attainable centroid location accuracy when feature width, offset of the feature relative to the CCD array, and the number of quantisation levels utilised are varied. The effect of CCD camera shot noise on centroid accuracy and the implications of smoothing may also be examined.

5.7.7.1 Overview of Numerical Simulation Technique

The numerical simulation employed enables a number of parameters to be set by the user. The left hand edge and right hand edge of the feature may be specified as any integral or non-integral value. The edges of the bounding box defining the region of centroid computation may be similarly specified. Intensity levels for the 'white' and 'black' level may be set. Additional parameters that may be adjusted include the level of the black threshold of the `warp` software below which intensity values may be set to zero, and the number of quantisation levels to employ. Smoothing may also be introduced, implemented by means of a simple blurring function.

For the results presented here, the configuration of Fig. 5.9 has been employed. The 'white' level is set to 16000 counts (half the CCD saturation value) while the 'black' level is set to 1000 counts. The size of the bounding box is at all times set to be twice that of the feature being examined. The position of the feature within the bounding box may be specified as a

certain *offset* value, as shown in Fig. 5.9. The values chosen for the simulation are intended to represent realistic operational conditions experienced while aligning the Doppler global velocimeter. In particular, the implementation of the `warp` code is such that bounding boxes are generally around twice the size of the features so contained.

In order to model closely the behavior expected in the real Doppler global velocimetry system, aspects of the `warp` image processing algorithm are introduced into the simulation. A simulated normalisation frame ('white' value throughout) and a simulated black frame ('black' value throughout) are utilised when computing centroid position.

In order to take CCD shot noise into account, values are sampled from a Gaussian distribution representing noise of known RMS value. Since uniformly distributed random values are required, the pseudo-random number generation method `ran1` proposed by Press *et al.* [Press *et al.* 1992] is used. A number of pseudo-random values are generated and the mean of these taken. According to central-limit theorem, the distribution of these mean values will be Gaussian in nature. It is straightforward to determine the variance of this Gaussian distribution and to relate it to the RMS value of the noise to be modelled. The RMS noise value used for the model is determined by application of the CCD shot noise analysis of Section 5.2. Noise is introduced into the simulated normalisation and black frames in the same way as for the simulated calibration frame itself.

Each run of the simulation code iterates through a number of candidate features, of both integer and non-integer widths. Non-integer features are formed by adding an additional quarter or half a pixel to the width of the integer features. For such features, the left hand edge is set to a integer pixel position and the right hand edge set to a non-integer position. The intensity profile of the feature is sampled into the appropriate number of one pixel wide bins, using linear interpolation for those bins which straddle the edge of the feature.

For the case of 32768 quantisation levels (the default for the CCD array in use at Oxford) quantisation is performed implicitly by means of the integer representation of the bins in the simulation. In order to simulate the use of fewer quantisation levels, it is necessary to adjust the bin values explicitly. If selected, smoothing is performed post-quantisation. In this implementation a simple three point blur function has been employed.

Calculation of centroid position in the presence of noise is performed 500 times for each configuration. It is trivial to compute the true centroid position, so the RMS value of the expected centroid error may be readily computed for each configuration modelled.

The feature widths modelled were 2.0, 4.0, 8.0, 10.0, 16.0, 20.0, 24.0, and 36.0 pixels. Non-integer feature widths, obtained by adding 0.25 or 0.50 pixels to the integer values, were also examined. A number of simulation runs were undertaken in order to study the effect of feature width, bounding box position, black threshold value, quantisation, and smoothing.

5.7.7.2 Results of the Numerical Simulation

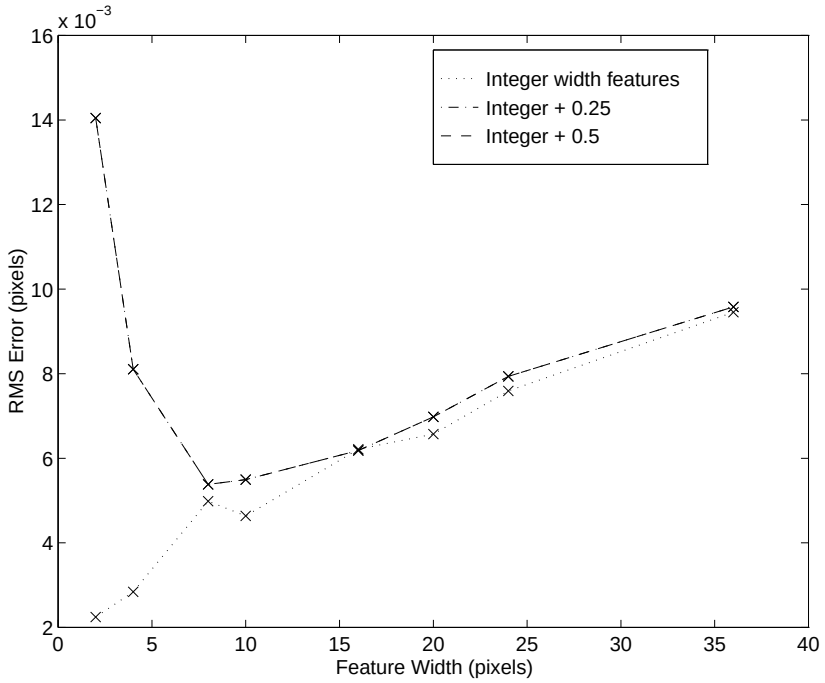


Figure 5.10: Simulation Results: RMS Centroid estimation error versus feature width for the case of a centralised bounding box. The black threshold is set to zero.

Having theoretically considered the influence of feature width on centroid estimation accuracy in Section 5.7.1, 5.7.2, and 5.7.3, it is interesting to model this relationship by means of the numerical simulation code. Fig. 5.10 shows the relationship between simulated feature width and estimated RMS centroid error for the case of a feature centred within its bounding box. The black level threshold has been set to zero in this case, and 32768 quantisation levels have been employed. The dotted line represents the estimated RMS centroid estimation error for integer width features. The dot-dashed and dashed lines represent the same analysis for the case of features with a non-integer width component of 0.25 and 0.50 respectively.

For integer width features, the RMS centroid estimate error in the presence of shot noise is seen to increase with respect to feature width. For a feature width of 2 pixels, the computed RMS centroid estimate error is around 0.002 pixels, rising to a value of around 0.008 pixels for a features width of 36 pixels. This observation is in accordance with the analysis of the effect of shot noise on centroid estimation presented in Section 5.7.3. That analysis was able to show that the variance of the computed centroid position increases with feature width

(eqn. 5.44).

It may be readily seen that the two sets of computed data for non-integer width features coincide. For the non-integer width case it appears that the estimated RMS centroid estimate error is far higher for small features than was the case for integer width features. Indeed, for a non-integer width feature of around 2 pixels in width, estimated RMS centroid error has been found to be 0.014 pixels. As the feature width is increased, the estimated RMS error is seen to diminish sharply, reaching a minimum of around 0.006 pixels when feature width is about 8 pixels. For feature widths in excess of this value, the non-integer data conforms quite closely to the results found for integer width features, rising gradually to around 0.008 pixel RMS error for a feature width of 36 pixels.

It is clear that for the case of non-integer feature widths, two effects are being superimposed. First, there is the effect of shot noise, which has been seen (Section 5.7.3) to lead to increased centroid estimation error with feature width. This effect dominates for features of width in excess of 8 pixels. For smaller features the effect of non-integer width itself (Section 5.7.2) leads to a sharp increase in centroid estimation error as the features are made smaller. The superimposition of these two effects means that there is an optimal feature width which may be employed in order to minimise the expected error in the computed centroid position. For the case examined in Fig. 5.10, the optimal value would appear to be around 8 pixels.

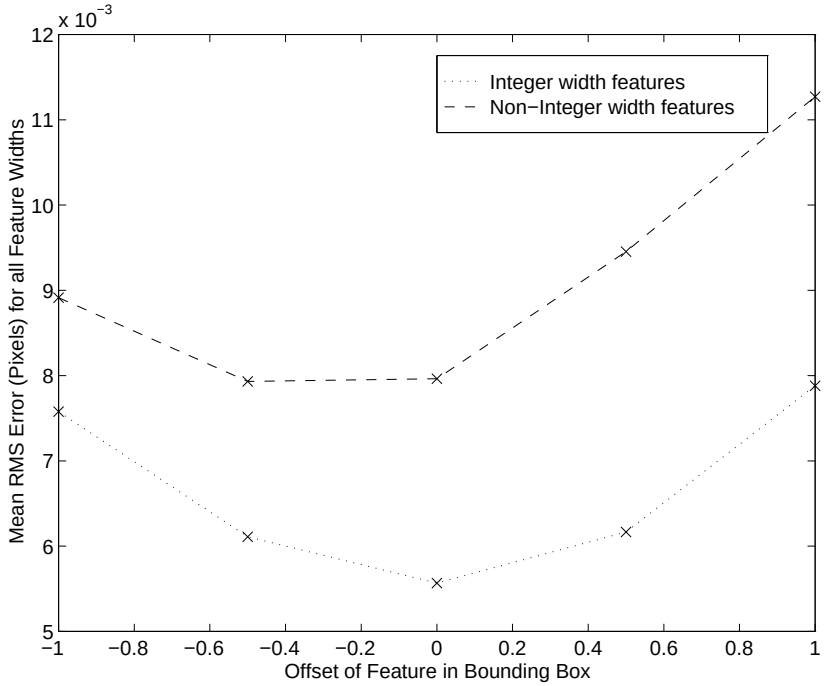


Figure 5.11: Simulation Results: Mean RMS error versus feature offset within bounding box. The black threshold is set to zero.

Having looked at the effect of feature width on predicted centroid estimation accuracy for the case of a feature in the centre of its bounding box, it is useful to examine the influence of the offset of the feature within the bounding box. Since the position of centroids is not known *a priori*, it is not possible to ensure that the alignment features will always be centred within the bounding box computed by the `warp` software. Ideally, it is hoped that off-centre features should be located with the same accuracy as those in the centre of their bounding box, since the black level image should remove any residual dark current which may bias the computation. In practice, however, both the calibration frame and the black frame are subject to shot noise. This means that dark current levels within the bounding box may not be eliminated entirely.

Fig. 5.11 shows the effect of feature offset, both for integer width features (dotted line) and non-integer width features (dashed line). The ordinate value is the mean RMS error computed for all feature widths under consideration. It is immediately apparent, once again, that non-integer width features exhibit a greater mean RMS centroid estimation error than integer width features. It is also clear that in the presence of shot noise, feature offset is significant in determining the accuracy to which centroids may be located. The general pattern is that centroid estimate error is lowest when the feature is in the centre of its bounding box (offset zero), and increases when the feature is moved off centre. For the case of integer width features, a minimum value of error for a centralised bounding box of around 0.006 pixels was determined. This value increases to around 0.008 pixels when the feature is at one extreme of the bounding box. For non-integer width features, these values were found to be 0.008 and 0.011 pixels respectively.

In order to attempt to overcome the effect of imperfect black level subtraction due to the presence of shot noise, a black threshold may be used. Any values which fall below this threshold may be set to zero. It is clear that the choice of such a value must be undertaken carefully. If the threshold is set too low, the effect of shot noise will still be noticeable in nominally black regions of the frame. If the threshold is set too high, on the other hand, the part full bins at the edge of the alignment feature may be set to zero, distorting the computed position of the centroid.

The effect on mean RMS error, for all the feature widths considered, of feature offset within the bounding box is shown in Figs. 5.12 and 5.13. The two figures differ only in the choice of black threshold used; 0.1 of the CCD saturation value for Fig. 5.12 and 0.05 for Fig. 5.13. The two figures are clearly very similar, showing a mean RMS error for all feature widths considered of around 0.006 for all offsets for the case of integer width features; and 0.008 for

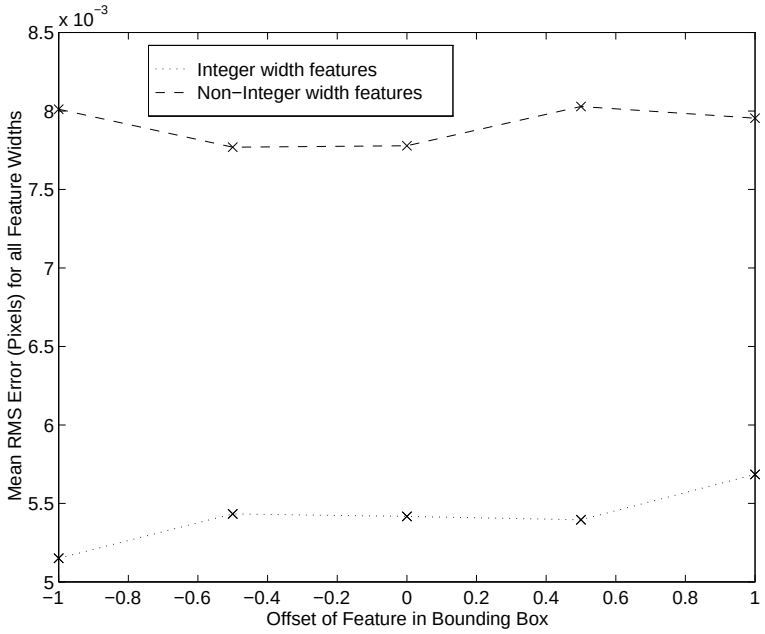


Figure 5.12: Simulation Results: Mean RMS error versus feature offset within bounding box. The black threshold is set to 0.1

all offsets for non-integer width features.

The fact that the two figures are so similar is a sign that even a black threshold value of 0.1 of saturation is insufficient to have a significant adverse effect on centroid estimation accuracy, and that a value of 0.05 of saturation is sufficient to compensate for the effect of shot noise on black level subtraction. The dependence of centroid estimation accuracy on offset of the feature within the bounding box is more or less entirely removed by this means. This suggests that the effect of shot noise on the nominally black region of the bounding box is indeed being overcome by means of the use of a black level threshold. The resulting centroid estimation accuracy determined is very similar to the optimal value for a centralised feature within the bounding box found when no black level threshold was in use.

It is possible to use an iterative scheme of centroid estimation, where the centroid estimate is determined for a certain initial bounding box, and a new bounding box is then computed centred on this position. The centroid estimate may then be computed again using this new bounding box. Such a scheme may be seen to offer improved immunity to the problem of features off-centre in their bounding boxes. The use of a black threshold value, however, is more straightforward and requires less computation. The fact that RMS centroid estimates have been found to be more or less insensitive to offset values when a black level threshold is specified suggests that there is little advantage in the use of an iterative scheme.

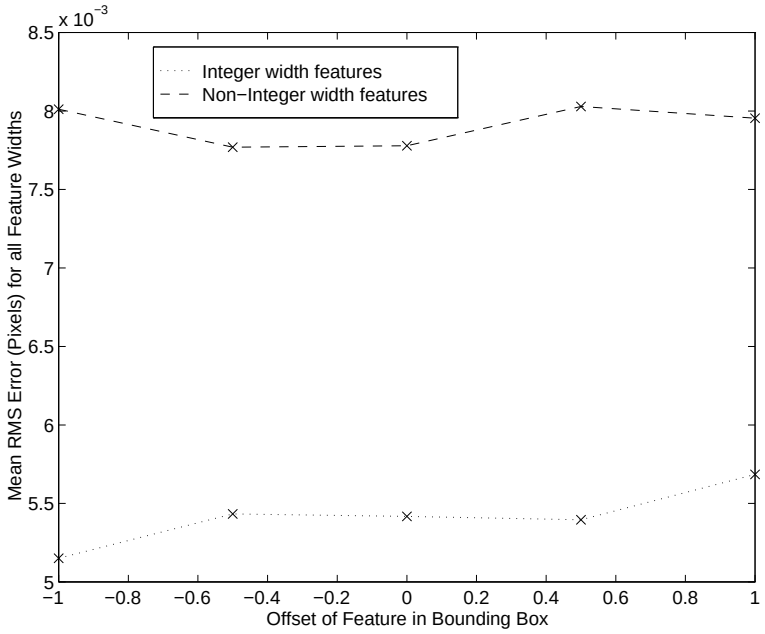


Figure 5.13: Simulation Results: Mean RMS error versus feature offset within bounding box. The black threshold is set to 0.05

The effect of the choice of black level threshold may be made clearer by plotting the mean RMS centroid estimate error for all widths and for all offsets against the value of the black level threshold, as shown in Fig. 5.14. It may be seen that setting the threshold to 0.05 of saturation markedly reduces error for both integer and non-integer width features. Increasing the black level threshold to 0.10 of saturation offers no further improvement, but has no adverse effect for the cases considered in the simulation. The improvement in worst case behaviour may be expected to be even more significant than that in mean behaviour.

All the numerical simulation data presented above has been determined using 32768 quantisation levels. However, the number of quantisation levels of the CCD camera system in use is also a siginificant parameter. The camera in use at Oxford offers 15 bit digitisation, representing 32768 quantisation levels, whereas video camera systems typically offer 6 to 8 bit digitisation, with 64 to 256 levels respectively. It is interesting to observe the effect of such an increase in quantisation on centroid estimation accuracy. It is also possible to choose to use a coarser quantisation in the image processing code than that offered by the camera, should this prove advantageous.

Fig. 5.15 shows the effect on mean RMS centroid estimation error for all feature widths and offsets of the number of quantisation levels employed. The black threshold has been set to zero in this case. Once again, it is seen that non-integer width features exhibit a greater centroid estimate error than that for integer width features. Using 100 levels,

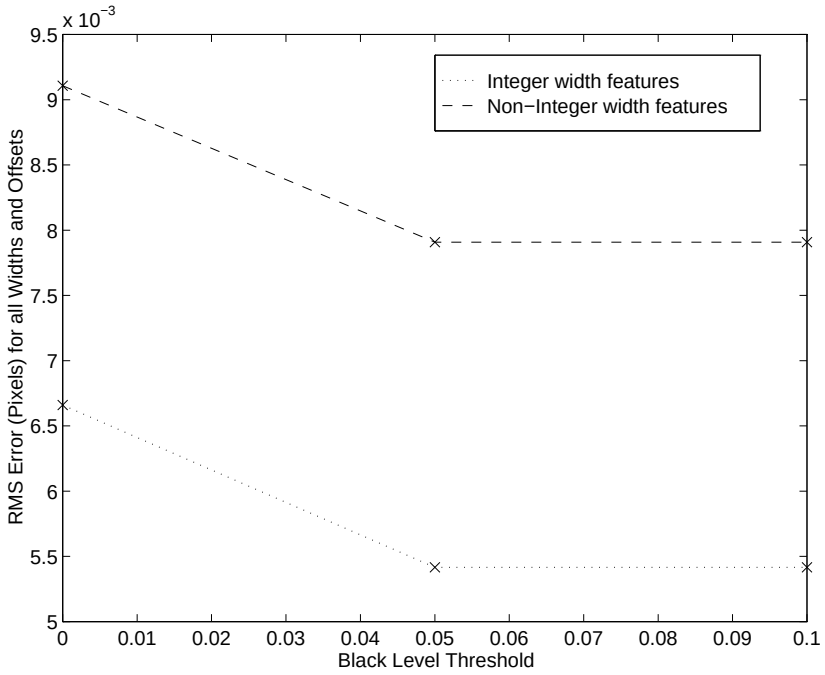


Figure 5.14: Simulation Results: Mean RMS error versus black level threshold setting.

both integer and non-integer width features offer a mean RMS centroid estimate error of around 0.02 pixels. This falls rapidly to around 0.007 and 0.009 pixels respectively using 250 levels. Increasing the number of quantisation levels beyond 250 has little effect on centroid estimation accuracy, although it may be a significant advantage in terms of acquiring actual flow images. The conclusion must be that at least 8 bit digitisation is required for accurate centroid estimation and that the use of a greater number of levels presents no significant drawback. The use of video cameras with fewer than 250 levels may lead to a significant increase in alignment error due to poor centroid estimation accuracy.

Finally, the effect of smoothing using a simple three-point blurring function has been modelled. Figs. 5.16 and 5.17 depict the RMS centroid estimate error versus feature width for the case of integer and non-integer width features respectively. A feature centred in its bounding box has been examined in these final simulation runs, with the black level threshold set to zero. Each figure shows the results obtained both while using smoothing (dotted line) and without smoothing (dashed line). Such smoothing may be seen to simulate the effect of imperfect focus in the optical system. Indeed, the use of such a modelling technique may be seen to mirror the analysis of Alexander and Ng [Alexander and Ng 1991] as presented in Section 5.7.6.

For integer features, as shown in Fig. 5.16, the unsmoothed data shows a continuous increase in RMS centroid estimation error with respect to feature width. However, for small features

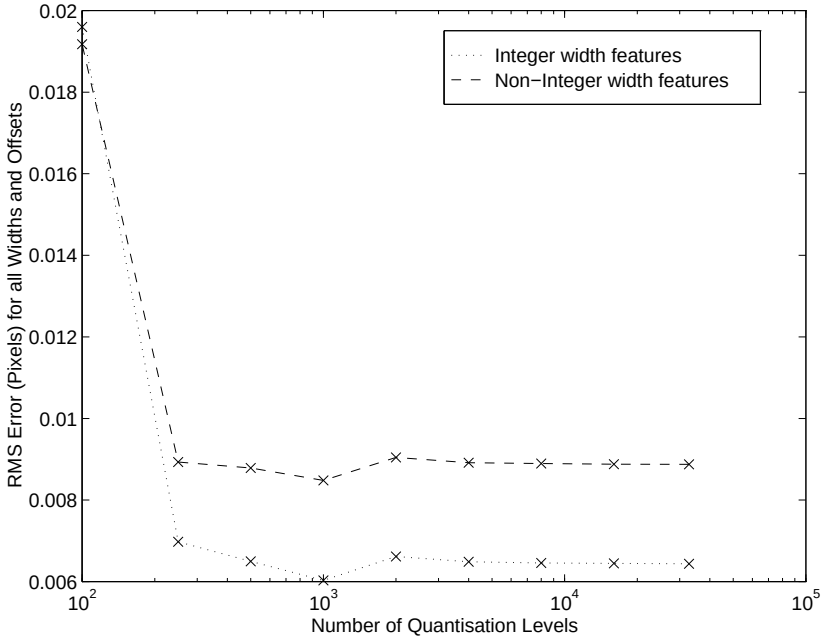


Figure 5.15: Simulation Results: Mean RMS error versus number of quantisation levels employed.

(2 and 4 pixels in this case), error was seen to increase markedly in the presence of smoothing. For larger features (above 8 pixels in width) the estimated error in the presence of smoothing is seen to converge with the results obtained in the absence of any smoothing.

For non-integer features, shown in Fig. 5.17, the pattern previously observed in the absence of smoothing where centroid estimate error diminishes and then increases with respect to feature width is seen. It may be seen, however, that the effect of smoothing for small features dominates over the effect of non-integer feature width, reaching an RMS error value of 0.09 pixels for a width of 2 pixels. The value of 0.09 corresponds well with the 0.1 pixel estimate presented by Alexander and Ng.

It is clear that some degree of blur will be present in the images obtained on the CCD array as a result of imperfections in the optical system . Features of sufficient width in order to maximise the centroid estimation accuracy attainable must therefore be chosen.

It has been seen that the effects of non-integer feature widths (which in general will be present) and of blurring are such that large errors may be encountered when the feature size is less that 8 pixels or so. The gradient of the RMS error curve is very steep to the left of this point, but much less steep to the right. It is thus prudent to choose a feature width for optimal centroid estimation somewhat in excess of 8 pixels. A small penalty may be incurred in doing so, but it is perhaps best to err on the side of caution, in case more blurring should be introduced by the optical system than is expected. The optimal feature

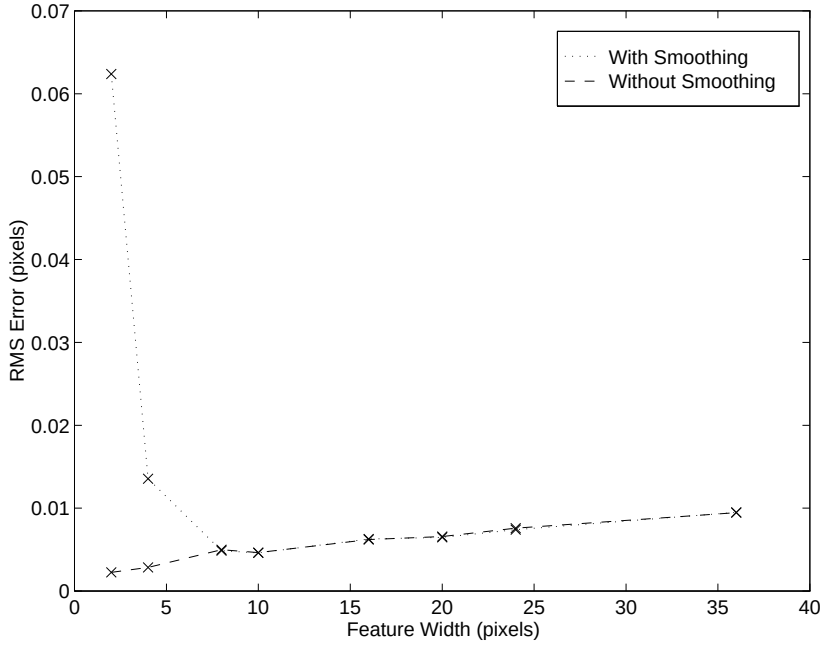


Figure 5.16: Simulation Results: Effect of smoothing on RMS error for the case of integer width features.

size is a trade-off between error introduced due to non-integer feature size and blurring on the one hand, and due to shot noise on the other. A feature size of 15 pixels would appear to be a good compromise. Accordingly, in experiments undertaken at Oxford, alignment cards were used in such a way that the features imaged onto to the CCD array occupied approximately 15 pixels.

Blurring has been seen to dominate over other sources of centroid estimation error in this simulation, leading to RMS error values close to 0.1 pixels for small features. In the case of the real Doppler global velocimetry system, rather than this numerical simulation, imperfections in the optical system must thus be expected to have a significant deleterious effect on centroid estimation accuracy. Systematic error is likely to be incurred at operating aperture settings. If suitable image processing steps are taken to reduce the error sources revealed by the numerical simulation, systematic error is expected to dominate, leading to an RMS centroid estimation error of around 0.1 pixel. For the chosen case of features 15 pixels in width, using the `warp` image processing software, this value has been seen to be realistically achievable.

5.7.7.3 Summary of the Causes of Centroid Estimation Error

The predominant causes of centroid estimation error may be thus characterised:

- Systematic error due to lens line-spread function and CCD sampling characteristics.
- CCD array shot noise

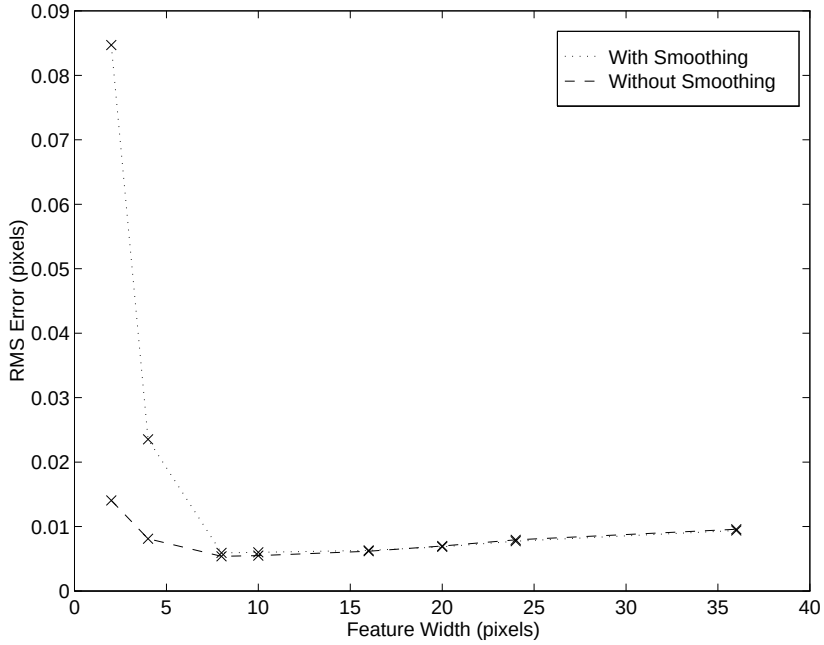


Figure 5.17: Simulation Results: Effect of smoothing on RMS error for the case of non-integer width features.

- Error due to non-integer width features
- The effect of choosing a threshold in order to exclude shot noise from the nominally 'black' region of the bounding box.
- The effect of quantisation

There is little that may be done to overcome errors due to non-integer width features, since it is impossible to set up the optical system so as to guarantee that features cover an integral number of pixels on the CCD array across the entire frame. Similarly, the variation in centroid estimation error with feature offset relative to the CCD array is unavoidable. It is necessary simply to attempt to minimise the worst case error with respect to both non-integer width and offset.

CCD shot noise is an unavoidable intrinsic phenomenon, albeit one which may be minimised by ensuring that multiple images of the calibration image, the calibration normalisation image, and the black level image are acquired and then averaged. As shown in Section 5.7.4, this will have the effect of reducing shot noise by a factor $1/\sqrt{n}$, where n is the number of frames acquired.

It has been seen from the results of the shot noise simulation that the use of a higher number of quantisation levels (such as the 32768 levels for the CCD system used at Oxford) does not necessarily improve centroid estimation accuracy. Using a number of quantisation

levels fewer than around 2000 should be avoided, since this has been seen to increase error markedly.

The results of the simulation have shown that for a noise-free signal, smoothing using a simple blur function has a significant deleterious effect. Any use of spatial smoothing techniques in the alignment software is thus to be avoided. The blurring function mimics the effect of the line-spread function of the lenses employed in the imaging system. Optics must be of the highest quality to minimise such systematic error. According to the analysis of Alexander and Ng [Alexander and Ng 1991], if the imaging system f number is set above a certain critical value, then no systematic error is present. For apertures larger than this critical setting (smaller value of the f number) a systematic error, largely insensitive to feature width, of around 0.1 pixel in the worst case is present. For the case of the Doppler global velocimetry optical system, choosing an aperture smaller than the critical value may prove impractical, since that would result in a marked increase in the exposure times necessary in order to obtain sufficient intensity. For larger apertures, the systematic error will dominate if adequate steps are taken to reduce other sources of error.

It is advantageous to choose an optimum feature width. From examination of the simulated results, it would appear that a feature width of around 15 pixels produces the best trade-off between the error due to non-integer width features and that due to shot noise.

A significant error may be introduced due to shot noise in the black region. In practice, the choice of tight bounding boxes and the effective elimination of noise in the region outside of the feature are required to minimise the problem.

Attempting to ensure that the feature is as close as possible to the centre of the bounding box will minimise any error due to the presence of poorly compensated black level or due to shot noise. Bounding boxes should be made as small as possible, so that they enclose just the feature and a minimal surrounding area. It is possible to implement an adaptive scheme where a bounding box is selected, its centroid estimated, and a new symmetrical bounding box computed. This scheme may be applied iteratively in order to reduce the effect of asymmetrical bounding boxes.

In order to reduce the adverse effect of any non-zero black level, it is sufficient to subtract a black level frame from the calibration and normalisation frames. A threshold may then be chosen below which the pixel values of the resulting adjusted normalised calibration frame may be set to zero. A value of 0.05 for this threshold has been found to be optimal in this simulation.

5.7.7.4 Techniques for Ensuring Optimal Centroid Estimates

Thus, finally, it is possible to present a number of recommendations for ensuring accurate centroid estimation for the Doppler global velocimetry application.

- The optical system should be accurately aligned and focused, with minimal blur across the entire frame, in order to ensure that the calibration features are sharply and consistently imaged on the CCD array. The highest practical value of f number should be used throughout in order to minimise systematic error.
- The target should consist of a regular array of square white dots on a matt black background. The size of the dots should be such that they occupy approximately 15 pixels of the CCD array as viewed. A number of different calibration targets will then be required for different configurations of the Doppler global velocimetry optical system. The optimal inter-dot spacing is determined by the requirements of the warping algorithm in use — this aspect will be considered in Section 5.8.
- A normalisation frame should be acquired in order to compensate for the effect of pixel sensitivity non-uniformity and for non-uniform illumination of the calibration target.
- A frame of a matt black target should be acquired and subtracted from the calibration frame and the normalisation frame in order to remove the black level, which may distort centroid location.
- Illumination and exposure time should be such that the dynamic range of the CCD array is well utilised. At higher intensities closer to CCD saturation, the signal to noise ratio is improved.
- Pixel values below the threshold of 0.05 should be set to zero in an attempt to eliminate the effect of shot noise from the black regions.
- Bounding boxes should be computed to ensure that features are located within the box as centrally as is possible, and that the box is no larger than necessary.
- Apertures and optical settings should not be adjusted between calibration and data acquisition, since this may alter the nature of optical aberrations. The restriction on the adjustment of the aperture between alignment and data acquisition is particularly irksome, since it disallows the use of small apertures (with resulting long exposure times) for alignment, and larger apertures for data acquisition. Thus the avoidance of systematic error proves to be difficult in practice. The use of high quality lenses

reduces this dependence on aperture, and the use of high intensity illumination helps reduce the constraints on the available apertures for data acquisition while allowing reasonably short exposure times.

- The calibration image must be acquired under laser illumination, to ensure that chromatic aberrations do not mar the alignment. Precautions must be taken (see Chapter 6) to reduce the effect of the laser speckle pattern.
- Averaging of multiple calibration, normalisation, and black images may be employed to reduce the effect of shot noise.
- Spatial smoothing techniques should not be employed.
- If it proves impossible to select an f number greater than the critical value, achievable centroid error should be close to the 0.1 pixel value given by the systematic error.

5.8 Image Transformation Techniques

While the expected accuracy of alignment feature centroid estimation has been evaluated, it is far from clear how this error may propagate through the warping algorithm. In order to appreciate fully the impact of this error on the final measured velocity, a numerical simulation was devised for both the bilinear and polynomial warp techniques.

5.8.1 Numerical Simulation of Warp Error

5.8.1.1 Overview of Numerical Simulation Technique

The simulation proceeds by setting up a pair of identical grids of regularly spaced points. In the case of the bilinear warp, only four pairs of points are required: they are positioned at the extreme corners of the frame. For the polynomial warp a greater number of points pairs are employed. A number of different inter-feature spacings may be modelled. One of the sets of points, representing those points determined by centroid estimation, is modified by the addition of noise of a Gaussian distribution of standard deviation 0.1 pixel in both the x and y directions; while the other set of features, representing the grid to which the alignment is to be made, is not subjected to any noise.

The warp coefficients necessary to transform the first set of features onto the second set are then computed, using the methods outlined in Chapter 3, and the warp itself is performed. The RMS misalignment is computed over the entire frame and compared to the RMS error in the points themselves. The ratio of the RMS misalignment over the frame to the RMS

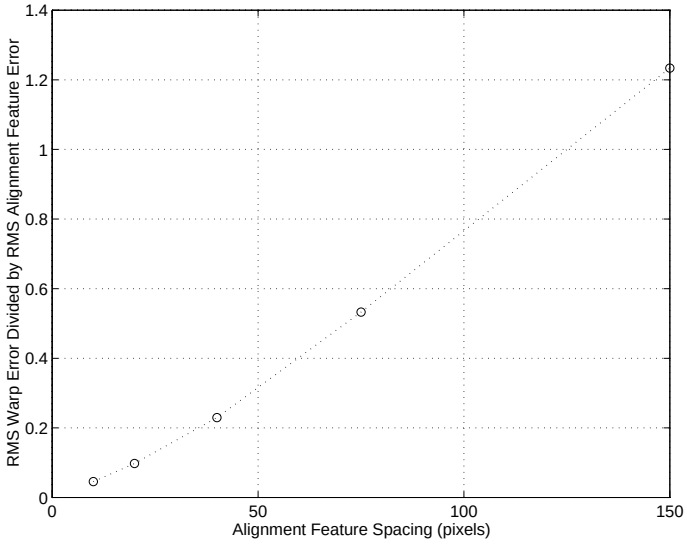


Figure 5.18: Results of numerical simulation of polynomial warp — the effect of alignment feature centroid estimation error on warp accuracy.

error of the features may be employed as a coefficient of centroid estimation noise immunity for the particular warping technique. For each warp configuration, the simulation is repeated ten times, the mean RMS errors being computed over all ten runs.

As stated in Chapter 3, the pseudo-inverse implementation of the polynomial warping technique is prone to significant numerical error. The reason for such error is that the method requires computation of the inverse $(\mathbf{W}^T\mathbf{W})^{-1}$. The condition number of $(\mathbf{W}^T\mathbf{W})$ is the square of the condition number of $\mathbf{W}$, so numerical errors are rapidly propagated.

A further simulation has been devised to compute the numerical error introduced by the pseudo-inverse implementation of the polynomial warp. The effect on the condition number of both the number of alignment features utilised for warp vector computation, and the positional regularity of those features, may be deduced.

This simulation proceeds, once again, by setting up two regular grids of features. One represents the detected feature centroids while the other represents the regular grid on to which the warp will attempt to transform the data images. The number of features modelled can be varied, subject to the requirement that there must be at least six pairs of features in order to specify the warp. In the simulation employed here, an N by N grid of features spaced at unit intervals is used. The position of those features in the 'source' grid may also be modified by adding a quantity of Gaussian noise in order to simulate the effect of a somewhat irregular array of features. When the features themselves have been defined, the matrix $\mathbf{W}$ may be determined. It is then a trivial matter to compute $\mathbf{cond}(\mathbf{W}^T\mathbf{W})$.

The values of the condition number are computed for a range of centroid estimate Gaussian noise levels, from zero through to a standard deviation of 0.45 pixel. The number of features is kept fixed as a ten by ten grid. For each Gaussian distribution standard deviation value, the computation is repeated fifty times and the mean value of the condition number is plotted. The effect of varying the number of features in the absence of Gaussian noise is also investigated. Grid sizes of three by three through to fifteen by fifteen are modelled.

Such a simulation should yield useful information about the situations in which the simple pseudo-inverse solution is adequate. High values of condition number indicate that there is probably a significant numerical error in the warp calculation. In such a case, the method of orthogonal polynomials should be employed.

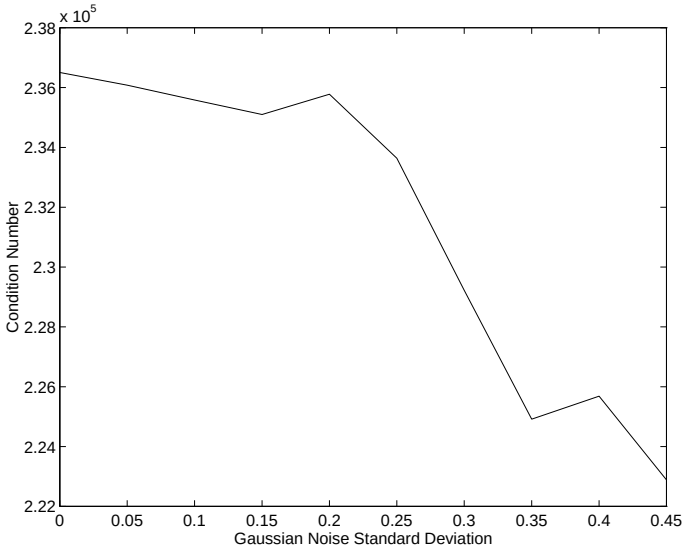


Figure 5.19: Results of the numerical simulation of the pseudo-inverse method of polynomial warp computation. This graph shows the effect of introducing noise to the otherwise regular feature arrays. A four-by-four array of features was used in this case. The standard deviation of the Gaussian noise introduced is plotted on the x axis, whereas the condition number of $(\mathbf{W}^T\mathbf{W})$ is shown on the y axis.

5.8.1.2 Results of the Numerical Simulation

Fig. 5.18 shows the results of the first numerical simulation applied to the polynomial warping technique. Inter-feature spacings of 10, 20, 40, 75 and 150 pixels were modelled. For the ten iterations of the simulation, the mean RMS warp alignment error divided by the mean RMS feature alignment error was found to range from 0.05 for a spacing of 10 pixels, to 1.23 for a spacing of 150 pixels. As demonstrated in Section 5.7, the expected error in centroid estimation is of the order of 0.1 pixels, so the expected alignment error under the polynomial warping technique is of the order of 0.005 to 0.123 pixels, depending on the spacing of the features chosen.

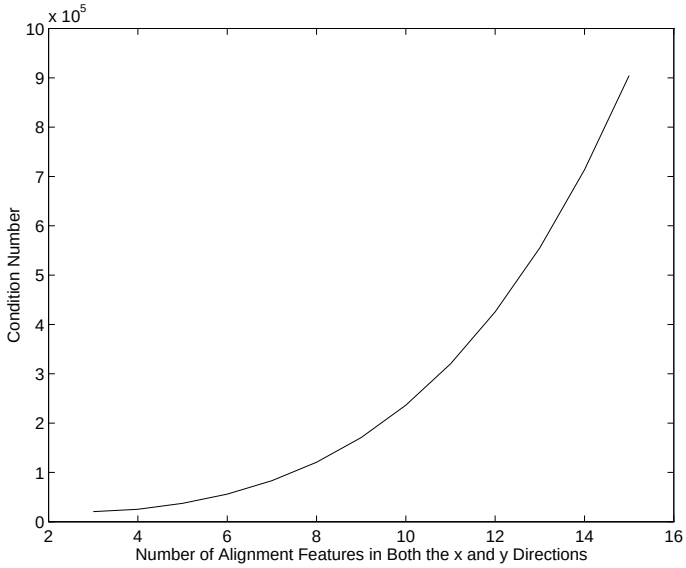


Figure 5.20: Results of the numerical simulation of the pseudo-inverse method of polynomial warp computation. This result shows the effect of varying the number of features used to define the warp. No Gaussian noise was introduced in this case. The number of features in each of the x and y directions is plotted on the x axis, whereas the condition number of $(\mathbf{W}^T\mathbf{W})$ is shown on the y axis.

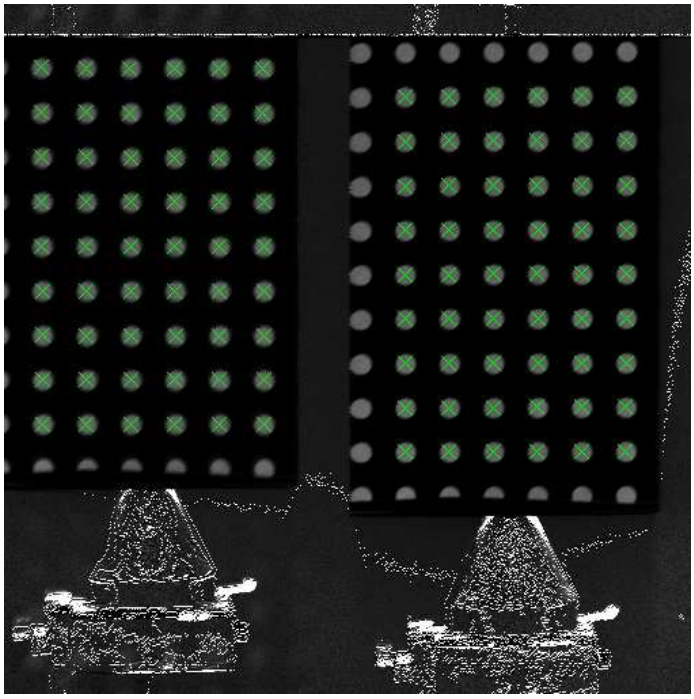


Figure 5.21: Alignment features prior to pseudo-inverse polynomial warp. The object visible at the bottom of each image is the nozzle discussed in Chaper 7.

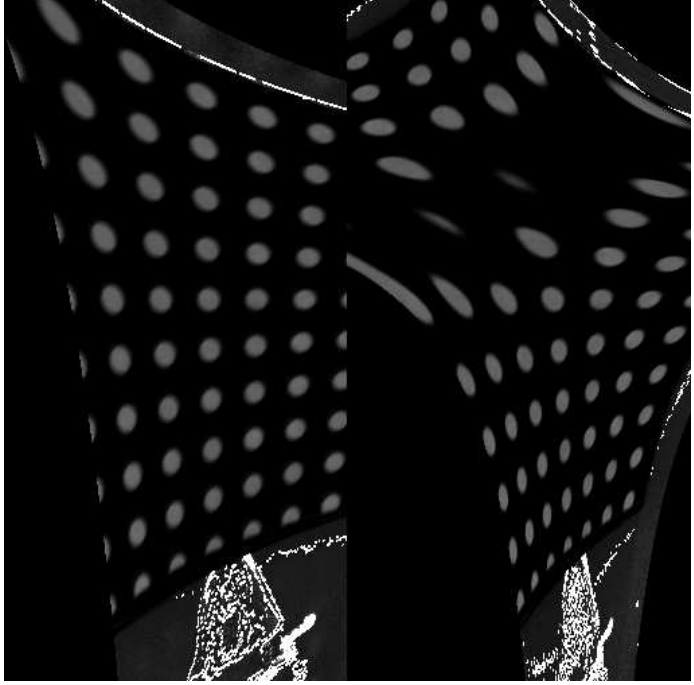


Figure 5.22: Alignment features after image transformation by pseudo-inverse polynomial warp, showing the effect of ill-conditioning on the warp calculation. Each of these images should have been warped to a regular grid.

It has been seen that for the polynomial warping technique, a smaller inter-feature spacing (implying the use of more pairs of points) markedly decreases the expected alignment error in the presence of a certain degree of centroid estimation inaccuracy. Clearly, then, as many features as possible should be employed in the computation of the warp.

As shown in Section 5.7.7, the optimal width of features, in order to minimise centroid estimation error, is around 15 pixels. It is necessary, in order to exploit the automatic location and processing of features provided by the `warp` software, that there should be sufficient space between features that horizontal and vertical dividing lines may be superimposed on the image in order to separate each feature into a bounding box. In practice, it has been seen that the space between features should be of the same order as the size of the features themselves in order for the image processing software to function reliably. This being the case, an inter-feature spacing of around 30 pixels appears to be the smallest spacing attainable. At such a spacing, the ratio of RMS warp error across the frame to RMS centroid estimate error is around 0.2. Thus an expected centroid estimation error of 0.1 pixels would imply an expected mean warp error of 0.02 pixels RMS across the frame.

For the case of the bilinear warping method, using four points positioned at the extreme edges of the frame in order to minimise error, the mean RMS warp alignment error divided by the mean RMS feature alignment error was found to be 0.7. This implies an expected alignment

error of 0.07 pixels RMS across the frame. The use of a bilinear warping technique will yield significantly higher mean RMS misalignment when compared to a polynomial technique, even if the distortion present in the optics can be modelled satisfactorily using a bilinear method alone. Using local methods of bilinear warping, or choosing four features further from the corners of the frame, will lead to more inaccurate alignment still.

Fig. 5.19 shows the results of the second numerical simulation — an investigation of the effects of feature grid regularity on pseudo-inverse polynomial warp computation accuracy. It may be seen from this figure that while increased grid irregularity leads to a lower value of the condition number of $(\mathbf{W}^T\mathbf{W})$, this reduction is very slight, and may be regarded as inconsequential.

The results of the simulation of the pseudo-inverse polynomial warp with varying numbers of features are shown in Fig. 5.20. Whereas the regularity of the grid of alignment features is seen to have little effect on condition number, the effect of the number of features utilised is more pronounced. The three by three array of features yielded $\mathbf{cond}(\mathbf{W}^T\mathbf{W}) = 2.1 \times 10^4$, whereas for a fifteen by fifteen array of features, the corresponding value was computed to be 9.0×10^5.

This investigation of the properties of the pseudo-inverse warp leads to the conclusion that this technique is best avoided. The orthogonal polynomials technique of warp computation avoids the squaring of the condition number which is implicit in the pseudo-inverse solution. The numerical accuracy of the latter technique is superior, especially in cases where there are a large number of features present.

In practice the expected numerical instability of the pseudo-inverse solution has indeed been seen to manifest itself when greater numbers of alignment features are employed. Figs. 5.21 and 5.22 show the pre-warp and post-warp results for a certain alignment grid, employing the pseudo-inverse method of polynomial warp computation. In this case the numerical instability may clearly be seen to cause extremely high levels of spatial distortion.

In order to model any non-linear distortion in the optical system effectively, a large number of features may be required. In addition, as shown above, the simulation of alignment error using the polynomial warp reveals that the greater the number of features employed for warp calculation, the smaller the overall alignment error resulting from inaccurate centroid estimation. Such results indicate that as many alignment features as possible should be employed but the use of the pseudo-inverse solution to the polynomial warp would lead to significant alignment errors in such a case. The bilinear warp and the pseudo-inverse solution

to the polynomial warp should thus be avoided. The orthogonal-polynomials solution is the preferred method.

5.8.2 Distortions Present in the Optical System

The sensitivity to inaccurately located centroids of the different warping techniques has now been established. There is also, however, the issue of how well a given technique may model the distortions inherent in the Doppler global velocimetry optical system. It has been seen, in Fig. 3.4, that the polynomial warp, unlike the bilinear warp, is capable of representing a class of non-linear distortions. The question which remains is whether the capabilities of the bilinear warp are sufficient to deal with the distortion inherent in the optics; and if not, whether the capabilities of the polynomial warp are sufficient.

In order to analyse the effect of different image processing parameters more systematically, it has been proposed that the processing should be applied to known computer generated images rather than to digitised images which have undergone unknown distortions in the optical chain. The use of such synthetic test images would appear to have no application to the determination of ideal warping techniques. In order to create a synthetic image, certain distortion characteristics must be assumed; yet such characteristics may not be representative of the optical chain in use. Having assumed the distortion properties, it is necessary to use a warping technique to create the synthetic image. If an nth order warping technique were used to create the synthetic image, then similar warping techniques of order greater than or equal to n could perfectly remove such distortion.

It is thus necessary to employ the Doppler global velocimetry system itself rather than merely using a numerical simulation as has been done for noise modelling. Each of image transformation algorithms presented in Chapter 3 have been utilised in order to obtain the results presented in Chapter 7. The reader is referred to that chapter for a discussion of the effect of the choice of technique in relation to the distortion present in the Oxford Doppler global velocimetry optical system.

5.9 Interpolation Methods

5.9.1 Sinusoidal Test Card

In Section 3.4 four different interpolation schemes were outlined: nearest neighbour, bilinear, cubic, and exponential. It was explained that these schemes possessed interpolation kernels that successively approached the ideal **sinc** function. The exponential technique is thus expected to be the superior method, followed by the cubic, bilinear, and nearest

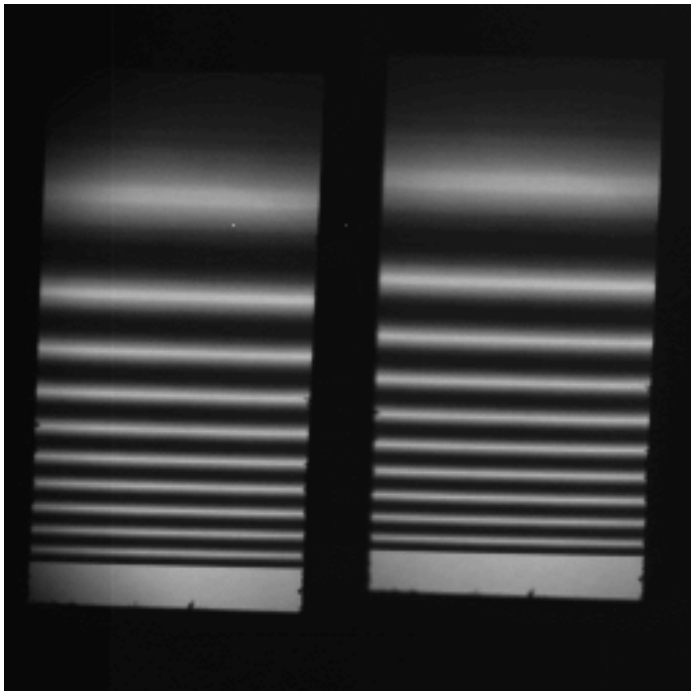


Figure 5.23: Sinusoidal Test Card.

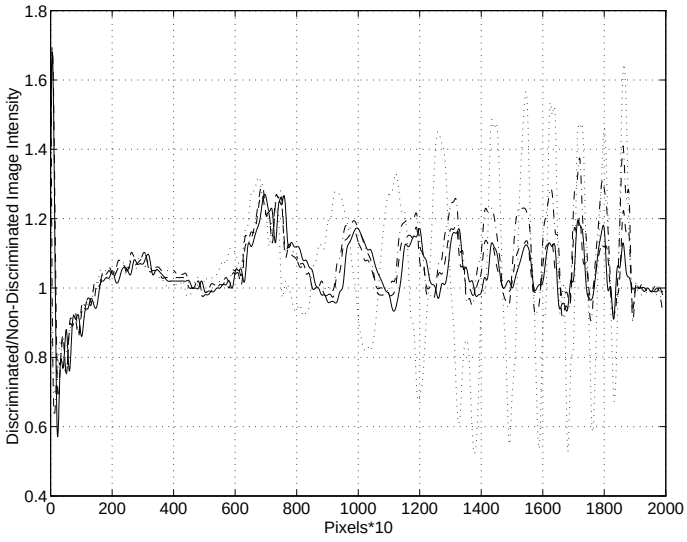


Figure 5.24: Effect of interpolation as revealed by the sinusoidal test card. Nearest neighbour is shown as the dotted line, bilinear interpolation as the dash-dotted line, cubic interpolation as the dashed line, and exponential interpolation as the solid line.

neighbour methods in that order. It is the case, however, that the nearest neighbour and bilinear schemes may be implemented at modest computational cost, whereas the cubic and exponential methods are time consuming to compute. It is necessary to investigate the quality of interpolation attainable with different interpolation techniques and also to ascertain whether the simpler methods are sufficient for Doppler global velocimetry image processing applications.

In order to evaluate the different interpolation schemes, a test-card of sinusoidally varying intensity was produced using a laser printer. The spatial frequency of the variations in intensity increases linearly across the card itself. The image of this test card, as seen through the Doppler global velocimetry optical system, is shown in Fig. 5.23. The use of such a test card should enable the expected error resulting from the use of the different interpolation techniques to be assessed for a range of spatial frequencies.

The Doppler global velocimetry system was set up and aligned as normal using a polynomial warp. The sinusoidal test card was then employed as a target. The previously computed warp coefficients were used to warp the discriminated and non-discriminated images into the same coordinate frame. A graph of the discriminated image divided by the non-discriminated image was plotted. In the absence of any velocity in the field of measurement, there should be no Doppler shift. Thus the discriminated image divided by the non-discriminated image should evaluate to unity throughout the frame.

In the presence of centroid estimation error, warp modelling error, and interpolation error, however, the ratio of the two images will not evaluate to unity. The results from the four interpolation techniques are presented in Fig. 5.24, allowing comparison between the different methods. It may clearly be seen, as expected from the analysis of Section 5.6, that the ratio departs further from the expected value as the spatial frequency of intensity variation increases.

As expected, when the interpolation kernels employed more closely resemble the optimum **sinc** function, the quality of the results obtained is seen to improve. Nearest neighbour is a poor technique, yielding unsatisfactory results for high spatial frequency image features — error is seen to increase markedly with frequency. In particular, for intensity fluctuations from a level close to zero to close to the CCD saturation value with a period of 20 pixels or less poor results may be expected from the use of this method. Bilinear interpolation is an improvement on the naïve nearest neighbour technique, but the performance of the cubic and exponential methods is significantly better. The computational cost of these methods is somewhat greater than for the simpler techniques, since more pixels must be considered

for each point. However, for a 512 by 512 image, even the exponential technique may be evaluated in around a minute on the laboratory workstation. For non real-time applications, then, computational cost alone need not preclude the use of complex interpolation methods.

By reference to Section 5.2, it may be seen that the bilinear interpolation scheme offered superior performance to the nearest neighbour method in the presence of pixel noise. The noise suppression performance of the cubic and exponential techniques was somewhat worse than that of the simpler bilinear approach. While cubic and exponential interpolation offer markedly better behaviour in the presence of high spatial frequency intensity variations, they are less advantageous in the presence of CCD array shot noise. The choice of interpolation routine is thus dependent on both the noise performance and the effect on image alignment of the scheme.

Intensity variations of such high spatial frequency are unlikely to be seen in typical flow fields. In many cases bilinear interpolation may well prove to be an adequate technique in terms of image alignment. The superior noise suppression properties of the bilinear interpolation method would suggest that it should be the preferred method for Doppler global image processing applications.

A number of different interpolation algorithms have been utilised in order to produce the results presented in Chapter 7. The reader is referred to that chapter for a discussion of the effect of the choice of technique in practice.

5.10 The Effect of Viewing Angle on Error

It may be the case, in certain circumstances, that the velocity component the experimenter is interested in does not lie in the direction $(\hat{o} - \hat{\imath})$. For example, in the case of the free jet (configuration I) measurements, discussed in Chapter 7 it is inconvenient to orientate the viewing system so as to observe the axial velocity component directly. Similarly, the convergent-divergent nozzle measurements discussed in Chapter 8 require that the system be orientated at a particular angle in order to reduce velocity sensitivity.

Very often, when the velocity is predominantly in a single direction, or the laser sheet is set up so that there are no out of plane velocity vectors, it is possible to scale the measured results by some constant k_{conf}, which may be determined from a consideration of eqn. 2.1. The magnitude of all the errors discussed here will then be scaled by the same factor k_{conf} as the velocity values themselves, leading to an increase in error in the resulting data.

5.11 Error Considerations of Two Component Measurements

Just as viewing angle must be taken into account when measuring single component velocities where the direction of the desired component differs from that to which the velocimeter is most sensitive (see Section 5.10), so the effect of observation and illumination directions must be taken into account when considering the influence of configuration on two component measurement accuracy.

In Section 2.9, three possible configurations were considered and the calculations required to convert the measured velocity components to the desired orthogonal velocity components were presented. In this section, each of these arrangements will be considered once again, this time in order to enable the expected error bounds to be deduced.

The first such configuration, shown in Fig. 2.7, uses a pair of imaging systems positioned in the same plane as the illumination vector. If the two desired orthogonal components of velocity are represented as v_x and v_y, it was concluded that

$$v_x = \frac{c(\Delta\nu_1 \sin\theta_2 - \Delta\nu_2 \sin\theta_1)}{\nu_o(\sin\theta_1 - \sin\theta_2 + \cos\theta_1 \sin\theta_2 - \sin\theta_1 \cos\theta_2)} \tag{5.54}$$

$$v_y = \frac{c(\Delta\nu_1(\cos\theta_2 - 1) - \Delta\nu_2(\cos\theta_1 - 1))}{\nu_o(\sin\theta_1 - \sin\theta_2 + \cos\theta_1 \sin\theta_2 - \sin\theta_1 \cos\theta_2)} \tag{5.55}$$

The first optical system may be seen to be measuring the value of $\frac{c\Delta\nu_1}{\nu_o}$. This quantity has units of ms^{-1} and is the value of the velocity component to which this system is most sensitive, as shown in Fig. 2.4. The quantity $\frac{c\Delta\nu_1}{\nu_o}$ may be denoted as v_{m1}. Similarly, for the second optical system, the quantity $\frac{c\Delta\nu_2}{\nu_o}$ may be denoted as v_{m2}. The values of v_{m1} and v_{m2} are subject to the errors resulting from camera pixel shot noise; iodine cell calibration error and instability; and misalignment discussed above. It remains to be seen how the error in v_{m1} and v_{m2} affects the accuracy of the desired orthogonal components v_x and v_y.

The equations above may be rewritten in terms of v_{m1} and v_{m2} in this way:

$$v_x = \frac{v_{m1} \sin\theta_2 - v_{m2} \sin\theta_1}{\sin\theta_1 - \sin\theta_2 + \cos\theta_1 \sin\theta_2 - \sin\theta_1 \cos\theta_2} \tag{5.56}$$

$$v_y = \frac{v_{m1}(\cos\theta_2 - 1) - v_{m2}(\cos\theta_1 - 1)}{\sin\theta_1 - \sin\theta_2 + \cos\theta_1 \sin\theta_2 - \sin\theta_1 \cos\theta_2} \tag{5.57}$$

It then remains to determine the partial derivatives of the desired values v_x and v_y with respect to the measured components v_{m1} and v_{m2}. The value of these partial derivatives may be used as a multiplying factor for the error found in the measured components. Thus

$$\varepsilon_{v_x} = \frac{\partial v_x}{\partial v_{m1}}\varepsilon_{v_{m1}} + \frac{\partial v_x}{\partial v_{m2}}\varepsilon_{v_{m2}} \tag{5.58}$$

$$\varepsilon_{v_y} = \frac{\partial v_y}{\partial v_{m1}}\varepsilon_{v_{m1}} + \frac{\partial v_y}{\partial v_{m2}}\varepsilon_{v_{m2}} \tag{5.59}$$

The partial derivatives may be expressed

$$\frac{\partial v_x}{\partial v_{m1}} = \frac{\sin\theta_2}{\sin\theta_1 - \sin\theta_2 + \cos\theta_1\sin\theta_2 - \sin\theta_1\cos\theta_2} \tag{5.60}$$

$$\frac{\partial v_x}{\partial v_{m2}} = \frac{-\sin\theta_1}{\sin\theta_1 - \sin\theta_2 + \cos\theta_1\sin\theta_2 - \sin\theta_1\cos\theta_2} \tag{5.61}$$

$$\frac{\partial v_y}{\partial v_{m1}} = \frac{\cos\theta_2 - 1}{\sin\theta_1 - \sin\theta_2 + \cos\theta_1\sin\theta_2 - \sin\theta_1\cos\theta_2} \tag{5.62}$$

$$\frac{\partial v_y}{\partial v_{m2}} = \frac{-\cos\theta_1 + 1}{\sin\theta_1 - \sin\theta_2 + \cos\theta_1\sin\theta_2 - \sin\theta_1\cos\theta_2} \tag{5.63}$$

The terms $\frac{\partial v_x}{\partial v_{m1}}$, $\frac{\partial v_x}{\partial v_{m2}}$, $\frac{\partial v_y}{\partial v_{m1}}$ and $\frac{\partial v_y}{\partial v_{m2}}$ may be seen to be constants for any given configuration. It is therefore necessary that care should be taken to consider the values of these partial derivatives when the system is first set up, prior to measurements being undertaken. The choice of the angles θ_1 and θ_2 should be made so as to ensure that the velocity error in the desired components is minimised. Poor choice of viewing angles may lead to extremely high velocity error in v_x and v_y. In general, it is not easy to recommend the optimal choice of angles, since it may be the case that one component is required at a higher accuracy than the other. It should be noted that the non-uniform scattering intensity with respect to observation angle implied by Mie scattering theory further complicates the choice of observation angle. It is important to ensure that sufficient intensity will be incident on the CCD array in order to minimise the effect of noise.

Turning to the second proposed configuration, as shown in Fig. 2.8, the error analysis may be repeated. If it is assumed that there is no out-of-plane velocity ($v_z = 0$), then

$$v_x = v_{m1} \tag{5.64}$$

$$v_y = \frac{v_{m1} - v_{m2}}{\sin\theta_2} \tag{5.65}$$

Hence

$$\frac{\partial v_x}{\partial v_{m1}} = 1 \tag{5.66}$$

$$\frac{\partial v_x}{\partial v_{m2}} = 0 \tag{5.67}$$

$$\frac{\partial v_y}{\partial v_{m1}} = \frac{1}{\sin\theta_2} \tag{5.68}$$

$$\frac{\partial v_y}{\partial v_{m2}} = \frac{-1}{\sin\theta_2} \tag{5.69}$$

It is necessary also to consider the final case presented in Section 2.9 where there is only one fixed observation direction but two different orthogonal illumination directions. Again assuming $v_x = 0$, for this arrangement

$$v_x = -v_{m1} \tag{5.70}$$

$$v_y = -v_{m2} \tag{5.71}$$

So

$$\frac{\partial v_x}{\partial v_{m1}} = -1 \tag{5.72}$$

$$\frac{\partial v_x}{\partial v_{m2}} = 0 \tag{5.73}$$

$$\frac{\partial v_y}{\partial v_{m1}} = -1 \tag{5.74}$$

$$\frac{\partial v_y}{\partial v_{m2}} = 0 \tag{5.75}$$

For this arrangement, $\varepsilon_{v_x} = \varepsilon_{v_{m1}}$ and $\varepsilon_{v_y} = \varepsilon_{v_{m2}}$. The technique of manipulating the light sheet propagation direction in order to obtain two measured velocity components, as described in Section 2.9, has the property that each of the two orthogonal components may be found with the same accuracy as a regular single component measurement. Unfortunately, such a technique is appropriate only for the case of repeatable flow fields.

5.12 Summary and the Development of an Error Modelling Code

Having examined a number of error sources individually, it is necessary to combine them to obtain an estimate of the expected error in Doppler global measurements using the Oxford

system. In the preceding sections each individual component of measurement error has been considered in turn and expressed either in terms of its effect on the resulting ratio of the discriminated to the non-discrimated images or in terms of its effect on the laser frequency determined.

To make use of such error estimates it is important to express each component in terms of its effect on the velocity measurement obtained. In this section each of the previously obtained error estimates will be recast in terms of its effect on velocity. Once this has been done, it is straightforward to combine the individual sources of error into a single expression for overall expected measurement error. The resulting error model has been implemented in `Matlab` code and applied to the measured data frames obtained from experimental work as explained in Chapters 6, 7 and 8.

The development of error estimation codes allows the expected error to be compared with the actual measurement error. While previous Doppler global velocimetry work has concentrated on a qualitative real-time approach, the research programme at Oxford has focussed on the development of an instrument for accurate quantitative measurements. In such an application, the ability to quantify measurement error is a key requirement. The approach adopted allows the experimenter to conveniently quantify the level of error present in measured data. In addition, the individual components of the estimated error may be estimated. Such analysis provides a good indication of the best procedure for reducing measurement error for a particular application.

The error sources have been categorised into three groups.

- Error due to noise introduced by the CCD array.
- Error introduced due to iodine absorption filter calibration uncertainty and operating instability.
- Error due to misalignment of the discriminated and non-discriminated images.

These categories will be examined in turn below.

Errors Introduced by the CCD Array	
1	CCD shot noise (eqn. 5.5) $\varepsilon_I = \left(\frac{II_{\max}}{N_{\text{fw}}}\right)^{\frac{1}{2}}$ grey levels
2	Effect of shot noise on adjusted data and normalisation images (eqns. 5.10, 5.11) $\varepsilon_{I_{\text{adjdata}}} \left(\varepsilon^2_{I_{\text{data}}} + \varepsilon^2_{I_{\text{black}}}\right)^{\frac{1}{2}}$ grey levels $\varepsilon_{I_{\text{adjnorm}}} = \left(\varepsilon^2_{I_{\text{norm}}} + \varepsilon^2_{I_{\text{black}}}\right)^{\frac{1}{2}}$ grey levels
3	Effect of shot noise on adjusted normalised data image (eqn. 5.16) $\varepsilon_{I_s} = \left(\left(\frac{\varepsilon_{I_{\text{adjdata}}}}{I_{\text{adjnorm}}}\right)^2 + \left(\frac{I_{\text{adjdata}}\varepsilon_{I_{\text{adjnorm}}}}{I^2_{\text{adjnorm}}}\right)^2\right)^{\frac{1}{2}}$
4	Effect of interpolation on shot noise (Section 5.2) factor k_i=1.00 for nearest neighbour factor k_i=0.58 for bilinear interpolation factor k_i=0.73 for cubic interpolation factor k_i=0.79 for exponential interpolation
5	Effect of shot noise on final ratio image (eqn. 5.18) $\varepsilon_R = \left(\left(\frac{k_i\varepsilon_{I_{s_d}}}{I_n}\right)^2 + \left(\frac{I_d k_i \varepsilon_{I_{s_n}}}{I^2_n}\right)^2\right)^{\frac{1}{2}}$
6	Effect on determined velocity $\frac{\partial\nu}{\partial\phi}\frac{\partial v}{\partial\nu}\varepsilon_R$ ms^{-1}

In Section 5.2 the effect of CCD shot noise was considered. The effect of shot noise on the ratio of the discriminated to the non-discriminated images was computed and expressed as ε_R. In the table above, the first five entries serve to summarise the key findings of Section 5.2. Taken together those relationships define the expected ratio error ε_R.

In summary, the properties of the CCD array potentially contribute thermal noise, read-out noise, and shot noise to the measured data. In the case of the Oxford system, it has been shown that the Peltier-cooled, slow-scan CCD camera employed exhibits levels of thermal and read-out noise which may be disregarded. Shot noise is the dominant source of error resulting from the use of a CCD array. The shot noise expected in terms of grey levels was calculated ([1] in the table above).

As was seen in Chapter 3, the acquired CCD frames are subjected to a number of image processing steps. In particular, the black frame is subtracted from both the data and normalisation frames to form the adjusted data and normalisation frames. Since shot noise is present on all CCD frames, it is necessary to consider the effect of shot noise on these adjusted frames ([2] in the table above).

The adjusted data frame is then divided by the adjusted normalisation frame in order to eliminate the effect of any vignetting or non-uniform pixel sensitivity. This processing step must also be considered when determining the effect of shot noise on the final ratio obtained.

An approximate formulation for the RMS error in the ratio of two images when each image is subject to noise was obtained and then applied to the particular case of the adjusted data and normalisation images. An expression was thus computed for the expected effect of shot noise on the adjusted normalised data image ([3] in the table above).

After image transformation has been undertaken, interpolation is a necessary step in order to map non-integer coordinates in the 'source' frame onto integer coordinates in the 'destination' frame. Interpolation may be regarded as a weighted mean of a number of pixels in the 'source' frame for each pixel in the 'destination' frame. The effect of this on the overall level of noise was examined, both analytically and numerically, for a number of interpolation schemes. A coefficient k_i representing the effect of interpolation on the overall level of noise was determined by means of numerical simulation (tabulated [4] above). The use of the numerical simulation technique revealed that the cubic and exponential techniques, which were previously assumed to have superior noise suppression properties, were in fact inferior to the bilinear method in this regard. This suggests that the implementation of these techniques is hampered by the limited numerical precision available during computation.

Finally, the discriminated image must be divided by the non-discriminated image, in order to obtain a final ratio image representing the absorption of the working iodine cell across the entire CCD frame. It has been seen that both the discriminated and non-discriminated images are subject to noise (as given by table entry [3] above). It was seen that the expression for the expected RMS error in the ratio of two images may be applied yet again in order to determine the RMS error ε_R in the final ratio ([5] in the table above).

In order to convert the estimated ratio error due to shot noise in the acquired frames into an estimate of velocity error it is necessary to take account of the transmission function of the iodine cell and the sensitivity of the Doppler global apparatus to the measured velocity component ([6] in the table above). To this end, the value of $\frac{\partial v}{\partial \nu}$ — which may be seen to represent the rate of change of velocity with respect to Doppler-shifted frequency — must be considered. $\frac{\partial v}{\partial \nu}$ is dependent on the position of the optical system and the illuminating light sheet, and may be obtained for the configuration in use by a consideration of eqn. 2.1. The term $\frac{\partial \nu}{\partial \phi}$, representing the rate of change of frequency with respect to iodine cell transmission ratio, must also be considered. A value for $\frac{\partial \nu}{\partial \phi}$ may be obtained from a consideration of the iodine cell transmission profile about the operating point.

Errors Introduced by the Iodine Cell	
1	Iodine cell stability (Section 5.4) 1.5 MHz
2	Iodine cell calibration (eqn. 5.20) $\left(2.9^2 + \left(\frac{\Delta\nu}{117}\right)^2\right)^{\frac{1}{2}}$ MHz
3	Laser no-shift frequency (Section 5.3) 2.9 MHz
4	Effect on determined velocity $\left(\left(1.5^2 + 2.9^2 + 2.9^2 + \left(\frac{\Delta\nu}{117}\right)^2\right)^{\frac{1}{2}} \times 10^6\right) c/\nu_o$ ms^{-1}

It has been seen that the use of an iodine absorption cell as a frequency-to-intensity conversion filter also results in the introduction of a number of error sources. Such errors result from inaccuracies in the cell calibration process, discussed in Section 5.3, and also from instability of the cell in day-to-day operation, examined in Section 5.4.

The transmission profile of the iodine absorption filter is dependent on temperature. The cold-finger temperature and the ambient temperature each have an effect on the transmission characteristics of the cell. The cold-finger of the cell is temperature controlled using a Peltier device. Although the transmission is very sensitive to cold-finger temperature, the precision of temperature control is sufficiently good that ambient temperature fluctuations affecting the temperature of the body of the cell have a far greater effect on the cell properties. Due to typical ambient temperature fluctuations, the RMS error in the Doppler-shift frequency determined by reference to the iodine cell transmission profile was estimated to be of the order of 1.5 MHz.

The remaining error sources associated with the iodine cell may be attributed to the mode-hop based calibration technique described in Section 4.1. The mode-hop spacing of the laser is specified as 143 MHz, with an uncertainty of ± 0.4 MHz. In addition the laser emission frequency is subject to a thermal drift of approximately 60 MHz/°C. These two effects, together with an estimated RMS error of 2.9 MHz due to the curve-fit approximation described in Section 4.1, result in an uncertainty in the iodine transmission curve obtained by means of the calibration procedure. An estimate of this error was computed ([2] in the table above).

Each time the transmission curve is employed to determine a velocity value during the processing of Doppler global data, the 2.9 MHz RMS error caused by the curve-fit approximation is incurred. The aggregate effect of iodine cell instability, calibration error and further curve-fit approximation error on the determined velocity may be determined. This may be done

by simply combining entries [1], [2] and [3] in the table above, using the method of eqn. 5.1. Since the error estimates were originally computed in units of MHz, it is necessary to multiply by 10^6. In order to convert the error in determined frequency into an error expressed in terms of velocity a conversion factor of c/ν_o must also also be taken into account. The resulting expression for overall error introduced by the iodine cell is shown as entry [4] in the table above.

Errors Introduced due to Misalignment	
1	Centroid estimation error (Section 5.7.6) 0.1 pixel systematic error
2	Effect of warp on alignment (Section 5.8.1.2) factor k_w=0.7 for bilinear factor k_w=0.2 for polynomial
3	Effect of misalignment on ratio (eqn. 5.26) $\varepsilon_R = \frac{\left(\frac{\partial I_n}{\partial x}\delta x+\frac{\partial I_n}{\partial y}\delta y\right)\phi(v)+\frac{\partial \phi}{\partial v}\left(\frac{\partial v}{\partial x}\delta x+\frac{\partial v}{\partial y}\delta y\right)\left\{I_n(x,y)+\frac{\partial I_n}{\partial x}\delta x+\frac{\partial I_n}{\partial y}\delta y\right\}}{I_n(x,y)}$
4	Effect on determined velocity (eqn. 5.27) $\varepsilon_V = \frac{\partial v}{\partial \phi}\varepsilon_R$ ms $^{-1}$
5	Effect on determined velocity $\frac{\partial v}{\partial \phi}0.1k_w\frac{\left(\frac{\partial I_n}{\partial x}+\frac{\partial I_n}{\partial y}\right)\phi(v)+\frac{\partial \phi}{\partial v}\left(\frac{\partial v}{\partial x}+\frac{\partial v}{\partial y}\right)\left\{I_n(x,y)+0.1k_w\left(\frac{\partial I_n}{\partial x}+\frac{\partial I_n}{\partial y}\right)\right\}}{I_n(x,y)}$ ms $^{-1}$

It is clear that precise image alignment is essential for accurate quantitative Doppler global velocimetry measurements to be obtained. Image processing methods have been developed in order to align the discriminated and non-discriminated images to sub-pixel accuracy. Such image processing methods, described in Chapter 3, consist of three stages: alignment feature centroid estimation, image transformation, and image interpolation. In Sections 5.6, 5.7, 5.8 and 5.9 the errors associated with each step in the image processing algorithms have been evaluated, allowing overall error to be estimated and different image transformation and interpolation methods to be compared.

Alignment feature centroid estimation has been studied both analytically and by means of numerical simulation, described in Section 5.7. A large number of factors contribute to the accuracy with which the position of features may be determined. The effect of CCD array shot noise, non-integer width features, non-zero black levels in off-centre bounding-boxes, sampling in a finite number of pixels, quantisation, and the so-called systematic error have all been discussed and evaluated.

It has been determined by means of numerical simulation that if appropriate image processing operations are performed, the effect of CCD shot noise, non-integer width features, non-zero black levels, sampling and quantisation on the determined centroid position may be reduced to a value of around 0.01 pixel. In this case, the systematic error will dominate. The presence

of systematic error dictates that the minimum centroid location error that may be achieved is 0.1 pixel. The use of suitable image processing algorithms effectively eliminates the other possible sources of centroid estimation error, so the RMS centroid error obtained in practice may be assumed to be 0.1 pixel.

In order to achieve this it is necessary to employ a cooled slow-scan CCD array in order to reduce noise, and to use a sufficient number of quantisation levels. The camera in use in the Oxford Doppler global velocimetry system allows the use of 32768 quantisation levels, and has the property that shot noise is the dominant noise source. It is thus ideally suited to the accurate estimation of alignment feature centroid locations. The use of video cameras, which have higher levels of thermal noise and read-out noise and a limited number of quantisation levels, would lead to appreciably higher levels of centroid estimation error.

The optimum size of alignment features has been determined to be 15 pixels, as imaged on the CCD array. The use of significantly larger features would increase the centroid estimation error due to the presence of shot noise, whereas the use of significantly smaller features would result in an increase in error due to the effect of non-integer width features.

The black level threshold of the `warp` software should be set to a value of around 0.05 of the saturation intensity in order to overcome the effect of shot noise present in the nominally 'black' segments of the bounding boxes. Lower values of threshold may allow error due to shot noise to distort centroid position, whereas higher values may introduce significant error due to the suppression of valid data at the extremes of each feature.

The optical system should be carefully focussed and aligned, and the lenses themselves should be of good quality, in order to reduce the extent of the line-spread function as far as possible. This will have the effect of minimising the systematic error present.

The degree of error in the centroid estimates themselves is not the sole determinant of the alignment error across the entire frame. The choice of warping algorithm is also significant in this regard. A numerical simulation is described in Section 5.8. Using this simulation, each alignment point under consideration is subject to noise, sampled from a Gaussian distribution. It was found that, averaged across the frame, the misalignment in the final image when the bilinear warp was performed was a factor of 0.7 of the RMS misalignment of the alignment points themselves. For the case of the polynomial warp, the degree of misalignment across this image was found to be dependent on the number of points used to specify the warp: the greater the number of points employed, the less sensitive the overall alignment error is to the centroid estimation error of the individual features. Since the

optimal feature size is around 15 pixels, and there is a need to ensure that sufficient space is present around each feature, a spacing of 30 pixels seems the minimum practical. At this value of spacing, the average misalignment across the image was found to be a factor of around 0.2 of the RMS misalignment of the alignment points themselves.

As described in Chapter 3, the polynomial warp may be computed by a simple pseudo-inverse method, or by means of orthogonal polynomials. A further numerical simulation was employed to investigate the stability of both methods of warp computation. It was found that the pseudo-inverse method becomes increasingly unstable as the number of alignment features is increased, possibly leading to very large errors in alignment. Since a large number of features is desirable in order to reduce overall alignment error due to centroid estimation inaccuracy, the use of the pseudo-inverse method is to be discouraged. The orthogonal polynomials method of computing the polynomial warp is thus preferred, using features (of size 15 pixels) spaced approximately 30 pixels apart.

The final image processing technique to be examined in terms of its error properties is that of interpolation. Section 5.9 describes a method of evaluation of the performance of different interpolation techniques by means of the use of a test card of sinusoidally varying intensity. Such a technique allows the frequency response characteristics of the different interpolation schemes to be assessed. It was found that while the nearest neighbour technique had very poor properties, leading to significant errors, the bilinear, cubic and exponential techniques were successively better. This is to be expected since, by reference to Chapter 3, it may be seen that these techniques were devised successively to approach the ideal **sinc** interpolation function. In practice, however, such high spatial frequencies of intensity variation are rarely encountered in a flow field to be measured. The choice of interpolation routine therefore becomes a trade-off between the high frequency performance and the technique's ability to suppress noise. It was seen above that the cubic and exponential methods offered less immunity to shot noise than the simpler bilinear approach. The bilinear interpolation scheme may thus be preferred for application to Doppler global velocimetry image processing.

Having made an assessment of the likely expected misalignment between the discriminated and non-discriminated images, it is necessary to determine how such misalignment may affect the resulting determined velocity. The effect of misalignment is highly dependent on the intensity and velocity gradients present in the flow field under consideration. Section 5.6 presents an analysis whereby the local intensity and velocity gradients may be taken into account to determine the overall error in ratio due to a misalignment $(\delta x, \delta y)$ (eqn. 5.26 [3] in the table above). The consideration of local intensity and velocity gradients, together with

the configuation-dependent value $\frac{\partial v}{\partial \nu}$ and the iodine-cell characteristic $\frac{\partial \nu}{\partial \phi}$, allow an estimate of velocity error due to misalignment to be determined for each pixel of the result image (eqn. 5.27 [4] in the table above).

Substituting in a value of $0.1k_w$ for the pixel alignment error expected due to centroid estimation systematic error and the effect of the image transformation algorithm, a final expression for the expected velocity error due to misalignment ([5] in the table above) may be obtained.

Other Error Sources		
1	Adverse reflections of light sheet (Section 5.5)	configuration dependent
2	Effect of viewing angle (Section 5.10)	factor of k_{conf} configuration dependent

It was seen in Section 5.5 that adverse reflections of the incident laser illumination may lead to velocity determination errors in the vicinity of reflective boundaries. Since this effect is wholly dependent on the geometry of the flow field and the angle of the illuminating light-sheet, no attempt has been made to quantify the resulting error for the general case.

The consideration that the velocity component of interest may differ from the measured component given by eqn. 2.1 was introduced in 5.10. Where this is the case, the overall level of velocity error will be multiplied by a constant k_{conf}, which may be determined from a consideration of eqn. 2.1 with regard to the current measurement geometry. As seen in Section 5.11, the consideration of the choice of viewing angle is of particular interest in the case of two component measurements.

The expressions for estimated velocity error derived in this chapter have been used to implement error modelling software. This modelling code may be applied to Doppler global velocimetry measurement frames. In this way, predicted levels of error resulting from shot noise, iodine cell calibration error, and misalignment may be obtained for all points in the measurement plane. Such error predictions may then be compared with the measurement error experienced when undertaking practical Doppler global measurements. Such error prediction calculations are presented alongside measured velocity results in Chapters 6, 7 and 8.

The error prediction and analysis undertaken in this chapter is significantly more detailed than that reported elsewhere. It is the belief of the author that such error analysis is a necessary foundation for the development of a velocimeter capable of accurate quantitative

measurements. The division of the error predictions into different categories allows the user of the Doppler global system to see readily which error source is dominant in any particular measurement configuration. In this way, effective precautions to ensure the best possible accuracy of velocity determination may be taken. The use of error prediction codes also enables error-bounds to be placed on the measured data throughout the measurement frame, as may be seen in the following chapters.

Chapter 6

Velocimetry Measurements of a Rotating Disc

6.1 Motivation

The first application of the Doppler global velocimeter was chosen to be the measurement of a single velocity component of a rotating disc. These results are presented as a elementary confidence test in both the Doppler global velocimetry system and the image processing software. The velocity of the rotating disc has been determined for all points on its surface. Since the angular velocity of the disc may be easily found, comparisons may be made between the measured and the expected velocity. A number of other workers (such as [Komine *et al.* 1991], [Meyers and Komine 1991], [Meyers 1992], [Morrison, Gaharan and DeOtte 1994], and also [Ford and Tatum 1994]) have published similar rotating disc measurements.

It will be shown that the alignment of the discriminated and non-discriminated images is not crucial in this case, due to the low intensity gradients present. Thus the rotating disc measurements may be obtained to a satisfactory level of accuracy using relatively crude image transformation techniques. For this reason the performance of the image processing code has not been evaluated in detail for the rotating disc. Rather, the processing of the disc data is used to illustrate the concepts expounded in Chapter 3. Each of the image processing steps is examined with reference to real measured data.

The rotating disc experimental arrangement consists of a cardboard disc of 70 mm diameter mounted on the spindle of a DC motor. The disc may be rotated at up to about 14000 r.p.m. An optical detector and a cathode ray oscillospcope are used to obtain a measurement of the angular velocity of the disc. For the measurements presented in this chapter, the angular velocity of the disc was found to be 13600 r.p.m., corresponding to a disc tip speed of 50 ms^{-1}.

Just as a seeded flow may be measured by determining the Doppler shift in the light scattered from the seeding particles, so the velocity of a solid body may be resolved from the light it scatters.

6.2 Configuration

The measurement configuration is shown in Fig. 6.1. The laser emission from the argon ion laser is expanded into a planar sheet of 80 mm width and 1.25 mm thickness, which is used

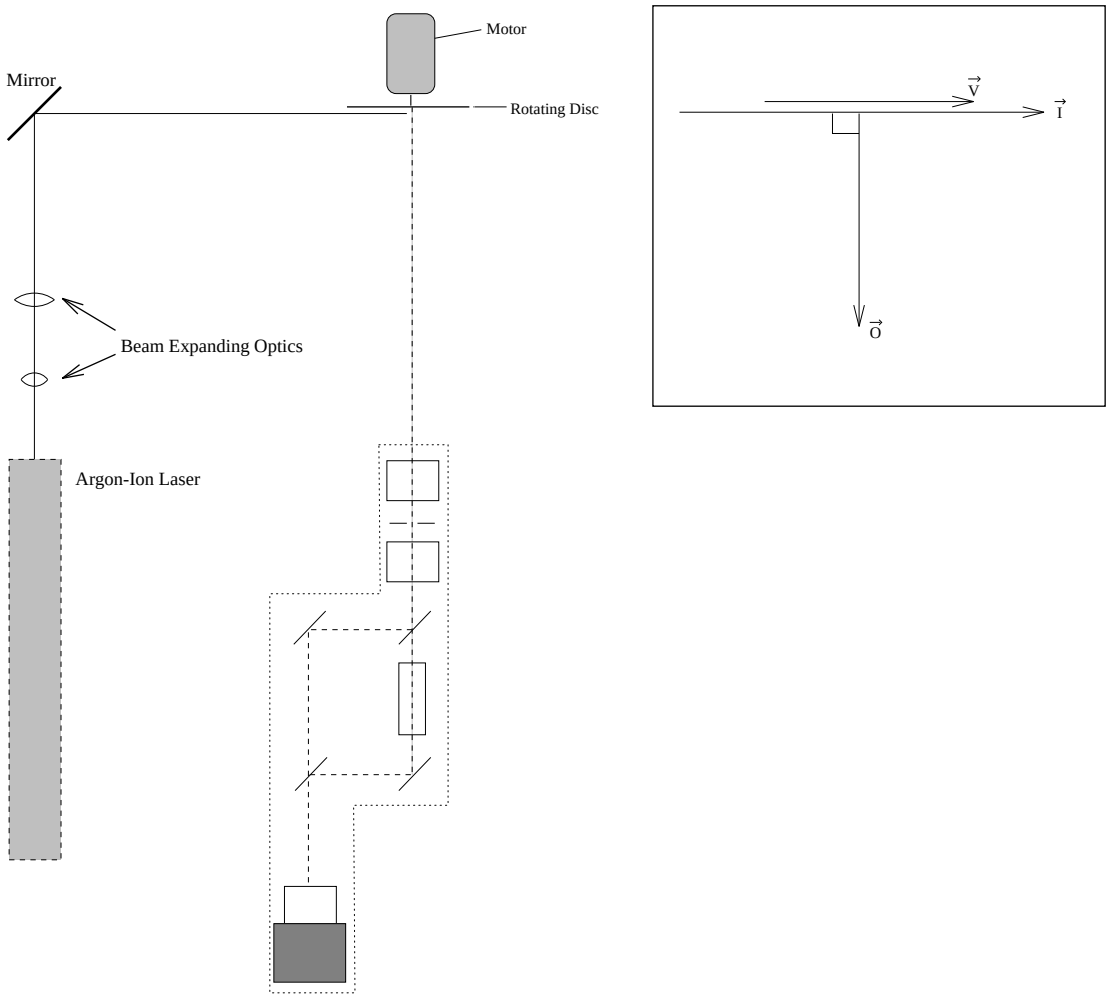


Figure 6.1: Rotating disc — viewing configuration.

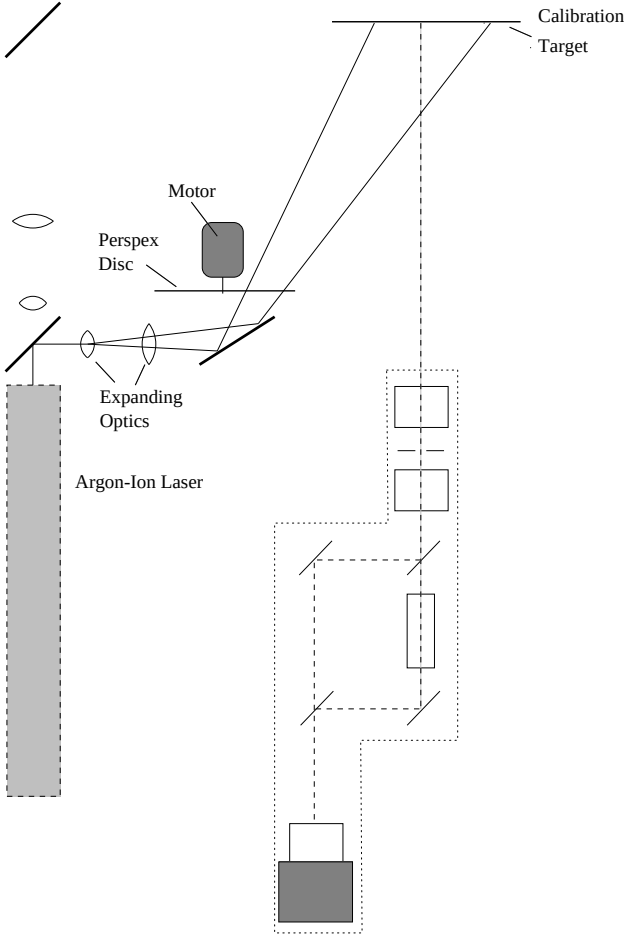


Figure 6.2: Illumination configuration for calibration.

to illuminate the disc from one side. The Doppler global velocimetry system receiving optics are then positioned to observe the disc itself from a direction normal to the face of the disc.

The viewing configuration is such that

$$\hat{\imath} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \tag{6.1}$$

$$\hat{o} = \begin{pmatrix} 0 \\ -1 \end{pmatrix} \tag{6.2}$$

Since there is no out-of-plane velocity expected, it may be said that

$$\vec{V} = \begin{pmatrix} v \\ 0 \end{pmatrix} \tag{6.3}$$

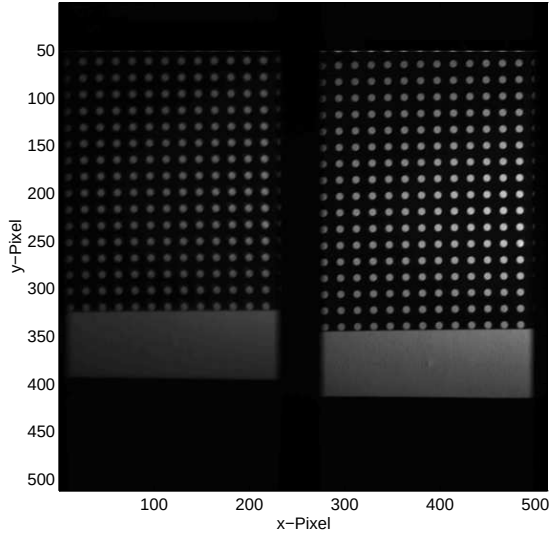


Figure 6.3: Calibration frame used for rotating disc measurement.

so, using eqn. 2.1

$$\Delta\nu = \frac{-\nu_o v}{c} \tag{6.4}$$

The system is sensitive both to velocity components in the plane of the disc, and to those perpendicular to it. In this application, however, all velocities are in the plane of the disc, and thus the measured values obtained may be interpreted as measurements of the velocity in the plane of the disc itself. The sensitivity of the system in this configuration is 1.94 MHz per ms^{-1}.

As a result of the rotation of the disc, both positive and negative Doppler shifts, representing positive and negative measured velocities, are expected. The magnitude of the shifts is expected to be comparatively small (from 0 MHz to a maximum of around 97 MHz), due to the relatively low velocities expected (from 0 ms^{-1} to around 50 ms^{-1}). Consequently, the laser may be tuned to the centre of one of the absorption wings of the iodine absorption filter profile. Such a setting of laser frequency will allow both positive and negative Doppler shift frequencies to be resolved. The low velocities and hence small Doppler shifts expected mean that there is no danger of running out of iodine cell bandwidth.

The `warp` application is employed, implementing the image processing techniques outlined in Chapter 3. Figs. 6.3 to 6.16 illustrate different stages of the image processing as applied to the rotating disc measurement. Each of these images is obtained directly from the `warp` display.

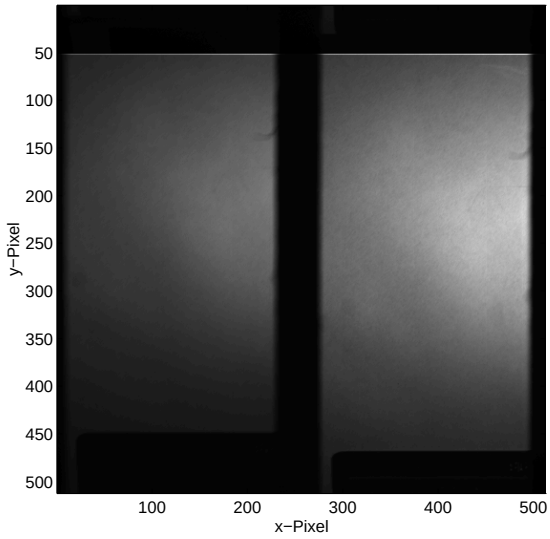


Figure 6.4: Calibration normalisation frame used for rotating disc measurements.

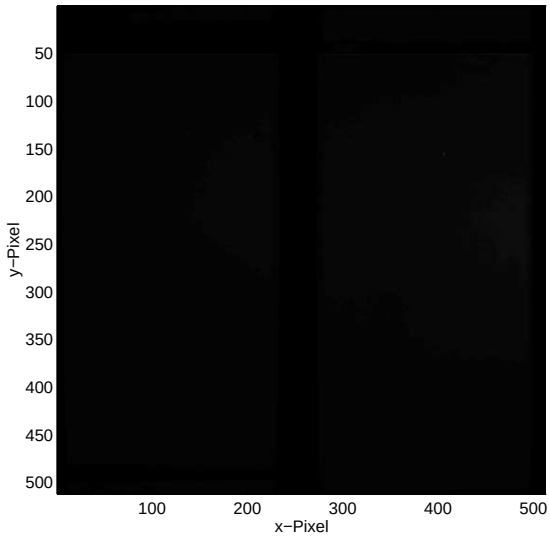


Figure 6.5: Black frame used for rotating disc measurements.

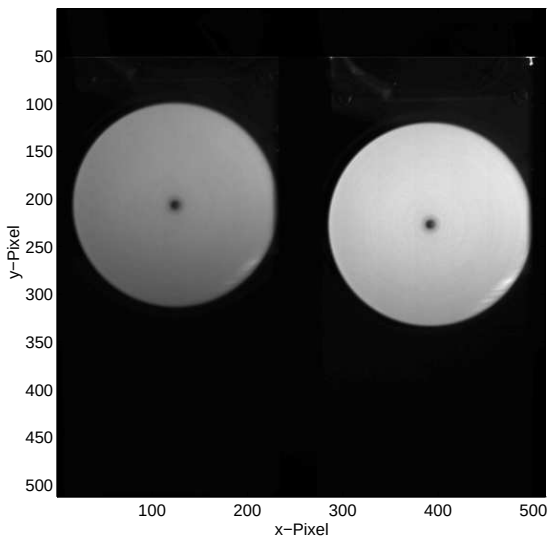


Figure 6.6: Data frame used for rotating disc measurements.

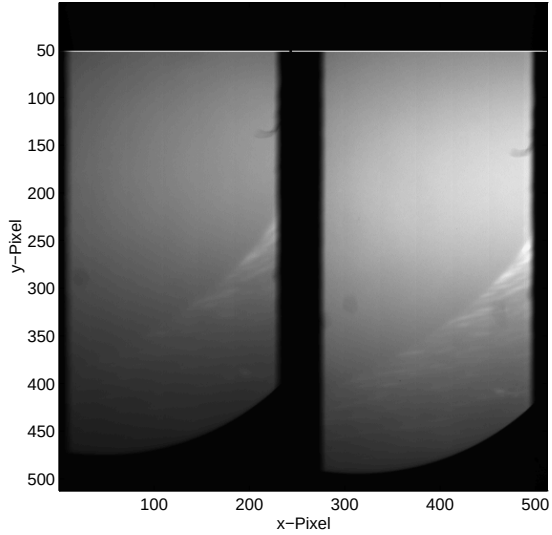


Figure 6.7: Normalisation frame used for rotating disc measurements.

6.3 Flow Measurement and Image Processing Procedure

An image of an alignment card, of regularly spaced white dots on a black background placed in the plane of the disc, is acquired as shown in Fig. 6.3. The alignment card is prepared in accordance with the findings of Chapter 5: white squares of such a size that they occupy around 10 pixels on the CCD frame are used; and the mark to space ratio of the dots is unity, so that a sufficient number of features (a twelve by seventeen array) may be viewed in the CCD frame allowing sufficient space to enable straightforward isolation into bounding boxes.

In order to avoid any chromatic distortion, whereby the position of the observed features varies with a change in the frequency of the illuminating light source, laser illumination is employed for the acquisition of this image. This presents an additional problem: the coherent nature of the laser illumination may result in interference patterns, manifesting themselves as a speckle pattern in the acquired image that may affect the intensity of the features and thus distort the estimated centroid positions.

The problem of interference patterns may be overcome by means of the illumination configuration of Fig. 6.2. A mirror is inserted into the path of the laser beam, diverting it through a pair of cylindrical lenses. The expanded beam is then passed through a rotating perspex disc, the surface of which has been roughened. The rotating disc has the effect of diminishing the coherence of the laser illumination and hence breaking up the speckle pattern. Thus, the calibration target may be illuminated with laser frequency light without incurring the problems associated with the speckle pattern.

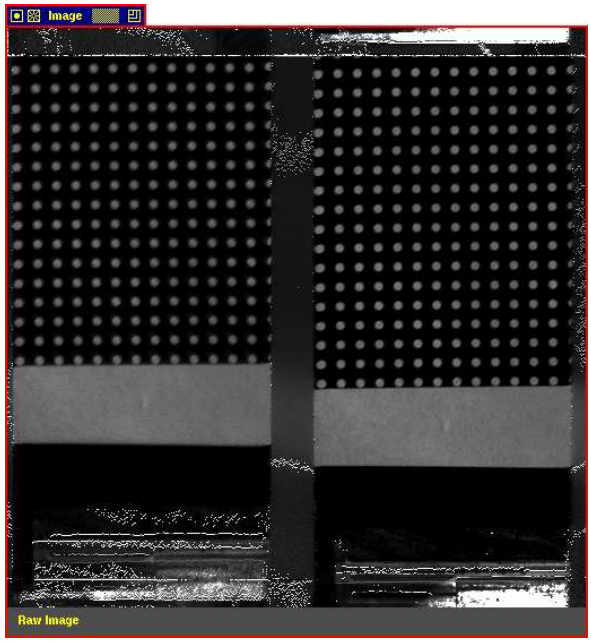


Figure 6.8: Normalised calibration frame with black level subtracted, as computed by the `warp` application.

A calibration normalisation frame and a calibration black frame, Figs. 6.4 and 6.5, are acquired by replacing the calibration target with a plain white card and a plain black card respectively. For all these frames, the camera aperture, the CCD gain value, and the exposure time are maintained at the same values.

The calibration frames having been obtained, the data frames themselves may be acquired. Fig. 6.6 shows the raw image of the rotating disc as acquired by the Doppler global velocimetry system. As explained in Chapter 3, a normalisation and a black frame for the data image are also required.

For each data frame acquired, the `IODINE` reference cell monitoring software described in Section 4.1 is utilised. The use of this program enables accurate mean reference cell transmission values to be obtained for each exposure. Thus the stability of the laser emission frequency may be assessed at the time of image acquisition, and any necessary laser frequency adjustments may be made to the normalisation and black images using the `norm_adjust` code after acquisition is complete. It is necessary to know the reference cell transmission ratio at the time of acquisition of the data images in order to compute quantitative velocity results using the `viewresults` code.

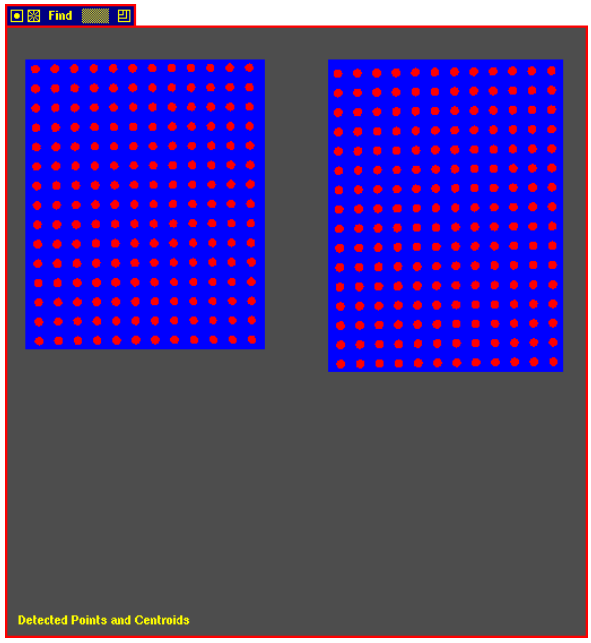


Figure 6.9: Amplitude threshold segmentation used to identify features, as computed by the `warp` application.

The image of a slowly rotating PTFE disc is used for the data normalisation frame. This disc is placed in the position previously occupied by the rotating disc for which measurements are to be taken. PTFE was found to have suitably uniform scattering properties. The disc should be set to rotate slowly in order to further suppress the effect of any speckle pattern. If the speed of rotation of this normalisation disc were too great, it would impart a significant Doppler shift of its own, and hence distort the measured results. The normalisation frame is shown in Fig. 6.7.

The data black frame is simply an image of the measurement configuration in the absence of laser illumination. This frame should enable the effect of stray illumination to be removed from the processed images.

Having acquired all six images, the `warp` application may be invoked in order to perform the required image processing. Initially, the calibration frame, the calibration normalisation frame, and the calibration black frame are loaded. The calibration black frame is subtracted from both the other frames, and the adjusted calibration frame divided by the adjusted calibration normalisation frame. The resulting frame is shown in Fig. 6.8. At this stage the effects of any illumination non-uniformity or any vignetting in the optical system should have been removed from the calibration frame data.

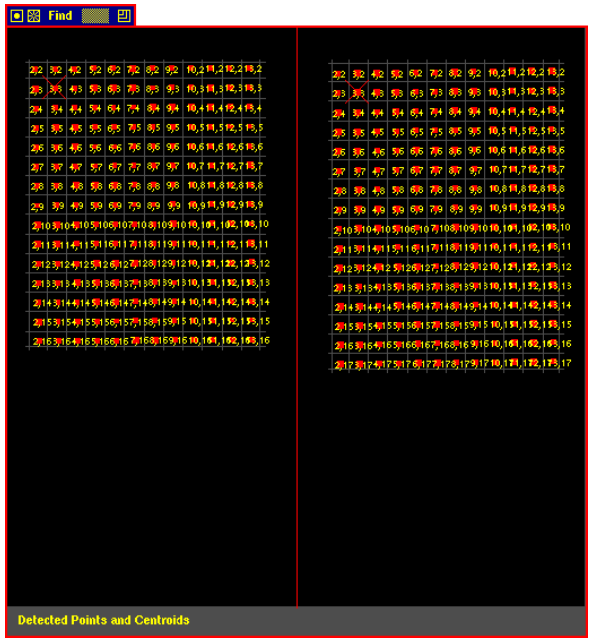


Figure 6.10: Bounding boxes computed to isolate features, as computed by the `warp` application.

Amplitude threshold segmentation is then performed in order to highlight the features of the calibration frame, as shown in Fig. 6.9. Bounding boxes may be computed by passing dividing lines through the highlighted features in the horizontal and vertical directions, as may be seen in Fig. 6.10. For each of these bounding boxes, the centroid of the feature so contained may be estimated. These computed centroid positions are shown in Fig. 6.11.

When the centroid positions for both the discriminated and non-discriminated images have been estimated, the correspondence between the two sets of features must be established. This may be performed by the `warp` application entirely automatically: the resulting correspondence is displayed in Fig. 6.12.

The determination of the correspondence of the two sets of estimated centroids is sufficient to enable the warp vectors to be computed. In this case an orthogonal polynomials warp has been employed. In Fig. 6.13, the transformation has been applied to the calibration image itself, in order to illustrate the effect of the warp in aligning the discriminated and non-discriminated images. The warp vectors may be saved to disk for future use, so that multiple measurements undertaken using the same geometry require only a single computation of the warp vectors.

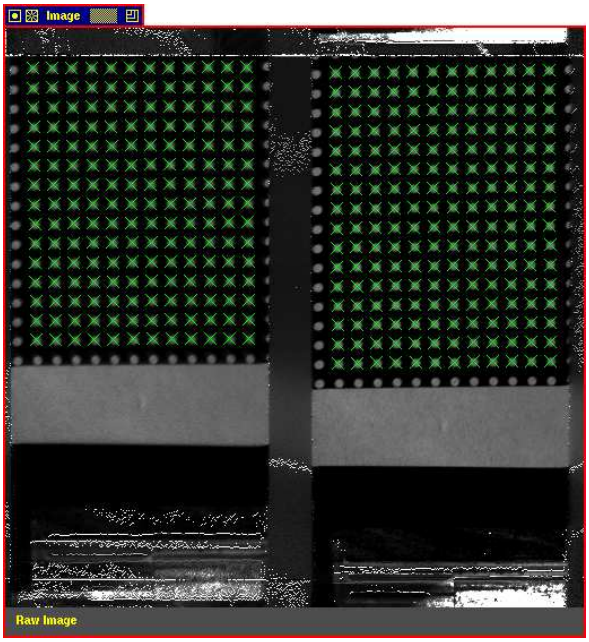


Figure 6.11: Computation of centroid position, as computed by the `warp` application.

The data images themselves may now be processed. Fig. 6.14 shows the rotating disc data frame with black level subtracted, divided by the normalisation frame also with black level subtracted. The previously computed warp vectors may then be applied to this data in order to effect alignment: the resulting frame is shown in Fig. 6.15.

Now that the discriminated and non-discriminated images have been accurately aligned, simply dividing the discriminated image by the non-discriminated image, as shown in Fig. 6.16, yields a representation of the transmission through the iodine absorption filter for all points in the measurement region. Subsequent processing, taking into account the iodine transmission profile using the application `viewresults` as described in Section 4.1, enables this data to be directly mapped into velocity values. A representation of the measured velocity is shown in Fig. 6.17.

6.4 Doppler Global Velocimetry Measurements

As seen in Section 6.2, consideration of eqn. 2.1 reveals that the measured component of velocity is such that the system is sensitive to velocity components both in the plane of the disc and perpendicular to the disc. The only velocity expected is in the plane of the disc, and the system is sensitive only to the component of that velocity which lies in the horizontal direction. Thus a series of horizontal bands is expected, representing a linear

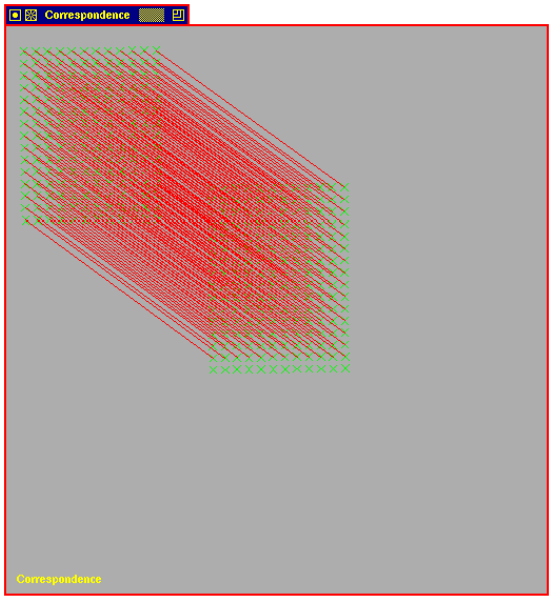


Figure 6.12: Establishing the correspondence between features in the discriminated and non-discriminated images, as computed by the `warp` application.

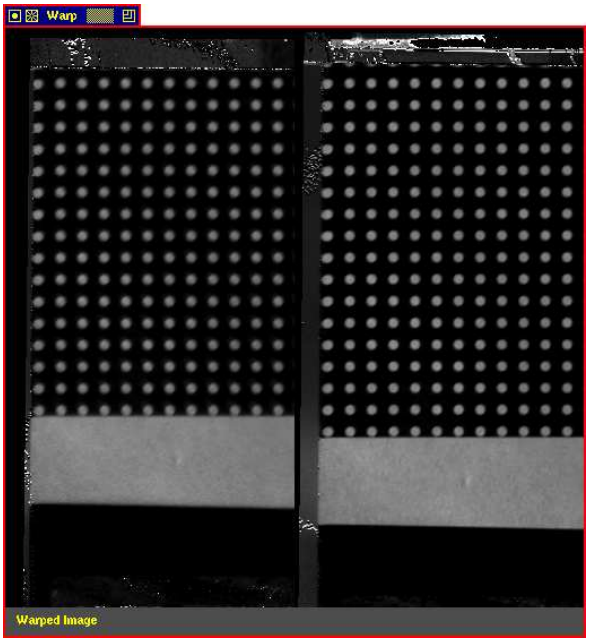


Figure 6.13: Warped calibration frame to illustrate the effect of the image transformation, as computed by the `warp` application.



Figure 6.14: Normalised rotating disc data frame with black level subtracted — as computed by the `warp` application.

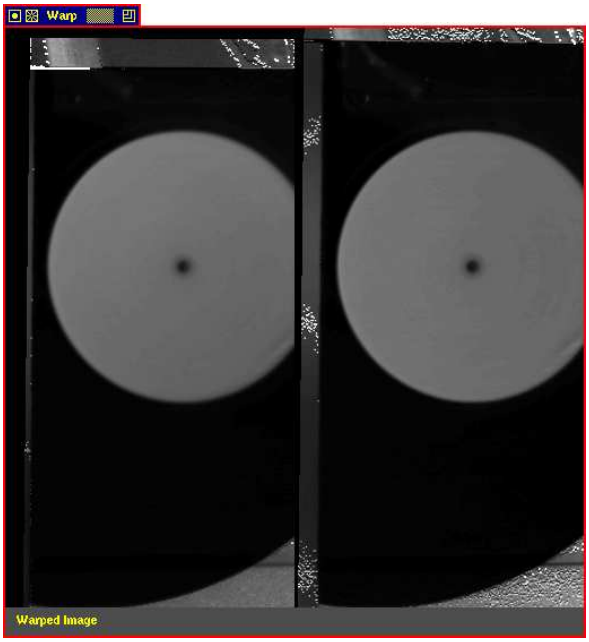


Figure 6.15: Warped rotating disc data frame — as computed by the `warp` application.

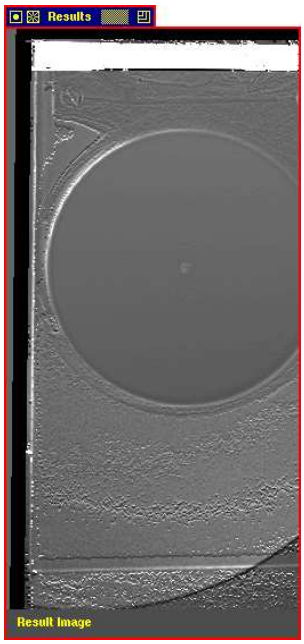


Figure 6.16: The discriminated image divided by the non-discriminated image, as computed by the `warp` application using an orthogonal polynomial warp with bilinear interpolation.

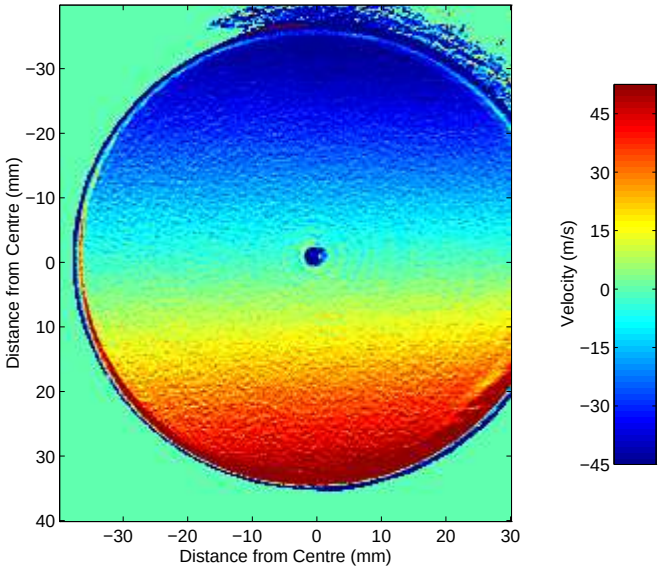


Figure 6.17: Rotating disc velocity, as measured by Doppler global velocimetry.

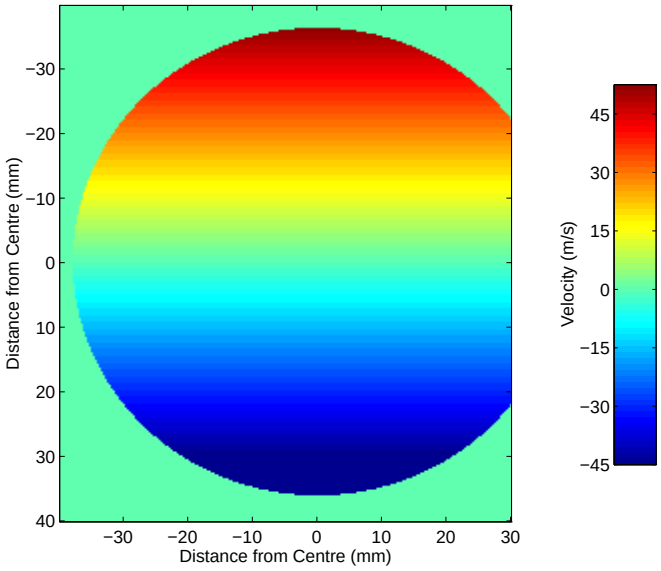


Figure 6.18: Synthetic disc velocity field.

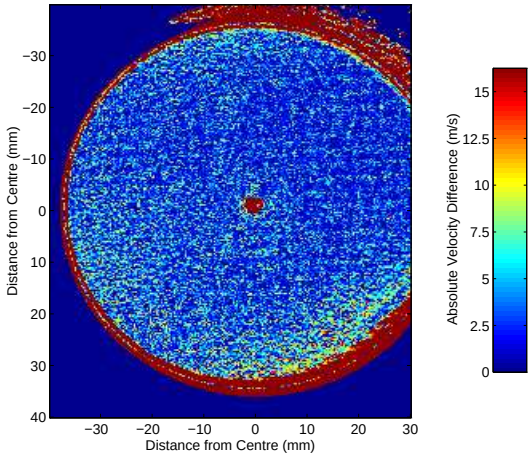


Figure 6.19: Absolute velocity difference between Doppler global velocimetry measurements of the rotating disc and the synthetic disc velocity field.

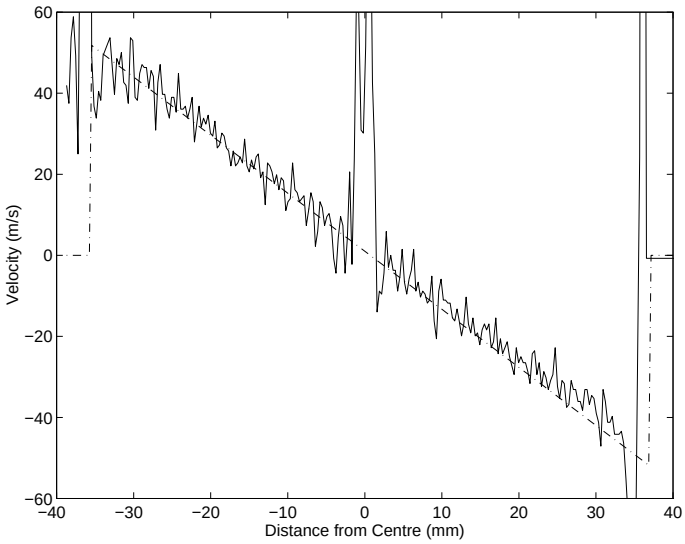


Figure 6.20: Centre-line velocity. The Doppler global velocimetry measurements for a single row of pixels are shown as a solid line. The synthetic disc velocity predictions are shown as a dotted line.

variation in velocity from zero at the hub to a peak value at the rim. As the disc is rotating, positive Doppler shifts will be seen to one side of the hub, and negative Doppler shifts on the other side. Fig. 6.17 shows a maximum absolute measured velocity of around 50 ms^{-1}. The rotating disc result thus demonstrates the ability of the Doppler global velocimetry system to resolve even relatively low velocity flow fields.

Fig. 6.18 shows a synthetic image of the expected disc velocity. Since the angular velocity of the rotating disc is known to be 13600 r.p.m., it is possible to recreate the entire expected velocity field so as to allow comparison with the Doppler global velocimetry measurements. Such a comparison is shown in Fig. 6.19, where the absolute velocity difference between the Doppler global velocimetry measurement and the predicted synthetic image is shown. The significant feature observed at the centre of the image is, in fact, the spindle at the centre of the disc itself. Even a small degree of misalignment will result in such a sharply defined intensity feature appearing as spurious noise on the results image.

Fig. 6.20 shows a line plot of velocity through the centre of the disc itself. The dotted line shows the corresponding velocity data for the synthetic disc image. Fig. 6.21, shows the same line plot, but this time a mean velocity over an area spanning ten pixels each side of the centre-line is computed.

The effect of different image processing techniques on the resulting final results may be reliably assessed only by implementing those techniques and applying them to real data. However, the case of the rotating disc is not a promising candidate for such experimenta-

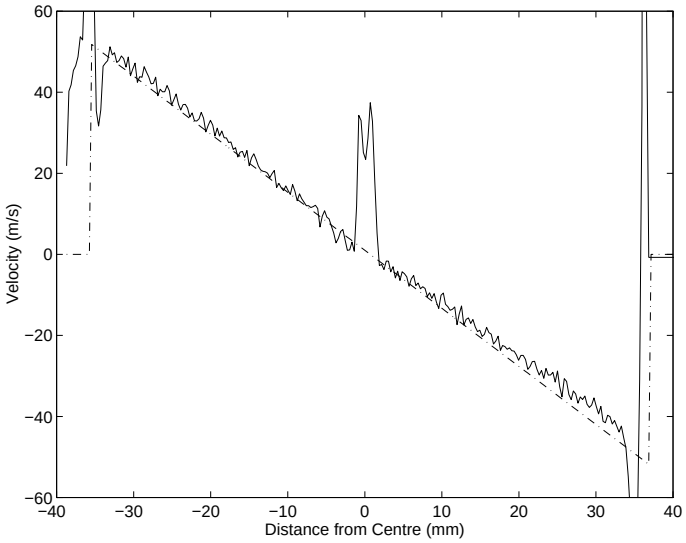


Figure 6.21: Centre-line velocity. The Doppler global velocimetry measurements, taken as a mean of ten pixel rows each side of the centre-line, are shown as a solid line. The synthetic disc velocity predictions, taken as a mean over the same region, are shown as a dotted line.

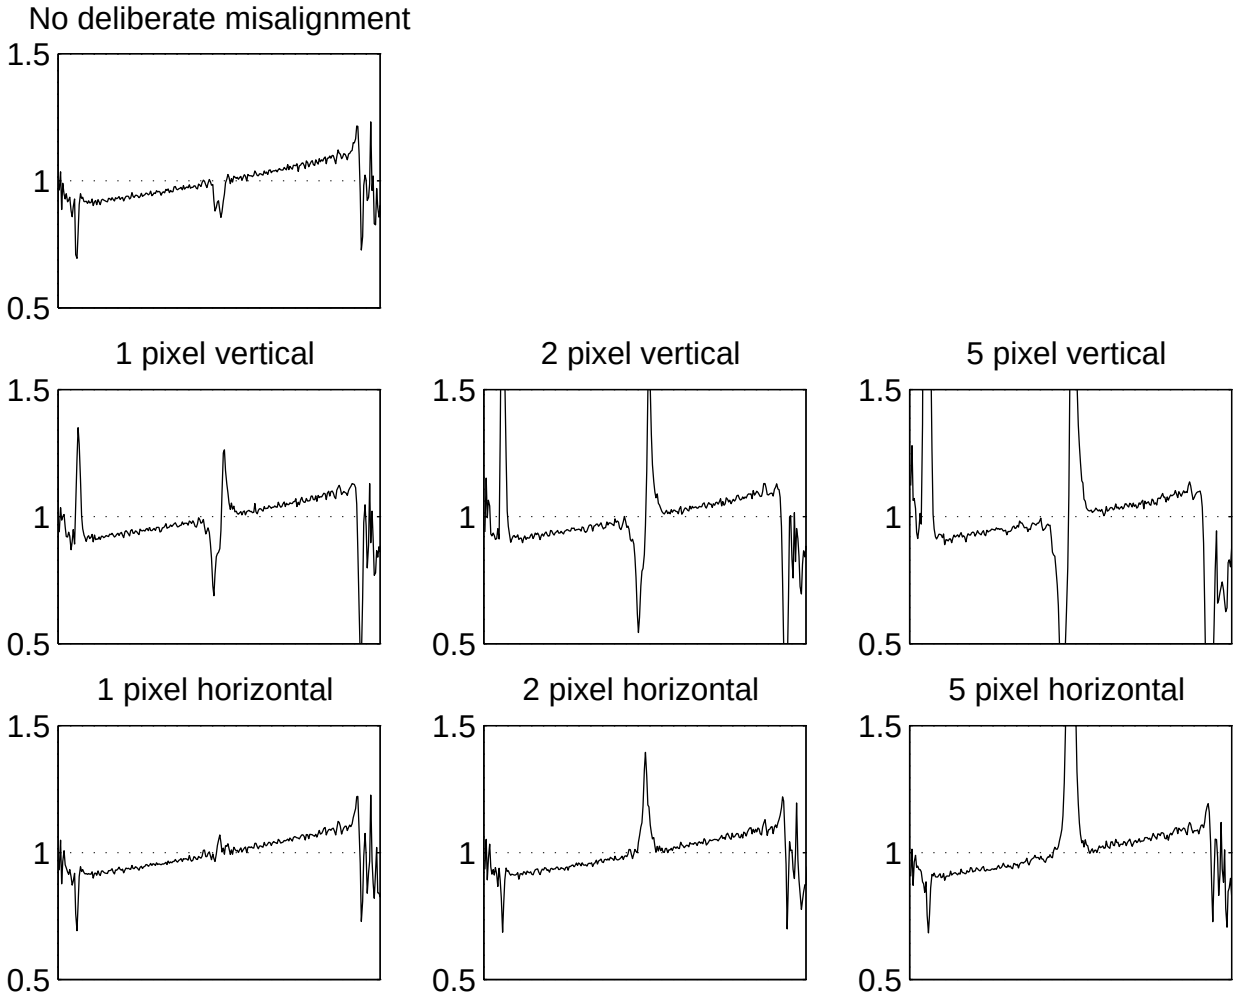


Figure 6.22: The effect on ratio of deliberately misaligning the discriminated and non-discriminated images by one, two and five pixels in the vertical and horizontal directions.

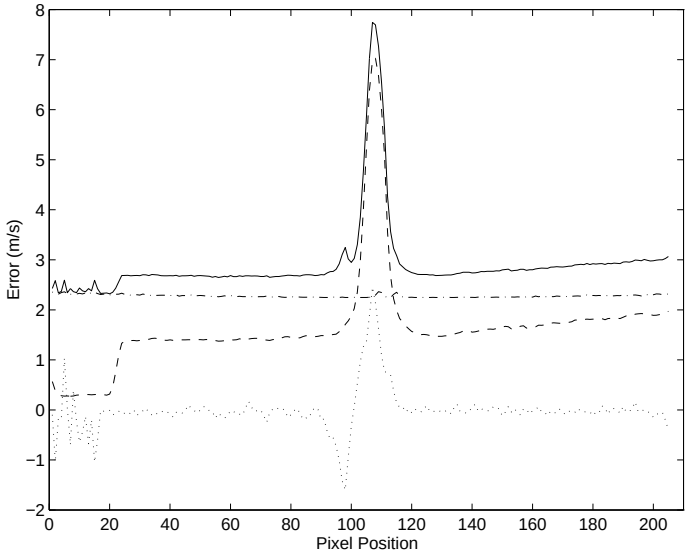


Figure 6.23: Predicted measurement error for the rotating disc measurement, along the centre-line of the disc, using the analysis of Chapter 5. The dotted line represents the predicted error due to misalignment, the dashed line the predicted error due to shot noise, the dot-dashed line the predicted error due to the iodine cell, and the solid line the overall RMS error.

tion. The low spatial frequency of intensity variation across the image is such that even extreme misalignment may be tolerated without significant deleterious effect being apparent in the resulting measurement. Fig. 6.22 shows the effect of deliberately misaligning the discriminated and non-discriminated images by one, two, and five pixels in each of the vertical and the horizontal directions (Here, horizontal is taken to mean parallel to the image x-axis, and vertical to mean parallel to the image y-axis.) In all cases, with the exception of a small region around the periphery of the disc, acceptable results are yielded. A flow field exhibiting a higher spatial frequency of intensity variation, such as the free jet of Chapter 7, must be employed in order to assess the effect of different image processing methods.

6.5 Error Analysis

The error estimates of Chapter 5 may be applied to the case of the rotating disc measurements. The expected errors due to misalignment, shot noise and iodine cell calibration may be computed for each pixel in the measurement frame, making use of the data, normalisation, and black images acquired. In computing the error estimates, the case of a polynomial warp with bilinear interpolation has been considered. The observed error may be readily deduced from a comparison of the Doppler global measurements and the synthetic disc image.

Fig. 6.23 shows the estimated error along the centre-line of the disc itself. All the quantities are shown as RMS values, with the exception of the misalignment error which is shown as a

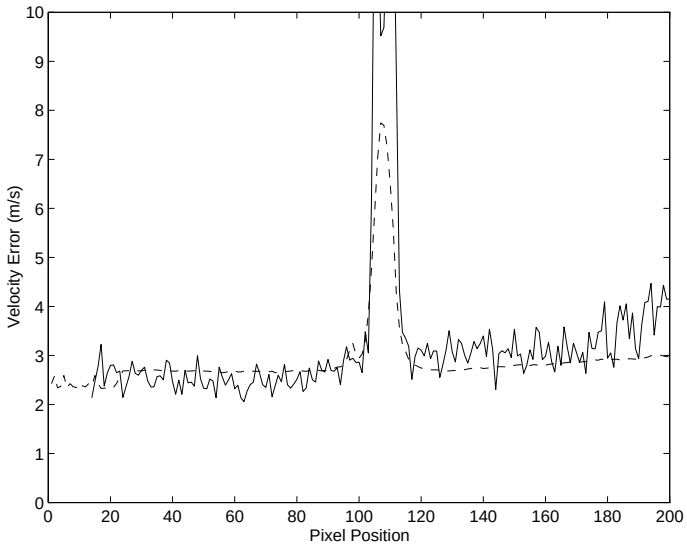


Figure 6.24: Comparison of predicted error with observed error, along the centre-line of the disc. The dashed line is the predicted RMS error and the solid line the observed measurement error.

signed quantity in order to allow the sense of the misalignment error to be readily seen. As expected, the error due to misalignment is negligibly small across the entire disc, with the exception of a region at the centre. The reason for the small influence of misalignment is the low spatial frequency of any intensity or velocity variations in the field under consideration. The larger value at the centre is due to the effect of a sharp intensity discontinuity in the middle of the disc, where the disc is attached to the motor spindle. The small effect of misalignment error is further evidence for the fact that alignment is relatively unimportant in this case. This adds weight to the previously stated conclusion that the rotating disc test case is inadequate for testing the different image processing algorithms conclusively.

The main error contribution derives from CCD array shot noise (approximately 1.5 ms^{-1} RMS) and iodine cell calibration errors (approximately 2.5 ms^{-1} RMS). It may be seen that the shot noise contribution increases gradually across the disc surface from around 1.0 ms^{-1} to a value of around 2.0 ms^{-1}. This is due to the slight intensity variation across the face of the disc. The shot noise contribution is significantly higher at the centre of the disc, due, once again, to the presence of the motor spindle. A low intensity region is present in both the data and normalisation images. A low value in the normalisation images leads to an amplification of the effect of the shot noise present when the adjusted data image is divided by the adjusted normalisation image. The iodine cell contribution, on the other hand, is more or less constant across the disc surface, increasing only slightly at positions further from the disc centre. The velocity of the disc is sufficiently low that there is no significant change in this error component within the frame. The overall RMS error prediction is also

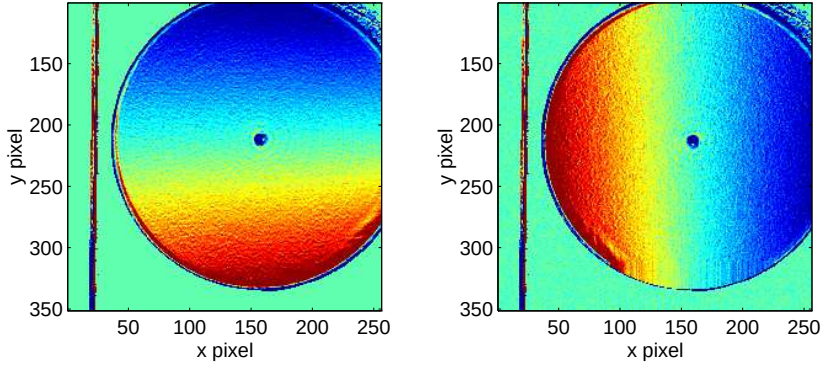


Figure 6.25: Two velocity representations of the spinning disc, using two different light sheet propagation directions $\hat{\imath}$ in order to perform the measurement.

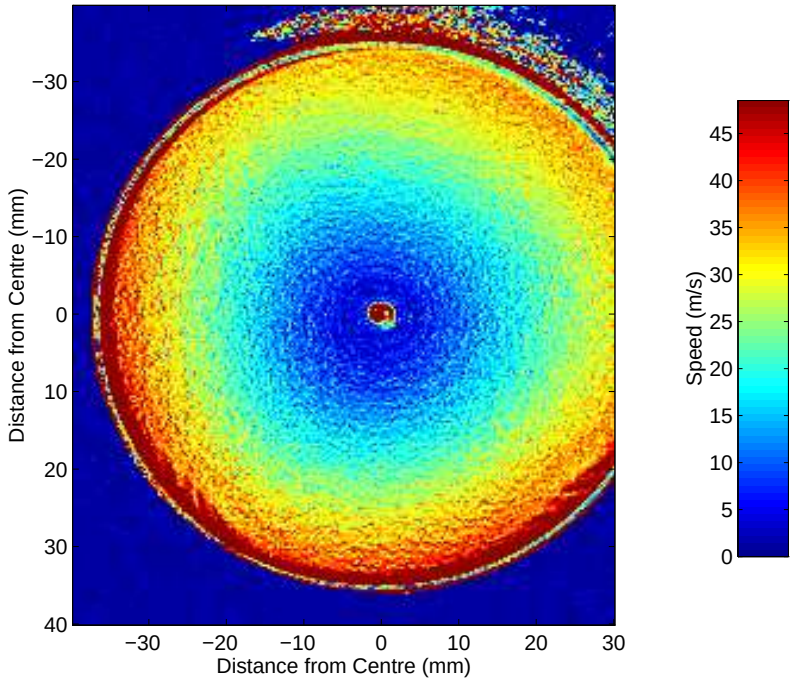


Figure 6.26: Representation of the speed of the rotating disc, computed by combining the two individual velocity components measured

shown in Fig. 6.23.

Fig. 6.24 allows a comparison of the RMS predicted error with the observed error, again along the centre-line of the disc. The observed error is computed as the absolute value of the discrepancy of all the pixels in the appropriate CCD row with the corresponding pixels in the synthetic image. It may be seen that the observed error corresponds with the expected error values extremely well. Both the expected and the observed errors may be seen to be around 3 ms^{-1} across the disc, representing a 6% error close to the top and bottom of the disc image where the measured velocity component is at its highest.

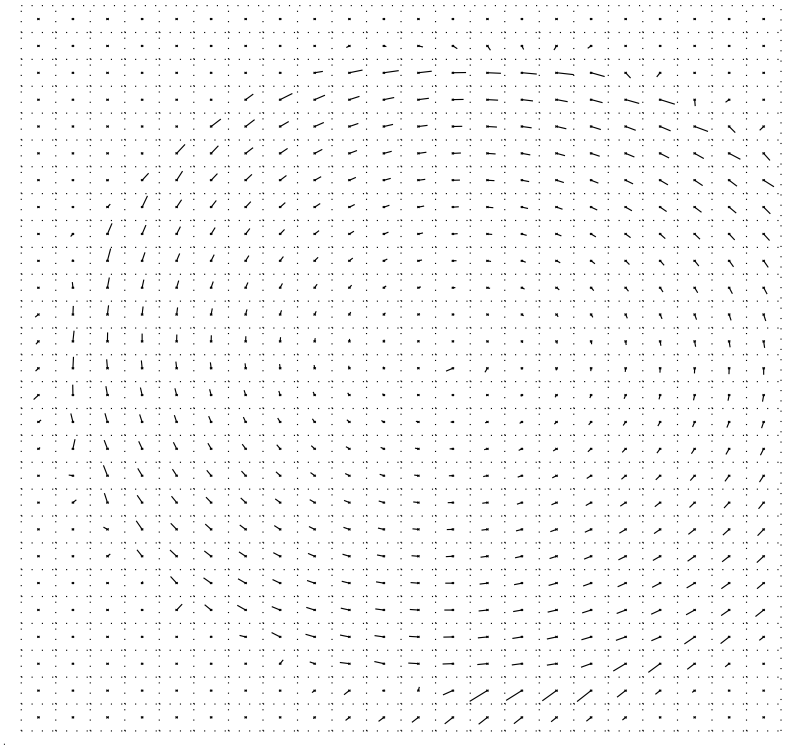


Figure 6.27: The two orthogonal velocity components of the spinning disc, represented in the form of a vector map

6.6 Two Component Doppler Global Velocimetry Measurement

A two component measurement of the rotating disc was undertaken in order to illustrate the principles outlined in Section 2.9 and to test the image processing software `2component` and `2component_vis` described in Section 4.2. Two orthogonal components of velocity were directly obtained using the method of manipulating the light sheet propagation direction $\hat{\imath}$ and performing a pair of single component Doppler global velocimetry measurements using the same configuration of the receiving optical system.

The first exposure was obtained using the configuration described in Section 6.2 above. In that way, the velocity component in the plane of the disc orientated horizontally was measured. For the second exposure, the direction of propagation of the laser light sheet was adjusted using an arrangement of three mirrors (as shown in Fig. 2.9) so that the vector $\hat{\imath}$ was orientated vertically, the light propagating upwards. It is crucial, when reconfiguring the light sheet in such a way, to ensure that the position of the plane of the light sheet is unaltered. In this way a further velocity component, also in the plane of the disc but vertically orientated, was obtained.

Once a pair of exposures, each allowing the measurement of a single component of velocity in the plane of the disc has been obtained, each must be processed, using the `warp` software in order to effect alignment of the discriminated and non-discriminated images, and the `viewresults` software to allow the determination of quantitative velocities. Since a single optical system is used, and it remains undisturbed between the two exposures, no consideration need be given to the problem of aligning the data resulting from each exposure with one another.

Fig. 6.25 shows the resulting velocity map images obtained from `viewresults` processing. The left hand image depicts the component of velocity in the horizontal sense, and the right hand image that in the vertical sense. The `2component` software may then be employed in order to process these two separate results images into a single image allowing both the orthogonal components of velocity to be combined. Fig. 6.26 depicts the speed across the surface of the disc, obtained from the `2component` output data files. Using `2component_vis` allows the velocity map to be visualised as a regularly spaced set of vectors. Such a representation may be seen in Fig. 6.27.

It is clear, as explained in Section 5.11, that the error expectation for each component of the measured velocity using this technique of light sheet manipulation is the same as that obtained for the single component measurements presented above. Taking the error expected in both the v_x component value and the v_y component value as ε_v, it may be stated that the error expected in the speed determined by such a technique, ε_s is

$$\varepsilon_s = \sqrt{(v_x + \varepsilon_v)^2 + (v_y + \varepsilon_v)^2} - \sqrt{v_x^2 + v_y^2} \tag{6.5}$$

Taking the error ε_v in each component to be around 3 ms^{-1} as obtained above, the error in the speed value obtained may be seen to vary across the measurement frame, reaching an estimated maximum value of around 4.3 ms^{-1}.

6.7 Summary

In this chapter, the results of the application of the Oxford Doppler global velocimetry system to the measurement of a single velocity component of a rotating disc have been presented. Each step of the image processing algorithm introduced in Chapter 3 has been shown, applied to real data frames.

It has been seen that due to the low velocity and intensity gradients present in the data, alignment error is non-critical in this particular application. The rotating disc velocity field

is thus not well suited to performing a rigorous evaluation of the merits of different warping and interpolation techniques.

Nonetheless, the error analysis of Chapter 5 has been applied to the data frames obtained and an estimate computed. The angular velocity of the rotating disc may be determined allowing a synthetic disc velocity field to be computed. Comparison of the measured disc velocity field with the synthetic image allows measurement error to be readily evaluated. This measurement error may be compared to the predicted error obtained from the analysis of Chapter 5.

It has been found that the main error contribution derives from CCD array shot noise (approximately 1.5 ms^{-1} RMS) and iodine cell calibration errors (approximately 2.5 ms^{-1} RMS). As expected, the error due to misalignment is negligible in this case. The overall predicted RMS error of around 3 ms^{-1} agrees well with the observed error.

A simple two component test case using a single optical system has been set up. The two component measurements obtained appear to be satisfactory. This test may be considered a proof of concept for the two component processing software `2component` and `2component_vis`.

In conclusion, the Doppler global velocimetry system has shown itself to be capable of accurate measurements of low velocity flow-fields. This is in itself a demanding test of the measurement technique, since Doppler shift values are so low for such a case. The error analysis techniques which have been developed have been seen to model the observed error faithfully.

Chapter 7

Velocimetry Measurements of a Free Jet Flow

7.1 Motivation

The Oxford Doppler global velocimetry system has been applied to the measurement of the velocity of a turbulent round free jet. The motivation for undertaking such measurements of a well known flow regime is to establish confidence in the system's ability to measure the velocity of a seeded air flow. The choice of a seeded free jet emerging into unseeded ambient air is an appropriate test-case since such a flow field is well documented; large intensity gradients are present in the frames acquired; and there is a significant variation in velocity across the profile of the jet.

The free jet flow is presented as a more demanding test case than the rotating disc described in Chapter 6. The free jet flow is easy to set up in the laboratory, and has been the subject of many previous investigations. A comparison of the Doppler global measurements with published data is thus possible. Whereas the results obtained from the rotating disc proved to be relatively insensitive to image alignment, this is not the case with the free jet. The large intensity and velocity gradients present mean that alignment must be performed to a high degree of accuracy if satisfactory results are to be achieved. With this in mind, the free jet is used as a test-case for the image processing software. The different warping and interpolation techniques discussed in Chapter 3 are compared and evaluated with application to real data. The expected error performance of these techniques is considered, and compared to the values obtained in practice.

Similar free jet Doppler global measurements have been presented by a number of other workers (such as [Komine and Brosnan 1991], [Komine *et al.* 1991], [Meyers and Komine 1991], [Morrison, Gaharan and DeOtte 1994], [Thorpe, Ainsworth and Manners 1995] and, in detail, by [Thorpe 1996]). Very few publications, however, have provided any detailed error analysis of the measured data. The results presented in this chapter are thus significant in that they represent one of the first attempts to correlate measured Doppler global velocimetry data with published results and to compare the apparent measurement error with a detailed model.

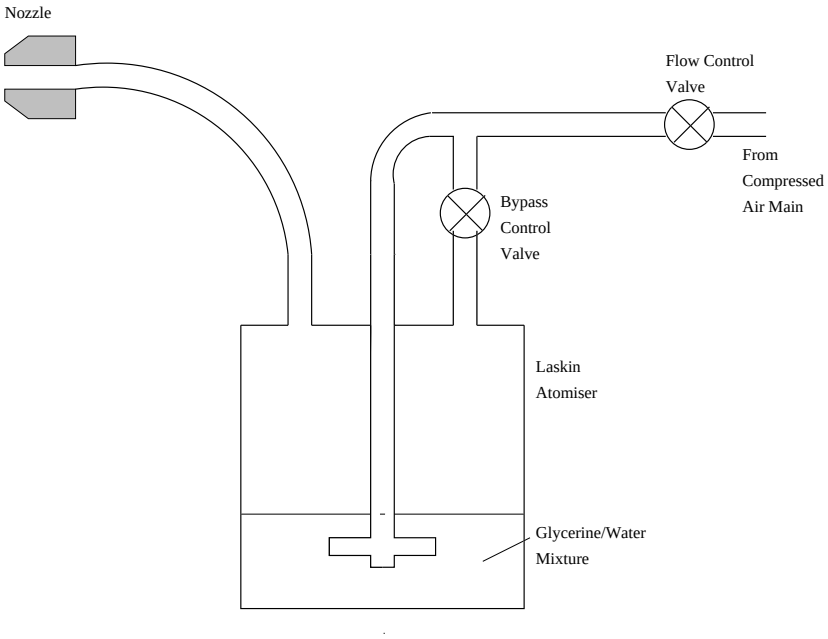


Figure 7.1: Laskin atomiser used to seed free jet.

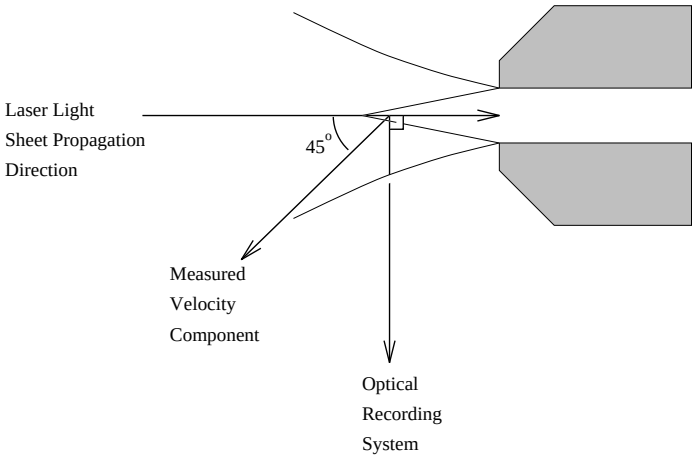


Figure 7.2: Doppler global velocimetry arrangement for free jet flow measurements — **Configuration I**.

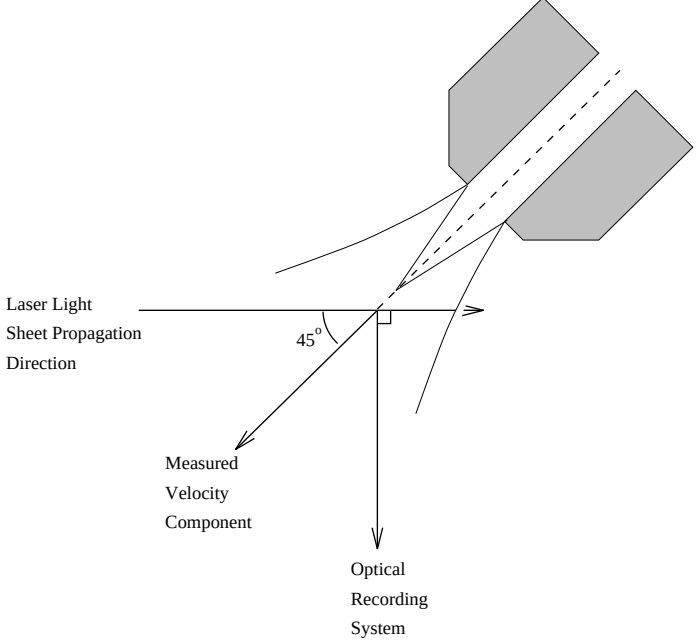


Figure 7.3: Doppler global velocimetry arrangement for free jet flow measurements — **Configuration II**.

7.2 Configuration

The jet nozzle of 12 mm internal diameter is set up as shown in Fig. 7.1. This nozzle contains flow straighteners and meshes to ensure a uniform exit flow, as described by Sheldrake and Ainsworth [Sheldrake and Ainsworth 1995]. A Laskin atomiser is used to seed the jet flow, the seeding medium being a mixture of water and glycerine. The nozzle is connected to the outlet of the atomiser using a flexible hose of 25 mm internal diameter, allowing the nozzle to be readily re-positioned. The flow into the atomiser is derived from the 400 psi compressed air main in the laboratory. The atomiser possesses two inputs, one of which passes to the seeding jets, the other being used as a bypass air control to set the quantity of seeding in the flow. A pair of valves are employed — one to set the overall flow, and the other to regulate the bypass flow, allowing the seeding density to be controlled by the experimenter. The Laskin atomiser provides glycerine/water particles of approximately 1 μm diameter. A mean exit flow velocity of 102 ms^{-1}, determined by the difference between measured jet exit total pressure and ambient pressure has been used in all the measurements described in this chapter.

Two different measurement configurations, as shown in Figs. 7.2 and 7.3, have been utilised. In configuration I, the jet is viewed by the Doppler global velocimetry optics so that the observation vector is at right angles to the jet axis, the laser light sheet itself being coincident with the axis of the jet. The measured component of velocity is then at 45° to the jet axis, as shown. This configuration corresponds with that employed for the measurement of the

rotating disc. Assuming that the light sheet is positioned accurately along the jet axis, there should be no velocity normal to the light sheet itself. Thus, although the system is sensitive to such a velocity, none should be observed in a symmetrical jet flow. The measured velocity may be interpreted as the axial velocity of the jet. A direct comparison between the measurements and the jet velocity correlations, which apply to the axial velocity component may therefore be made. Equation 2.1 gives an expected Doppler shift of 1.94 MHz per ms^{-1} of jet axial velocity.

Configuration II is shown in Fig. 7.3. In this case, the laser light sheet cuts across the flow at an angle of 45° to the jet axis, allowing cross-sectional profiles of the free jet to be measured. In configuration II, the Doppler global velocimetry system is sensitive to the component of velocity that coincides with the jet axis. For this case, eqn. 2.1 gives a sensitivity of 2.75 MHz per ms^{-1} of axial velocity. Even though the light sheet cuts the flow at 45°, the cross-sectional results obtained are devoid of any perspective effect. This is because the image processing transformation employed (see Chapter 3) has been designed to remove such effects.

For each of these configurations, the maximum expected axial jet velocity, of the order of 100 ms^{-1}, is such that the laser emission frequency ν_o may be set to a point midway along one of the wings of the iodine absorption profile. At such a frequency setting, there is no danger of the Doppler-shift being sufficiently great to cause problems with the limited iodine absorption profile bandwidth available. The maximum expected Doppler shift is of the order of 194 MHz for configuration I, and 275 MHz for configuration II.

In order to accomplish the change from configuration I to configuration II, it is necessary only to move the position of the nozzle itself: the nozzle is rotated by 45° while both the Doppler global velocimetry optical system and the laser light sheet are left in the same position. Furthermore, in configuration II, the nozzle, mounted on a linear optical rail, may be moved backwards and forwards, allowing cross-sections to be measured at varying distances from the nozzle exit plane. In this way it is possible to use the same computed alignment for both configuration I and configuration II and to readily change the position of the measured cross-section without undertaking any re-alignment.

In each case, the laser beam was expanded to form a light sheet of width 50 mm and thickness 1.25 mm. This ensures that the intensity of the laser emission is maximised in the region of interest.

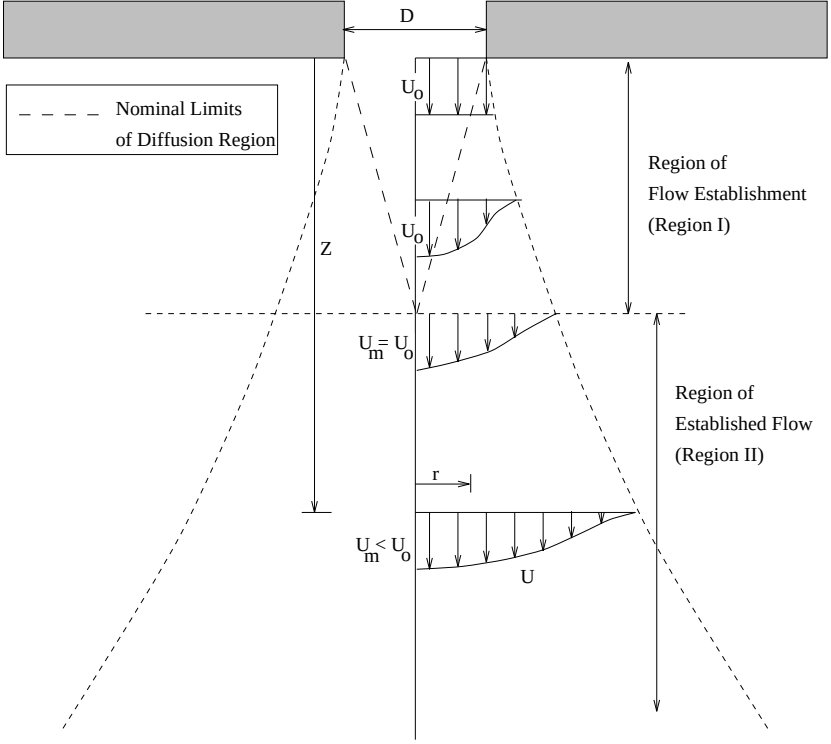


Figure 7.4: Diagram showing features of the free jet flow in the near-field region (after Gauntner, Lingwood and Hrycak [Gauntner, Lingwood and Hrycak 1970])

In accordance with the methods outlined in Chapter 3 and the analysis of Chapter 5, an alignment test card with square features of such a size that they occupy approximately 15 pixels in the CCD camera frame was used. The mark to space ratio of these features was unity, allowing a sufficient number of features to be viewed in the CCD frame (an eight by seventeen array of features for both the discriminated and non-discriminated images). The considerations of Chapter 5 indicate that the preferred image processing method is to use a polynomial warp with orthogonal polynomials together with a bilinear interpolation scheme. In the processing of the open jet images, a range of different image processing techniques has been employed, with the aim of verifying the theoretical considerations presented in Chapter 5 and determining the significance of the choice of algorithm in practice.

7.3 Description of the Free-Jet Flow Field

Turbulent free jet flows have historically received much attention from experimentalists. This is in part due to the fact that they are amongst the easiest flows to produce in the laboratory. There are a number of descriptions of this well understood flow field, such as that presented by Gauntner, Lingwood and Hrycak [Gauntner, Lingwood and Hrycak 1970]. Fig. 7.4 presents a diagrammatic representation of the near field flow pattern for a turbulent round free jet. Such a flow may be characterised by the variation in axial and radial velocity components.

The jet flow may be considered to be made up of a potential core region, and a spreading shear layer where the high velocity of the jet meets the stationary ambient air. The potential core represents a cone of constant axial flow velocity and zero radial velocity. Within the shear layer, the axial velocity has been observed experimentally to have a Gaussian variation as a function of the radial distance from the edge of the potential core region. This variation is expressed by Gauntner, Lingwood and Hrycak [Gauntner, Lingwood and Hrycak 1970] as

$$\frac{U}{U_o} = \exp\left(-\frac{\left(r + Kz - \frac{D}{2}\right)^2}{2(Kz)^2}\right), \tag{7.1}$$

where U is the local axial velocity component, U_o is the exit velocity of the jet, r is the radial distance from the jet centre-line, z is the axial distance from the jet exit plane, D is the diameter of the jet exit, and K is the spreading rate of the jet.

In order to compare Doppler global velocimetry data with this model, it is necessary to establish a value for the spreading rate K. This is given by

$$K = \frac{D}{2z_c} \tag{7.2}$$

where D is again the jet exit diameter and z_c the potential core length of the jet. For the jet in use in Oxford, $D = 12$ mm and z_c may be readily obtained from the measured results themselves, as outlined in Section 7.5.

Equations 7.1 and 7.2 may be taken to be valid throughout the region between the exit plane and the end of the potential core. Beyond approximately ten exit diameters from the exit plane, the velocity profiles become self-similar: that is, they have the same variation with radius when radius and velocity are suitably normalised. This part of the flow is known as the established or self-preserving region. Only the flow region upstream of the end of the potential core is considered in the present work.

7.4 Doppler Global Velocimetry Measurements

As explained above, the free jet measurements are well suited to the task of evaluating the performance of the different warping and interpolation techniques described in Chapter 3. To this end, the Doppler global velocimetry results obtained from configuration I have been processed using each of the different algorithms.

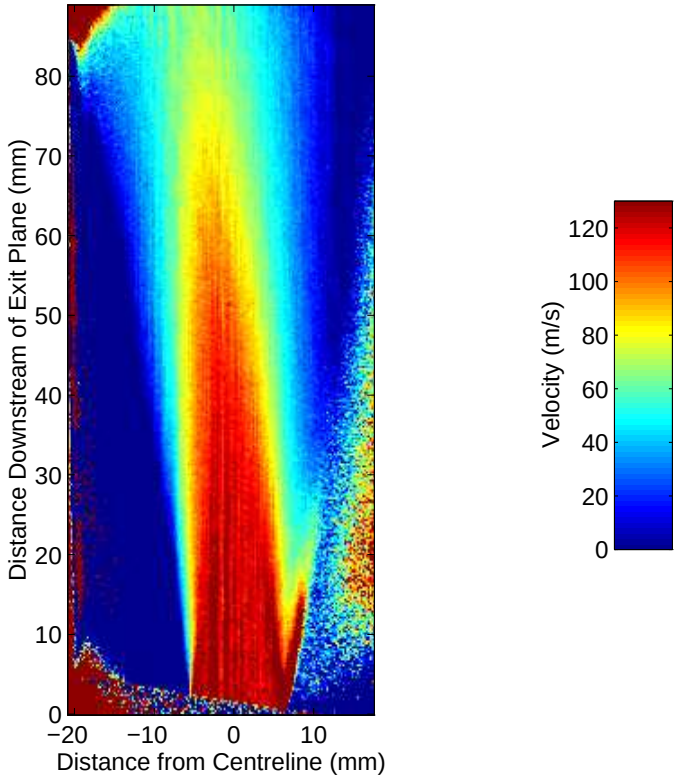


Figure 7.5: Free jet flow as measured by Doppler global velocimetry, employing bilinear warp method with bilinear interpolation. **Configuration I.**

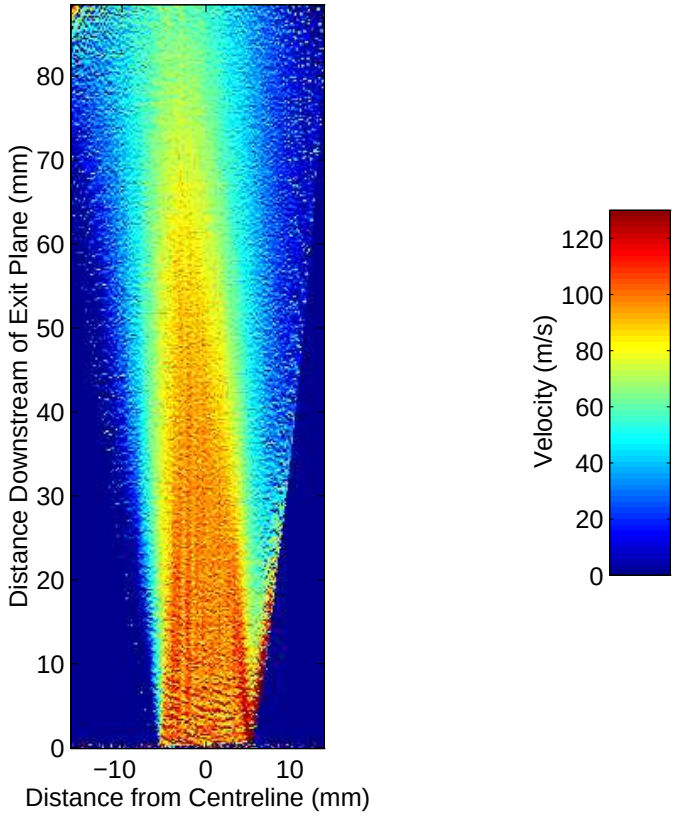


Figure 7.6: Free jet flow as measured by Doppler global velocimetry, employing pseudo-inverse polynomial warp method with bilinear interpolation. **Configuration I.**

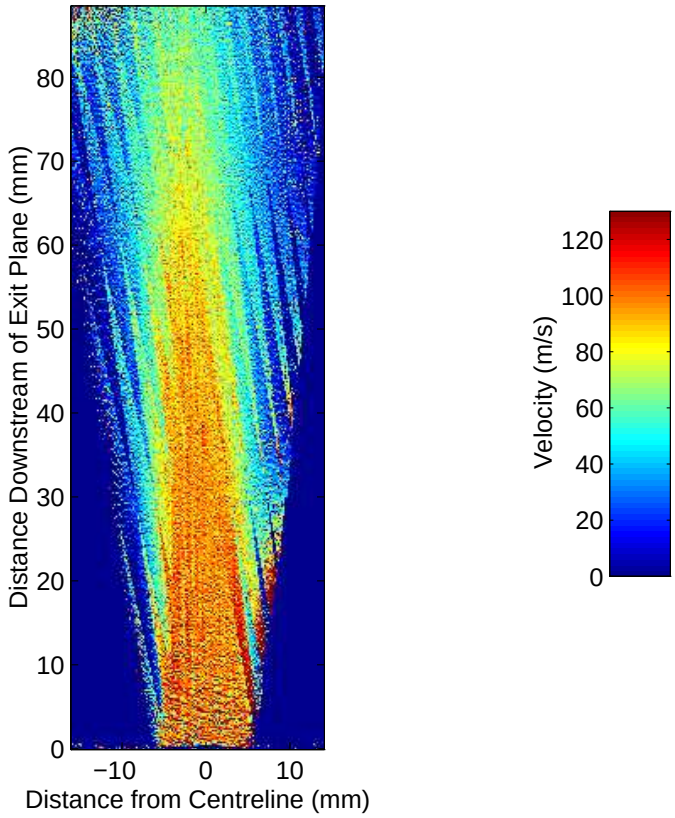


Figure 7.7: Free jet flow as measured by Doppler global velocimetry, employing orthogonal polynomials warp method with the nearest neighbour interpolation scheme. **Configuration I.**

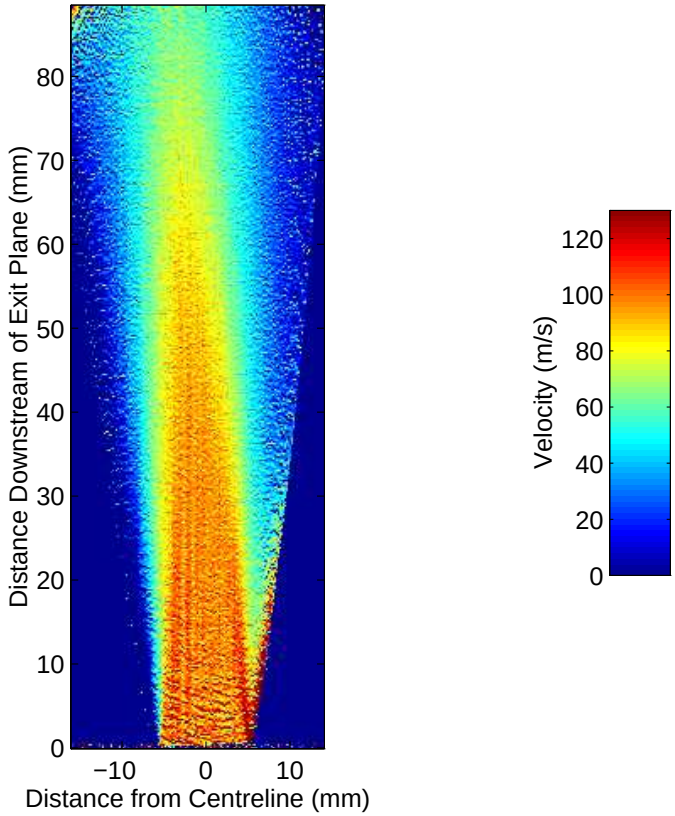


Figure 7.8: Free jet flow as measured by Doppler global velocimetry, employing orthogonal polynomials warp method with the bilinear interpolation scheme. **Configuration I.**

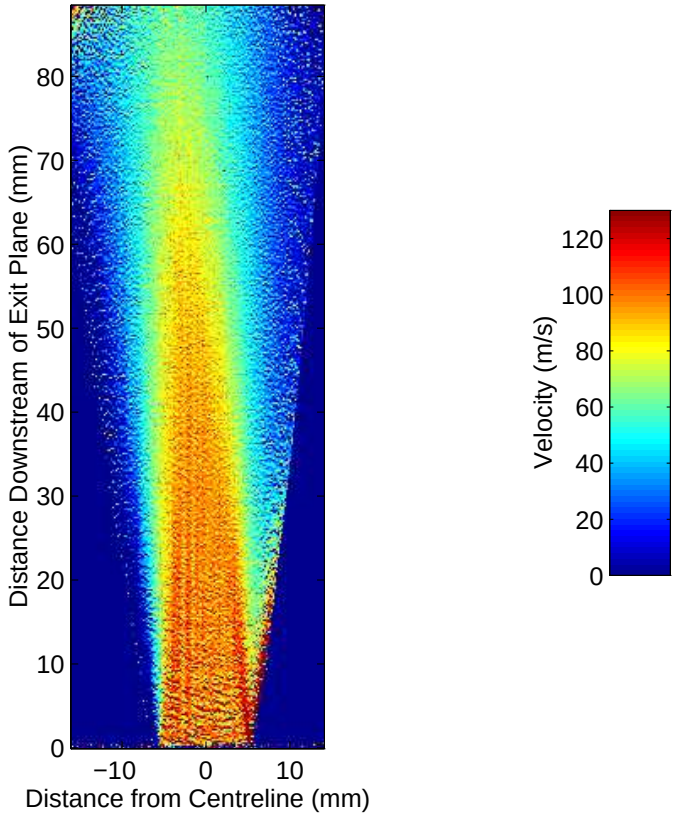


Figure 7.9: Free jet flow as measured by Doppler global velocimetry, employing orthogonal polynomials warp method with the cubic interpolation scheme. **Configuration I.**

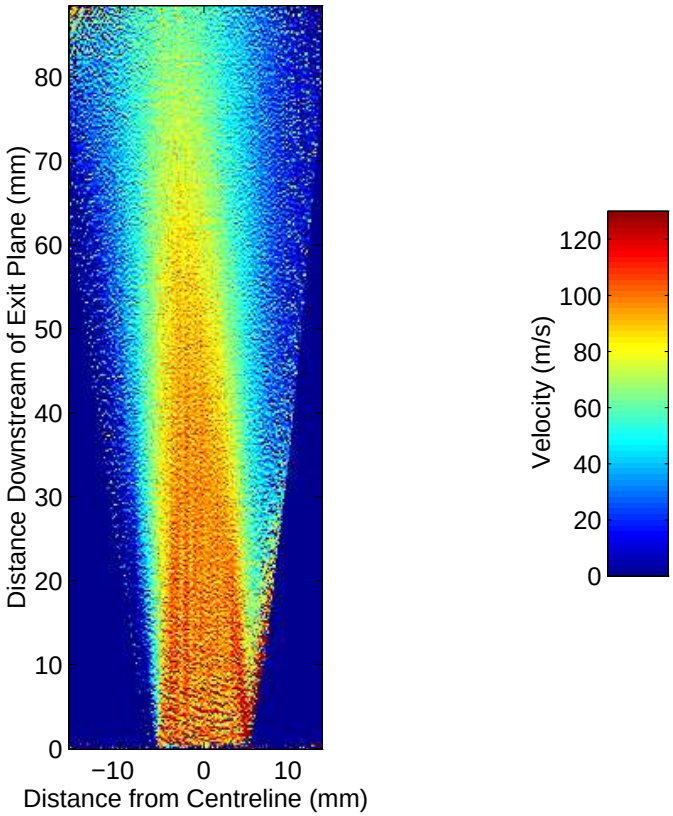


Figure 7.10: Free jet flow as measured by Doppler global velocimetry, employing orthogonal polynomials warp method with the exponential interpolation scheme. **Configuration I.**

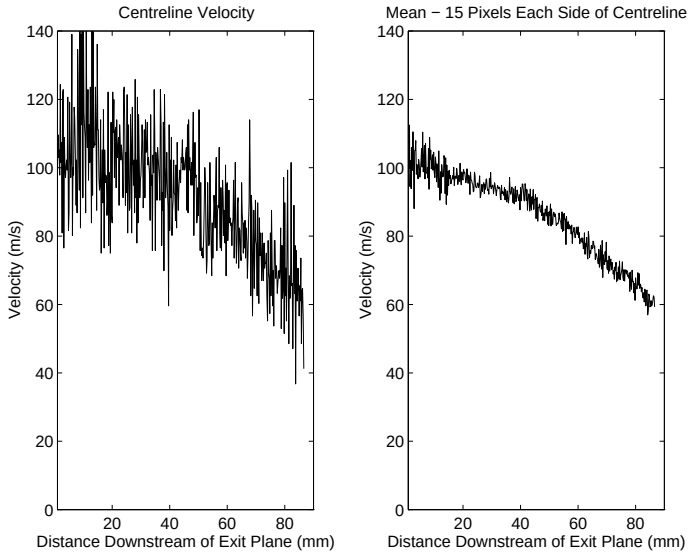


Figure 7.11: Line plot of axial velocity along free jet centre-line, as computed by Doppler global velocimetry, using the orthogonal polynomials warp method with the nearest neighbour interpolation scheme. The left hand image shows the raw measurement, whereas the right hand image shows the mean value of velocity for the region spanning 15 pixels each side of the centre-line. **Configuration I.**

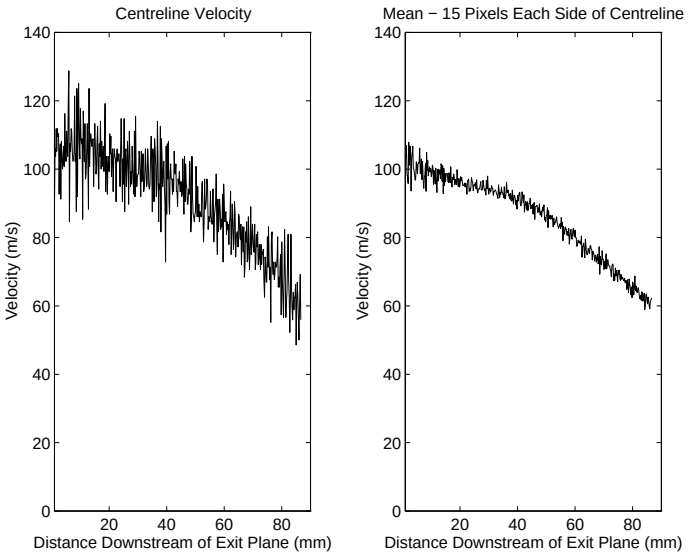


Figure 7.12: Line plot of axial velocity along free jet centre-line, as computed by Doppler global velocimetry, using the orthogonal polynomials warp method with the bilinear interpolation scheme. The left hand image shows the raw measurement, whereas the right hand image shows the mean value of velocity for the region spanning 15 pixels each side of the centre-line. **Configuration I.**

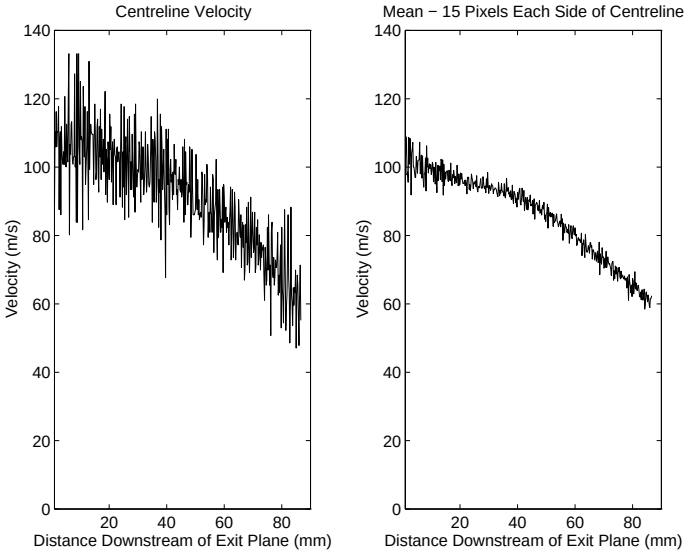


Figure 7.13: Line plot of axial velocity along free jet centre-line, as computed by Doppler global velocimetry, using the orthogonal polynomials warp method with the cubic interpolation scheme. The left hand image shows the raw measurement, whereas the right hand image shows the mean value of velocity for the region spanning 15 pixels each side of the centre-line. **Configuration I.**

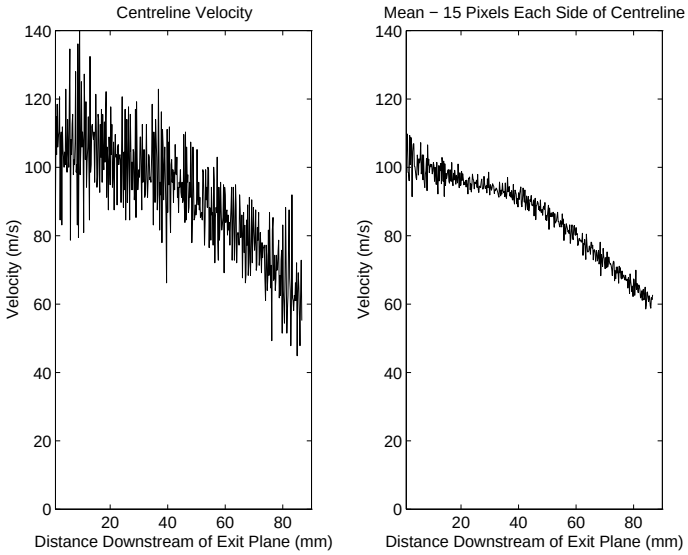


Figure 7.14: Line plot of axial velocity along free jet centre-line, as computed by Doppler global velocimetry, using the orthogonal polynomials warp method with the exponential interpolation scheme. The left hand image shows the raw measurement, whereas the right hand image shows the mean value of velocity for the region spanning 15 pixels each side of the centre-line. **Configuration I.**

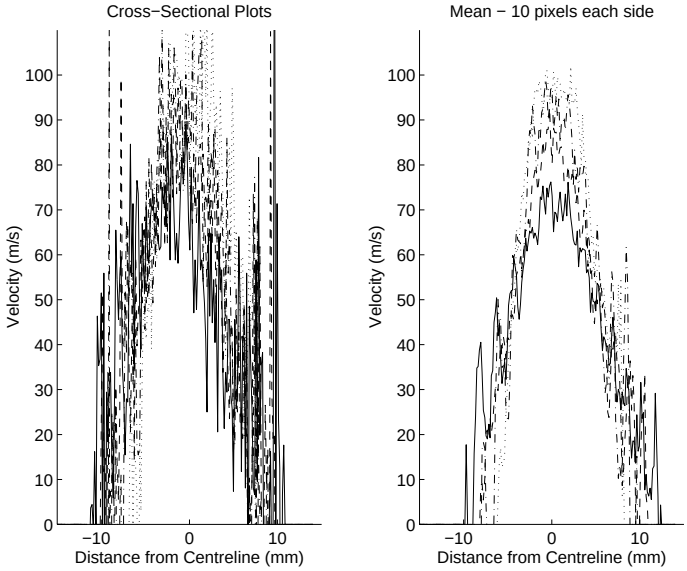


Figure 7.15: Line plot of the velocity cross-section of the free jet, as computed by Doppler global velocimetry, using the orthogonal polynomials warp method with the nearest neighbour interpolation scheme. The dotted, dot-dashed, dashed, and solid lines represent cross-sections taken at 30, 45, 60 and 75 mm downstream of the exit plane respectively. The left hand image shows the raw measurement, whereas the right hand image shows the mean value of velocity for the region spanning 10 pixels each side of the appropriate plane. **Configuration I.**

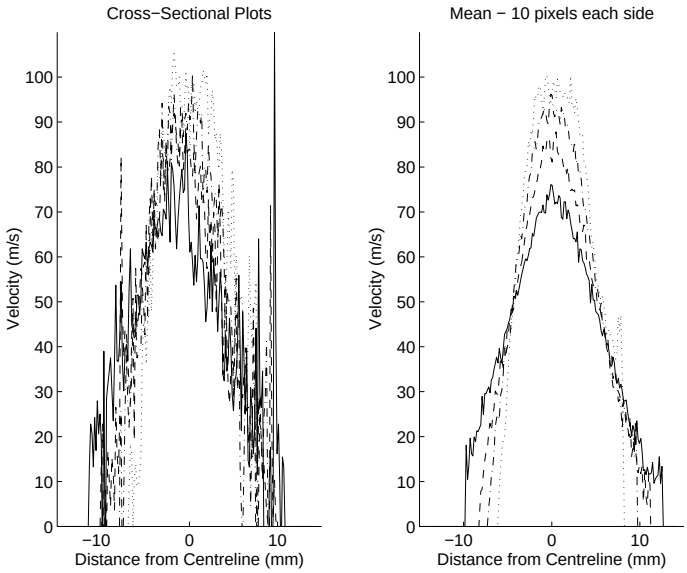


Figure 7.16: Line plot of the velocity cross-section of the free jet, as computed by Doppler global velocimetry, using the orthogonal polynomials warp method with the bilinear interpolation scheme. The dotted, dot-dashed, dashed, and solid lines represent cross-sections taken at 30, 45, 60 and 75 mm downstream of the exit plane respectively. The left hand image shows the raw measurement, whereas the right hand image shows the mean value of velocity for the region spanning 10 pixels each side of the appropriate plane. **Configuration I.**

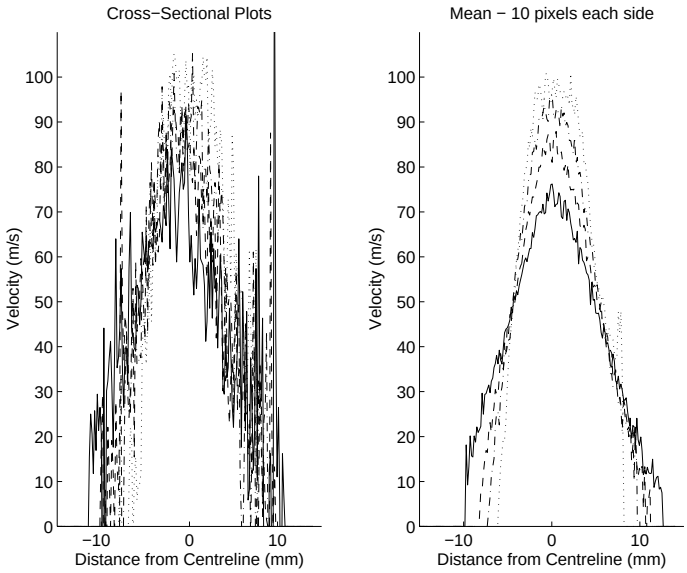


Figure 7.17: Line plot of the velocity cross-section of the free jet, as computed by Doppler global velocimetry, using the orthogonal polynomials warp method with the cubic interpolation scheme. The dotted, dot-dashed, dashed, and solid lines represent cross-sections taken at 30, 45, 60 and 75 mm downstream of the exit plane respectively. The left hand image shows the raw measurement, whereas the right hand image shows the mean value of velocity for the region spanning 10 pixels each side of the appropriate plane. **Configuration I.**

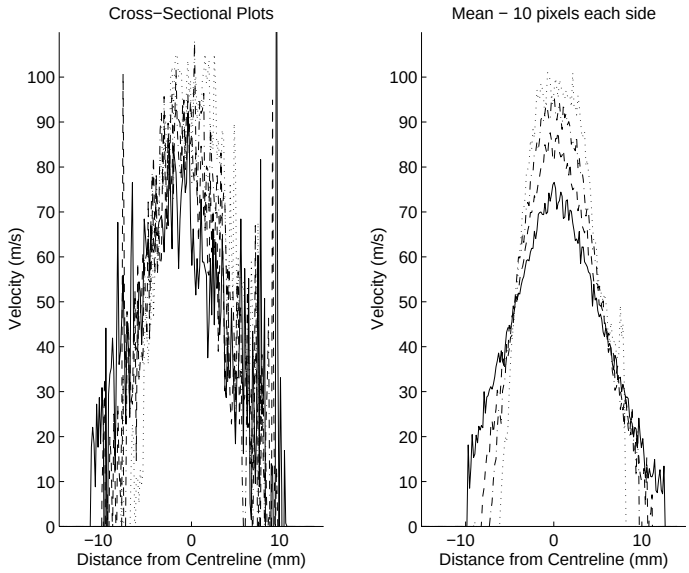


Figure 7.18: Line plot of the velocity cross-section of the free jet, as computed by Doppler global velocimetry, using the orthogonal polynomials warp method with the exponential interpolation scheme. The dotted, dot-dashed, dashed, and solid lines represent cross-sections taken at 30, 45, 60 and 75 mm downstream of the exit plane respectively.The left hand image shows the raw measurement, whereas the right hand image shows the mean value of velocity for the region spanning 10 pixels each side of the appropriate plane. **Configuration I.**

Three different warping methods — bilinear, polynomial (pseudo-inverse) and polynomial (orthogonal polynomials) — have been experimentally evaluated. Fig. 7.5 shows a colour-map representation of the free jet flow, computed using a bilinear warp with a bilinear interpolation method. Figs. 7.6 and 7.8 show the results of processing the same data using the polynomial (pseudo-inverse) and the polynomial (orthogonal polynomials) warp methods. In all cases the bilinear interpolation technique is used.

It was seen in section 5.8.1.1 that the pseudo-inverse method was subject to ill-conditioning, which may cause extreme distortion. It appears from Fig. 7.6 that such ill-conditioning is not an issue in this case, since no extreme distortion is evident. In general, however, the use of the pseudo-inverse method is to be discouraged because of the dangers of ill-conditioning. Since the orthogonal polynomials method is equivalent, and is less subject to such errors, the pseudo-inverse method will be given no further consideration.

The four different interpolation algorithms discussed in Chapter 3 — nearest neighbour, bilinear, cubic, and exponential — are also considered. It is convenient to evaluate the performance of the interpolation routines with respect to a single warping method: polynomial warping with orthogonal polynomials has been chosen for this purpose.

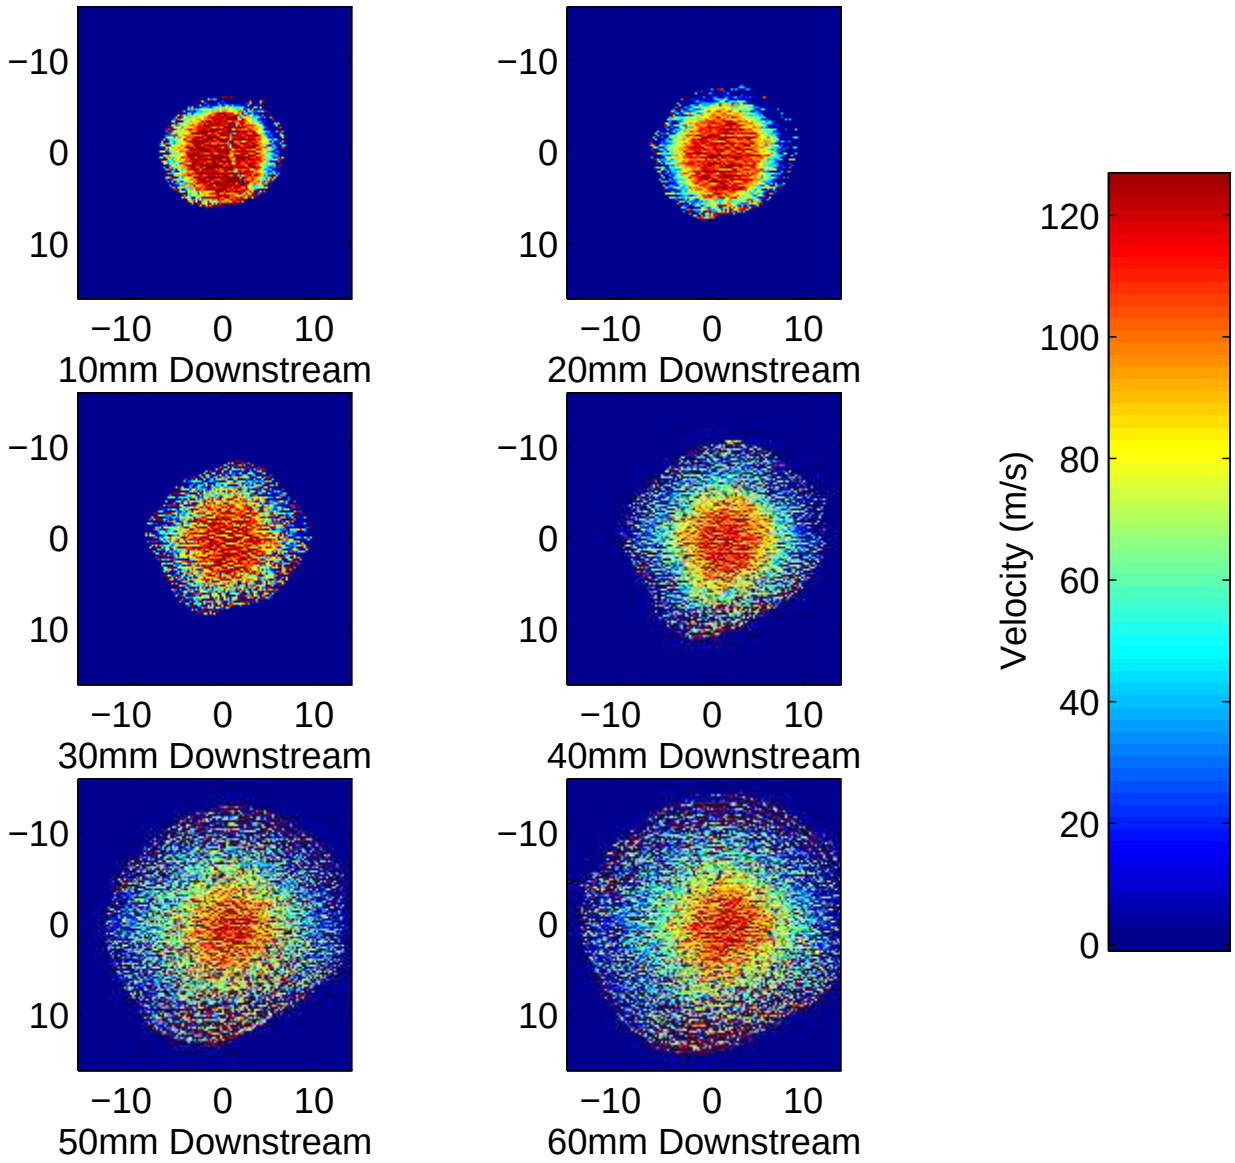


Figure 7.19: Doppler global velocimetry measurements of the free nozzle flow. Each of these colour-map representations shows the velocity throughout a cross-sectional slice of the flow at a certain distance from the nozzle exit plane. The axes are calibrated in millimetres from the centre-line of the nozzle. The orthogonal polynomials warp method was employed, using bilinear interpolation. **Configuration II.**

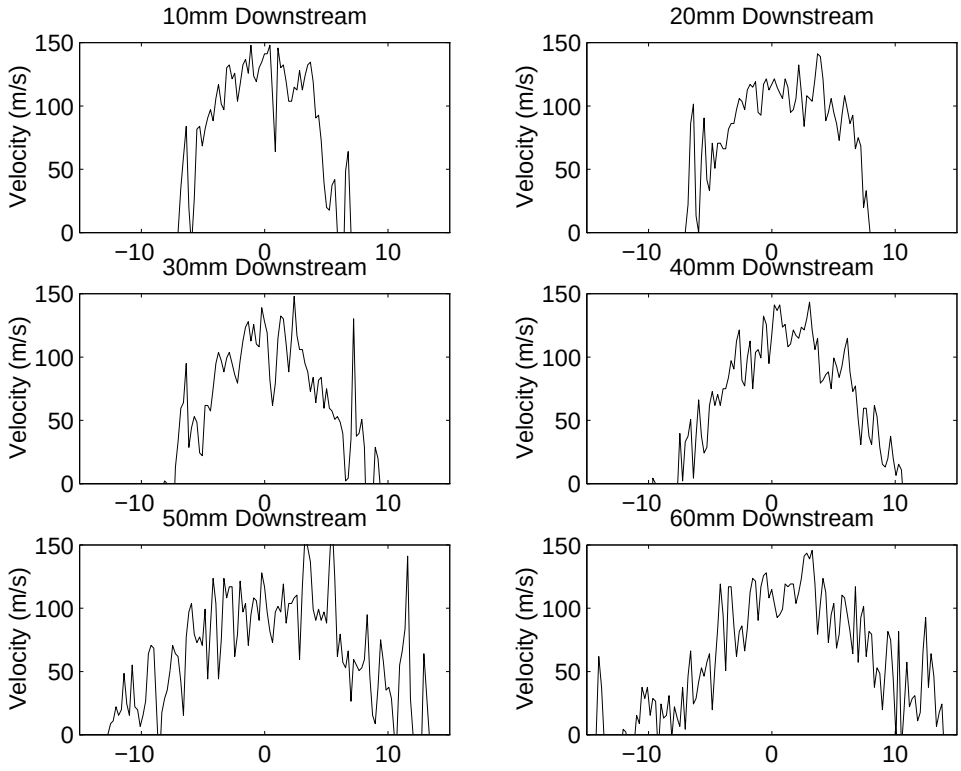


Figure 7.20: Doppler global velocimetry measurements of the free jet flow. Each graph shows the velocity across the centre of a cross-sectional slice at a certain distance from the nozzle exit plane. The orthogonal polynomials warp method was employed, using bilinear interpolation. **Configuration II.**

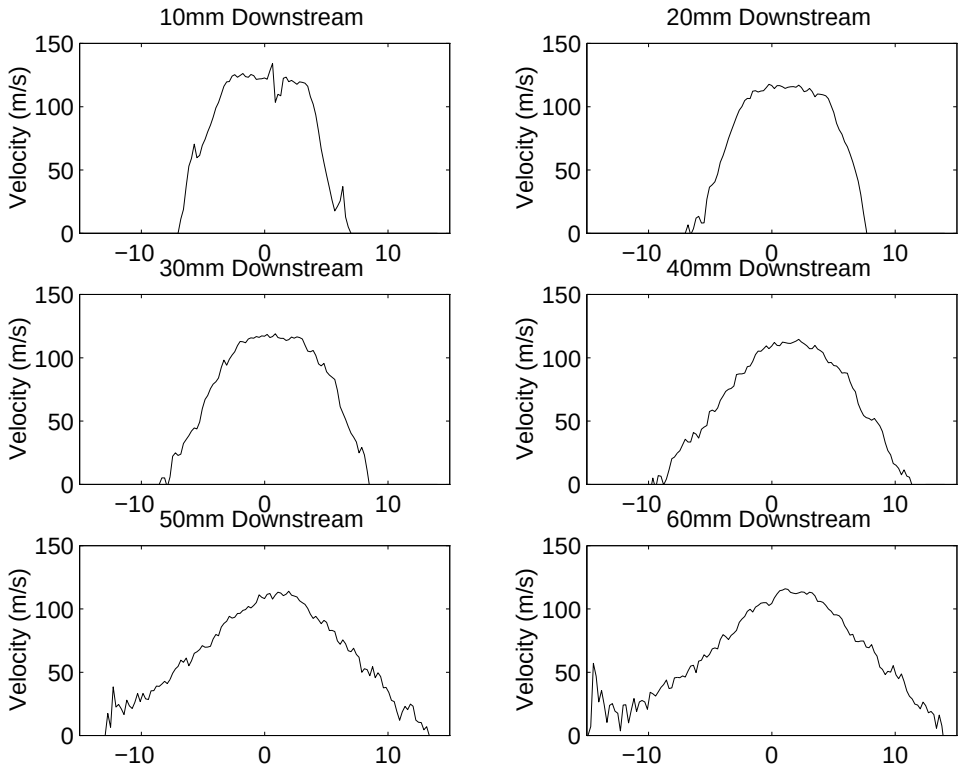


Figure 7.21: Doppler global velocimetry measurements of the free jet flow. Each graph shows the velocity across the centre of a cross-sectional slice at a certain distance from the nozzle exit plane. In this figure, the velocities shown are the mean of ten pixel rows each side of the line under consideration. The orthogonal polynomials warp method was employed, using bilinear interpolation. **Configuration II.**

Figs. 7.7, 7.8, 7.9 and 7.10 show the free jet velocity, as a colour-map representation, for the case of an orthogonal polynomial warp with nearest neighbour, bilinear, cubic, and exponential interpolation respectively.

The velocity data may also be represented in line plot form. Figs. 7.11, 7.12, 7.13 and 7.14 show the axial velocity along the centre-line of the jet for each of the interpolation techniques mentioned above. The raw data is plotted in each case, together with the velocity data obtained by taking a mean of fifteen pixels each side of the centre-line.

Similarly, cross-sectional velocity profiles may be plotted. Figs. 7.15, 7.16, 7.17 and 7.18 show the axial velocity across 'slices' 30, 45, 60 and 75 mm downstream of the exit plane of the jet, for each of the interpolation techniques in turn. Once again, raw data is shown, together with that obtained by taking a mean of ten pixels each side of the line under consideration.

All of the measurements discussed above were undertaken using configuration I. The results obtained using configuration II are represented by Figs. 7.19, 7.20 and 7.21. The first of these figures shows, as a colour-map representation, the velocity measured throughout cross-sectional slices of the flow 10, 20, 30, 40, 50 and 60 mm downstream of the jet exit plane. Figs. 7.20 and 7.21 show the same data, this time as velocity plots across the centre-line of the cross-sections of Fig. 7.19. Since the data obtained using configuration II has been shown to be in close agreement with that of configuration I, only the configuration I results have been employed in order to evaluate the image processing algorithms.

Before detailed analysis of the image processing performance is undertaken, it is important to consider the degree to which misalignment of the discriminated and the non-discriminated images may distort the measured velocity obtained from the Doppler global technique in this case. Fig. 7.22 shows, in colour-map form, the effect of misaligning the images deliberately by one, two, and five pixels in the vertical and horizontal directions. It may readily be seen that the velocity field obtained is very sensitive to any misalignment in the horizontal direction, but that the effect of misalignment in the vertical sense is less pronounced since the intensity gradient is far lower in that direction. From Fig. 7.22 it is also clear that the region closest to the nozzle exit plane is more subject to errors due to misalignment than the region further downstream. This may be attributed to the presence of high velocity and intensity gradients at the edge of the jet, close to the exit plane.

The importance of accurate alignment in the horizontal direction is reinforced by the graph of Fig. 7.23, where horizontal misalignment is deliberately imposed on the data while examining the resulting velocity profile at pixel row 400 of the image. Misalignment in the horizontal

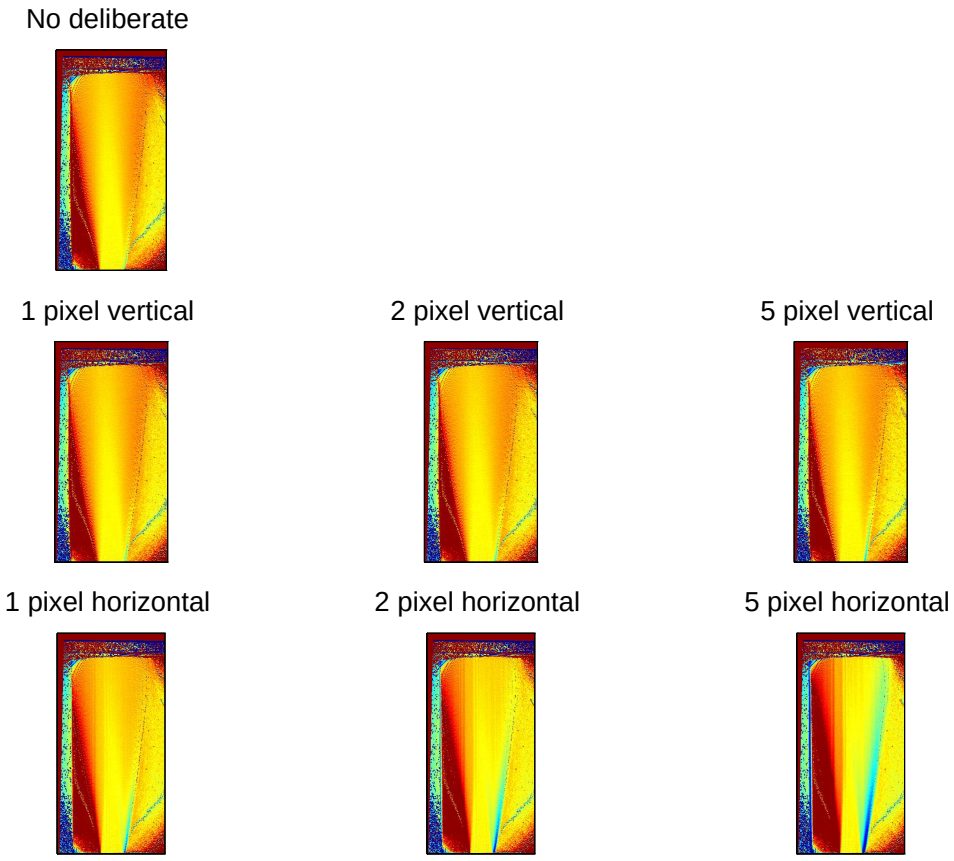


Figure 7.22: The effect on ratio of deliberately misaligning the discriminated and non-discriminated images by one, two, and five pixels in the vertical and horizontal directions. **Configuration I.**

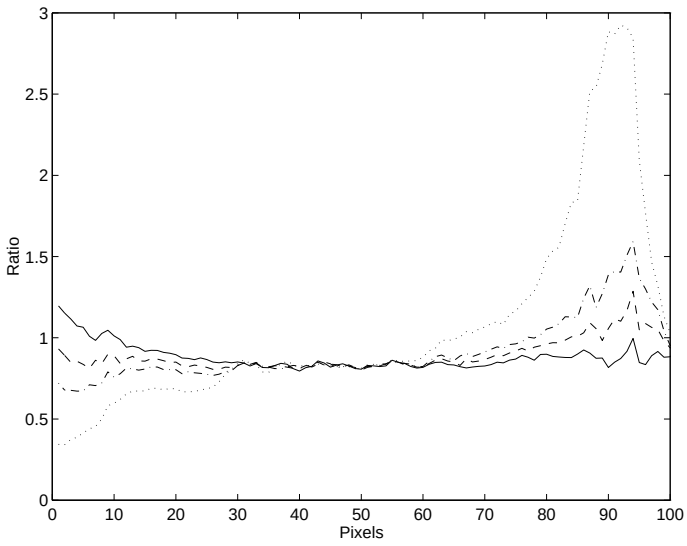


Figure 7.23: The effect on ratio (at pixel row 400) of deliberately misaligning the discriminated and non-discriminated images by one, two, and five pixels in the horizontal direction. The ratio obtained with no deliberate misalignment is shown as a solid line, whereas the ratio resulting from a deliberate misalignment of one, two, or five pixels is shown as dashed, dot-dashed, and dotted lines respectively. **Configuration I.**

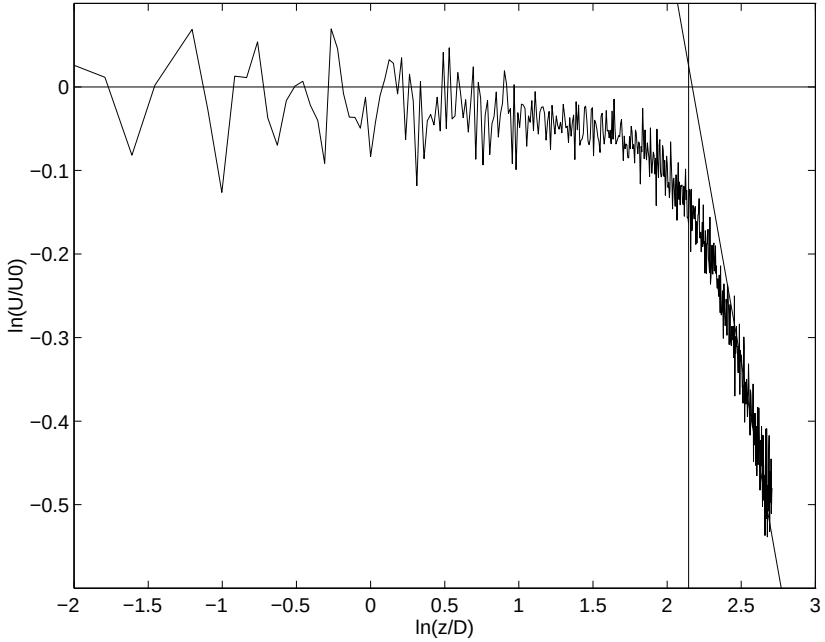


Figure 7.24: Method of establishing length of potential core from Doppler global measurements.

direction may be seen to lead to significant distortion in the velocity profiles obtained, especially towards the edges of the flow field.

7.5 Correlation of Published Free Jet Flow Field Characteristics with Doppler Global Velocimetry Measurements

In order to assess the capabilities of the different image processing algorithms, it is necessary to compare the results obtained using the different algorithms to the established correlations for the axial velocity of the free jet flow, as outlined in Section 7.3. The established model, described by eqn. 7.1, relates to the axial velocity component of the jet. The correlation is applicable to the region from the jet exit plane to the end of the potential core. Comparisons between Gauntner's model and the measured velocities obtained by Doppler global velocimetry are performed for configuration I, in the form of line plots of the data at positions one, two, and three nozzle diameters downstream of the exit plane. In performing the correlation, the jet exit velocity U_o has been assumed to be 102 ms^{-1}.

To fit the model to the measured data, it is necessary to calculate the length of the potential core. This may be determined from the Doppler global measurements using a technique described by Gauntner. Using this method, the jet centre-line velocity data is normalised by the exit velocity U_o. That data is then plotted along logarithmic axes as a function of distance from the exit plane, which is normalised by the jet exit diameter D. Such a graph is shown in Fig. 7.24. Beyond the potential core, the axial velocity of the jet decays as

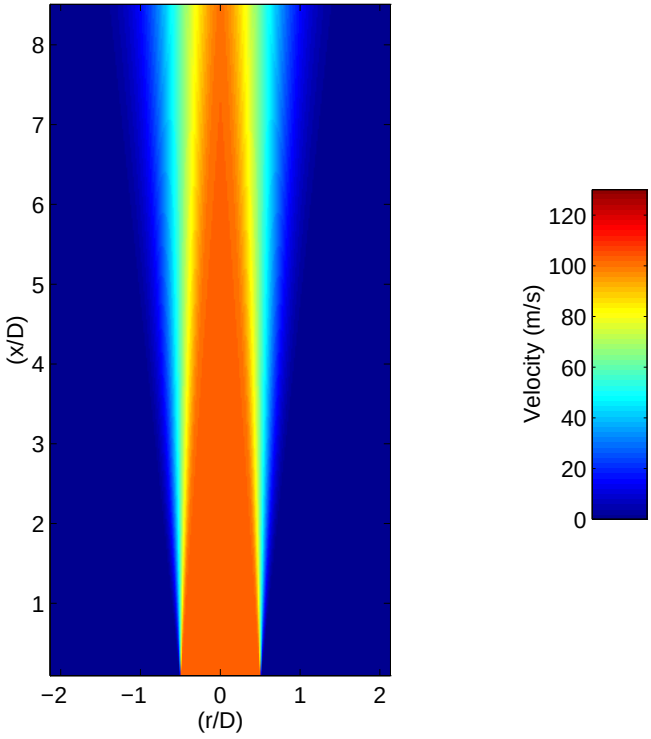


Figure 7.25: Expected free nozzle flow as predicted by the work of Gauntner, Lingwood and Hrycak [Gauntner, Lingwood and Hrycak 1970]

the inverse of axial distance, while in the core the axial velocity is constant. The data can thus be represented by two straight lines in the axial plot of Fig. 7.24, one of gradient zero for the core region and the other of gradient -1 downstream of the core. The location of the intersection of these lines determines the potential core length. From the data obtained (using an orthogonal polynomial warp with bilinear interpolation, with configuration I) the potential core length has been found to be 9.0 exit diameters, equivalent to 108 mm.

For circular air jets such as that under consideration here, reported values of potential core length vary from 4.7 to 7.7 exit diameters [Gauntner, Lingwood and Hrycak 1970]. For example, Albertson *et al.* [Albertson *et al.* 1950] report a value of 6.2 exit diameters, Pai [Pai 1954] a value of 6.5, and Tani and Komatsu [Tani and Komatsu 1966] a value of 4.8. The length of the potential core is thought to be independent of jet Reynolds number based on the exit diameter [Abramovich 1963], with the reported variations in core length attributed to differences in the exit flow turbulence level and velocity profile. The value of core length determined in the current work has thus been found to be slightly greater than that determined by previous workers. It is possible that this discrepancy may be attributed to the level of turbulence in the free jet flow.

Using 9.0 exit diameters as the value of the potential core length, and taking U_o to be 102 ms^{-1}, the model may be evaluated. The result of such a computation is shown in

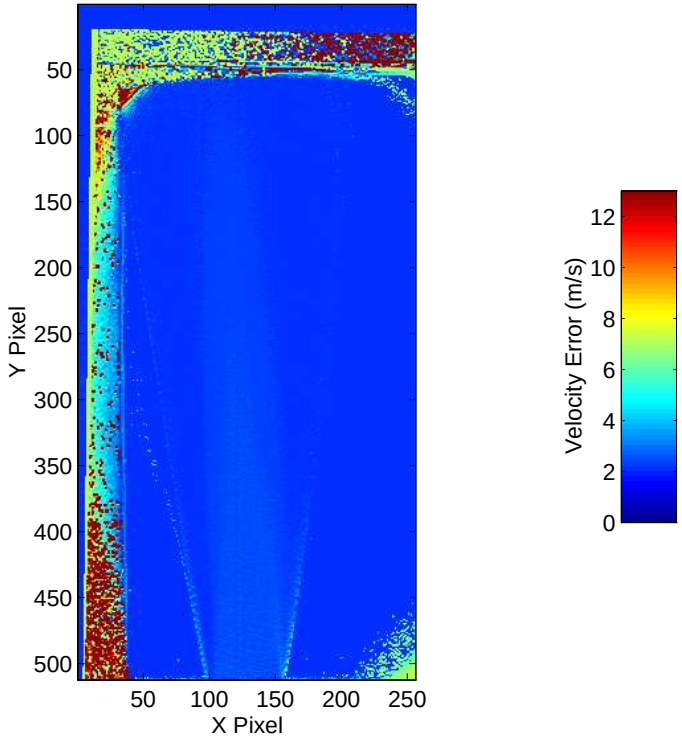


Figure 7.26: Predicted error due to iodine calibration error, as given by the application of the analysis of Chapter 5.

Fig. 7.25.

7.6 Error Analysis

In Chapter 5 it was explained that the error expected in Doppler global measurements may be seen to be composed of three components: error due to iodine calibration error, error due to the presence of shot noise in the acquired data and error due to misalignment. Expressions were developed for each of these error sources, enabling them to be evaluated across the entire measurement field. These expressions may be applied to the free jet flow field and global error estimates hence obtained.

The expected error due to iodine cell calibration does not vary with the choice of warping or interpolation algorithm. This error component has been computed and is shown in Fig. 7.26. It may be seen that the iodine cell error component is more or less uniform across the measurement field, at an RMS value of around 2 ms^{-1}.

The expected error due to misalignment, on the other hand, is dependent on the choice of warping algorithm in use. In Chapter 5 it was seen that a parameter k_w representing the sensitivity of overall frame misalignment to calibration feature misalignment could be determined for each warping method. Accordingly, the expected value of misalignment error throughout the entire frame has been computed for both the bilinear warping technique,

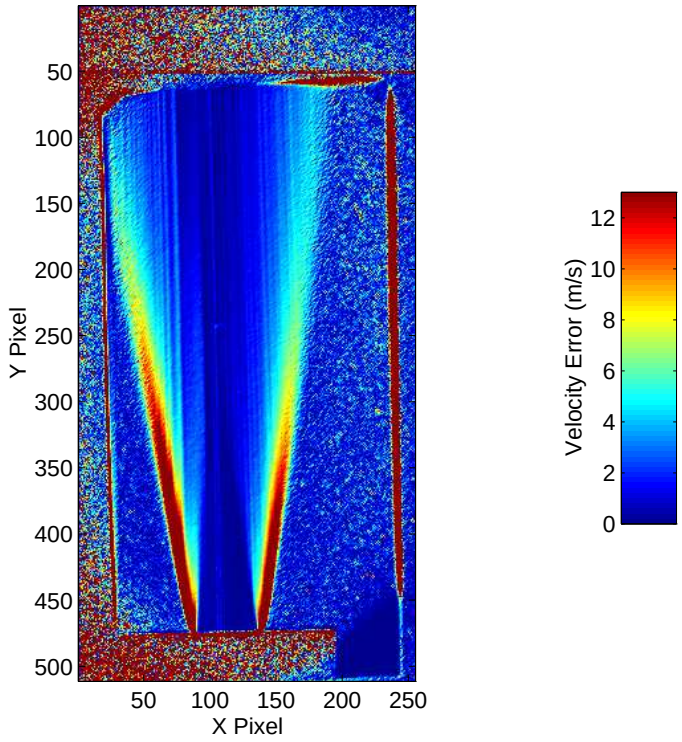


Figure 7.27: Predicted error due to misalignment, for the case of the bilinear warping technique, as given by the application of the analysis of Chapter 5.

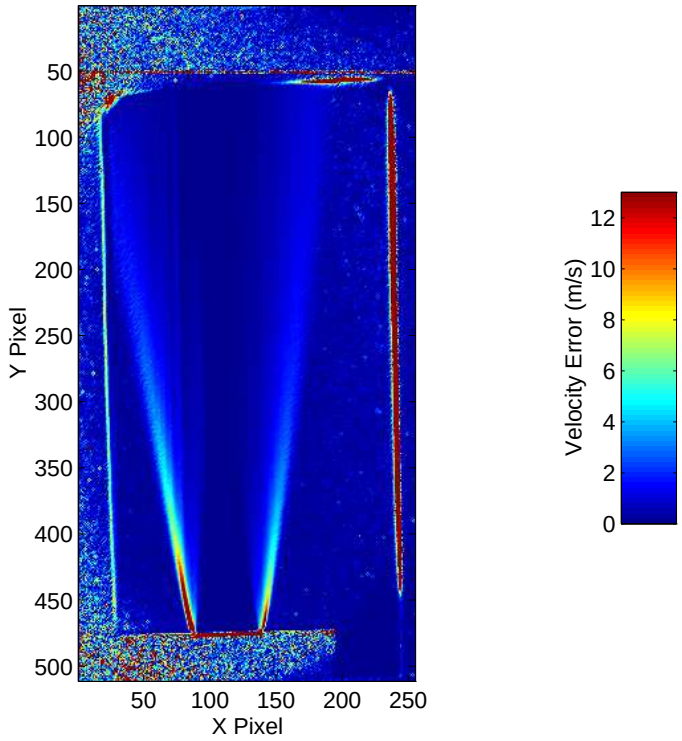


Figure 7.28: Predicted error due to misalignment, for the case of the polynomial warping technique, as given by the application of the analysis of Chapter 5.

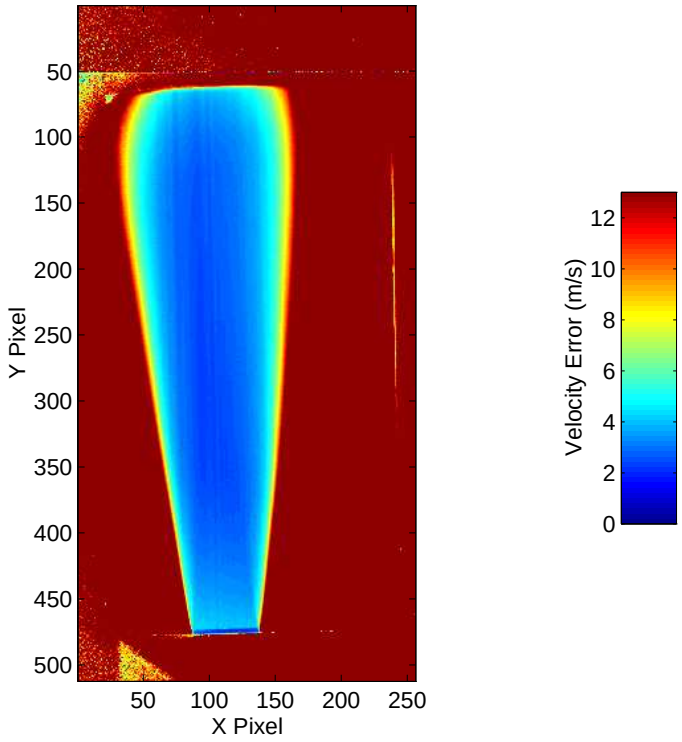


Figure 7.29: Predicted error due to shot noise, for the case of the nearest neighbour interpolation technique, as given by the application of the analysis of Chapter 5.

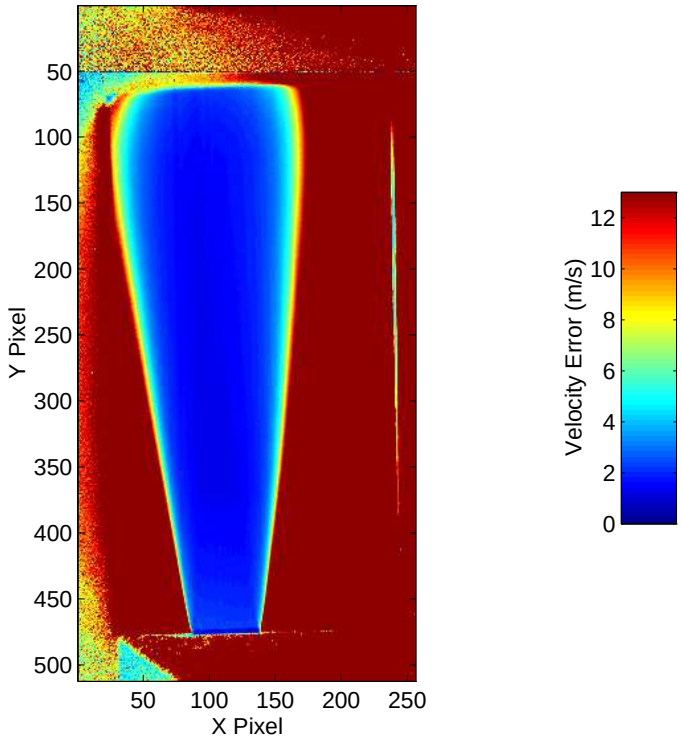


Figure 7.30: Predicted error due to shot noise, for the case of the bilinear interpolation technique, as given by the application of the analysis of Chapter 5.

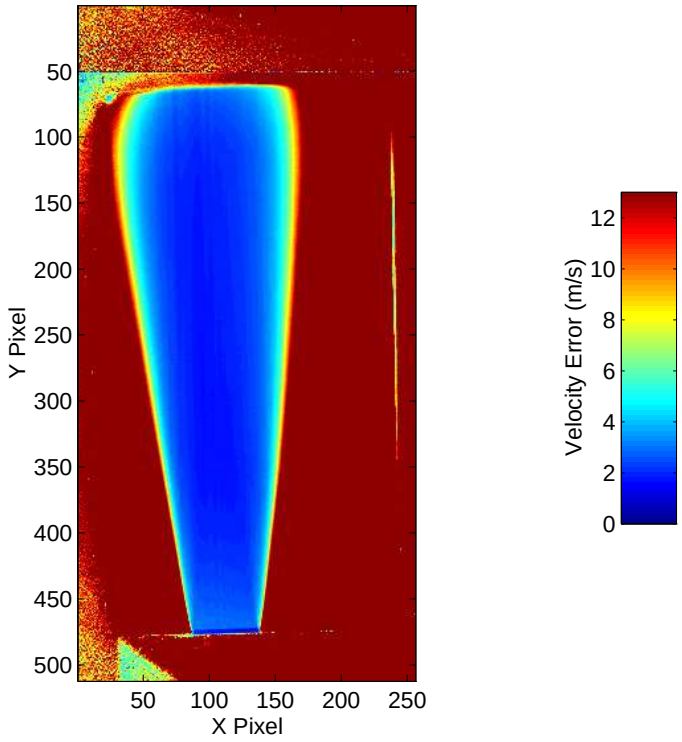


Figure 7.31: Predicted error due to shot noise, for the case of the cubic interpolation technique, as given by the application of the analysis of Chapter 5.

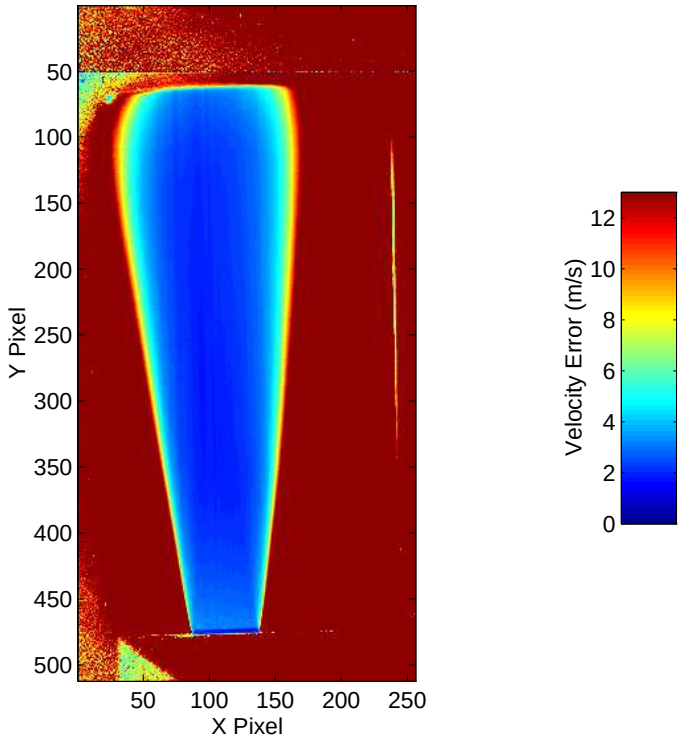


Figure 7.32: Predicted error due to shot noise, for the case of the exponential interpolation technique, as given by the application of the analysis of Chapter 5.

Fig. 7.27 and the polynomial warping technique, Fig. 7.28. An estimated value of centroid estimation error of 0.1 pixel RMS has been assumed throughout. From these two figures it may readily be seen that the expected error due to misalignment is significantly more severe in the case of bilinear warping than in the case of the polynomial transformation. This is due to the improved immunity to centroid estimate noise provided by the least-squares approach of the polynomial warping algorithm.

Error due to misalignment is most severe at the edges of the jet flow itself, where intensity and velocity gradients are at their highest. The error at the margins of the jet flow is very sensitive to small misalignments in the horizontal sense, as shown in Figs. 7.22 and 7.23. This being the case, it would be expected that the polynomial warp would yield significantly better results than the bilinear warp in this region.

The value of the expected error due to misalignment is approximately zero at the centre of the jet flow, where intensity gradients are small and the velocity gradient (within the potential core) is zero. Close to the margins of the jet, however, the expected error due to misalignment for the case of the polynomial warp is of the order of 10 ms^{-1} at a distance of one nozzle diameter from the exit plane, falling to around 5 ms^{-1} three nozzle diameters from the exit plane. For the bilinear warp, the expected error due to misalignment is again close to zero at the centre of the jet flow, rising to around 15 ms^{-1} at the edge of the jet at one nozzle diameter from the exit plane, 7 ms^{-1} at three nozzle diameters from the exit plane.

It seems to be the case that the choice of warping algorithm effects not so much the maximum expected misalignment error value, but rather the extent of the region within which this error is significant. Employing the polynomial warp rather than the bilinear method has the effect of reducing the width of the region within which misalignment error is significant by a factor of around three. The magnitude of the expected misalignment error is significantly higher for the bilinear method even far downstream of the nozzle. For example, five nozzle diameters from the exit plane, the peak estimated misalignment error for the bilinear warping algorithm is around 8 ms^{-1}. The polynomial technique exhibits a worst case behaviour of around 3 ms^{-1} misalignment error at this point.

Further, it may be observed from Fig. 7.27 that a number of vertically orientated striations are present in the estimated misalignment error; these are absent from the polynomial warped estimates of Fig. 7.28. Such variations are due to non-uniform intensity in the original acquired images resulting from dirty lenses or mirrors in the illuminating system. The polynomial warp offers superior suppression of errors due to such intensity discontinuities,

thanks to the lower overall misalignment which may be achieved with such a technique.

Just as the choice of warping technique proved to be significant in estimating the expected error due to misalignment, the level of expected error due to CCD array shot noise is dependent on the interpolation method employed. Figs. 7.29, 7.30, 7.31 and 7.32 show the estimated error due to shot noise, evaluated across the entire frame, for the nearest neighbour, bilinear, cubic, and exponential interpolation algorithms respectively.

In Chapter 5 a parameter k_i, dependent on interpolation method, was computed. This parameter relates the overall shot noise error expected to the shot noise at each pixel of the CCD array. It was found that the exponential and cubic schemes may be seen as performing a weighted mean of a greater number of pixel values than the bilinear interpolation method. The nearest neighbour scheme, on the other hand, performs no such weighted mean operation. It is expected that the operation of calculating such a mean should reduce the overall level of shot noise error in the final interpolated frame. Numerical precision was found to be a factor, however, leading to the conclusion that the bilinear scheme could be expected to perform better than the more complex cubic and exponential methods.

Each of the shot noise error estimates exhibits the characteristic that the level of shot noise error increases in regions of low intensity. For the case of the free jet, this manifests itself as a large estimated error at the edges of the jet flow, falling to a lower constant value within the body of the jet itself.

As expected, the nearest neighbour error estimates of Fig. 7.29 show poor performance by comparison with the other methods. Within the body of the jet, an error value of around 5 ms^{-1} is observed, compared with around 2 ms^{-1} for the bilinear interpolation technique. In addition, the region of significant error (around 10 ms^{-1}) at the edges of the jet is somewhat larger for the case of nearest neighbour than for the more sophisticated methods.

Comparing Figs. 7.30, 7.31 and 7.32, it is clear that while it is expected that bilinear interpolation should offer noise immunity superior to the cubic and exponential methods, this is less pronounced than the predicted difference in performance between the nearest neighbour and the bilinear algorithms. Within the body of the jet, the estimated shot noise error for the case of the bilinear interpolation method is 2 ms^{-1} RMS, as opposed to 3 ms^{-1} RMS for the more complex techniques. There appears to be little to choose between these three interpolation algorithms with regard to their performance near the edges of the jet flow.

In Section 7.4, Doppler global velocimetry measurements of the free jet flow were introduced. Sections 7.3 and 7.5 introduced the free jet model proposed by Gauntner, Lingwood

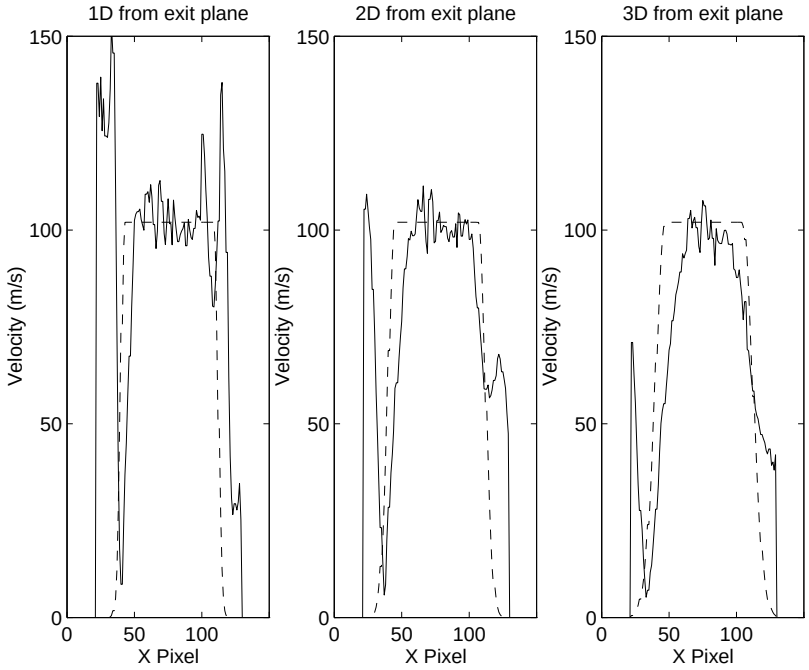


Figure 7.33: Comparison of Doppler global measurements (solid line) with expected data (dashed line) for the case of the bilinear warp with bilinear interpolation. Cross-sections one, two, and three nozzle diameters from the exit plane are considered.

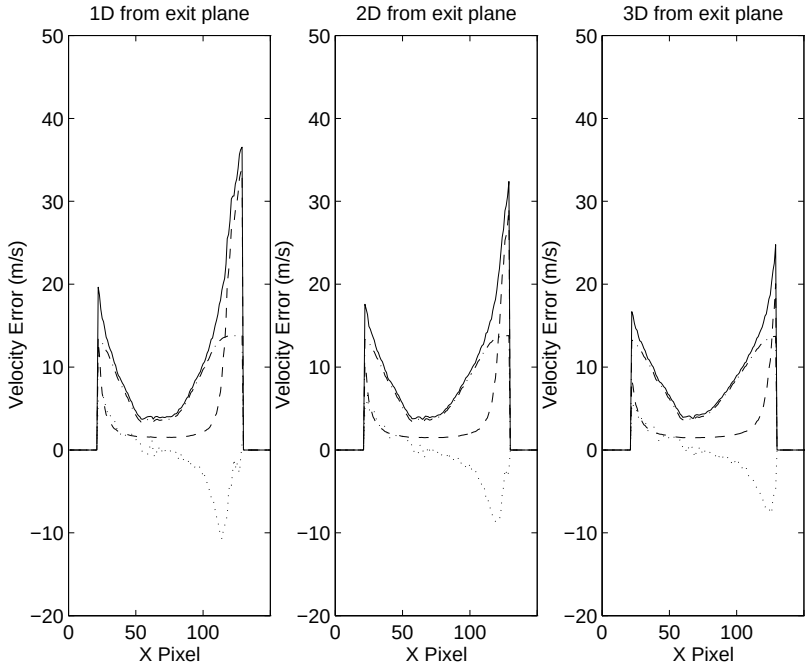


Figure 7.34: Predicted measurement error using the analysis of Chapter 5 for the case of the bilinear warp with bilinear interpolation. The dotted line represents the predicted error due to misalignment, the dashed line the predicted error due to shot noise, the dot-dashed line the predicted error due to the iodine cell, and the solid line the overall RMS error.

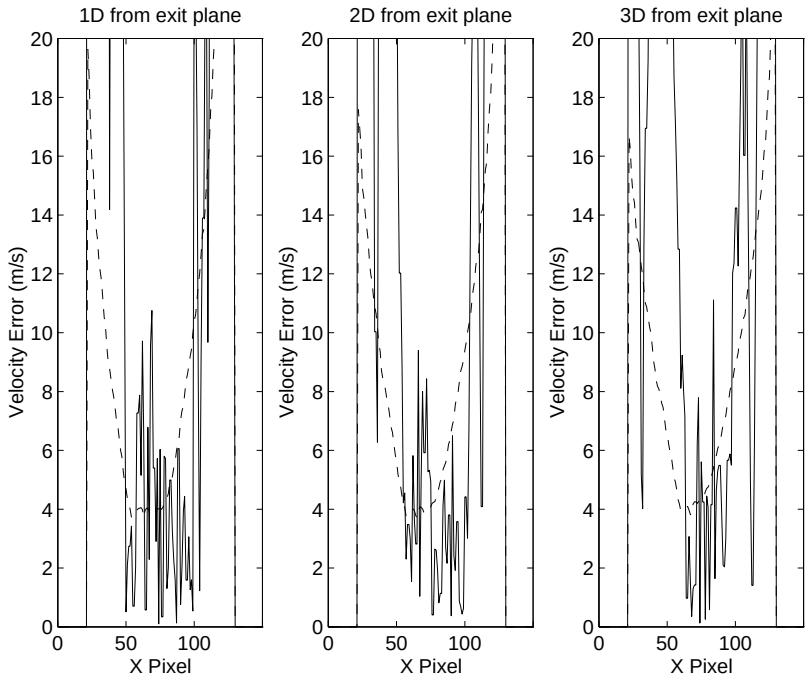


Figure 7.35: Comparison of predicted RMS error with observed error for the case of the bilinear warp with bilinear interpolation.

and Hrycak, and described the procedure for correlating those results with the Doppler global measurements obtained. The analysis of Chapter 5 has been applied to the free jet measurements as described above, yielding error estimates across the entire measurement frame. It is now possible to evaluate the performance of the different image processing algorithms as applied to the free jet flow measurements.

The procedure for assessing image processing performance is as follows. The Doppler global measurements may be compared to the theoretical results of Gauntner in order to obtain an estimate of measurement error. The predicted measurement error due to iodine cell calibration, CCD shot noise and misalignment have also be determined. The estimated error may then be compared to the observed error. This method is applied at a number of positions in the flow field, at cross-sectional slices one, two, and three nozzle diameters downstream of the exit plane. Such analysis enables the theoretical predictions of Chapter 5 to be tested in a practical environment.

The first case for analysis is that of the bilinear warp together with bilinear interpolation. Fig. 7.33 shows the Doppler global measurements of the free jet as determined by such an image processing technique, at positions one, two, and three nozzle diameters downstream of the exit plane. Superimposed on the measured data is the theoretical free jet axial velocity distribution obtained from the Gauntner model. Fig. 7.34 shows the results of computing the predicted error at these same flow field locations. The dotted line represents the predicted

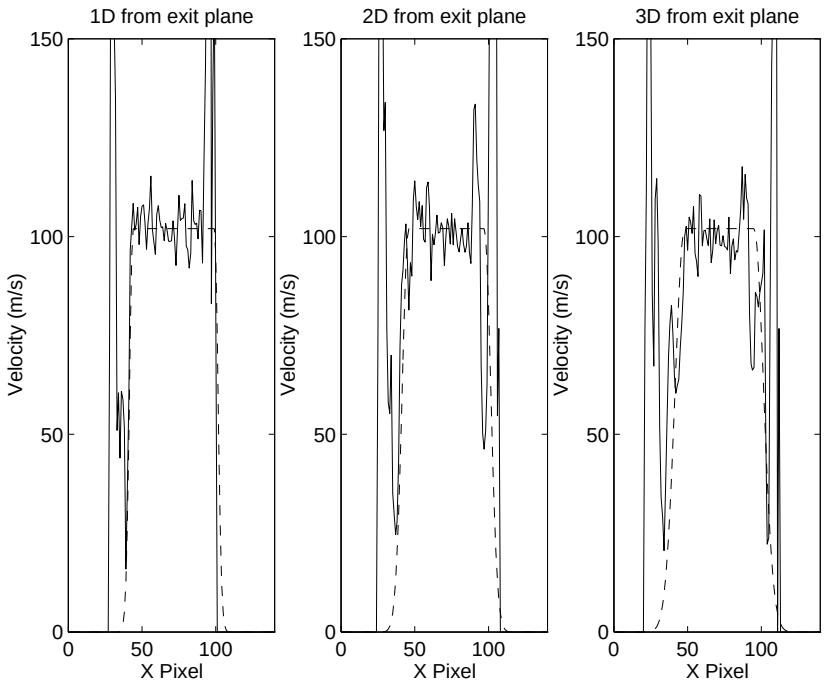


Figure 7.36: Comparison of Doppler global measurements (solid line) with expected data (dashed line) for the case of the orthogonal polynomials warp with nearest neighbour interpolation. Cross-sections one, two, and three nozzle diameters from the exit plane are considered.

error due to misalignment, the dashed line the predicted error due to shot noise, the dot-dashed line the predicted error due to the iodine cell, and the solid line the overall RMS error. For every predicted error graph in this chapter, all the quantities are shown as RMS values, with the exception of the misalignment error which is shown as a signed quantity in order to allow the sense of the misalignment error to be readily seen. The observed error (that is, the discrepancy between the measured data and the correlation) and the predicted RMS error are plotted together in Fig. 7.35.

In order to be able to assess the effect of both the choice of warping algorithm and the choice of interpolation method, it is necessary to perform the same analysis for the orthogonal polynomials warping technique, using each of the possible interpolation routines: nearest neighbour, bilinear, cubic and exponential.

Figs. 7.36 to 7.38 present the same analysis as discussed above, this time for the case of orthogonal polynomials using nearest neighbour interpolation. Figs. 7.39 to 7.41 show the results for bilinear interpolations. Figs. 7.42 to 7.44 and 7.45 to 7.47 allow comparison of the measurement error with the predicted error for the cases of cubic and exponential interpolation respectively. This data is examined in order to discern the effect of the choice of warp and interpolation method in Sections 7.7.1 and 7.7.2 below.

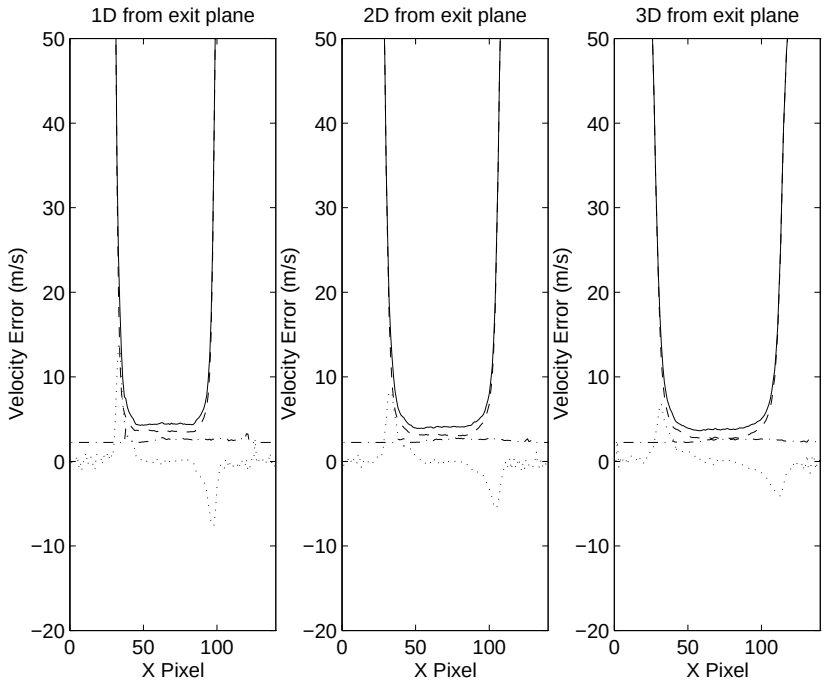


Figure 7.37: Predicted measurement error using the analysis of Chapter 5 for the case of the orthogonal polynomials warp with nearest neighbour interpolation. The dotted line represents the predicted error due to misalignment, the dashed line the predicted error due to shot noise, the dot-dashed line the predicted error due to the iodine cell, and the solid line the overall RMS error.

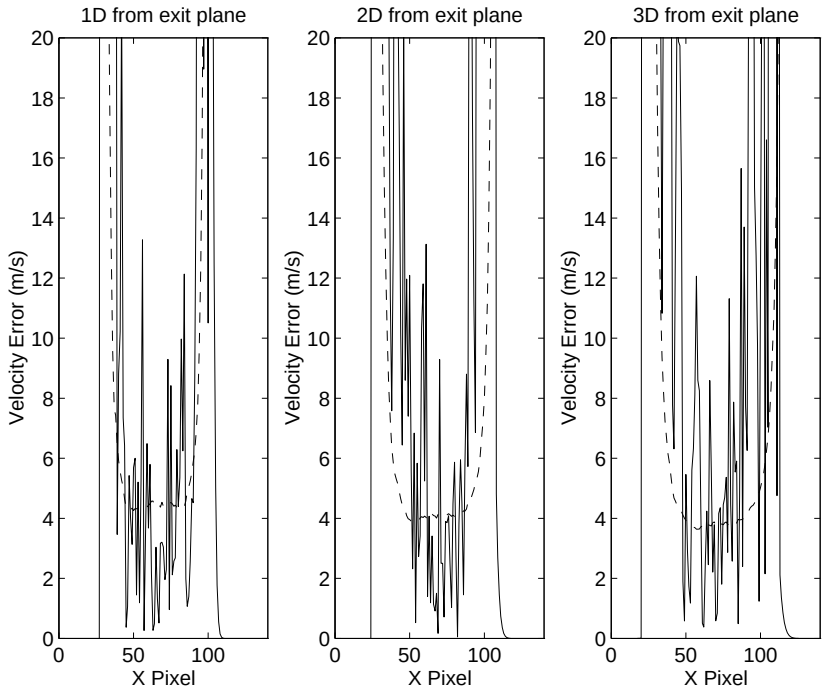


Figure 7.38: Comparison of predicted RMS error with observed error for the case of the orthogonal polynomials warp with nearest neighbour interpolation.

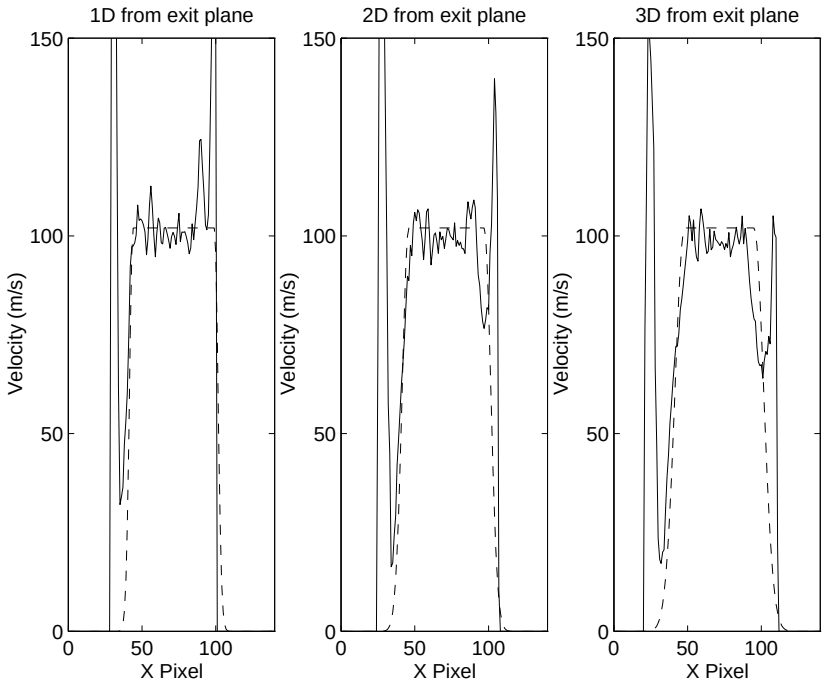


Figure 7.39: Comparison of Doppler global measurements (solid line) with expected data (dashed line) for the case of the orthogonal polynomials warp with bilinear interpolation. Cross-sections one, two, and three nozzle diameters from the exit plane are considered.

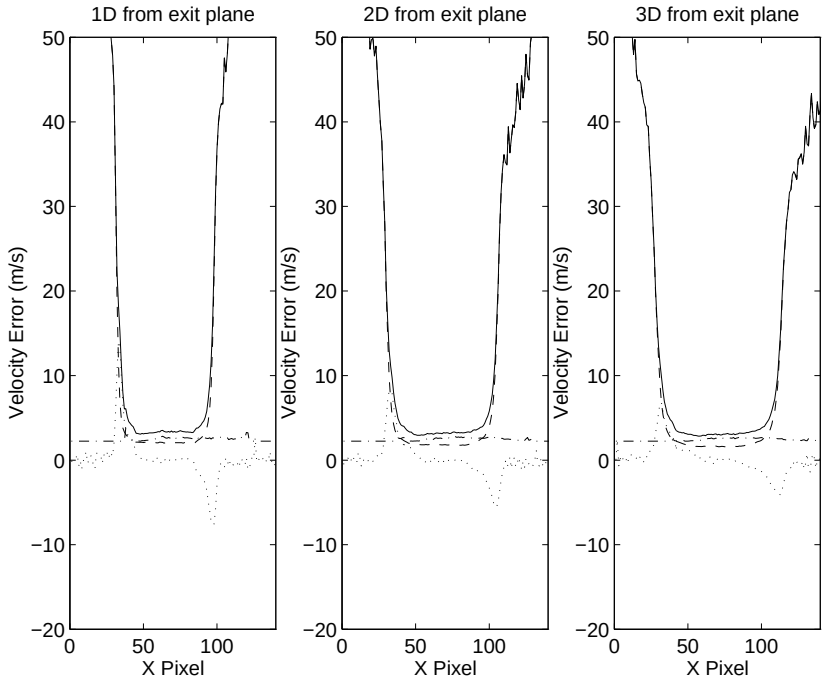


Figure 7.40: Predicted measurement error using the analysis of Chapter 5 for the case of the orthogonal polynomials warp with bilinear interpolation. The dotted line represents the predicted error due to misalignment, the dashed line the predicted error due to shot noise, the dot-dashed line the predicted error due to the iodine cell, and the solid line the overall RMS error.

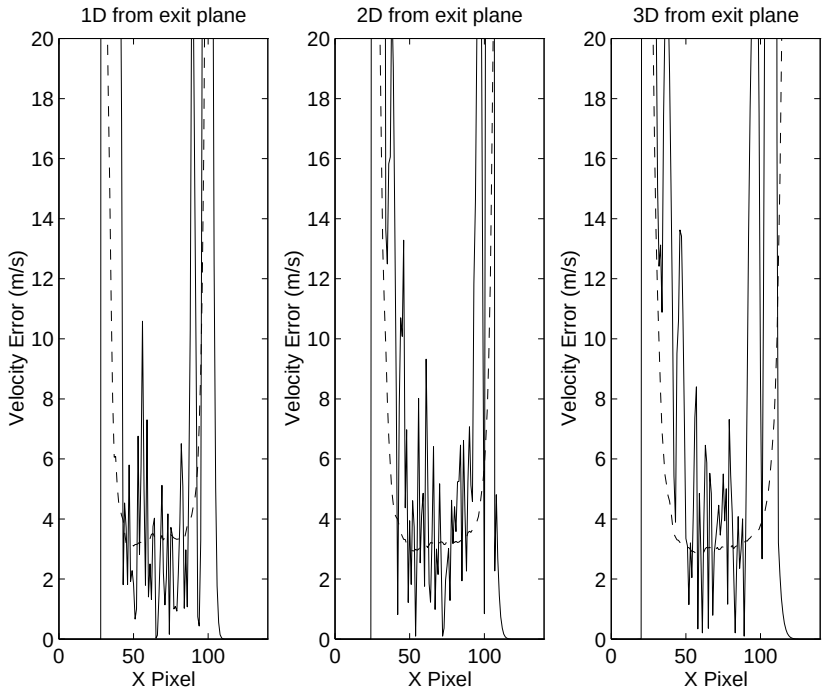


Figure 7.41: Comparison of predicted RMS error with observed error for the case of the orthogonal polynomials warp with bilinear interpolation.

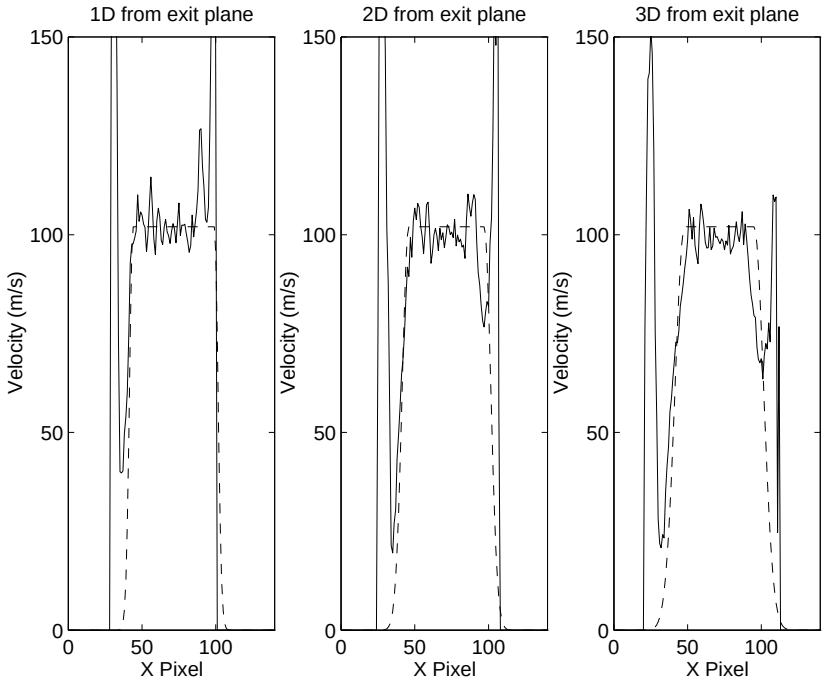


Figure 7.42: Comparison of Doppler global measurements (solid line) with expected data (dashed line) for the case of the orthogonal polynomials warp with cubic interpolation. Cross-sections one, two, and three nozzle diameters from the exit plane are considered.

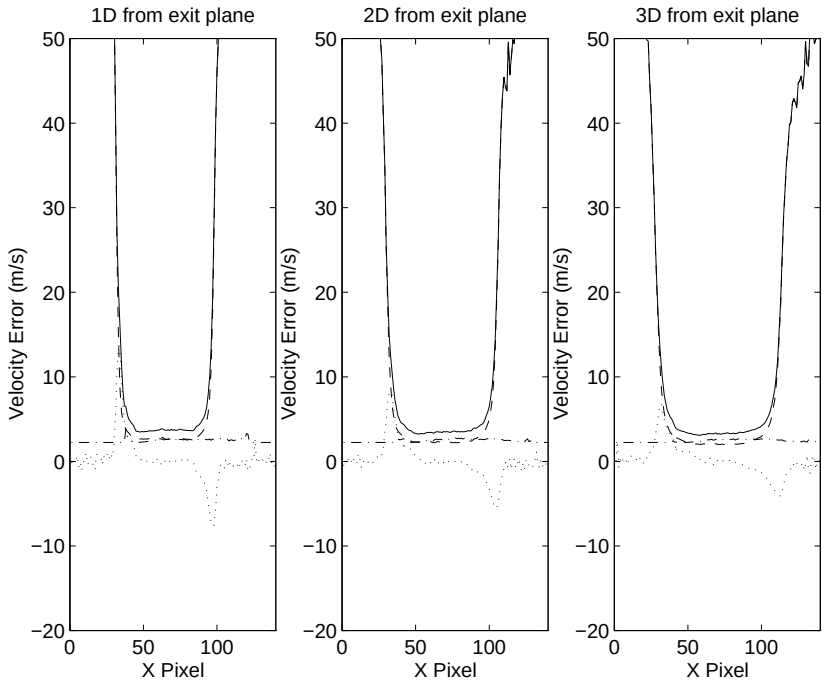


Figure 7.43: Predicted measurement error using the analysis of Chapter 5 for the case of the orthogonal polynomials warp with cubic interpolation. The dotted line represents the predicted error due to misalignment, the dashed line the predicted error due to shot noise, the dot-dashed line the predicted error due to the iodine cell, and the solid line the overall RMS error.

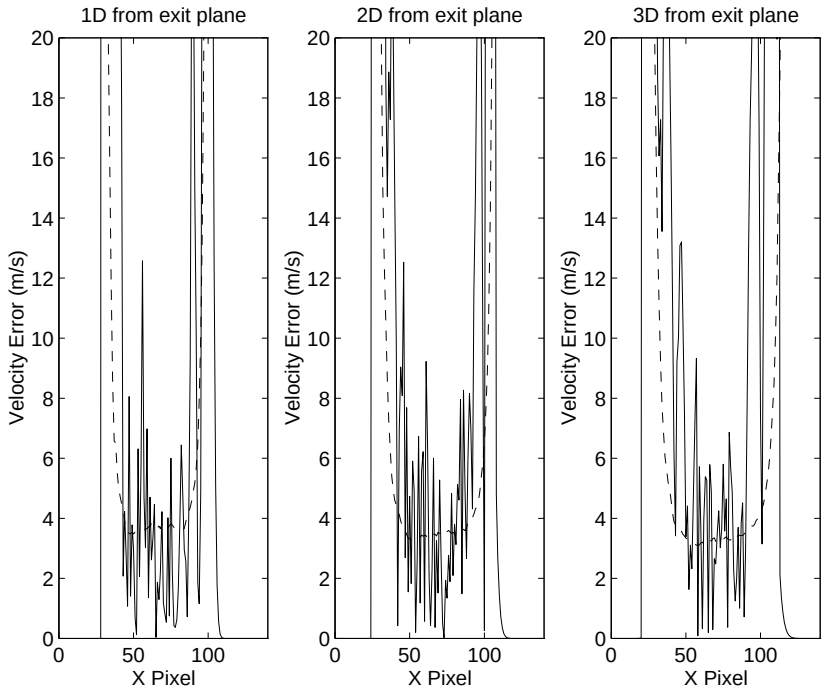


Figure 7.44: Comparison of predicted RMS error with observed error for the case of the orthogonal polynomials warp with cubic interpolation.

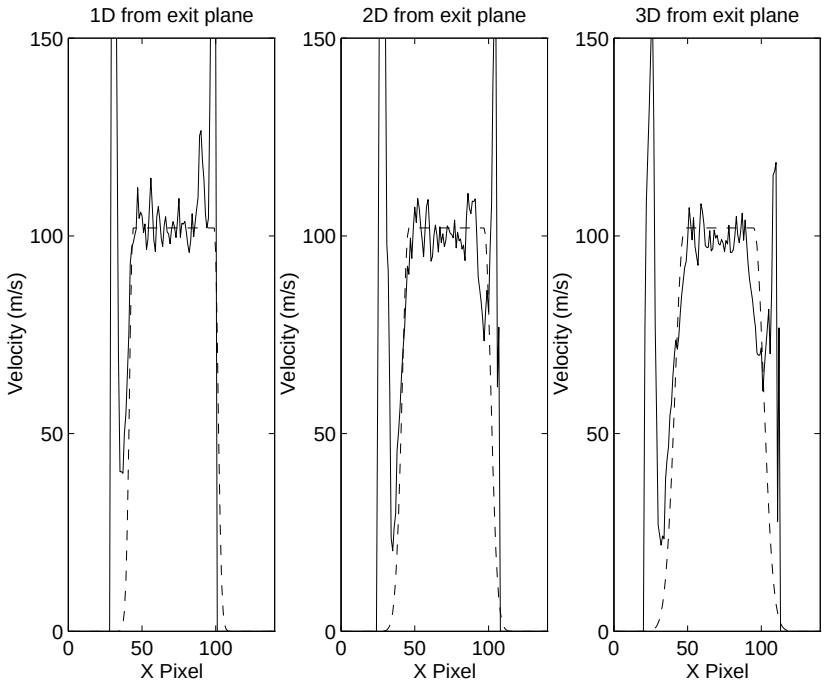


Figure 7.45: Comparison of Doppler global measurements (solid line) with expected data for (dashed line) the case of the orthogonal polynomials warp with exponential interpolation. Cross-sections one, two, and three nozzle diameters from the exit plane are considered.

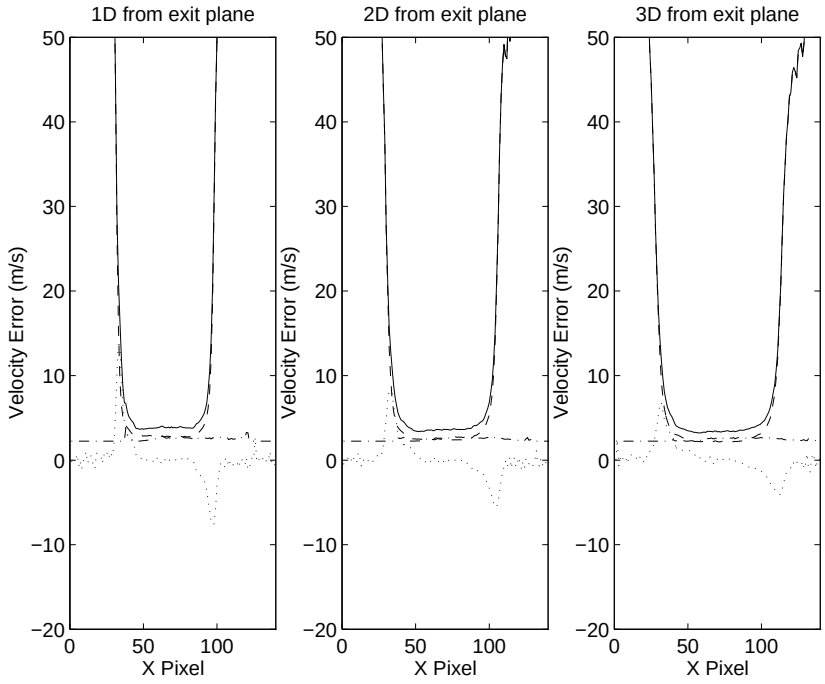


Figure 7.46: Predicted measurement error using the analysis of Chapter 5 for the case of the orthogonal polynomials warp with exponential interpolation. The dotted line represents the predicted error due to misalignment, the dashed line the predicted error due to shot noise, the dot-dashed line the predicted error due to the iodine cell, and the solid line the overall RMS error.

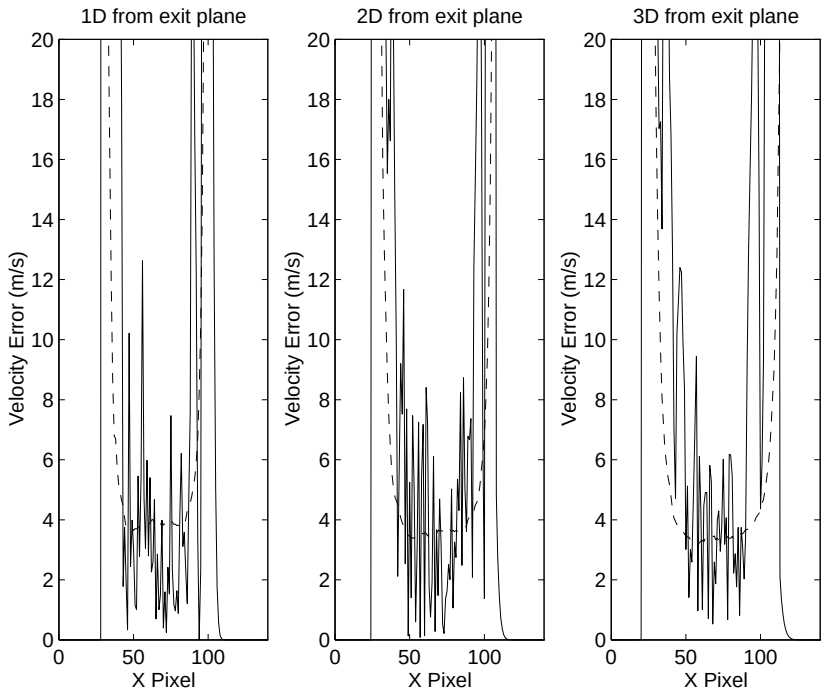


Figure 7.47: Comparison of predicted RMS error with observed error for the case of the orthogonal polynomials warp with exponential interpolation.

7.7 Effect of the Choice of Image Processing Technique

7.7.1 Choice of Warp Method

It is now possible to assess the effect of the choice of warping technique by comparing the results obtained by use of the bilinear warp with bilinear interpolation (Figs. 7.33 to 7.35) to those obtained using the orthogonal polynomials warp with bilinear interpolation (Figs. 7.39 to 7.41).

As explained in Chapter 5, the choice of warping technique affects the overall alignment error across the entire frame. This misalignment results from two different aspects of the warping procedure. First, for a given centroid estimation error (here assumed to be 0.1 pixels) the bilinear warp is expected to yield a greater overall misalignment across the frame than that obtained by use of the polynomial warping technique. This effect is caused by the least-squares fit performed on the centroid positions in the computation of the polynomial warp. The second aspect which must be considered is the ability of the warp to model the distortions present in the optical system. In this regard the polynomial warping technique offers a greater degree of flexibility than the bilinear method. Since optical aberrations are difficult to assess, only the first of these contributions to misalignment is modelled in the analysis of Chapter 5. The pseudo-inverse and the orthogonal polynomials implementations of the polynomial warp should yield identical results. However the danger of ill-conditioning when using the pseudo-

inverse method may result in significant error, as shown in Fig. 5.22. Although this appears not to be the case for the free jet measurements, in general the pseudo-inverse method is best avoided. The choice is thus between the orthogonal polynomials implementation of the polynomial warp and the simpler bilinear transformation.

Fig. 7.34 shows the predicted error using the analysis of Chapter 5 for the case of the bilinear warp with bilinear interpolation. This may be compared with Fig. 7.40 which shows the same predictions for the case of the orthogonal polynomials warp. The predictions due to shot noise and iodine cell errors are the same regardless of the warp method chosen, so the area of interest is the predicted error due to misalignment. For the case of the bilinear warp the peak error due to misalignment at the extreme edge of the jet is around 10 ms^{-1} at one jet diameter from the exit plane, falling to around 6 ms^{-1} three exit diameters downstream. The orthogonal polynomials warp, on the other hand, exhibits a peak predicted misalignment error of around 13 ms^{-1} one exit diameter downstream, falling to around 5 ms^{-1} at three exit diameters. While these figures do not make for a convincing argument in favour of the polynomial warp, it is evident from a comparison of Figs. 7.34 and 7.40 that the region within which the misalignment error is significant is far larger for the case of the bilinear warp. Thus although the peak error values at the extreme edge of the jet may not be significantly reduced by choosing the polynomial warp over the bilinear, the error outside of this region may be improved. This effect may be seen clearly by reference to Figs. 7.27 and 7.28.

It is now necessary to study the actual measured data. Examination of the measured values obtained, as shown by Figs. 7.33 and 7.39, shows that the polynomial warped data exhibits a lower level of apparent noise than that produced by means of the bilinear warp. It is also evident that the large region of high misalignment error which is predicted for the bilinear warp manifests itself as a significant inaccuracy towards the edges of the jet itself. The use of the bilinear warp as shown in Fig. 7.33 results in measured velocity falling off close to the edges of the jet, even though the model due to Gauntner shows no such expected behaviour. The use of the polynomial warp, on the other hand, leads to measurements in closer accordance with the model.

Referring now to Figs. 7.35 and 7.41, it may be seen that the predicted error for the polynomial warp is in close accordance with that observed. As expected, it has a mean value of around 3 ms^{-1} RMS at the centre of the jet flow, and rises sharply in a narrow region at each side. For the case of the bilinear warp, however, the observed data does not fit the expected pattern so closely. The error values at the centre of the jet are erratic due to the

effect of misalignment, but the mean level falls below that expected. Significantly greater errors than those expected occur closer to the edges of the jet flow. The region of extremely high error due to misalignment is even larger than that predicted by the error analysis of Chapter 5.

It has been seen that the free jet flow field has proved to be a stiff test of alignment accuracy and that the actual measured results reinforce the predictions of Chapter 5. The bilinear warp shows even worse performance by reference to the Gauntner model than expected. The reason why the observed performance of the bilinear warp was inferior to that predicted may be the presence of distortions in the optical system that the bilinear warp is unable to represent. The polynomial warp allows better modelling of such non-linear distortion in the optical system. The ability of the bilinear warp to model non-linear distortion is not considered in the error analysis predictions of Chapter 5. In conclusion, the measured performance of the warping techniques provides additional weight to the argument that the orthogonal polynomial warp is to be preferred over the bilinear method.

7.7.2 Choice of Interpolation Method

The impact of the choice of interpolation method may be evaluated from a comparison of all four techniques using the orthogonal polynomials warp (Figs. 7.36 to 7.47).

As has been seen in Chapter 5, the choice of interpolation method may affect the results obtained in two ways. First, each interpolation method may be viewed as a weighted mean of a number of pixels. The greater the size of the interpolation kernel, the more pixels are included in this mean. This means that higher order interpolation schemes should offer superior immunity to errors due to pixel shot noise. It was found, however, that the higher order (cubic and exponential) schemes suffered from numerical precision limitations in their implementation, counteracting any sought-for improvement in noise performance. The second aspect of the choice of interpolation routine is the effect on the misalignment of the discriminated and non-discriminated images. This effect was seen to be most pronounced where high spatial frequency variations in intensity were present. Only the former aspect of the interpolation algorithm is modelled in the error analysis.

Fig. 7.36 shows the results of applying nearest neighbour interpolation. A colour-map representation of the resulting measured flow field may be seen in Fig. 7.7. It is clearly the case, by reference to the latter figure, that the misalignment introduced by the nearest neighbour technique (as much as half a pixel in the worst case) is detrimental in the extreme to the fidelity of the measurements obtained. The theoretical predicted error for this case is shown

in Fig. 7.37. By reference to Fig. 7.38 this prediction may be compared with the observed error one, two, and three nozzle diameters from the exit plane. At the centre of the jet, an error of 4 ms^{-1} is predicted. This error estimate is far greater at the edges of the jet flow, due to a combination of misalignment and shot-noise error contributions. The observed error agrees reasonably well with the prediction close to the centre of the jet, but is seen to exceed that predicted close to the sides of the jet. This effect may be due to the fact that the estimated error ignores the additional misalignment introduced by the interpolation method, which may be significant in this case.

Bilinear interpolation may be employed, yielding the results of Fig. 7.8. It is immediately apparent that the use of bilinear rather than nearest neighbour interpolation has the effect of eliminating the structure in the results data which arises due to misalignment. In order to assess the performance of bilinear interpolation fully, the cross-sectional plots of measured data must be compared with the theoretical model of Gauntner as shown in Fig. 7.39. The error prediction is recalculated for the case of bilinear interpolation and is plotted in Fig. 7.40. The comparision of observed error with predicted error is shown in Fig. 7.41. Whereas for nearest neighbour, the predicted error was around 4 ms^{-1} at the centre of the jet, for bilinear the predicted error is only 3 ms^{-1} at the same point. Once again, predicted error rises sharply at the edges of the jet itself. The error in the measured data fits the predicted error satisfactorily yet exceeds it towards the edges of the jet. The variance of the error in this case is somewhat reduced in comparison with the case where nearest neighbour interpolation is employed.

It seems clear that the bilinear technique is to be preferred over the nearest neighbour method for a number of reasons. The structure evident in the nearest neighbour interpolated data of Fig. 7.7 is not present in the bilinear interpolated data of Fig. 7.8. A lower estimate of error is obtained from the error analysis technique if bilinear interpolation is used rather than the nearest neighbour approach. This lower estimate seems to be justified by the observed error values, which are correspondingly lower at the same point in the measurement frame when compared to those arising from the use of the nearest neighbour interpolation method.

Cubic and exponential interpolation techniques were used to obtain the results images of Fig. 7.9 and 7.10 respectively. The corresponding line plots and correlations may be seen in Figs. 7.42 and 7.45. Performing error estimation at measurement planes one, two and three nozzle diameters from the exit plane produces the plots of Figs. 7.43 and 7.46. Finally, Fig. 7.44 shows the observed and predicted errors plotted together for the case of cubic interpolation. Fig. 7.47 shows the same results for the case of exponential interpolation.

In both cases the theoretical predicted error is close to 3.5 ms^{-1} at the centre of the jet flow. It is fair to say from comparison of the line plots of Figs. 7.44 and 7.47 and also from comparison of the colour-map representations 7.9 and 7.9 that there is no discernable difference in performance between the cubic and the exponential methods.

Comparing these results with those arising from the use of a bilinear interpolation technique reveals a number of interesting points. First, the estimated error within the jet itself is 3.5 ms^{-1} for the cubic and exponential techniques, and only 3 ms^{-1} for the bilinear method. By means of comparison, nearest neighbour manages a 4 ms^{-1} RMS error at this point in the flow. The observed data too is seen to follow the same trend, lending credence to the error predictions of Chapter 5.

This being the case, it is straightforward to say that the bilinear approach is superior to the nearest neighbour method. It is less easy to say whether cubic or exponential schemes are to be preferred to a bilinear method, however.

It has been seen that the higher order schemes have performed slightly less well that the simpler bilinear algorithm on the free jet data. The error introduced by the limited machine precision has dominated over any expected improvement in shot noise immunity offered by the cubic or exponential schemes. In the case of the flow field under consideration, there are no very high spatial frequency variations in intensity. Were such features present, the higher order methods would be expected to perform better, as shown in Chapter 5.

In general, for most practical flow fields under consideration, such high spatial frequency intensity variations are not to be expected. This being the case, it seems hard to justify the use of the cubic or exponential schemes, since their performance is so often inferior to that of the bilinear method.

A final consideration in the choice of interpolation scheme may be made with the aid of the line plots of axial velocity along the jet centre-line shown in Figs. 7.11 to 7.14. It may be seen at a glance from these graphs that the observed variance of the measurement error is greatest for the nearest neighbour technique and least for the bilinear interpolation technique. The cubic and exponential schemes are once again seen to offer very similar performance to one another, superior to that of the nearest neighbour method, but inferior to that of the bilinear scheme in this application. Thus the choice of the bilinear interpolation method may be seen to offer not only the least measurement error with respect to the predictions of Gauntner but also the lowest error variance. Accordingly the bilinear interpolated colour-map representations of the measured velocity superficially appear to be the least noisy.

7.8 Summary

In this chapter, Doppler global velocimetry measurements of the seeded flow field of a free jet have been presented. The apparatus has been configured in two different ways. The first configuration allows measurements of axial velocity to be obtained as seen from one side of the jet, whereas the second configuration allows a cross-sectional view of the jet velocity to be observed at a number of positions downstream of the exit plane.

In order to correlate this model to the measured data, a model of the expected flow field resulting from such a free jet (taken from Gauntner, Lingwood and Hrycak) has been outlined and a method for determining the potential core length has been employed. The divergence between the model and the measured values may then be taken as an estimate of Doppler global velocimetry measurement error. The error analysis of Chapter 5 has been utilised to predict the expected measurement error due to iodine cell calibration, misalignment, and shot noise.

It has been seen that unlike the error in the rotating disc experiment of Chapter 6, the error obtained in measurements of the free jet flow is extremely sensitive to misalignment because of the high intensity and velocity gradients at the edge of the jet flow itself. The free jet thus makes an ideal test case, enabling the relative performance of the different warping and interpolation techniques to be assessed. This practical assessment has been compared to the predicted performance of the techniques obtained in Chapter 5.

In evaluating the relative performance of the bilinear and the orthogonal polynomials warping techniques, it has been found that for a significant region close to the edge of the jet flow the measured error due to the use of the bilinear method was far in excess of that due to the polynomial algorithm. The analysis of Chapter 5 predicted that the bilinear warp would perform less well than the polynomial warp, due to its greater susceptibility to centroid estimation error. It was expected that the bilinear warp would exhibit poor performance in the region of high intensity and velocity gradients at the edge of the jet flow. This prediction was found to be valid, although the extent of the error due to the use of the bilinear warp was found to be significantly greater than that predicted. The orthogonal polynomials warp, on the other hand, performed in close accordance with the predicted behaviour.

Measured values were also used to test the validity of the predicted error values due to the choice of interpolation method. By comparing with the Gauntner model the results obtained using the nearest neighbour, bilinear, cubic and exponential schemes, all implemented with the orthogonal polynomials warp, an impression of the performance of the techniques was

obtained. As expected, the bilinear interpolation method offered superior error performance to the nearest neighbour method, the latter method introducing significant structure into the computed velocity field. The more complex cubic and exponential schemes were also found to be superior to the nearest neighbour method, but were inferior to the bilinear method, as a result of the limited machine precision used for computation. In conclusion then, for flow fields such as the free jet, where extremely high spatial frequencies of intensity or velocity variation are not expected to be present, the bilinear interpolation method is the preferred technique.

With the use of bilinear interpolation and an orthogonal polynomial warp, measurements may be undertaken with errors of around 4 ms^{-1} within the body of the jet flow itself. This error rises sharply within a small region at the edge of the jet flow due to the misalignment sensitivity in this region discussed above, and also due to increased shot noise in the low intensity region.

Such accurate measurements of the free jet flow have been shown to be attainable only with accurate alignment. Thus, the use of the orthogonal polynomials warp has been found to be necessary for adequate performance in such a case. The error analysis of Chapter 5 has been seen to model the effect of the overall error with a good degree of accuracy, lending additional confidence to the theoretical predictions with regard to image processing techniques. The error attainable using appropriate image processing techniques has been seen to be dominated by the unavoidable effects of shot noise in the CCD array.

Chapter 8

Velocimetry Measurements of a Convergent-Divergent Nozzle

8.1 Motivation

The Doppler global velocimetry measurements of a convergent-divergent nozzle presented in this chapter are the next step towards the eventual aim of developing a Doppler global system that may be applied to a turbomachinery flow field. The nozzle is the first transonic flow field to which Doppler global velocimetry has been applied. Whereas the rotating disc experiment tests the low velocity resolution of the instrument, the convergent-divergent nozzle presents new problems in high velocity flow field determination.

The convergent-divergent nozzle was chosen as a further test case for the Doppler global velocimetry system for a number of reasons. The nature of the flow in such a nozzle is well known, allowing a comparison of measured data with theoretical predictions to be made. Computational fluid dynamics codes have been run using the same nozzle geometry. Those predictions are presented here, and compared to the measured values obtained. This test case allowed the ability of the technique to measure seeded flows in the transonic regime, rather than those of lower velocity such as the free jet discussed in Chapter 7, to be assessed.

The future proposed application of the Doppler global velocimetry method is the measurement of complex turbomachinery flows. Such an application will require that fluid flows in confined passages be imaged. The measurements presented in this chapter are the first attempt at Doppler global velocimetry for an enclosed flow. In addition, any attempt at measuring turbomachinery flows will require measurements to be made through another transparent medium, such as glass or perspex. Again, the convergent-divergent nozzle experiment is the first attempt at such a measurement.

8.2 Configuration

Fig. 8.1 shows the apparatus employed in the convergent-divergent nozzle test rig. A vacuum tank of approximate volume 2 m^2 is evacuated using a small rotary vacuum pump. This vacuum tank is connected to the nozzle itself using 2 inch diameter piping, incorporating a pair of ball valves. The downstream ball valve is set to a partially open position in order to act as a throat. To ensure time-invariant flow conditions in the nozzle, this throat should be choked during the measurement time. The upstream ball valve, which is simply used to

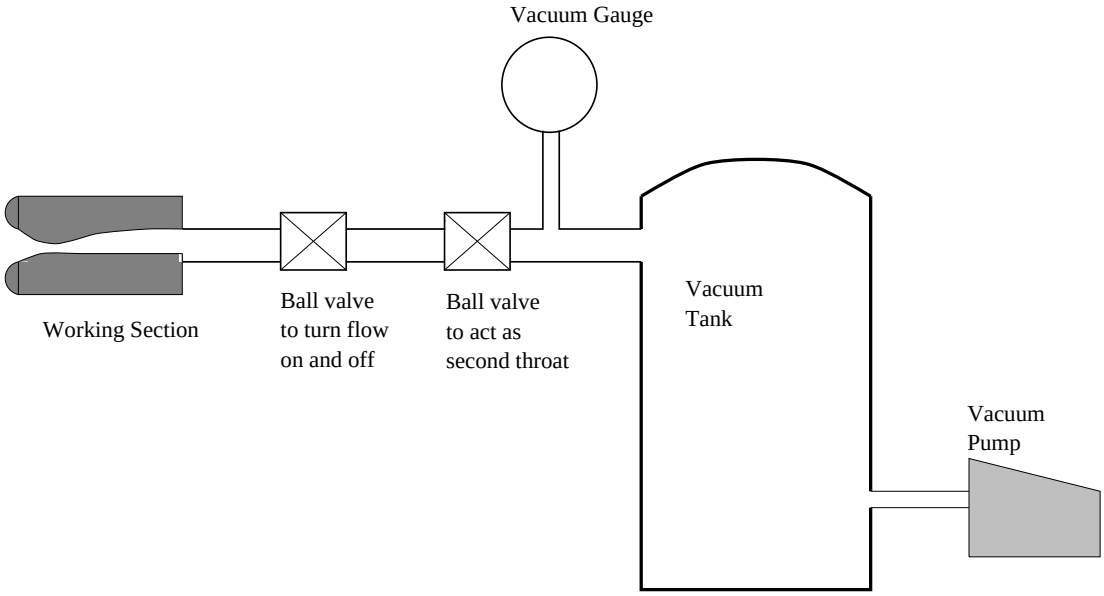


Figure 8.1: Convergent-divergent nozzle apparatus.

turn the flow on and off, is either fully opened or fully shut.

The mouth of the nozzle is subject to atmospheric pressure, whereas the vacuum tank is typically evacuated to around -0.95 bar relative to atmospheric pressure, as indicated by a Bourdon vacuum gauge. Upon opening the upstream ball valve, flow commences, and the downstream ball valve is choked for 4 to 5 seconds, during which time measurements may be undertaken. The working fluid is, of course, air.

Since the valve remains choked only for a short period, this is the limiting factor which determines the duration of exposure which may be employed. The use of shorter exposures clearly has the effect of reducing the image intensity obtainable. In order to counteract this, the intensity of illumination, the seeding density, or the aperture size may be increased. Since the laser was operating at full output power, and the seeding density was as great as the atomiser employed would allow, the only option for increasing image intensity was to employ larger apertures. This has the effect of reducing depth of field. Since the flow itself is viewed at an angle of 30º from the normal, a certain depth of field is required in order to allow the entire region of interest to be in focus. This constraint limits the aperture size which may be employed. The use of larger apertures for alignment may result in an increase in systematic error and hence may be detrimental to alignment accuracy, as explained in Section 5.7.6. The use of a more powerful laser, such as a Nd:YAG pulse laser would help to overcome the problems arising from low image intensity. The use of such a pulse laser would allow truly instantaneous, rather than time-averaged, velocity data to be acquired.

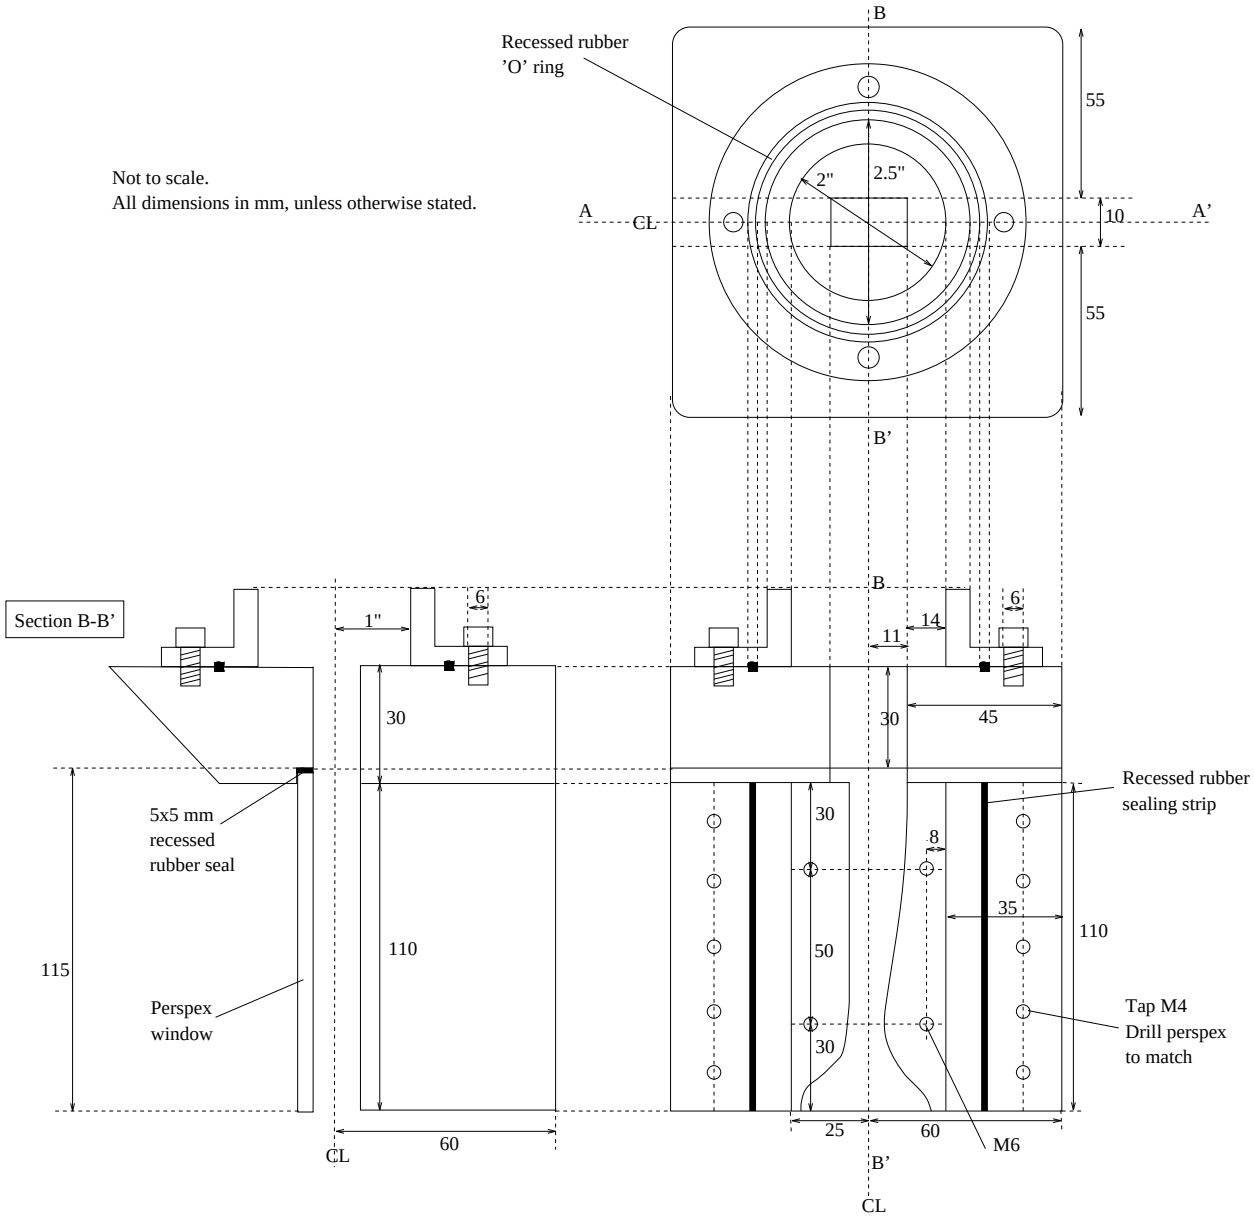


Figure 8.2: Convergent-divergent nozzle working section.

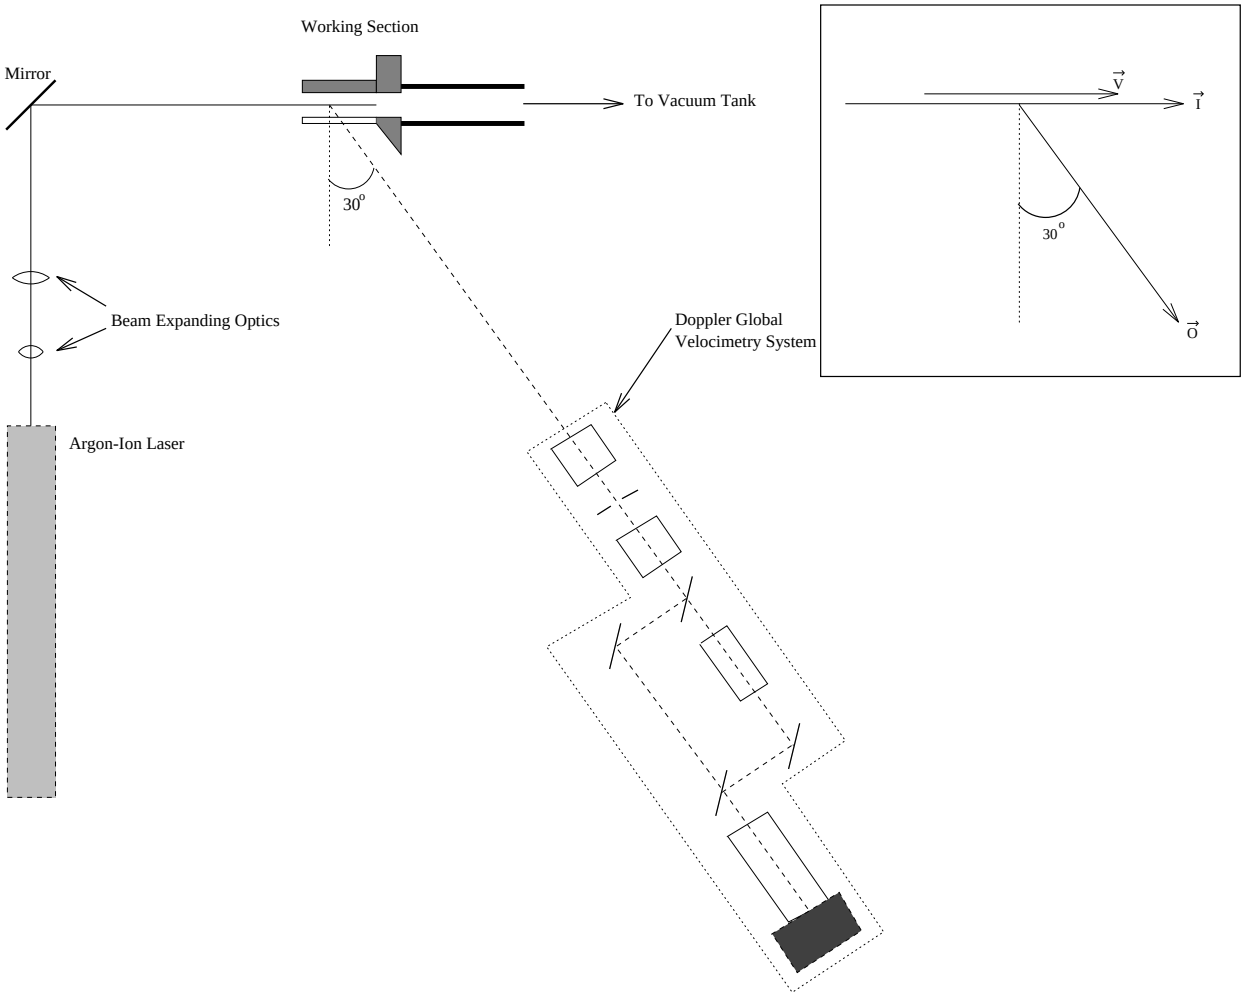


Figure 8.3: Convergent-divergent nozzle – viewing configuration.

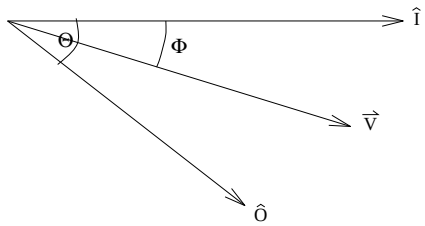


Figure 8.4: General Viewing Configuration.

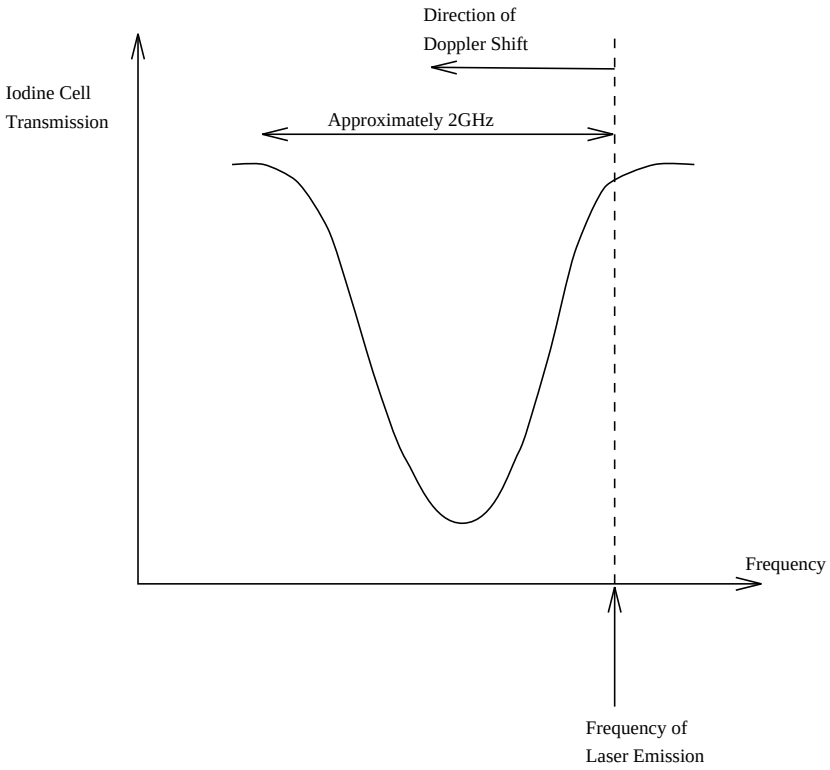


Figure 8.5: Laser Frequency Setting for Convergent-Divergent Nozzle Application.

The design of the convergent-divergent nozzle itself is shown in Fig. 8.2. The nozzle is made up of a pair of liners, one flat and the other profiled, which form the convergent-divergent passage itself. The design of the liners is copied from an existing supersonic nozzle experiment in use at the laboratory. The pair of liners is mounted in an aluminium block, allowing connection to the vacuum pipe. The front face of the nozzle is made from perspex, allowing optical access, and the downstream face of the aluminium housing has been cut away to allow unimpeded viewing. The entire assembly has been sprayed with matt black paint in an attempt to minimise the effect of reflections within the nozzle itself. A bellmouth of radius 25 mm is fitted to the entrance of the nozzle to minimise the effect of separation. The nozzle itself is 100 mm in length, and of breadth 10 mm at the throat. The channel depth is 10 mm throughout.

Fig. 8.3 shows the configuration of the Doppler global velocimetry system relative to the nozzle. The emission of the argon ion laser is expanded into a light sheet using a cylindrical lens and a spherical lens. This light sheet is diverted through a right angle and hence into the mouth of the nozzle itself. The light sheet is positioned at the centre of the channel depth. The illuminating optics are adjusted so that a narrow light sheet width is obtained, of such a width that the sheet occupies the entire flow area of the nozzle. Care must be taken to ensure that the light sheet does not impinge excessively on the side walls of the

nozzle.

The Doppler global velocimetry system is positioned to observe the flow at an angle of 30° from the normal. The image processing techniques described in Chapter 3 are employed, so the perspective effect is eliminated from the measured frames, resulting in the flow field appearing as if viewed head-on.

The reason for the choice of observation angle may be seen by reference to Section 2.4 where the dependency of the measured component on illumination and observation directions is described.

A general viewing configuration is shown in Fig. 8.4. The velocity vector is defined at an angle ϕ to the incident illumination direction, whereas the observation vector is defined at an angle θ to the incident illumination direction. For this configuration, let

$$\hat{\imath} = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \tag{8.1}$$

$$\hat{o} = \begin{pmatrix} \cos\theta \\ -\sin\theta \end{pmatrix} \tag{8.2}$$

$$\vec{V} = \begin{pmatrix} v\cos\phi \\ -v\sin\phi \end{pmatrix} \tag{8.3}$$

so, using eqn. 2.1

$$\Delta\nu = \frac{\nu_o v}{c}((\cos\theta - 1)\cos\phi + (\sin\theta.\sin\phi)) \tag{8.4}$$

If $\phi = 0^o$,

$$\Delta\nu = \frac{\nu_o v}{c}(\cos\theta - 1) \tag{8.5}$$

The maximum velocity expected in the convergent-divergent nozzle flow is of the order of 550 ms^{-1}. It has been seen in section 2.5 that the usable bandwidth of the iodine absorption feature is of the order of 1.5 GHz. Only half of the allowable frequency span of the iodine absorption profile may be employed, so that there is no ambiguity in determining the frequency of the scattered radiation. Were the entire iodine cell bandwidth employed, it would not be possible to discriminate between the two wings of the absorption curve.

In addition, because of the position of the argon ion laser resonant modes and the small gradient of the absorption features at its extremes and at its centre, it is best to further restrict the portion of the absorption profile in use. The Doppler shift of the scattered light is expected to be negative; in other words, the scattered light is expected to be of lower frequency than the incident laser radiation. Assuming a maximum allowable value of $\Delta\nu$ of 0.65 GHz, setting ν_o/c to $1/(514.5 \times 10^{-9})$ m^{-1} and assuming a maximum velocity of 550 ms^{-1}, it may be seen that

$$-0.65 \times 10^9 \leq \left(\frac{550.0}{514.5 \times 10^{-9}}\right)(\cos\theta - 1) \tag{8.6}$$

So,

$$\phi \leq 67.0^{\mathrm{o}} \tag{8.7}$$

In other words, the observation angle must be greater than $90 - 67.0 = 23.0^{\mathrm{o}}$ from the normal. An observation angle 30^{o} from the normal was thus chosen, as shown in Fig. 8.3.

As discussed above, the sense of the Doppler shift in the scattered light is expected to be negative. The frequency of the scattered light is thus lower than that of the incident laser illumination. Accordingly, the laser emission frequency is tuned to the high frequency side of the iodine absorption spectrum, as shown in Fig. 8.5. In order to ensure that sufficient frequency range may be obtained without encountering velocity ambiguity it is essential that the laser frequency should be set as close as possible to the extreme high frequency side of the feature. Accurate setting and monitoring of the laser frequency for this purpose was achieved by means of the reference cell apparatus and the `IODINE` code.

It has been seen that, in the presence of the large velocities expected in the convergent-divergent nozzle, it has been necessary to choose an oblique viewing angle. In this configuration, it is necessary to choose a small aperture in order to obtain sufficient depth of field, further exacerbating the problems presented by low image intensity. For such situations, it would be useful if the working iodine cell profile had a greater bandwidth, allowing the flow field to be viewed head on. If small quantities of impurities are introduced into the iodine cell during fabrication, the bandwidth of the absorption feature may be increased, the resulting cell being known as a *broadened cell*. Future work should examine the fabrication of such cells, and investigate the application of an appropriate broadened cell to the measurement of the convergent-divergent nozzle.

In addition to the problem of excessive Doppler shift, the measurement of flow within the convergent-divergent nozzle presents a number of other difficulties. The nozzle is the first flow field contained within a channel for which Doppler global velocimetry measurements have been attempted. If the incident light sheet impinges on the walls of the channel excessively, a significant proportion of the light is scattered back directly into the optical system. The application of matt black paint to the nozzle walls helped to reduce this problem, but failed to eliminate it entirely. The intensity due to scatter from the side walls, imaged onto the CCD array of the camera may easily lead to saturation. Since the CCD camera in use makes no provision for draining the charge from saturated pixels, large areas of the frame may be rendered unusable by excessive glare. The only effective remedy to this problem which has been found is to reduce the effective light sheet width by adjusting the illuminating optics until the glare is of sufficiently low magnitude to avoid CCD saturation.

As explained in Section 5.5, the measured component of velocity is dependent on the direction of propagation of the incident light — thus reflections from the side walls of the nozzle may result in errors in the determined velocity in that region. In the configuration used here, even with a reduced light sheet width, it is inevitable that some of the light sheet will be scattered from the side walls of the nozzle. The intensity profile of the light sheet may be represented as a Gaussian curve. Hence, even reducing the width of the sheet will not eliminate scatter from the side walls altogether.

There are a number of avenues that may be explored in order to overcome the problem of excessive scatter from the side walls of an enclosed passage. It may appear attractive to employ polarising filters in order to absorb the light which is scattered from the side walls. Such a strategy is unwise, however, since the polarisation state of the light scattered from the seeding particles within the flow is itself unknown. Great care has been taken to ensure that the Doppler global velocimetry optical system is non polarisation sensitive, so that measured velocities may not be distorted because of the polarisation state of the radiation scattered from the seeding particles.

Rather than painting the liners matt black in an attempt to absorb the laser emission, it may be possible to employ certain fluorescent paints which absorb light at one frequency and re-emit it at another. It would then be straightforward to filter out this re-emitted radiation. Another approach is to manufacture nozzle liners from a translucent material, so that the incident laser radiation passes into the liners where it is absorbed, rather than being reflected at the surface.

It has been proposed [McKenzie 1996] that an additional iodine cell be introduced into the Doppler global velocimetry receiving optical system in order to eliminate the effect of stray reflections from channel walls. Such an cell would be mounted in front of the entire optical system, so that it has an effect on both the discriminated and the non-discriminated images. By maintaining this cell at an elevated temperature, the iodine absorption profile could be sharpened. In this way, the cell could be configured so that light at the laser emission frequency ν_o was entirely absorbed, but light that had been Doppler shifted by a sufficient amount would be transmitted. Unshifted light scattered from the nozzle itself may thus be suppressed, avoiding the problem of CCD saturation. The addition of this extra cell would not, however, solve the problem of velocity determination error due to stray scatter. The scattered radiation has been Doppler-shifted in frequency, so is no longer at the laser emission frequency, ν_o. A further drawback is the inevitable loss of intensity caused by the introduction of the extra cell.

Latex spheres of 0.5 μm diameter in a suspension of isopropanol, available commercially for use with particle image velocimetry techniques, were employed as the seeding medium. As has previously been explained, the exposure time was limited to five seconds. In order to avoid severe shot noise error, it was important to ensure sufficient seeding density in order to allow significant light levels to reach the CCD. The volume flow rate through the nozzle is significant, around 0.033 $\mathrm{m^2 s^{-1}}$. In order to obtain a satisfactory concentration of seeding, a commercial seeding generator, the T.S.I. Six Jet Atomiser, was used. Seeding was delivered through a flexible tube of one inch diameter to the mouth of the nozzle itself. The seeding was propelled by the laboratory 400 psi compressed air main.

An analysis of one-dimensional flow in a channel with negligible friction, with particular regard to the case of the convergent-divergent nozzle, may be found in Massey [Massey 1989].

8.3 Doppler Global Velocimetry Measurements

The Doppler global velocimetry measurement of the convergent-divergent nozzle, in the form of a colour-map representation, is shown in Fig. 8.6. In accordance with the theoretical findings of Chapter 5 and the practical confirmation (for the case of the free jet) presented in Chapter 7, the orthogonal polynomials warp with the bilinear interpolation algorithm has been used for processing the convergent-divergent nozzle measurements. The shape of one of the nozzle liners is discernable on the left hand side of this image — the nozzle liner on the other side of the channel is flat. Within a region approximately 8 mm wide centred within the nozzle itself, the laser light sheet is of sufficient intensity to allow accurate measurements to be made. Outside of this region, the light sheet intensity is too low to allow any interpretation

of the measured velocity. The threshold for rejection of data on the grounds of low non-discriminated image intensity was set to zero to enable the profile of the nozzle to be discerned in the final image. The dark spot on the image approximately 12 mm downstream of the throat is due to the presence of a blanking plate on the back wall of the nozzle itself, which has the effect of scattering laser light into the optical system, causing saturation of the CCD locally.

As expected, an acceleration is seen through the throat of the nozzle, reaching Mach 1.0 at the throat itself, and continuing to accelerate to approximately 430 ms^{-1}, 20 mm downstream of the throat.

A computational fluid dynamics prediction of the flow in the convergent-divergent nozzle was undertaken, using the commercial package CFDS CFX Flow3D. A colour-map representation of the velocity predictions from this simulation for the entire nozzle (rather than just the region measured using Doppler global velocimetry) is shown in Fig. 8.7.

To enable comparison of the computational fluid dynamics predictions with the measured velocity, the two representations must be aligned with one another by reference to the shape of the nozzle itself. Figs. 8.8 and 8.9 show the two sets of data suitably aligned for comparison. Fig. 8.10 shows the absolute velocity difference between the prediction and the measured flow field, whereas Fig. 8.11 represents the percentage error of the measurement relative to the prediction. While it is envisaged that an eventual application of the Doppler global velocimetry technique will be to provide experimental measurements to facilite the validatation of CFD codes, in this chapter a computation solution is used to establish confidence in the Doppler global velocimetry results.

It is immediately apparent that across the main part of the channel downstream of the throat, there is only a small difference between the measured velocity and the predictions, representing an error of around five percent or 20 ms^{-1}.

Closer to the side walls, however, the error increases to as much as 40 or even 50%. This error may be attributed to scattered light from the side walls influencing the measured result as outlined above. It may be seen that the Doppler global velocimetry measurements are lower than those predicted in this region — this is consistent with a lower observed velocity, caused by scattering of the light sheet from the side walls.

The other region where significant error is observed is upstream of the throat. In this region, the Doppler global velocimetry technique shows a significantly higher velocity than that predicted by the computational fluid dynamics code. The code used was two-dimensional

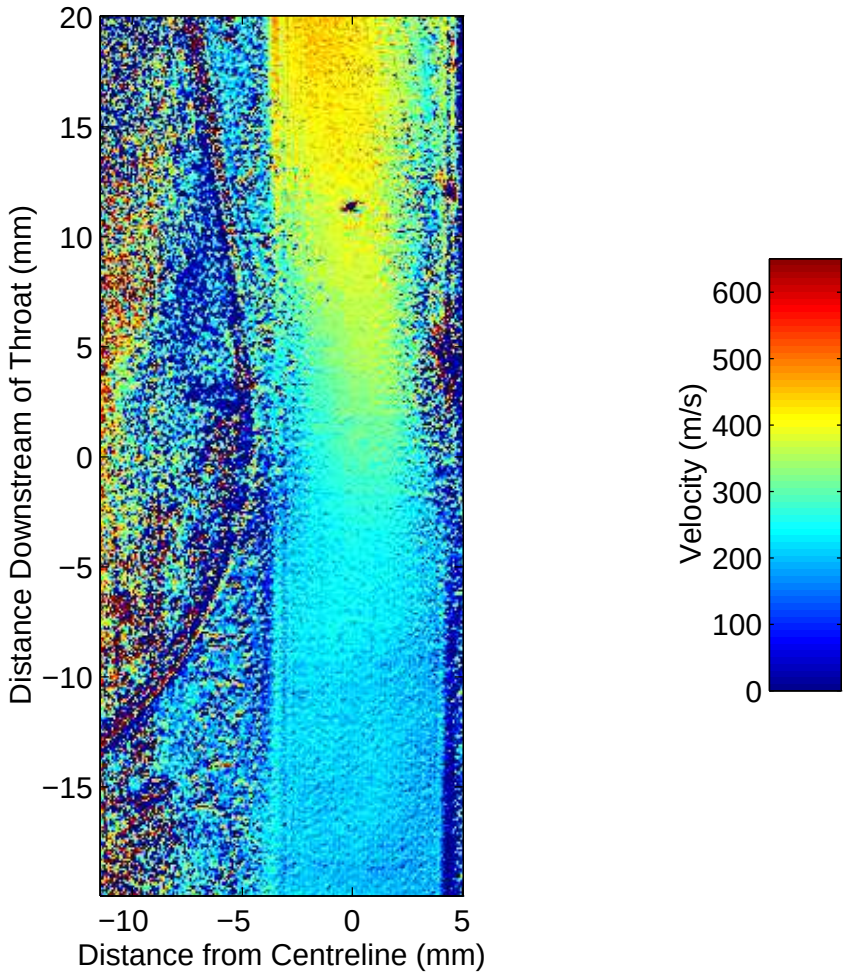


Figure 8.6: Convergent-divergent nozzle flow as measured by Doppler global velocimetry.

only, and is thus unable to model separation effects across the depth of the channel. It is probable that the flow separates from the channel walls at entry, leading to a higher velocity at the centre of the channel than expected.

Fig. 8.12 shows a line plot of measured centre-line velocity against the predicted velocity along the same line. It may be seen that the measurements and the predictions are in close agreement with one another downstream of the throat, although the Doppler global velocimetry measurements are, of course, significantly noisier than the Flow3D predictions. Fig. 8.13 is a similar line plot, but this time the mean of twenty pixel rows each side of the centre-line has been taken for both the measurement and the prediction. The use of such a mean enables a significant proportion of the measurement noise to be eliminated. It may be seen that from a position 5 mm upstream of the throat the measured and predicted velocities are in good agreement.

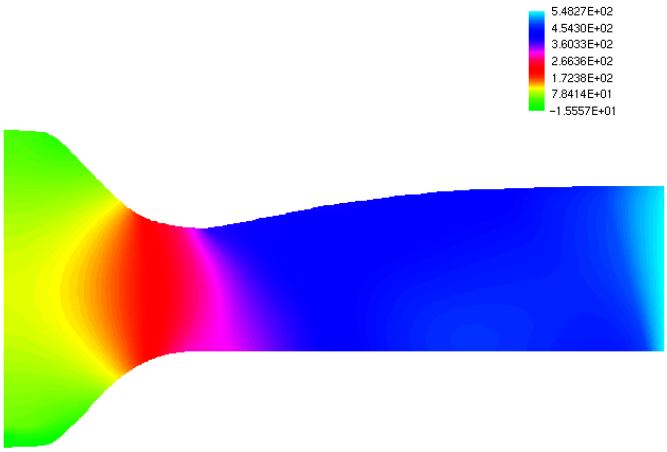


Figure 8.7: Convergent-divergent nozzle flow as predicted by FLOW3D computational fluid dynamics code.

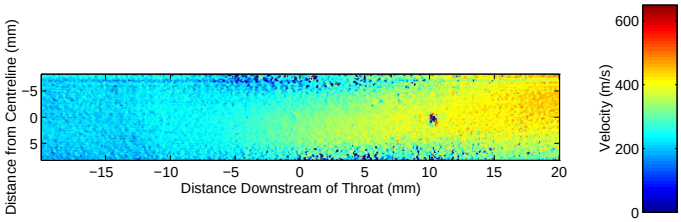


Figure 8.8: Convergent-divergent nozzle flow as measured by Doppler global velocimetry. The region of interest has been aligned to allow comparison with computational fluid dynamics results.

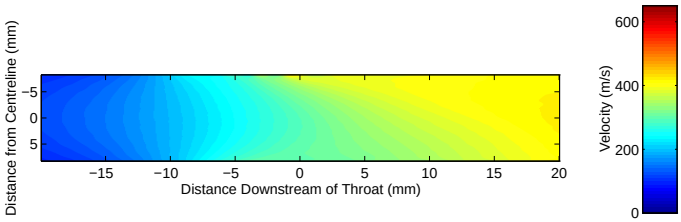


Figure 8.9: Convergent-divergent nozzle flow as predicted by FLOW3D computational fluid dynamics code. The region of interest has been aligned to allow comparison with Doppler global velocimetry results.

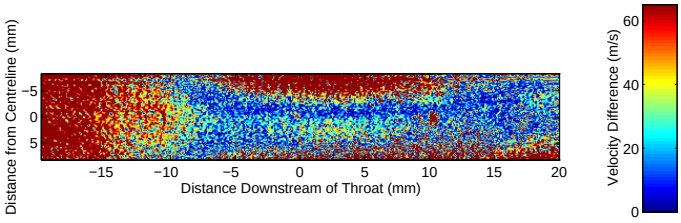


Figure 8.10: Absolute velocity difference between computational fluid dynamics prediction and Doppler global velocimetry measurement.

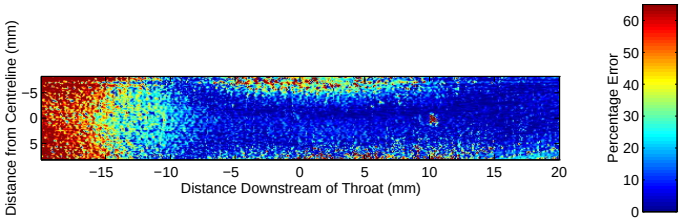


Figure 8.11: Percentage error between computational fluid dynamics prediction, and Doppler global velocimetry measurement.

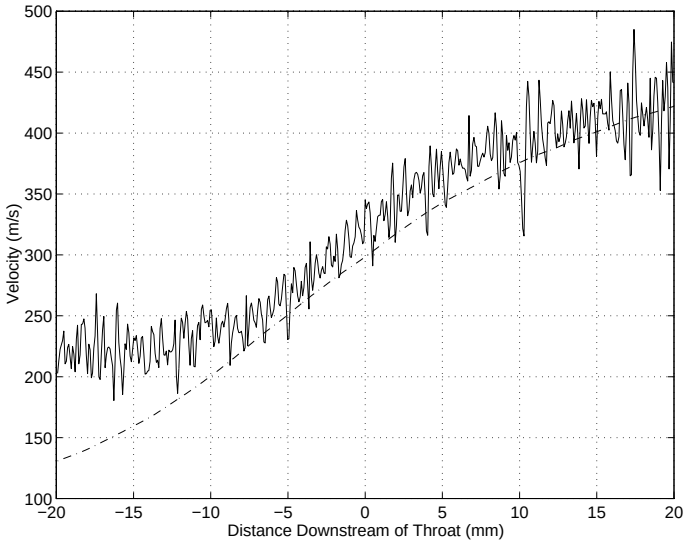


Figure 8.12: Centre-line velocity. The Doppler global velocimetry measurements for a single row of pixels are shown as a solid line. The computational fluid dynamics predictions are shown as a dotted line.

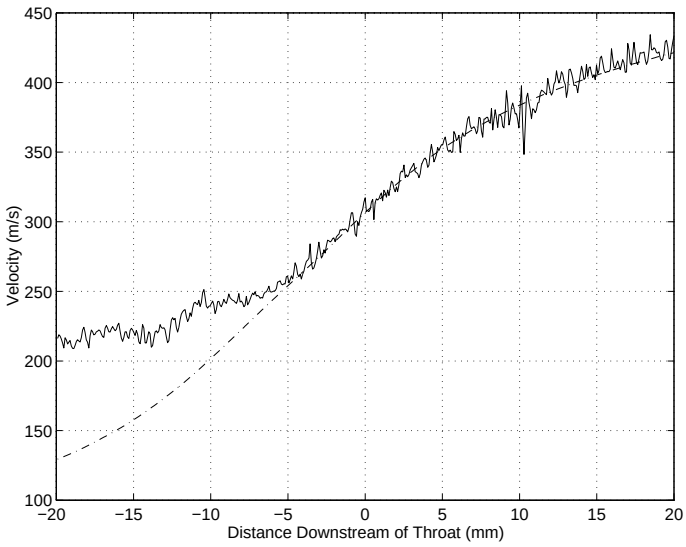


Figure 8.13: Centre-line velocity. The Doppler global velocimetry measurements taken as a mean of twenty pixel rows each side of the centre-line are shown as a solid line. The computational fluid dynamics predictions, taken as a mean across the same region, are shown as a dotted line.

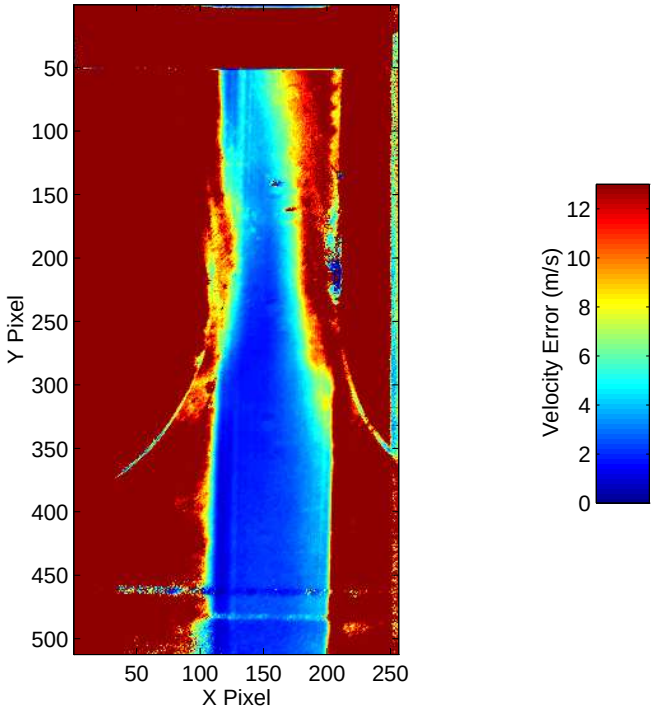


Figure 8.14: Predicted error caused by imaging system shot noise, using the analysis of Chapter 5 (shown as a colour-map representation).

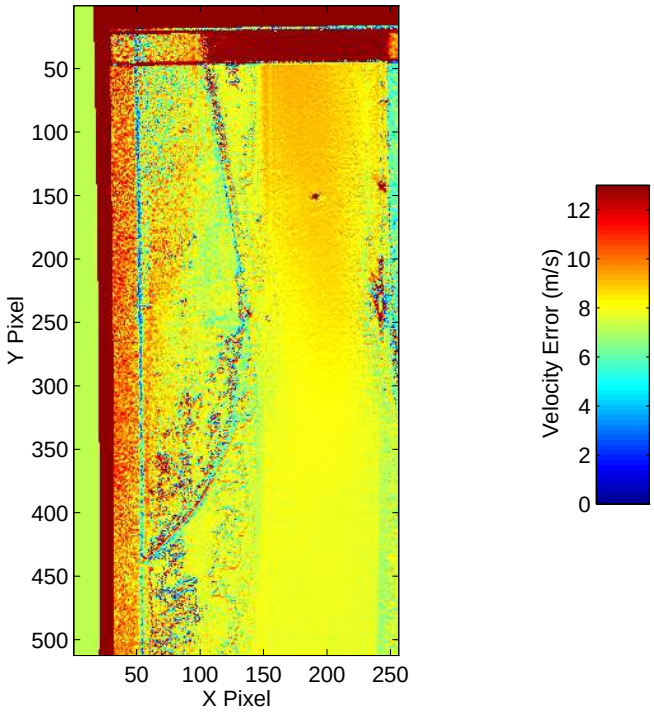


Figure 8.15: Predicted error caused by iodine cell calibration, using the analysis of Chapter 5 (shown as a colour-map representation).

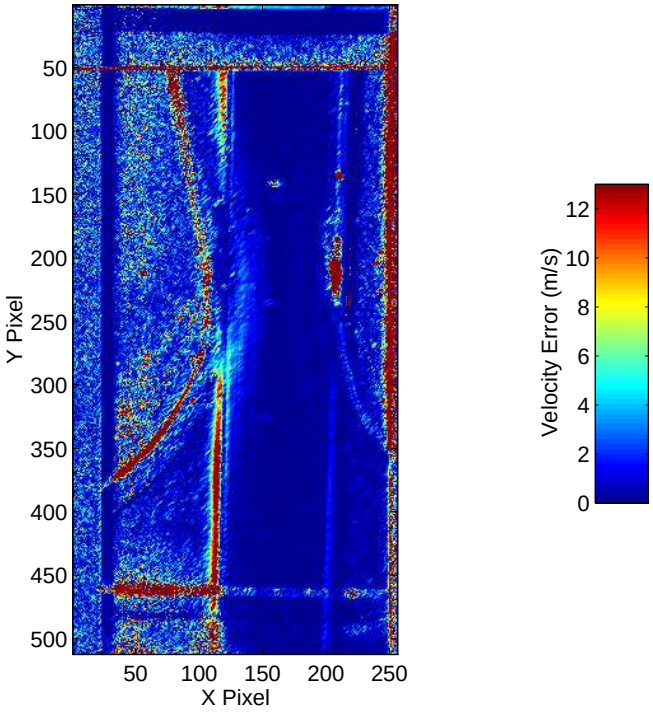


Figure 8.16: Predicted error caused by misalignment, using the analysis of Chapter 5 (shown as a colour-map representation).

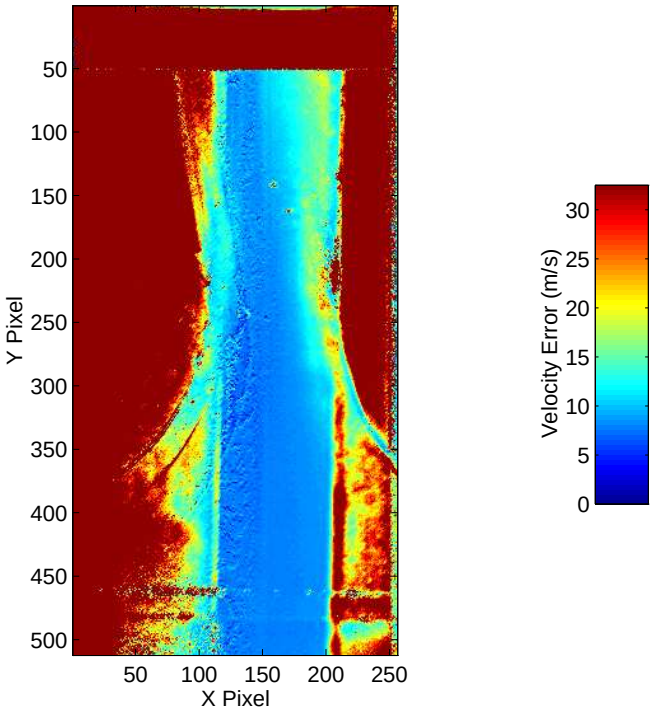


Figure 8.17: Predicted RMS error casued by shot noise, iodine cell calibration and misalignment, using the analysis of Chapter 5 (shown as a colour-map representation).

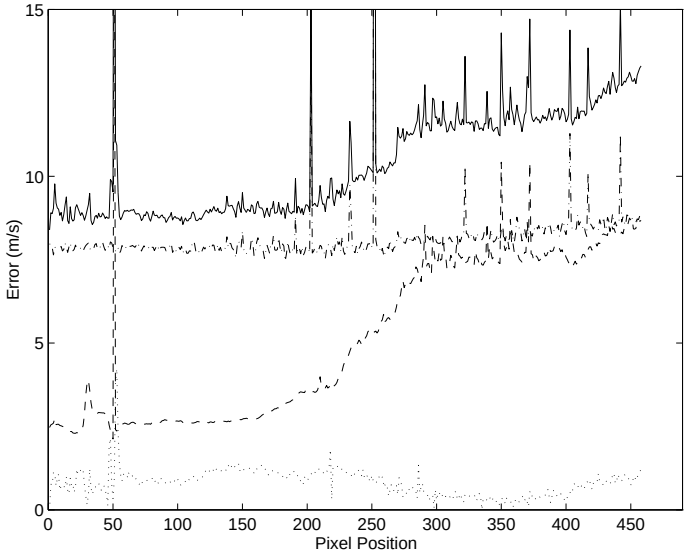


Figure 8.18: Predicted measurement error for the convergent-divergent nozzle along the centre-line of the flow, using the analysis of Chapter 5. The dotted line represents the predicted error due to misalignment, the dashed line the predicted error due to shot noise, the dot-dashed line the predicted error due to the iodine cell, and the solid line the overall RMS error.

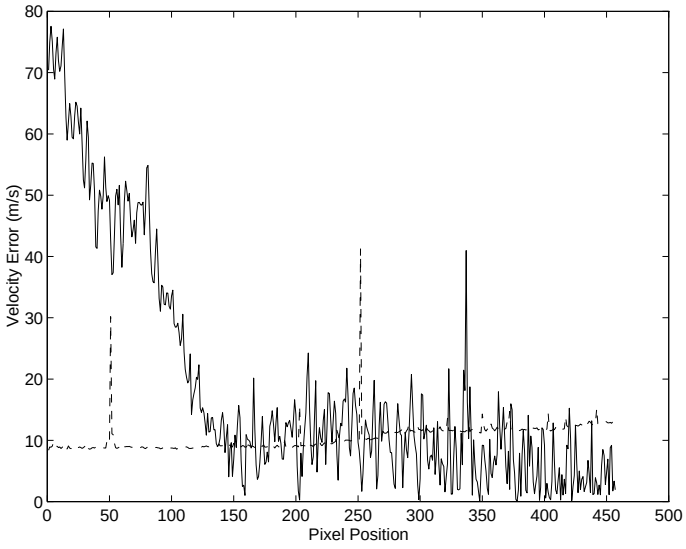


Figure 8.19: Comparison of predicted error with observed error along the centre-line of the flow. The dashed line is the predicted RMS error and the solid line the observed measurement error.

8.4 Error Analysis

In Chapter 7 a detailed analysis of the practical performance of the different warping and interpolation algorithms as applied to the measurement of a free jet was undertaken. This analysis required the fitting of the measured data to an established model, so that the difference between the theoretical model and the measured data could be interpreted as measurement error. The expected measurement error was calculated according to the analysis of Chapter 5. In this way, predicted and observed error values were compared.

The discussions of Chapter 5, reinforced by the practical findings of Chapter 7 suggest that the preferred image processing approach is to employ a polynomial warp (using orthogonal polynomials) together with bilinear interpolation. Accordingly, such a method has been applied to the determination of the convergent-divergent nozzle flow field.

It is not worthwhile to devote any further space to a consideration of the relative merits of the different warping and interpolation algorithms. It is, however, interesting to apply once again the error analysis of Chapter 5 in order to compare the predictions thus made with the observed error. In the case of the convergent-divergent nozzle, observed error is taken to be the difference between the measured flow field and the computational fluid dynamics prediction.

Once again assuming orthogonal polynomial warping, bilinear interpolation, and a 0.1 pixel RMS error in centroid estimation, a number of error predictions may be made across the entire field of measurement.

Fig. 8.14 shows the predicted error due to shot noise across the entire frame, as a colour-map representation. It may be seen that throughout the bulk of the flow field, the estimated error due to CCD array shot noise is low, at around 2 ms^{-1}. However, error due to shot noise becomes significantly higher (around 12 ms^{-1}) in the region close to the side wall of the nozzle itself. It has been seen that, in order to avoid excessive glare from the side walls, as well as to minimise the effect of adverse light reflections, the extent of the light sheet was restricted as far as possible. This results in low intensity close to the side walls, and hence an increase in the level of error due to CCD array shot noise.

The effect of iodine cell calibration error is shown in Fig. 8.15. As was the case with the free jet and the disc measurements, there is little variation in the level of iodine cell calibration error across the frame. The overall level of error is higher, at around 8 ms^{-1}, due to the higher flow velocities (and consequent greater Doppler shifts) present in this case. It is

notable that downstream of the throat, where flow velocities are higher, the predicted iodine cell calibration error increases to around 10 ms^{-1}.

The final component of error which may be considered is that due to misalignment. This prediction is presented in Fig. 8.16. Misalignment error seems to be less of a problem for the convergent-divergent nozzle flow than was the case for the free jet flow. The reason for this is the lack of high velocity or intensity gradients within the region to be measured. An overall level of predicted misalignment error of around 2 ms^{-1} is observed throughout the flow, increasing sharply at the very edges of the region of interest due to the high intensity gradient at the edge of the laser light sheet.

Fig. 8.17 shows the combined predicted error due to shot noise, iodine calibration, and misalignment. It may be seen that within the flow region itself, the overall level of predicted error is relatively constant: around 9 ms^{-1} upstream of the throat and 13 ms^{-1} downstream of the throat. Only close to the side walls is the predicted error seen to increase sharply as a result of the presence of increased shot noise in less well illuminated regions.

For the purposes of error comparison, a line passing along the centre of the nozzle has been examined. The predicted error along this line, together with the components that contribute to it, is illustrated in Fig. 8.18: the dotted line represents the predicted error due to misalignment, the dashed line the predicted error due to shot noise, the dot-dashed line the predicted error due to the iodine cell, and the solid line the overall RMS error. All the quantities are shown as RMS values, with the exception of the misalignment error which is shown as a signed quantity in order to allow the sense of the misalignment error to be readily seen.

Finally, Fig. 8.19 allows a ready comparison to be made, along the centreline of the nozzle, between the observed error (difference between computational fluid dynamics predictions and the measured values) and the predicted error. It may be seen that upstream of the throat the computational model and the measured results diverge by as much as 75 ms^{-1}, whereas the error is typically less that 10 ms^{-1} downstream of the throat. The error predictions of Chapter 5 are unable to explain the divergence upstream of the throat. As explained above, it is believed that this error arises due to separation at the mouth of the convergent-divergent nozzle. This effect is not modelled in the computational model. Downstream of the throat, however, the predicted error values are in good agreement with those observed from the actual measurement. Indeed, at the extreme downstream end it may be said that the measured values exhibit somewhat less error than that predicted theoretically.

8.5 Summary

This experiment is the first attempt at applying Doppler global velocimetry to an internal transonic flow field. The problem of adverse reflections has been partially overcome by restricting the extent of the illuminating light sheet so as to avoid excessive impingment on the walls of the nozzle itself. 0.5 μm diameter latex spheres were selected as a suitable seeding medium.

A computational fluid dynamics model of the flow geometry has been made and suitable codes run in order to yield a predicted flow field. This prediction has been compared with the measured Doppler global results in order to obtain an estimate of the measurement error. This estimated measurement error has been compared with that error which may be predicted using the techniques explained in Chapter 5.

It has been found that at the higher flow velocities observed in the nozzle, iodine calibration error becomes a significant source of measurement inaccuracy. In this case the precise alignment of the discriminated with the non-discriminated images is somewhat less crucial than was the case with the free jet experiment described in Chapter 7. In fact, both the error due to the presence of CCD array shot noise and the error due to iodine cell calibration dominate over that due to misalignment in this case. The major divergence between the predicted and the measured flow fields is in the region upstream of the throat, where the measured flow is significantly higher in velocity than that predicted by computational methods. This is attributed to separation occuring at the mouth of the nozzle itself. Downstream of the throat, however, the computational model fits the measured data to a degree of accuracy which is consistent with the error estimates of Chapter 5.

The Doppler global system has proved itself capable of measuring a short-duration transonic flow field in an enclosed channel with an accuracy better than 10 ms^{-1}. Encouraging though these results are, there is scope for significant further Doppler global velocimetry work within the convergent-divergent nozzle apparatus. Two problems in particular require further work. The first is that of low image intensity, which may be conveniently remedied by the use of a more powerful Nd:YAG pulse laser. The second problem — that of scatter from the side walls of the nozzle — is more difficult to overcome. Further work should be undertaken to investigate the effects of translucent nozzle liners, fluorescent paint and the application of an additional iodine cell maintained at an elevated temperature in order to reduce the problem of stray scatter. The measurement of transonic flows in enclosed channels is an important application for Doppler global velocimetry, and such a capability will be required if measurements are to be made of turbomachinery flow fields.

Chapter 9

Summary and Conclusions

9.1 Summary

Doppler global velocimetry is a technique for acquiring instantaneously a measurement of fluid velocity across an entire plane. Whereas previous workers have concentrated on the development of real-time visualisation systems, the work presented in this thesis is directed towards the development of a single-shot instrument. The relaxation of the requirement for real-time velocity data visualisation allows the use of image processing software to ensure the accurate alignment of the discriminated and non-discriminated images, and also to take account of the profile of the iodine absorption curves. Such image processing operations have been seen to be a necessary prerequisite if accurate quantitative velocity measurements are to be obtained.

The fundamental technologies required in order to implement the Doppler global technique are presented in Chapter 2. It was seen that the development of a successful velocimeter was not possible without an understanding of laser physics (to ensure a stable and controllable emission frequency), absorption spectroscopy (for the construction of suitable iodine vapour cells), image acquisition and processing (to allow accurate alignment of images) and fluid dynamics (for successful flow-field seeding).

The image processing technique which may be employed in order to align the discriminated and non-discriminated images was considered in Chapter 3. The three stages of image processing — centroid estimation, image transformation and interpolation — were considered from a theoretical standpoint. The manner in which these techniques are used as part of the overall image processing methodology implemented in the `warp` software was also explained.

Further image processing techniques, required in order to take account of the non-linear profile of the iodine cell, were introduced in Chapter 4. The procedure for calibrating the working and reference iodine cells was presented and the operation of the `viewresults` quantitative visualisation software was discussed. The codes `2component` and `2component_vis` are required in order to compute and visualise two component flow data. These programs and the associated measurement procedure were discussed in detail in the same chapter.

Since the objective of the current programme of research was to develop a Doppler global velocimeter capable of accurate quantitative velocity measurements, it is of key importance

to consider each of the sources of error in the system. In Chapter 5, each error source, including the effects of camera pixel shot noise, iodine cell transmission uncertainty, and image misalignment, was considered in detail. The individual error sources were modelled mathematically, allowing a combined analytical error estimate to be developed for the entire system. This estimate was implemented as a computer code, allowing each component of error to be estimated throughout the measured image. Such error estimates are useful to enable bounds to be placed upon measured data, and also to help in the consideration of where the system may be improved to reduce overall error.

The next three chapters, 6, 7, and 8, present practical measured results that have been obtained using the Oxford system. The measurement of the velocity of a rotating disc was employed as a simple test case to allow the operation of the system to be tested. The image processing stages required in order to take raw CCD images and produce velocity data were illustrated step by step. A satisfactory image of this low velocity data field was obtained, both as a single component and as a two component measurement. Application of the error model developed in Chapter 5 revealed that the observed error fitted the expected error closely, having a peak RMS value of around 3 ms^{-1}.

The measurements of a subsonic free jet proved to be a more demanding test of the velocimeter. The presence of large intensity and velocity gradients within the flow field meant that the measured velocity was acutely sensitive to image alignment. This data was therefore used, in conjunction with the error model, to test the practical effect of the different image transformation and interpolation algorithms. The use of an established model of the flow field allowed an estimated measurement error to be calculated at each point in the flow, which may be compared with the predicted error obtained from the model of Chapter 5. It was found that the orthogonal polynomials form of the polynomial warp was the prefered image transformation method, and the bilinear technique was favoured as an interpolation scheme. For a peak velocity of around 100 ms^{-1}, RMS error values were found to be around 3.5 ms^{-1}, and the velocity measurements obtained were judged to be satisfactory.

The final measurements to be presented, in Chapter 8, were those of a transonic enclosed flow. It is proposed that the velocimeter shall be eventually applied to the measurement of turbomachinery flow fields so the ability to measure transonic flow fields in enclosed channels is important. The convergent-divergent nozzle flow was found to be difficult to measure due to the problem of adverse reflections of the laser light sheet from the nozzle itself, low image intensity and limited iodine cell bandwidth. Future work with such a nozzle, employing broadened iodine cells and a more powerful Nd:YAG laser should help to eliminate some of

these difficulties. Despite these problems, measurements of the transonic flow were obtained. Comparison with computational predictions showed that downstream of the nozzle throat, the measured data fitted that expected closely. Within this region, the RMS error was again found to be in good agreement with the model of Chapter 5 at around 10 ms^{-1}. It was seen that the Doppler global velocimetry system was capable of accurate measurements of velocities of up to around 425 ms^{-1} in an enclosed channel.

The key points of this thesis may be thus summarised:

- Doppler global velocimetry is a technique for acquiring instantaneously a measurement of fluid velocity across an entire plane. Although it may be considered to be a developing technology at this time, it has already proved to be a useful measurement tool, offering significant advantages over other whole field techniques, such as particle image velocimetry.

- In order to obtain accurate quantitative measurements, alignment of the discriminated and non-discriminated images must be performed to a high degree of accuracy. The requirement for such alignment precludes the development of a real time instrument of acceptable accuracy. In the current work a single-shot instrument has been developed and an extensive suite of image processing software has been written in order to effect alignment.

- The alignment operation may be considered to be composed of three independent tasks — centroid estimation, warping and interpolation. A number of candidate image processing algorithms have been evaluated both theoretically and experimentally.

- The application of analytical and numerical error analysis suggests that the preferred algorithms for Doppler global velocimetry image processing are the orthogonal polynomials warp method combined with bilinear interpolation.

- For the case of a free jet flow, comparison of measured velocimetry data with a theoretical model allows the observed measurement error to be compared with that predicted. The agreement between the error model and the observed error is good. The measured results thus add further weight to the conclusion that the orthogonal polynomials warp method combined with bilinear interpolation is the preferred method.

- The Doppler global velocimetry technique has demonstrated great versatility — both low velocity fields (the rotating disc) and high velocity fields in a confined passage (the convergent divergent nozzle) have been successfully measured.

- Doppler global velocimetry is capable of measurements of two or three velocity components. Preliminary work on two component measurements applied to the case of the rotating disc has been encouraging.

- The Oxford Doppler global velocimetry system has been seen to be capable of accurate whole field velocity measurement for a range of applications. The expected error resulting from all components within the velocimetry system has been assessed and quantified. Such systematic error analysis is necessary in order to optimise the accuracy of the system. In addition the error analysis code allows error bounds to be assigned to measured velocity data.

9.2 Suggestions for Future Work

During the course of the programme of research described in this thesis, a number of difficulties arose which suggest that further work may be beneficial in terms of expanding the range of applications to which the Doppler global velocimetry technique may be applied. In particular, the measurement of the transonic flow within the confined passage of the convergent-divergent nozzle proved to be quite difficult to obtain due to reflections within the channel, limited iodine cell bandwidth, and low scattered intensity.

Since enclosed transonic flows are of great interest in turbomachinery applications, it is worthwhile to undertake further work with this nozzle configuration. Reflections within the channel present two independent problems. First, the high intensity of such reflections causes saturation of the CCD array leading to a loss of data over a large part of the frame. Second, the reflected rays may lead to error in the determined velocity, as explained in Section 5.5. Both of these effects may be minimised by reducing the overall level of scatter from the channel itself. The intensity of the laser illumination is such that using a matt black finish for the nozzle liners proved to be insufficient to eliminate the effects of scatter from the side walls. The use of polarising filters to remove the light scattered from the walls is not to be recommended, since the polarisation state of the radiation scattered from the seeding particles is itself unknown. The use of fluorescent paints, capable of absorbing the laser radiation and re-emitting it at a different frequency should be investigated. Alternatively, the use of translucent liners which are able to absorb the laser radiation and internally dissipate the energy may prove to be a possible solution to this problem. A third iodine cell, maintained at an elevated temperature, may be placed in front of the entire optical system in order to prevent any unshifted laser emission from entering the receiving optics — this will prevent the problem of CCD saturation, but not that of erroneous velocity measurement

due to reflected rays.

It was seen in Chapter 8 that the receiving optical system had to be positioned at an angle away from the normal in order that the Doppler shifted radiation was not such as to exceed the allowable bandwidth of the working iodine cell. For such high velocity applications, it would be desirable to have a working iodine cell of greater bandwidth than that usually employed. Such a cell may be produced by introducing small quantities of an impurity into the iodine contained within the cell. The use of such a broadened cell would eliminate the requirement for the extreme viewing position employed and allow a greater intensity of scattered radiation to be imaged onto the CCD array.

In order to further improve the intensity recorded on the CCD array, it is attractive to use a more powerful laser. In particular, the Nd:YAG laser has been used for Doppler global velocimetry applications in the past. The high intensity of such a laser, compared with the argon ion device used for the present work means that transient transonic flows may be more reliably imaged. The Nd:YAG laser is a pulse rather than a continuous wave device so an instantaneous rather than a time-averaged view of the flow may be acquired.

In this thesis a number of stategies for the acquision of two component measurements have been outlined, and two component measurements of the rotating disc have been presented. It would be desirable to construct a second Doppler global velocimetry optical system to enable further two component work to be undertaken. A third system could also be constructed, and suitable software written, to allow all three components of velocity to be simultaneously measured within the plane of the laser light sheet.

The eventual objective of this work is to obtain an instrument suited to the measurement of turbomachinery flow fields. A key aim in any future work must therefore be the integration of the instrument into an existing experimental turbomachinery experiment such as the Rotor Facility at Oxford.

9.3 Conclusions

A comprehensive image processing strategy for application to Doppler global velocimetry has been devised and implemented. A number of different techniques for image transformation and interpolation have been examined and applied to the task of aligning the discriminated and non-discriminated images. The error resulting from both physical and the image processing aspects of the Doppler global velocimetry technique has been examined and quantified, allowing an analytical model of the expected error to be developed.

The Oxford Doppler global velocimetry system has been applied to the measurement of three test cases — the rotating disc, the free jet flow, and the transonic flow within an enclosed convergent-divergent nozzle. The practical data so-obtained has been compared with existing accepted models or with computational predictions, allowing the observed error to be determined throughout the measurement plane. This observed error was compared to the predicted error obtained from the analytical model developed. In every case, the error predicted was found to closely fit that observed, demonstrating that the Doppler global velocimetry system was performing as well as could be expected in the presence of certain fundamental error sources.

In every case the Doppler global velocimetry system was seen to be capable of accurate determination of the flow field under consideration. The different image processing techniques implemented were compared and assessed. The assertion that accurate quantitative Doppler velocimetry measurements could not be obtained without the use of extensive image-processing operations was confirmed. The software implemented in the current work represents the most comprehensive attempt to undertake data processing for Doppler global velocimetry applications. Without the use of such software, the measurement accuracy achieved in the current work would not have been realised.

References

[Abramovich 1963]

Abramovich G.N.

The Theory of Turbulent Jets,

MIT Press.

[Ainsworth and Thorpe 1994]

Ainsworth R.W., & Thorpe S.J., 1994

The Development of a Doppler Global Velocimeter for Transonic Turbine Applications,

Presented at the International Gas Turbine and Aeroengine Congress and Exposition, The Hague, Netherlands, June 13-16. ASME Paper Number ASME 94-GT-146.

[Ainsworth, Thorpe and Manners 1996]

Ainsworth R.W., Thorpe S.J., & Manners R.J., 1996

A New Approach to Flow-Field Measurement — A Review of Doppler Global Velocimetry Techniques,

Invited Paper presented at the International Conference on Turbulent Heat Transfer, San Diego, U.S.A., organised by the Engineering Foundation, New York, U.S.A.

[Ainsworth, Thorpe and Manners 1997]

Ainsworth R.W., Thorpe S.J., & Manners R.J., 1997

A New Approach to Flow-Field Measurement — A View of Doppler Global Velocimetry Techniques,

International Journal of Heat and Fluid Flow, Vol. 18, pp. 116-130.

[Albertson *et al.* 1950]

Albertson M.L.,Dai Y.B., Jenson R.A., & Rouse H., 1950

Diffusion of Submerged Jets,

Transactions of the American Society of Chemical Engineers, Vol. 115, pp. 639-97.

[Alexander and Ng 1991]

Alexander B.F & Ng K.C, 1991

Elimination of Systematic Error in Subpixel Accuracy Centroid Estimation,

Optical Engineering, Vol. 30, No. 9.

[Astromed 1992]

Astromed Ltd., Innovation Centre, Cambridge Science Park, Milton Road, Cambridge, UK.

The Role of Charge Coupled Devices in Low Light Imaging.

[Bryanston-Cross and Epstein]

Bryanston-Cross P.J. & Epstein A., 1990

The Application of Sub-micron Particle Visualisation for PIV at Transonic Speeds,

Progress in Aerospace Science, Vol. 27, pp. 237-65.

[Chan *et al.* 1995]

Chan V.S.S., Heyes A.L., Robinson D.I., & Turner J.T., 1995

Iodine Absorption Filters for Doppler Global Velocimetry,

Measurement Science and Technology, Vol. 6., pp. 784-94.

[Chartier, Robertsson and Fredin-Picard 1991]

Chartier J.M., Robertsson L., & Fredin-Picard S., 1991

Recent Activities at B.I.P.M. in the Field of Stabilised Lasers — Radiations Recommended for the Definition of the Meter,

I.E.E.E. Transactions on Instrumentation and Measurement, Vol. 40, No. 2, pp. 181-4.

[Coleman and Steele 1995]

Coleman H.W., & Steele W.G., 1995

Engineering Application of Experimental Uncertainty Analysis,

AIAA Journal, Vol. 33, No. 10, pp. 1888-96.

[DANTEC 1984]

Dantec Electronic Ltd.

55L10 Seeding Generator, Instruction and Service Manual

[Echols and Young 1963]

Echols W.H., & Young J.A., 1963

Studies of Portable Air-Operated Aerosol Generators,

U.S. Naval Research Laboratory report Number 5929.

[EEV 1987]

The English Electric Valve Company Ltd., Waterhouse Lane, Chelmsford, UK.

CCD Imaging III

[Foley *et al.* 1990]

Foley J., van Dam A., Feiner S. & Hughes J., 1990

Computer Graphics,

Addison-Wesley.

[Ford and Tatum 1994]

Ford H.D. & Tatum R.P., 1994

Real-time Doppler Global Velocimetry,

Proceedings of Engineering Applications of Optical Diagnostic Techniques, Cranfield University, Bedford, UK..

[Fredin-Picard 1989]

Fredin-Picard S., 1989

A Study of Contamination in $^{127}I_2$ *Cells Using Laser Induced Flourescence*,

Metrologia, Vol. 26, No. 4, pp. 235-44.

[Gauntner, Lingwood and Hrycak 1970]

Gauntner J.W., Lingwood J.N.B., & Hrycak P., 1970

Survey of Literature on Flow Characteristics of a Single Jet Impinging on a Flat Plate,

NASA Technical Note 5652.

[Gerstenkorn and Luc 1978]

Gerstenkorn S. & Luc P., 1978

Atlas du Spectra d'Absorption de la Molecule de l'Iodine 4800-20000 cm^{-1},

C.N.R.S. Paris.

[Gill and Thomson 1986]

Gill P., and Thompson R.C., 1986

The Preparation and Analysis of Iodine Cells,

Metrologia, Vol. 23, pp. 161-66.

[Invent GmbH 1991]

Invent GmbH, AmWeichselgarten 21, D-8520 Erlangen-Tennenlohe, Germany.

STREU: A Computational Code for the Light Scattering Properties of Spherical Particles.

[Jenkins and White 1976]

Jenkins F.A., & White H.E., 1976

Fundamentals of Optics, 4th Edition,

McGraw-Hill, New York, U.S.A.

[Komine and Brosnan 1991]

Komine H., & Brosnan S.J., 1991

Instantaneous, Three-Component, Doppler Global Velocimetry,

Presented at the 4th International Conference on Laser Anemometry, Advances and Applications.

[Komine *et al.* 1991]

Komine H., Brosnan S.J., Litton A.B., & Stappaerts E.A., 1991

Real-Time Doppler Global Velocimetry,

AIAA Paper number AIAA-91-0337.

[Komine *et al.* 1994]

Komine H., Brosnan S.J., Long W.H.,& Stappaerts E.A., 1994

Development of a Flight Research Intrumentation System for Application to Non-intrusive Measurements of the Flow Field,

NASA Contractor Report 191490.

[Laskin 1948]

Laskin S., 1948

Submerged Aerosol Unit,

AEC Project Quarterly Report UR-38, University of Rochester.

[Manners, Thorpe and Ainsworth 1996]

Manners R.J., Thorpe S.J., & Ainsworth R.W., 1996

Image Processing Techniques for Doppler Global Velocimetry,

Proceedings of Optical Methods and Data Processing in Heat and Fluid Flow, City University, London, U.K.

[Massey 1989]

Massey B.S., 1989

Mechanics of Fluids,

Van Nostrand Reinhold.

[Mathworks 1993]

The Mathworks Inc., Cochituate Place, Prime Park Way, Natick, Mass., U.S.A.

Matlab High-Performance Numeric Computation and Visualisation Software User's Guide.

[McKenzie 1996]
McKenzie R.L., 1996
Measurement Capabilities of Planar Doppler Velocimetry Using Pulsed Lasers,
Applied Optics, 35, 948-964.

[Melling and Whitelaw 1973]
Melling A. & Whitelaw J.H., 1973
Seeding of Gas Flows for Laser Anemometry,
DISA Information, Dantec Electronic Ltd.

[Melling 1986]
Melling A., 1986
Seeding Gas Flows for Laser Anemometry, AGARD Conference Proceedings No. 399.

[Metrologia 1984]
Documents Concerning the New Definition of the Meter,
Metrologia, Vol. 19, No. 4, pp. 163-77.

[Meyers and Komine 1991]
Meyers J.F. & Komine H., 1991
Doppler Global Velocimetry: A New Way to Look at Velocity,
ASME 4th International Conference on Laser Anemometry, Advances and Applications.

[Meyers, Lee and Cavone 1991]
Meyers J.F., Lee J.W & Cavone A.A., 1991
Signal Processing Schemes for Doppler Global Velocimetry,
IEEE 14th International Congress on Instrumentation for Aerospace Simulation Facilities, Rockville, MD, pp. 321-328.

[Meyers 1992]
Meyers J.F., 1992
Doppler Global Velocimetry — The Next Generation,
AIAA Paper number AIAA 92-3897.

[Meyers 1995]
Meyers J.F., 1995
Development of Doppler Global Velocimetry as a Flow Diagnostics Tool,
Measurement Science and Technology, Vol. 6, pp. 769-83.

[Mie 1908]
Mie G., 1908
Optics of Turbid Media,
Ann. Phys., Vol. 25, pp. 377-445.

[Morrison, Gaharan and DeOtte 1994]
Morrison G.L., Gaharan A.G., & DeOtte R.E., 1994
Doppler Global Velocimetry: Problems and Pitfalls,
A.S.M.E., Fluids Engineering Division, Vol. 191, Proceedings of Laser Anemometry Advances and Applications.

[Nye 1990]
Nye A., 1990
Xlib Programming Manual,
O'Reilly & Associates, Inc.

[Pai 1954]
Pai S., 1954
Fluid Dynamics of Jets,
D. Van Nostrand.

[Pratt 1991]
Pratt W.K., 1991
Digital Image Processing. 2nd edition.
Wiley Interscience.

[Press *et al.* 1992]
Press W., Teukolsky S., Vetterling W. & Flannery B., 1992
Numerical Recipes in C,
Cambridge University Press.

[Reiss and Radin 1988]
Reiss L. & Radin J., 1988
X Window Inside and Out,
McGraw Hill.

[Roehle and Schodl 1994]
Roehle I., & Schodl R., 1994
Evaluation of the Accuracy of the Doppler Global Velocimetry Technique,
Proceedings of Optical Methods and Data Processing in Heat and Fluid Flow, pp. 155-61.

[Schlichting 1979]

Schlichting H., 1979

Boundary Layer Theory,

McGraw-Hill, London.

[Sheldrake and Ainsworth 1995]

Sheldrake C.D., & Ainsworth R.W.A., 1995

The Use of Hot Wires Applied to Aerodynamic Measurements in a Model Turbine Stage,

First European Conference on Turbomachinery Fluid Dynamic and Thermodynamic Aspects, pp. 149-73, Nuremburg, Germany.

[Spectra-Physics 1987]

Spectra-Physics Ltd., Boundary Way, Hemel Hempstead, Hertfordshire, UK.

Literature relating to single frequency model 165 argon ion laser products.

[Spectra-Physics 1995]

Spectra-Physics Ltd., Boundary Way, Hemel Hempstead, Hertfordshire, UK.

1995 Scientific Product Catalogue.

[Stroustrup 1991]

Stroustrup B., 1991

The C++ Programming Language,

Addison Wesley.

[Tani and Komatsu 1966]

Tani I. & Komatsu Y., 1966

Impingment of a Round Jet on a Flat Surface,

Proceedings of Eleventh International Congress of Applied Mathematics. Springer-Verlag, pp.672-6.

[Thorpe and Ainsworth 1994]

Thorpe S.J., Ainsworth R.W., 1994

Doppler Global Velocimetry — Application to a Round Free Jet,

Proceedings of Engineering Applications of Optical Diagnostic Techniques, Cranfield University, Bedford, UK.

[Thorpe, Ainsworth and Manners 1995]

Thorpe S.J., Ainsworth R.W., & Manners R.J., 1995

The Development of a Doppler Global Velocimeter and its Application to a Free Jet Flow,

Proceedings of the ASME/EALA 6th International Conference on Laser Anemometry, Advances and Applications, Hilton Head Island, South Carolina, U.S.A.

[Thorpe, Ainsworth and Manners 1996]

Thorpe S.J., Ainsworth R.W., & Manners R.J., 1996

Time Averaged Free Jet Measurements using Doppler Global Velocimetry,

Proceedings of the Symposium on Experimental and Numerical Flow Visualisation and Laser Anemometry, San Diego, California, U.S.A.

[Thorpe 1996]

Thorpe S.J., 1996

D.Phil Thesis,

University of Oxford, Department of Engineering Science.

[Quinlan *et al.* 1997]

Quinlan N.J., Ainsworth R.W., Bellhouse B.J., Manners R.J., Thorpe S.J., 1997

Application of Doppler Global Velocimetry to Supersonic Gas-Particle Flows in Drug Delivery,

Proceedings of Laser Anemometry Advances and Applications.

[Usry, Meyers and Miller 1992 **(1)**]

Usry J.W., Meyers J.F., Miller L.S., 1992

Doppler Global Velocimeter Measurements of the Vortical Flow Above a Thin Delta Wing,

AIAA 92-0005.

[Usry, Meyers and Miller 1992 **(2)**]

Usry J.W., Meyers J.F., Miller L.S., 1992

Assessing the Capability of Doppler Global Velocimetry to Measure Vortical Flow Fields,

IMechE Optical Methods and Data Processing in Heat and Fluid Flow, 1992, pp. 103-114.

[Yariv 1985]

Yariv A., 1995,

Introduction to Optical Electronics. Third Edition,

Holt, Rinehart and Winston.

[Watt 1989]

Watt A., 1989

Fundamentals of Three-Dimensional Computer Graphics,

Addison-Wesley.

[Wolberg 1991]

Wolberg G., 1991

Digital Image Warping,

IEEE Computer Society Press.

[Wright 1993]

Wright Instruments Ltd., Unit 10, 26 Queensway, Enfield, Middlesex, UK.

TK-512 CCD Camera Specification

[Yea and Cummins 1964]

Yea Y., & Cummins H.Z., 1964

Localised Fluid Flow Measurements with a He-Ne Laser Spectrometer,

Applied Physics Letters, Vol. 4, pp. 176-8.