Three-Dimensional Particle Image Velocimetry

Olalekan Owodunni
Department of Engineering
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

September 2006

Fitzwilliam College

Declaration

This dissertation is the result of my own work and includes nothing which is the outcome of work done in collaboration except where specifically indicated in the text. No part of this dissertation has already been, or is being currently reviewed for any other degree, diploma or qualification. This Dissertation contains approximately 30,000 words and 77 Figures.

……………………………………………………

Olalekan Owodunni
September 2006

Keywords

- Flow measurement
- Fluid velocity
- Particle Image Velocimetry
- Particle Tracking Velocimetry
- Correlation
- Resolution
- Accuracy
- Valid data yield
- Measurement Speed

Abstract

Digital Particle Image Velocimetry (DPIV) is now a standard technique for global velocity measurement of fluid flows.

In most applications, the velocity information is obtained over a plane in the flow field and the major challenge in PIV measurements today is the efficient implementation of PIV to give all three components of velocity at many points in a volume of interest.

The PIV system arrangements and analysis techniques that have been developed to date to give three-component velocity information throughout a given volume are severely limited in important aspects of their performance. Their measurement speed, accuracy, spatial and temporal resolution are low. Also associated with these arrangements are complex measurement procedures and high equipment cost which have further limited their widespread use.

This dissertation describes the development of a relatively simple high-performance and cost-effective system which uses a single 3CCD colour camera and multi-coloured illumination to identify and track the three-dimensional locations of the particles suspended in a flow. By comparing particle positions from two flow images taken a short time apart using a hybrid cross-correlation/particle tracking technique, all three velocity components of each particle within the volume of interest were determined.

This novel system was applied to measure laminar vortex ring flow.

The measurements obtained were in good agreement with expected values, thus, confirming the potential of the system to give accurate, high resolution, 3-component velocity measurements in a volume within a practical time frame.

Acknowledgements

I wish to begin by thanking Dr. Tim Nickels, for giving me the opportunity to work on this exciting project. It has been a privilege to study under his supervision and I thank him for his guidance, generosity and support throughout the course of the project.

My thanks and appreciation also go out to Dr. Holger Babinsky and Dr. Stuart Dalziel for their interest and suggestions which have resulted in a better research outcome.

I also wish to thank all the CUED technicians who have been involved in this project. I note in particular, Michael Underwood, Alistair Ross and Dave Pittock who ensured a prompt and skillful construction of some of the equipment that was used in this project. Many thanks also to the Librarians for all their assistance.

Contents

4 The 3D colour PIV-PTV system: Image Analysis 113

List of Figures

List of Tables

Nomenclature

Roman Symbols

I	pixel/voxel intensity
L	interrogation window/block size along the x-direction
M	interrogation window/block size along the y-direction
N	interrogation block size along the z-direction

Greek Symbols

α	orientation angle of the calibration plate with the incident light beam
Δ	maximum correlation search distance
γ	3D grid compression ratio
μ	dynamic viscosity of working fluid
ω	vorticity
ρ	density of the working fluid

Subscripts

i	pixel/voxel x-coordinates
j	pixel/voxel y-coordinates
k	voxel z-coordinates

x	x-direction
y	y-direction
z	z-direction

Other Symbols

$\bar{U}_p$	average piston velocity
$\Delta\vec{s}$	particle image displacement vector
δ_y	calibration point grid spacing along the y-direction
δ_z	calibration point grid spacing along the z-direction
Δt	time between successive image exposures
Δ_i	correlation pixel/voxel search distance along the x-direction
Δ_j	correlation pixel/voxel search distance along the y-direction
Δ_k	correlation voxel search distance along the z-direction
$\vec{V}$	particle(local fluid)velocity
D_o	orifice Diameter of the vortex ring cylinder
L_p	piston stroke length
M_i	image magnification
N_x	CCD array size (measured in pixels) along the x-direction
N_y	CCD array size (measured in pixels) along the y-direction
Re	Reynold's Number
S_x	number of grid voxels along the x-direction
S_y	number of grid voxels along the y-direction
x'	x-image coordinate of a point

x_p	x-coordinate of a point in real-world space
y'	y-image coordinate of a point
y_p	y-coordinate of a point in real-world space
z_p	z-coordinate of a point in real-world space

Acronyms

CCD	Charge-Coupled Device
$DCCC$	Direct Cross-Correlation of a Compressed grid
$DPIV$	Digital Particle Image Velocimetry
FFT	Fast Fourier Transform
FOV	Field Of View
PIV	Particle Image Velocimetry
PTV	Particle Tracking Velocimetry
ROI	Region Of Interest

Chapter 1

Particle Image Velocimetry: Advances in performance

1.1 Introduction

In experimental fluid mechanics, the need to acquire reliable qualitative and quantitative information about the behaviour of flows has led to the development of many types of flow measurement and visualization techniques over the years. An understanding of the flow around and within engineering devices is of critical importance in improving the design of these devices and can lead to substantial improvements in performance and efficiency. The measurement of the characteristic flow property — its velocity — is of great importance in helping to further improve the understanding of the structure and behaviour of many complex flows. Such velocity measurements also provide experimental data, which can be used to validate and improve CFD codes used to predict and explain the behaviour of many practical flows.

Several well-established quantitative measurement techniques exist today for the measurement of flow velocities. Instruments such as Pitot-Static tubes, hot-wire anemometers and Laser Doppler Velocimeters (LDV) are now routinely used in research to measure flow velocities. However, these are single point techniques, which allow the velocity at only one point in the flow field to be taken at an instant in time. Furthermore, the measurement procedure associated with Pitot and hot wire techniques requires the insertion of a physical probe into the flow thus introducing the risk of disturbing the flow.

LDV on the other hand has the important advantage of being almost non-intrusive. LDV uses optical probes to obtain local flow velocity information at a point by measuring the velocity of small particles suspended in the flow as they pass through the point of intersection from one to three pairs of coherent laser beams — one pair for every velocity component being measured.

Although LDV is non-intrusive, it is still a single point measurement technique. However, the nature of many practical flows of interest require that several instantaneous global velocity measurements be taken if their spatial structures and their development in time is to be better understood. Furthermore, it is difficult to perform whole field measurements with single-point probes. The effort involved in scanning the flow is enormous and is often not practical.

Faced with such limitations associated with single point probes, a lot of research has focused on the development of non-intrusive, multi-point flow measurement systems and techniques based on pulsed illumination.

In the last decade, the strong development in the technology underlying pulsed illumination and image recording has allowed decisive progress in the development of velocity field measurement methods derived from visualization techniques. Pulsed Light Velocimetry (PLV) enables velocity measurements to be made over global domains by observing the motions of small marker particles, suspended in the flow, over a predetermined time. It is assumed that the particles faithfully follow the flow thus allowing the local flow velocities to be inferred from the particle velocities. The distance moved by the particles between observations is measured and divided by the observation time to give an estimate of the particle velocities.

In PLV, the locations of the markers at various instants are recorded optically by very short pulses of light that freeze the marker images on an optical recording medium such as photographic film or a digital video recorder. The time interval between successive illuminations is typically very short — usually in the order of a millisecond or less — and small compared to the flow time scales, thus making it possible to obtain near-instantaneous velocity maps, typically, in a plane of the flow.

Today, Pulsed Light Velocimetry (PLV) has successfully evolved to become a widely used technique for multi-point, near-instantaneous velocity field measurements. It offers much higher accuracy and resolution than Streak Photography — its predecessor — where the pulse duration is long enough to produce streaks of the particle image motion. When applied to fluid flow measurements, PLV has two modes of operation, depending on the particle density and choice of analysis technique[1]:

- Particle Tracking Velocimetry (PTV) for low seeding density and
- Particle Image Velocimetry (PIV) for high seeding density.

In the particle-tracking mode, the distance travelled by each individual marker particle between illuminations is measured to obtain an estimate for the local velocity at each particle location. The particle displacements are determined by first identifying the images of the same particle in two or more recordings of the flow field and then measuring the difference between their positions. The higher the number of particles that can be successfully tracked, the higher the spatial resolution of the measurements.

The fundamental assumption made when tracking particles directly is that the particle image displacements are small compared to the mean particle spacing so that nearest neighbouring particle images are likely to belong to the same particle. This implies that to ensure high pairing success rates, the particle concentration must be low.

The low particle concentration (i.e. seeding density) associated with the particle tracking technique however means that the extracted velocity vectors may be spaced too far apart relative to the scales of motion, thus limiting particle tracking applications to the study of relatively large flow structures. Furthermore, in situations where it may be desirable to interpolate the randomly located velocity vectors onto a regular grid, for example, when comparing the experimental results to the results obtained from simulations, the interpolations

---

[1] A third mode called Laser Speckle Velocimetry (LSV) involves higher particle concentration than PIV. However it has been argued (see Adrian (1984)) that for practical flow measurements with tracers, the concentration of the illuminated particles cannot be high enough to produce speckle patterns in the image plane.

could be too coarse and as a result, give very poor estimations of velocities at the desired points in the flow field.

The PIV mode, on the other hand, allows for high tracer particle concentrations so that high-resolution flow measurements can be realised. In this mode, the particle concentration in the flow is not limited by the anticipated maximum tracer displacements since the displacement of each individual particle image is not sought. Instead, the particle images are grouped into clusters and the average displacements of each cluster is determined by analysing the cluster motion between successive recordings of the particle images.

We note that the system hardware arrangement for image acquisition in PIV is similar to that for PTV. The major difference between the two modes lie in the method employed to extract the velocity information from the acquired images. PIV and PTV can therefore be viewed as different modes of operation of the same instrument. In other words, PTV can be regarded as the low-density mode of PIV.

1.1.1 Research focus

PIV is fast emerging, as a popular choice for global flow velocity measurements and the advancement of the PIV measurement technique is the focus of this dissertation.

PIV's ability to provide a high resolution global view of the instantaneous spatial structure of flows in a qualitative and quantitative manner together with the continuous technological advances and reduction in costs of the underlying hardware continue to drive its ever-growing acceptance as a practical flow velocity measurement technique. Digital PIV (DPIV)[1] measurements are now routinely performed in research laboratories and industry using systems consisting of high speed computers, high repetition-rate, double-pulsed, solid-state lasers[2] as illumination sources, interline-transfer digital CCD (Charge-Coupled Device) cameras and correlation-based image-interrogation software as standard.

---

[1] Digital PIV here, refers to an arrangement where the images of the flow field is digitally recorded using a CCD camera.

[2] A common illumination source is the Nd:YAG laser, which is capable of producing hundreds of milli-joules of energy per pulse at repetition rates in tens of Hz.

While the same basic principle of recording successive images of particles suspended in flows and tracking their movements within the flow remain common to all PIV measurements, the way in which the illumination, recording and analysis of the resulting particle images are performed during PIV measurements have evolved since the early published measurements by Barker and Fourney (1977); Dudderar and Simpkins (1977) where the parabolic velocity profile of laminar tube flow was revealed using a planar laser light sheet for illumination and recording two exposures of the flow field on the same photographic film. Other notable early work include Meynart (1980, 1983a,b) and for DPIV, Willert and Gharib (1991). Today, there are now several ways in which the images of the flow field are acquired and analysed with every advance either in PIV hardware arrangement or image-interrogation technique ultimately aimed at improving one or more aspects of PIV performance.

Improving the performance of PIV, in particular, advancing the technique to make it a simple, yet practical and powerful tool capable of giving high quality three-component velocity information throughout a volume of interest is the objective of this project. Here, we map out a technique where a single multi-spectral camera is combined with multi-colour illumination of the flow field to give accurate, reliable and high resolution, three-dimensional velocity information within a practical time frame.

However, before we present the details of this novel measurement technique, it is appropriate to first introduce the reader to how the simplest PIV measurements are conducted and also explain the issues surrounding PIV performance including a review of other PIV arrangements and analysis techniques which have served to advance PIV performance.

1.2 The classical PIV system arrangement

In its simplest form, PIV gives two-component velocity field information in a plane of a flow field. A typical planar PIV arrangement and measurement chain are illustrated schematically in Figure 1.1 and Figure 1.2 respectively. A cross-section of a seeded flow field is illuminated, in the direction of the prevailing flow velocity, by a sheet of light pulsed typically twice in quick succession. During each illumination, the particles within the light sheet volume are projected onto the image plane of a recording medium whose optical axis is aligned normal to the illuminated plane in the flow. An instantaneous 2D velocity vector field map for the recorded plane in the flow field is subsequently obtained from the displacements of local clusters of particle images between two recordings of the flow field.

The illuminated flow field can be recorded on photographic film which is then developed and scanned (e.g. with a film scanner) and the digital information stored on a computer ready for analysis (see Prasad (2000a)). Today, the more common way of acquiring particle images of the illuminated flow field is to digitally record the particle images with a CCD camera thereby bypassing the film-development and scanning processes, hence, making the image data readily available for analysis. See Willert and Gharib (1991).

To extract the velocity field information from the recorded flow field images, a grid of interrogation windows is defined over the digital images as shown in Figure 1.3. Each window would typically contain enough particle images so that a correlation process can successfully be used to determine the 'average' displacement of the particle image cluster within it.

Pairs of corresponding interrogation windows from two image recordings are interrogated to determine the average displacement of the particle image cluster within each pair of interrogation windows, over the time separating the two recordings.

The local velocity of the region in the flow field represented by each window is subsequently approximated by the expression:

$$\vec{V} = \frac{\Delta \vec{s}}{\Delta t} M_i \tag{1.1}$$

Where $\Delta \vec{s}$ is the particle-image cluster displacement vector and is measured in image coordinates (in pixels), M_i is the image magnification and Δt is the time separating the two recordings of the flow field. A typical PIV velocity field obtained from such window interrogations is shown in Figure 1.4.

Before the images are interrogated, some pre-processing is sometimes carried out to remove image noise such as background reflections. Also, after

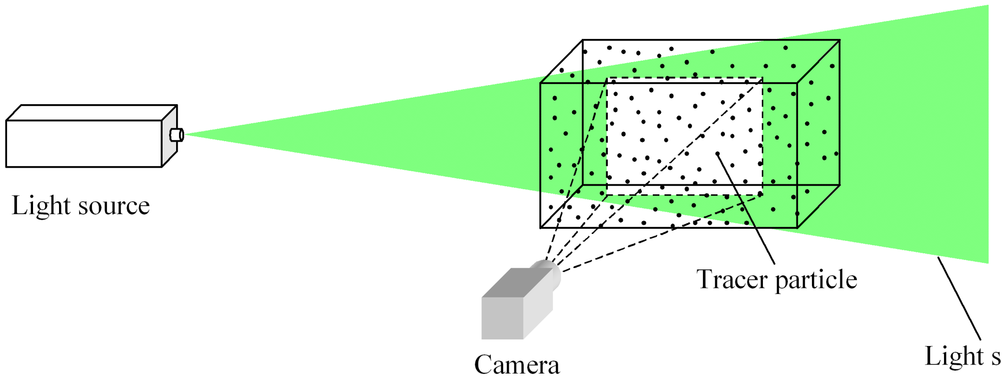

Figure 1.1: A schematic representation of the classical 2D PIV experimental arrangement.

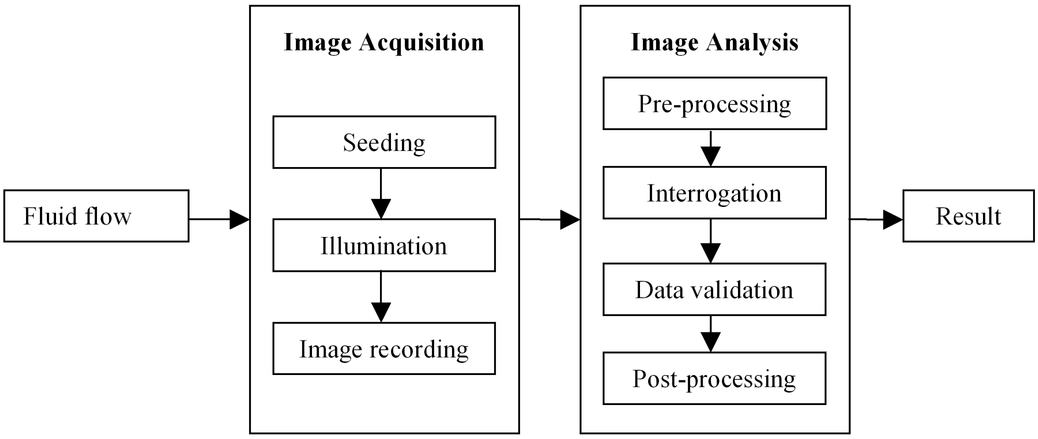

Figure 1.2: The classical 2D-PIV measurement chain.

interrogation, the resulting displacement vector field is often validated to remove and possibly replace vectors that may not accurately represent the local flow movement. Typical validation and replacement methods include comparing each vector with the mean or median of the nearest neighbours and substituting an ‘invalid’ vector with an average of the neighbouring vectors. The velocity field is obtained from the displacement vectors using equation (1.1). Derived flow quantities such as vorticity may subsequently be calculated from the velocity data.

We now explain, in greater detail, how the velocity vector field is obtained from the digital camera images.

1.2.1 Interrogation of digital PIV images

The standard means of extracting displacement information of the particles within the interrogation windows is by correlating corresponding windows in two image recordings of the flow field.

The correlations of the windows is a search for the best pattern match between the cluster patterns formed by the same particles at different instants in time. Given two interrogation regions, the correlations can be described as moving one region over the other, and at each search coordinate (in whole pixels), evaluating the degree to which the cluster patterns match. The most probable cluster displacement in the time separating the image recordings, is taken to be the search coordinates that corresponds to the best pattern match.

The best pattern match is taken as the highest value in a correlation table. The correlation table shows the degree of correlation of the image patterns over the whole range of search coordinates where the mean particle cluster displacement is assumed to be given by the coordinates of the highest value (i.e. the peak) in the correlation table.

The values in the correlation table are obtained by evaluating, at each search coordinate, the discrete cross-correlation function expressed as:

$$C(\Delta_i, \Delta_j) = \sum_{i=1}^{M} \sum_{j=1}^{N} I_1(i, j) \cdot I_2(i + \Delta_i, j + \Delta_j) \tag{1.2}$$

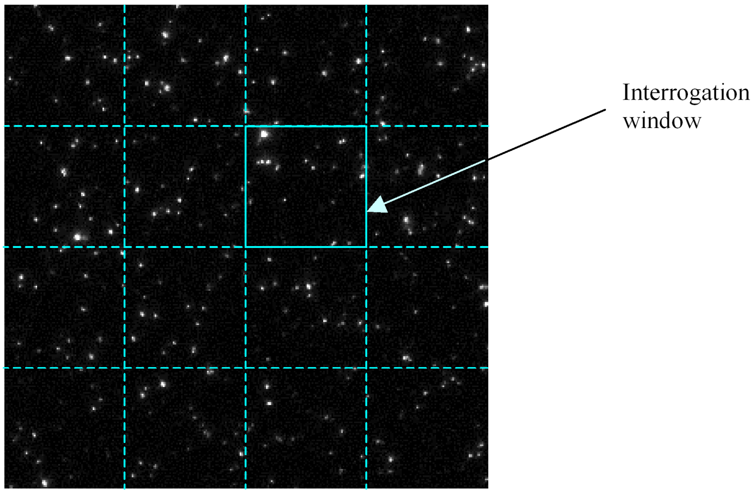

Figure 1.3: A 296 × 296 pixel sample of a singly-exposed digital image divided into interrogation windows of size 74 × 74 pixels. The particle images appear as white spots on a dark background.

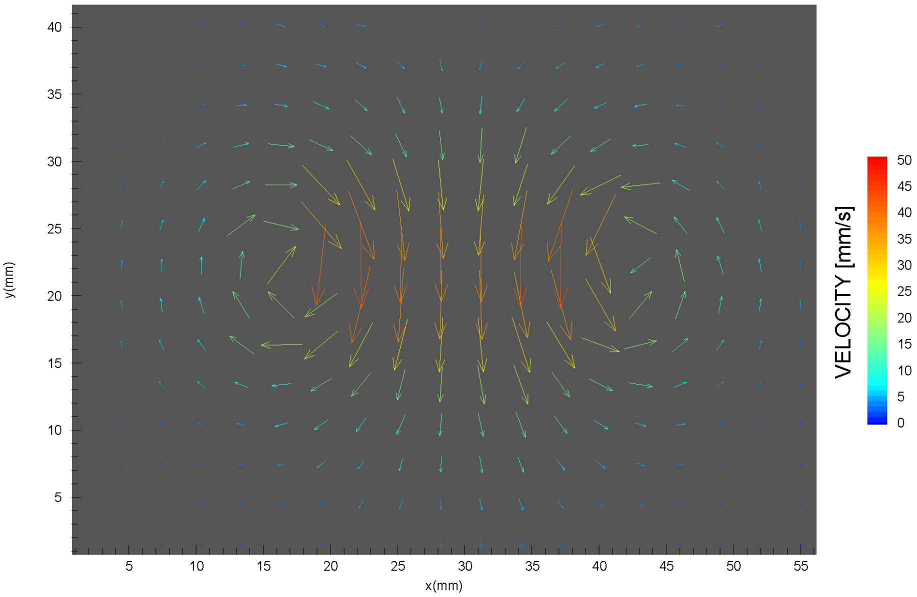

Figure 1.4: Velocity vector field of a laminar vortex ring flow obtained by analysing a pair of 1392 × 1040 pixel flow images representing a 57mm×42mm area of the flow field. Each vector represents the local velocity over a 3mm×3mm (i.e. 74×74 pixel interrogation) area of the recorded flow field. Part of the analysed image is shown in Figure 1.3. Images analysed with the author's 2D PIV software package.

where the size (in pixels) of the interrogation areas are $M \times N$, I_1 and I_2 are the pixels intensities in interrogation area 1 and interrogation area 2 respectively and the Δ_i and Δ_j are the correlation search coordinates. See Figure 1.6.

The correlation function can be computed directly or via FFT[1]. Figure 1.5 illustrates how the discrete cross-correlation function is evaluated directly, for a corresponding pair of interrogation windows. At the search coordinate (Δ_i, Δ_j), only the overlapping pixels contribute to the value of the correlation function. We note that if the second interrogation window is effectively 'cut-out' from the digital image and the correlation function is computed directly, then at larger search coordinates, the number of overlapping pixels that contribute to the value of the correlation function will be less. This will result in the estimated mean cluster displacement being biased towards lower values.

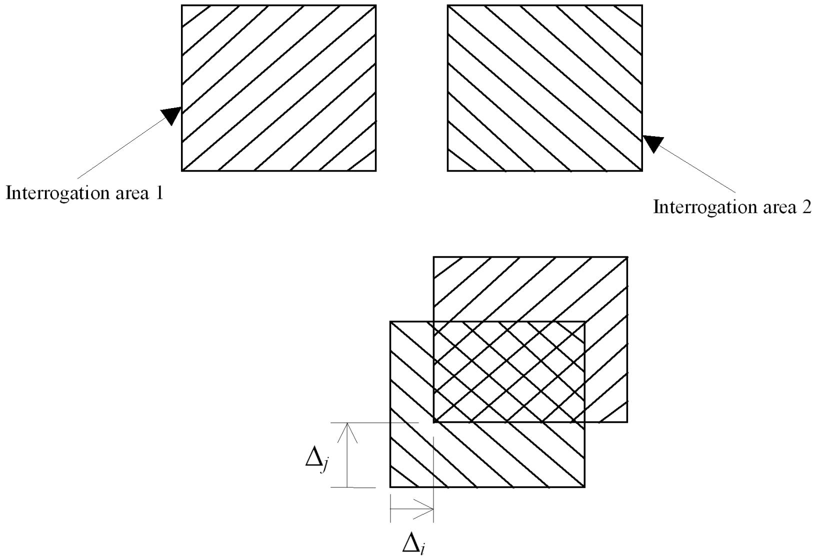

Figure 1.5: An illustration of how two interrogation windows are correlated at the search coordinate (Δ_i, Δ_j). For the pair of cut-out interrogation windows, only the intensities of the pixels in the overlapping region contribute to the value of the correlation function when the function is computed directly.

---

[1] For the radix-2 FFT implementation, the correlation window sizes will be a power of 2, e.g. 32 pixels, 64 pixels, 128 pixels.

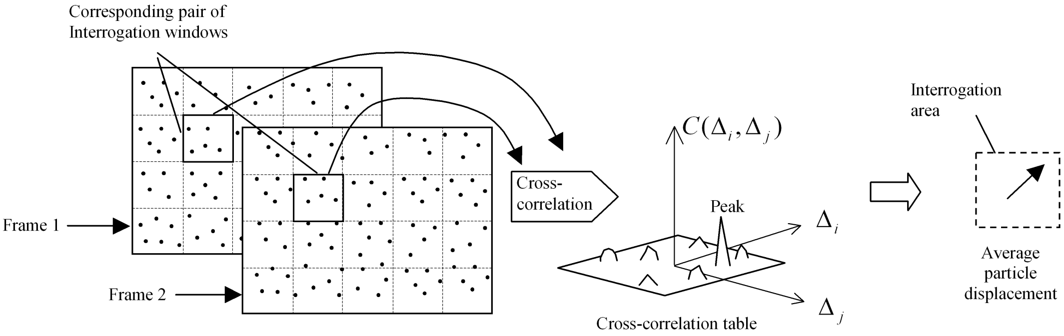

Figure 1.6: The steps involved in the cross-correlation of two singly-exposed images to give the average displacement of the cluster of particle images in the interrogation window. The coordinates of the peak in the cross-correlation table give the cluster displacement.

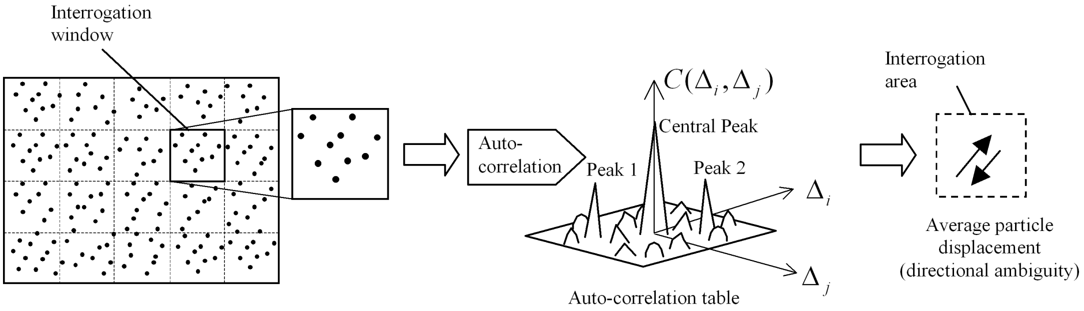

Figure 1.7: The steps involved in the auto-correlation of an interrogation area of a doubly-exposed image to give the average displacement of the cluster of particles in the interrogation window. The cluster displacement is given by the coordinates of one of the two symmetrically located peaks in the auto-correlation table.

1.2.1.1 Cross-correlation vs. auto-correlation

When the images of the moving particles in the flow field are recorded on separate camera frames, a pair of corresponding windows from two single exposed images (frames) is cross-correlated, and the average window particle displacement is indicated by the location of highest peak in the resulting table of evaluated correlation function. However when the two recordings of the flow field are registered on the same image frame, we have a doubly-exposed image. In this case, each interrogation window from the image is auto-correlated and the average window particle image displacement is indicated by one of a pair of dominant symmetrical peaks in the resulting auto-correlation table. The cross-correlation and auto-correlation processes are depicted in Figure 1.6 and Figure 1.7 respectively. For a review of both correlation methods, see Adrian (1991); Keane and Adrian (1990, 1991).

The choice between auto-correlation and cross-correlation is influenced by the way in which the images are recorded which in turn is influenced by the camera hardware.

Historically, auto-correlation has been the more common method because two illuminated images of the flow field could not be recorded onto separate camera frames in a time interval short enough to give adequate temporal resolution of moderate to high-speed flows. However, the advent of interline transfer technology in digital cameras has driven the change from the auto-correlation of double exposures to the cross-correlation of singly exposed image recordings. The technology allows two sequential images to be recorded, in rapid succession, onto separate camera frames thus making it possible to use cross-correlation analysis and exploit its advantages over auto-correlation. These advantages are discussed next.

The cross-correlation and auto-correlation of an artificially generated pair of uniform flow images and their resulting correlation tables are shown in Figure 1.8 and Figure 1.9 respectively. The simulated doubly exposed image in Figure 1.9 is simply a combination of the simulated two singly exposed images in Figure 1.8.

The resulting auto-correlation table from the interrogation of the doubly exposed image consists of two dominant peaks appearing almost symmetrically on both sides of a central self-correlation peak[1]. In the auto-correlation table, the presence of the central correlation peak makes it increasingly difficult to distinguish the symmetrical peaks from the central peak as the particle image displacements increasingly become less than about a particle image diameter. Therefore it becomes more difficult to resolve the velocities in regions where the particle image displacement is less than the particle image diameter. As a result, the dynamic range of measurements that can be achieved with auto-correlation is relatively low. We discuss the dynamic range in section 1.4.

The existence of the two symmetrical dominant peaks in the auto-correlation table also causes directional ambiguity problems since it becomes impossible to determine which of the two sets of peak coordinates represents the average particle image displacement—unless the flow is uni-directional and the sign of the displacement is known a priori. Elaborate techniques such as image shifting Adrian (1986); Landreth et al. (1988) and the use of different colour illumination Goss et al. (1991) can be used to remove the directional ambiguity during the multiple recordings of particle images on the same camera frame. Such methods however add more complexity to the PIV system arrangement and image analysis procedure.

However, in cross-correlation analysis of particle images recorded on separated frames, the directional ambiguity problem is not present. Instead, the direction of the flow is determined automatically by the order of the correlating images. Furthermore, the absence of a central self-correlating peak means that much smaller particle image displacements can be resolved, therefore allowing a larger dynamic range of measurements to be realised.

Other advantages that cross-correlation has over auto correlation include the possibility of higher valid-data-yield — we discuss valid-data-yield in section 1.6 — because the signal-to-noise ratio in the correlation table will in general be higher due to lesser number of random mismatch of particle images. Furthermore, for the same seeding density, cross-correlation gives better accuracy due to lesser number of overlapping particle images in the recordings. In other words, cross-

---

[1] The self-correlation peak as discussed here refers to that due to the correlation of the particle images. However, we note that in real experiments, the auto- or cross-correlation of flow images where images of non-moving boundaries are present will also give a self-correlation peak.

correlation allows greater seeding density, before the same degree of particle image overlaps occurs; hence it offers higher spatial resolution measurements (see section 1.7 for discussion of spatial resolution).

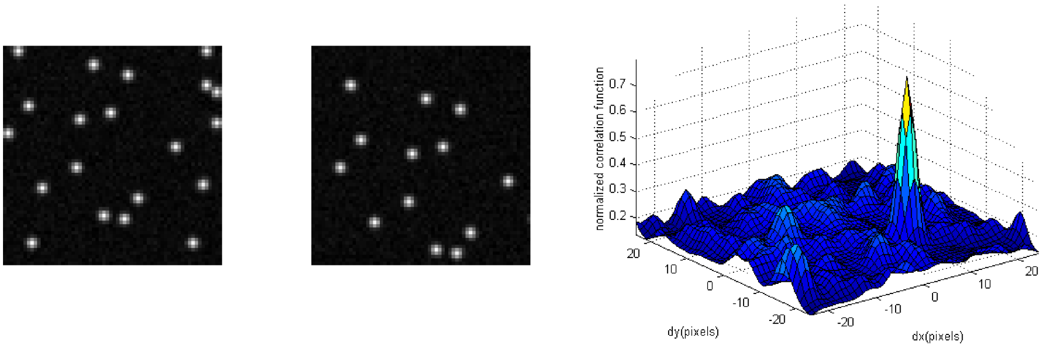

Figure 1.8: A pair of simulated 64 × 64 pixel singly-exposed interrogation windows and their cross-correlation table. The particle image displacement is (7, -10) pixels. The highest peak in the correlation table is located at the coordinate (7, -10).

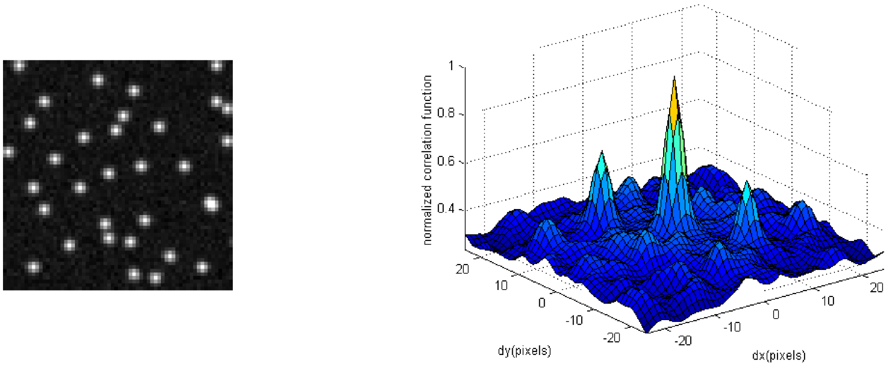

Figure 1.9: A simulated 64 × 64 doubly-exposed interrogation window and its auto-correlation table. The image is a composite of the two images shown in Figure 1.8. The true particle image displacement is (7, -10) pixels. Note the central peak and the two symmetrical peaks at the correlation search coordinates (7, -10) and (-7, 10).

1.3 The Measures of PIV performance

Beyond the basic PIV system arrangement, several other PIV systems have been developed to address one or more aspects of PIV performance. The performance of PIV systems is typically assessed in terms of the following criteria:

- Dynamic velocity range
- Image acquisition speed and image processing speed
- Accuracy and valid-data-yield
- Spatial resolution
- Number of velocity components revealed by each vector
- Spatial range of the measurements (i.e. planar or volumetric)

In the next sections, we define these measures of performance and discuss how the hardware arrangement and experimental parameters influence them.

1.4 Dynamic velocity range

The Dynamic Velocity Range is defined as the maximum measurable velocity divided by the minimum measurable velocity.

For good accuracy, PIV experiments would normally be arranged so that the minimum pixel displacement of the particle images between any two successive image recordings is greater than the uncertainty in the measurements. Also, the maximum particle image displacements must be kept low to prevent large relative particle displacements — a major source of correlation error.

In many instances, the larger the imaged region of the flow, the wider the range of velocities that will exist across the region. Therefore, as the imaged region is increased, a PIV system with a larger dynamic range will be required to keep the smallest particle image displacements sufficiently larger than the measurement uncertainties.

A higher dynamic range system will also permit more accurate measurements of a wider range of flow scales of motion.

High dynamic range PIV systems are characterized by small particle displacement measurement capability and low measurement uncertainties.

Camera noise is a major source of measurements uncertainties and high dynamic range PIV systems will typically be comprised of low-noise cameras. The choice of correlation method also plays a major role in determining the affordable dynamic range. The absence of a central self-correlation peak in the cross-correlation table means particle displacements smaller than the particle image diameter can be measured. For cross-correlation based systems, the lower limit of the dynamic range is determined by the uncertainty in the measurements instead of the particle image diameter, as is the case for auto-correlation based systems.

The upper limit of the dynamic range, on the other hand, is restricted by factors that do not favour large particle displacements. These factors include the local flow gradients, approximations of the true flow path by straight trajectories of the particles and movements of particles across the boundaries of the interrogation windows. We shall discuss these factors in section 1.6 where we explain the accuracy of PIV systems. However, here, it is sufficient just to note that the time interval between image recordings is typically set to give particle image displacements no greater than 10-20 pixels which translates to a dynamic range in the order of hundreds for cross-correlation based image analysis.

1.5 PIV measurement speed

An important area of performance of a PIV system is the speed at which the velocity information about a flow field can be obtained. We define the measurement speed as the time required to acquire and interrogate the particle images.

Advances in camera technology have had a significant impact on the preferred choice of PIV image recording and hence, the measurement speed. Since its early use for PIV image recording (see Westerweel (1993); Willert and Gharib (1991)), the spatial and temporal resolution of digital video recording has advanced to make it the firm choice over photographic image recording. The high image throughput offered by the digital imaging process has been decisive in investigating flows such as turbulence where hundreds or even thousands of images need to be recorded and analysed to obtain reliable estimates of the statistics of the structures.

In the digital approach, the flow is imaged directly onto a CCD camera, thus bypassing the time consuming procedure associated with the processing of photographic film such as film development and scanning of the negatives. The immediate availability of the digital images also means that image interrogations could be performed in real time thus enabling experimentalists to make adjustments to experimental parameters in order to improve the quality of the measurements. Today, image acquisition is more commonly performed digitally. Digital cameras with framing rates of over 1000 frames per second (improved temporal resolution) at high resolutions are available today.

The time required to interrogate the images is also important to the performance of a PIV system. Such is the power of modern day computers that several hundred high-resolution flow images can be analysed in a couple of hours. Today a typical 2000×2000 pixel image can be interrogated in about 10 seconds at a rate of over 400 vectors per second using Fast Fourier Transform (FFT)[1] window sizes of 32 pixels on a PC with a Pentium III 1GHz processor with 512 MB of RAM. This is significant progress from the days when a single typical 100×100 vector field image could take hours to interrogate and some researchers (see Meinhart et al. (1993)) employed parallel processors to improve the image interrogation speed.

However, while the data storage speed, hard disk capacity and processing power of computers will continue to grow, the increasing resolution and framing rates of cameras is likely to make the need for faster image transfer, larger storage capability and image processing speed remain just as great as it is today. We may begin to see frequent use of image compression techniques to reduce image data hence speed up image transfer and conserve disk space. The compression of images however could compromise the accuracy and valid data yield of the

---

[1] Correlation based interrogation techniques typically use the FFT to achieve fast correlations since the operation count for correlations via FFT is much less than that for the direct computation of the correlation function. The computational cost involved in computing the correlation function for an interrogation window of size $N \times N$ via FFT is $O\left[N^2 \log_2 N\right]$ while the computational cost is $O\left[N^4\right]$ when the correlation function is computed directly.

measurements, as some particle image information is lost. The effects of lossy and losselss image compression techniques on the quality of the velocity field information is examined by Cenedese et al. (2000).

In order to speed up image interrogation, the use of multiple processors may become more frequent. Also, the development of faster image interrogation algorithms will be important. An example of a fast image interrogation algorithm is the Sparse Array correlation algorithm whose potential to significantly speed up correlations was demonstrated by Hart (1998b). Key to such sparse correlationalgorithm is the omission of the low intensity pixels that contribute insignificantly to the value of the correlation function. The gain in correlation speed versus the quality of the velocity field information obtained is, however, dependent on how many of the true particle image pixels are considered during the evaluation of the correlation function.

In Appendix F, we describe a high-speed image processing algorithm where the particle images are first identified and then reconstructed on a coarse grid on which the correlations are subsequently performed.

1.6 Accuracy and valid-data-yield

Two of the most important performance indicators are the accuracy and valid-data-yield of the measurements. The accuracy and valid-data-yield are closely related and we discuss both in this section. The accuracy describes how well the velocity vectors represent the local flow velocities while the valid-data-yield indicates the percentage of the vectors that can be accepted as valid representations of the local flow velocities. In other words, the valid-data-yield is a measure of the reliability of the measurements.

We note that high accuracy is required in measurements such as turbulence where reliable turbulence statistics are sought and the mean flow is often subtracted from the results to reveal the smaller structures in the flow, see Adrian et al. (2000); Christensen (2004). Furthermore, the accuracy in the estimation of derived flow quantities such as vorticity depends on the accuracy of the raw velocity data. For PIV and PTV, the maximum particle displacements between successive image illuminations would typically be between 10 to 20 pixels and the uncertainties in estimating the particle image displacements would typically not exceed about 5% of the particle image displacement.

In PTV, the accuracy to which each particle displacement is resolved depends on the accuracy to which its particle image centroid is determined. The valid-data-yield on the other hand, depends fundamentally on the effectiveness in correctly identifying the true particle images from the flow image. Poor particle identification techniques could drastically reduce the valid-data-yield. Background pixels that are erroneously identified as particles could present themselves as nearest neighbours. Also, a particle image whose true pair has not been identified as may still end up being matched with another particle image.

In correlation-based PIV, the accuracy and valid-data-yield largely depend on the size of the interrogation window and the local flow gradients within the area represented by the window. Smaller interrogation windows will typically give a better representation of the areas of the flow field they represent since the variation of velocities across smaller windows is less. Smaller interrogation windows also translate to more windows hence more vectors to represent the same imaged area. In other words, smaller interrogation windows translate to higher spatial resolution. Defining smaller interrogation areas may, however, affect other correlation parameters in a way to make the eventual vector from the correlations a poor representation of the particle cluster displacement. The window size trade-off is discussed in greater detail shortly.

In order for the vectors obtained from the correlations to be valid, the highest peak in the correlation table must come from the correlations between true particle image pairs and not any of the peaks formed as a result of random correlations between the particle images. Peaks from random correlations are termed as noise peaks. When the experimental parameters are properly specified, the height of these peaks will be much smaller than the true displacement peak. Such noise peaks are visible as small peaks in the correlation tables shown in Figure 1.8 and Figure 1.9.

When the peak in the correlation table is identified, to improve the accuracy of the displacement vectors beyond half integer values, the coordinates of the peak is typically located to sub-pixel accuracy by either by fitting a curve through the peak and its surroundings or evaluating the centroid of the peak profile. In particle tracking, the particle image centroid is similarly determined to sub-pixel accuracy by fitting a curve through the highest-intensity pixel of the particle image and its neighbours or using the zero-order moment equation to determine its centroid. Commonly used curve-fitting estimators include the three-point Gaussian estimator and the parabolic peak estimator. The expressions for these sub-pixel estimators are given in Appendix E. With these estimators, sub-pixel accuracies in the order of $\frac{1}{10}$th of a pixel are typically achieved today during PIV/PTV measurements, see Huang et al. (1997) and Stuer et al. (1999).

1.6.1 The types of errors in PIV

Determining the correlation peak or the particle image centroid to sub-pixel accuracy does not guarantee highly accurate or valid vectors. There are several experimental factors that influence the accuracy and the validity of the vectors. In this section, we explain the type of errors in PIV and their sources.

1.6.1.1 Bias error

When the displacement peak in the correlation table or the particle image centroid is determined to sub-pixel accuracy, there is a systematic and sometimes-significant error associated with the sub-pixel fit. This error is commonly known as the mean bias error and as its name suggests, is characterised by a bias of the estimated sub-pixel peak/centroid coordinates towards the nearest integer pixel values.

Bias errors arise as a result of the discretization of the particle images due to the finite resolution of the camera.

The magnitude of the bias error is influenced by the particle image size and the choice of sub-pixel estimator and according to Prasad et al. (1992), the choice of the best estimator depends on the particular flow being investigated. However, the same bias trend is observed for all the aforementioned estimators. Figure 1.10 shows the typical trend in the bias error, over a range of sub-pixel particle image displacements, for three particle image sizes. The bias error typically decreases with an increase in the particle image size. However, larger particle images would generally give larger random errors when noise is present in the images. We discuss random errors in section 1.6.1.3.

When the bias error is the dominant error in a PIV data, a histogram of the particle/window displacement data would reveal peaks around the integer displacement values. This bias phenomenon is commonly referred to as 'peak-locking'.

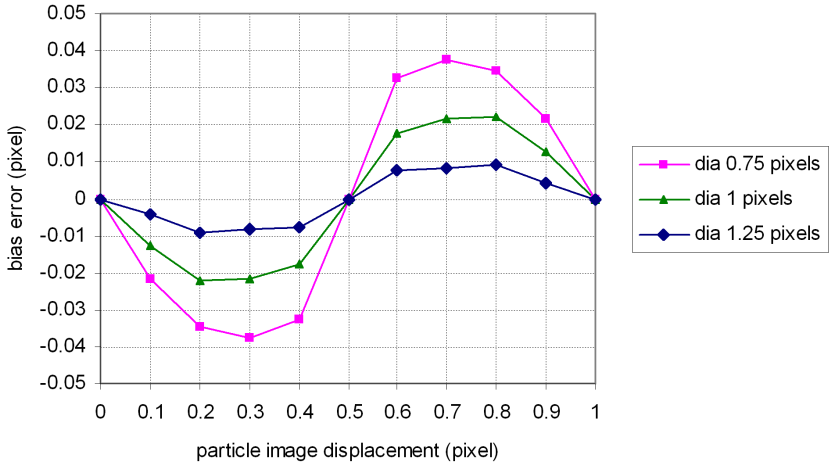

Figure 1.10: Bias error versus particle image displacement for three particle image sizes. The simulated images contain no noise and the peak in the correlation table was resolved to sub-pixel accuracy using the Gaussian sub-pixel estimator.

1.6.1.2 Errors due to local flow velocity gradients

Another type of error introduced in PIV measurements is that resulting from the velocity gradients that exist in the flows. In PIV, a single vector that is typically placed at the centre of the interrogation window approximates the velocity of the flow field region that is represented by the interrogation window. Placing the vector at the centre of the interrogation window however does not account for the nature of the velocity variation across the interrogation area or the distribution of the particle images in the area.

We note that when the relative particle displacements are high, then less particles correlate with their pairs at the same correlation search distance Δ.When this is the case, the resulting velocity vector from the correlations is likely to come from the regions within the interrogation window where the particles are mostly concentrated and the relative particle displacements are smaller. The deviation of the measured velocity from the true particle velocityat the centre of the interrogation area will therefore depend on the variation of the velocities across the interrogation area and the distribution of the particles within it.

While the velocity gradient error is reduced when smaller interrogation windows are defined, we note that the number of particles contained within the interrogation window decreases. The probability of achieving a valid correlation outcome however decreases with the number of particle images in the interrogation windows, see Keane and Adrian (1992). The super-resolution techniques which are discussed later in section 1.7 serve to reduce the gradient error (while maintaining a high valid-data-yield) because the number of velocity vectors obtained using such image analysis techniques are denser thereby effectively representing smaller interrogation areas.

Higher pixel-resolution cameras can also help to reduce the gradient error because they allow the flow field to be divided into smaller interrogation areas. For example, given the same size Field of View (FOV), a 32×32 pixel interrogation window on a higher resolution camera represents a smaller region of the FOV. Smaller particles at higher seeding can then be used to maintain the desired particle image size and a sufficiently high number of particles per window as required for a high correlation valid-data-yield.

1.6.1.3 Random errors

The image distortion caused by flow gradients is somewhat systematic in the sense that it is these distortions that reveal the flow characteristics. However there are several experimental factors that cause random image-intensity pattern distortion and as a result introduce uncertainties in the location of the correlation peaks or the location of the particle image centroid.

Factors that contribute to random particle image distortion include the camera noise, particle image overlaps, lens aberrations, uneven illumination of the particle images, saturation of particle image pixels, boundary reflections, reflection of particles onto other particles, the change in the intensity profile of the particle image from one image to another due to the tumbling effect of a non-spherical tracer particle, and the movement of particles across the boundaries of the illumination plane.

Here, we refer to those factors that unsystematically change the particle image intensity profile from one captured image to another as image noise.

The uncertainties introduced in the measurements due to the random change in the particle image intensity profiles and their cluster patterns are generally referred to as random errors.

The random errors could either be large or small depending on the severity of the particle image distortions. Sources of random image distortions such as particle image overlaps, camera noise, boundary reflections and non-spherical particle shapes typically give small measurement uncertainties. Camera noise and boundary reflections add to the intensities of the pixels, therefore, they change the particle image intensity profile from one image to another. Also, when the particles are not spherical, the particle image intensity profiles will typically change from one captured image to another as the particles tumble between image recordings of the flow field. Flow velocity gradients also cause particle images that are initially separated in one image to overlap in another image hence leading to a change in the image profile of the particles involved. Overlaps affect the accuracy to which the particle image centroids are determined during the particle tracking process and would typically reduce the valid data yield as overlapping particles are difficult to identify, split and pair correctly. Severe distortions could also hinder the ability to correctly identify and pair particle images or reduce the probability that the highest peak in the correlation table is the displacement peak. When the highest correlation peak is not the displacement peak, the displacement vector is said to be spurious since it does not represent the true particle displacements.

The presence of spurious vectors in PIV data is obviously not desirable. The more spurious vectors there are in a measurement, the lower the valid data yield, and the less reliable are the measurements. Spurious vectors typically show up in velocity field maps as those vectors that deviate extensively from their surrounding neighbouring vectors, assuming the neighbours themselves are not spurious. Unvalidated PIV data will typically contain a number of such vectors and the presence of a least a few of these spurious vectors in PIV data is almost unavoidable, even in carefully designed experiments.

Spurious vectors or outliers, as they are sometimes called, occur in correlation based PIV processing when the aforementioned experimental factors severely distort the particle image patterns to the extent that the correlations between true particle image pairs are not sufficient to produce a peak which is higher than those peaks due to random particle image correlations.

The major causes of spurious vectors in correlation based PIV measurements are:

- large relative particle displacements (i.e. large flow velocity gradients) compared to the particle image diameter,
- excessive number of particles moving across the boundaries of the interrogation windows and illumination plane,
- insufficient number of particle images in the interrogation windows to ensure that the correlations give a valid vector as the previous two factors become more prominent.

High local velocity gradients and the movement of particle images out of interrogation windows all contribute to a loss of correlation pairs.

As the velocity gradients become severe, the relative dispersion of particle displacements increases and less particles match at the same correlation search coordinates (Δ_i, Δ_j). As a result, the height of the displacement peak is reduced.

When the local gradient is severe enough, an otherwise strong displacement peak will be replaced by smaller peaks, each representing the correlations of subsets of the particle clusters.

The movement of the particle images across the illumination plane and the interrogation window boundaries also have the effect of reducing the displacement peak. The larger the particle image displacements, the greater the number of particles entering and leaving the interrogation windows and the illumination plane. As a result, more particles in the first interrogation window will lose their pairs in the second window, hence leading to a reduction in the height of the displacement peak.

In PTV, factors such as particle image overlaps due to excessive seeding, high camera noise, uneven illumination of the particles and reflections off boundary surfaces will generally make the identification of true particle images in low contrast flow images more difficult. The valid-data-yield will typically be reduced as more cases of missing particle image pairs occur and greater number of background pixels (identified as particle images) are matched with true particle images.

1.6.1.4 Perspective error

Another very important error in the measurements is that which is introduced in the measured in-plane displacements of the particle images when the particles are imaged through a lens with a significant angle-of-view (or perspective).

We note that in classical single-camera 2D PIV arrangement, the planar illumination assumption is an approximation. The light sheet has a finite thickness (although it is a small thickness relative to the planar dimensions) and the particle images are distributed over this thickness.

When particles at the same x- and y-coordinates but different z-coordinates are imaged with a lens of finite Angle-of-View (or perspective), their images will not be recorded at the same position on the CCD. The exception being when the particles lie on the camera's optical axis. Also, when particles, separated along the z-direction, move by the same amount in the same direction, their image displacements will not be the same due to the variation of image scale with imaging distance.

Therefore, in the classical PIV arrangement, there is a perspective error in the measurements when:

- The same imaging scale is applied to all particles, which in real space are separated by some distance along the z-direction.
- There is z-directional movement of the imaged particles.

The perspective error can be significant when the out-of-plane velocity components have a non-negligible magnitude relative to the in-plane components. Using long focal length lenses will help reduce the perspective error because the variation in the image magnification with distance along the camera optical axis is smaller. Later in section 3.3.1.2, when we discuss the 3D colour PIV system, we explain how to correct for the perspective error when the z-coordinate of the imaged particles is known.

1.6.1.5 Tracking error

We note again that the particle velocities, as determined from their displacements over a given time interval, are only approximations of the instantaneous flow velocities. Therefore, a vector that accurately represents the motion of the particles may still be considered as spurious if the particles do not faithfully follow the flow or if the approximation of the local Eulerian flow velocity from the Lagrangian motion of tracer particles is poor.

1.6.2 Reducing the errors in PIV measurements

1.6.2.1 Reducing the tracking error

The errors associated with the response of the particles and the approximations of local instantaneous velocities by particle trajectories could be reduced by careful selection of particles and the time between illuminations respectively.

Suitable particles for liquid and gas flows are reviewed by Melling (1997); Merzkirch (1987); Raffel et al. (1998). In general, the particles should be able to follow the flow faithfully without excessive slip. They should be spherical, highly reflective, chemically inert, neutrally buoyant and smaller than the length scales of the structures present in the flow. The particle size distribution should also be narrow to prevent the saturation of the particle images belonging to the larger particles in the distribution.

Smaller particles follow the motion of the flow better than larger particles. Smaller particles also provide an opportunity for higher resolution measurements because the seeding density can be increased while still maintaining a low number of particle image overlaps and a high optical transparency of the medium.

However, the benefits of using small particles in the experiments is tempered by the need to have a large particle images to withstand high flow velocity gradients and for good light scattering to give a high particle image intensity-to-noise ratio. A trade-off is therefore typically made between the need for large particles for improved light scattering and improved valid correlation outcome and the need for small particles to improve particle response.

The difficulty involved in the particle size trade-off can be reduced by noting that the response of the particles in a flow also depends on how closely the particle density and fluid density match.

Assuming very low Reynolds number flow, the gravitationally induced velocity (settling velocity) U_g of a spherical particle of diameter d_p and density ρ_p in a fluid of density ρ and dynamic viscosity μ is given by the expression:

$$U_g = d_p^2 \frac{(\rho_p - \rho)}{18\mu} g \tag{1.3}$$

where g is the acceleration due to gravity, see Raffel et al. (1998).

The particle lag in a continuously accelerating fluid is also given by a similar expression with the g replaced by the fluid acceleration a. See Raffel et al. (1998).

To minimize the discrepancy between the fluid and particle motion, the particle properties should be selected so that U_g is negligible compared to the actual flow velocities. A particle approaches neutral buoyancy as U_g approaches zero and from equation (1.3), a very close approximation to neutral buoyancy can be achieved either by using extremely small particles or using particles whose densities are close to that of the working fluid. However, good light scattering requirement suggests that the better option to approaching neutral buoyancy is by carefully selecting particles whose densities match the density of the working fluid. We note that for fluid flows, matching particles with the working fluid is not too difficult as there is a wide range of particles from which to choose, but this may be relatively difficult to achieve in air.

Particle sizes used in PIV experiments typically range from about 5 microns to about 200 microns.

To reduce errors associated with approximating the local velocity by the Lagrangian motion of the particles, the illumination time interval, Δt, should be chosen carefully. The illumination time interval should be set to maintain an optimum balance between the need to have a much larger particle image displacement compared to the expected measurement errors and the need to have a small illumination time interval to prevent large relative particle displacements, excessive movement of particle image motions across the interrogation window boundaries and a large particle trajectory approximation for the flow.

1.6.2.2 Reducing the bias and random errors

The reduction of the bias and random errors and the detection of spurious vectors or, better still, the minimization of their occurrence in PIV has been a subject of active research. The work of several researchers, often consisting of theoretical analysis as well as computer simulations of flows have provided valuable information about how some crucial experimental parameters affect the accuracy of PIV measurements.

The bias error and schemes aimed at reducing it have been discussed by Huang et al. (1997, 1993, 993b); Lourenco and Krothpalli (1995); Nogueira et al. (1997, 2001); Prasad (2000a); Prasad et al. (1992); Raffel et al. (1995). Prasad et al. (1992) studied the effect of image resolution on the bias and random errors and suggested that for correlations, the total error (i.e. bias plus random error) is minimized when the particle image diameter is about 2 pixels[1].

Critical to obtaining valid vectors from the correlations is the choice of the interrogation window size. Smaller interrogation windows are desired because higher resolution measurements can be achieved thus allowing smaller flow structures to be revealed. In other words, the interrogation windows should be smaller than the scales of motion. However, the smaller the interrogation window, the less the number of particles contained within it and the greater the probability of particle images in neighbouring areas in the flow field correlating better with the particles in the window to give a spurious vector.

The obvious solution to keeping the number of particles in the interrogation windows high when the interrogation window sizes are reduced is to increase the particle seeding density. However, there is a limit on the extent to which one can seed the flow. High seeding can result in excessive number of particle image overlaps and a great reduction in the optical transparency of the medium.

---

[1] The particle image diameter is defined as the number of pixels that covers 68.3% of the light intensity recorded over the central horizontal or vertical line of pixels.

On the other hand, if the interrogation window is too large, then the relative particle image displacement across it may be too high to expect a valid or highly accurate velocity vector from the correlations[1].

A careful balance must therefore be maintained between using too small a window size and using large interrogation windows.

Keane and Adrian (1990, 1992) studied the effects of PIV correlation parameters such as interrogation window particle concentrations, particle displacements and local flow gradients in the windows and gave a set of guidelines aimed at maintaining a high valid data yield from the correlations of the windows. To achieve more than a 95% probability of valid cross-correlation outcome, they recommended the following:

- The number of particle images in each interrogation window should be at least 7 and
- the relative displacements of the particle images within the windows should be less than a particle image diameter.

It is also generally accepted that the particle image displacement should be less than half of the interrogation window size to reduce the movement of particles across the illumination plane and interrogation window boundaries.

1.6.2.3 Advanced interrogation and validation algorithms

To minimize the movement of particles out of the interrogation windows, the second window in the pair can be made sufficiently larger so that all the particles in the first window are contained in the second window, assuming there is no particle movement out of the illumination plane. Note that the first window will be 'zero-padded' up to the size of the second window.

If the correlation function is to be computed directly, then an alternative to using a larger second window, is to initiate a search for the best pattern match

---

[1] Larger interrogation windows also mean lower spatial resolution measurements since the number of vectors representing the flow is less — unless overlapping windows are defined or advanced interrogation techniques such as the iterative correlation (to be discussed in section 1.7) are used.

in the second frame instead of explicitly defining a corresponding window in the second image, see Hart (1999).

Defining a larger second window or ‘floating’ the second window also ensures that every pixel in the first window is correlated with a pixel in the second window for all correlation search coordinates. This significantly reduces the bias of the displacement peak coordinates towards smaller displacements.

The correlation displacement bias can also be minimized by normalizing the correlation function so that the function is weighted according to the number of pixels involved in the correlations. Normalizing the correlation function also reduces the errors associated with the uneven particle image intensities (see Huang et al. (1997)).

Higher accuracy and valid data yield can therefore be achieved if the correlation function is redefined to take a form such as that expressed by Hart (1999):

$$C(\Delta_i, \Delta_j) = \frac{\sum_{i=1}^{M}\sum_{j=1}^{N}(I_1(i,j)) \cdot (I_2(i+\Delta_i, j+\Delta_j))}{\sqrt{\left[\sum_{i=1}^{M}\sum_{j=1}^{N}(I_1(i,j))^2\right] \cdot \left[\sum_{i=1}^{M}\sum_{j=1}^{N}(I_2(i+\Delta_i, j+\Delta_j))^2\right]}} \tag{1.4}$$

To help detect the almost inevitable occurrence of spurious vectors a vector-validation criteria termed ‘detectability’ and defined as the ratio between the highest peak in the correlation plane and the second highest peak, is widely used. See Keane and Adrian (1992). We note that this is also the definition of the correlation signal-to-noise ratio. Using ‘detectability’ as validation criteria is based on the observation that the correlation tables that give spurious vectors are characterized by low signal-to-noise ratios typically less than about 1.1, so a detectability value between 1.1 and 1.5 is typically specified. However, it has been argued by Westerweel (1994) and Hart (1998a), who also suggested spurious vector detection alternatives, that using the ‘detectability’ as a method of validating PIV data is not very efficient.[1]

---

[1] The efficiency of a validation procedure is defined as the fraction of displacement vectors that are correctly identified as either valid or spurious data.

The vector validation algorithms described by Westerweel (1994) involve comparing each displacement vector with nearby vectors with the hope that the spurious vectors will deviate largely from the neighbours. However, such validation methods, though popular, rely heavily on the assumption that the neighbours themselves are not spurious. In other words, the efficiency of such 'neighbour-validation' methods relies on obtaining an initial high valid data yield so that the neighbour vectors are mostly valid. These methods also miss vectors that are 'spurious' but are not vastly different from their neighbours.

Hart (1998a) also presented an algorithm aimed at reducing the occurrence of spurious vectors where adjacent correlation tables of overlapping interrogation windows are multiplied with the aim of reducing the height of the noise peaks relative to the displacement signal peak. This method however is not robust to high gradients, which is a major cause of spurious vectors in the first instance.

The popularity of vector validation methods that use detectability as validation criteria and neighbouring vector information may be attributed to the lack of robust alternatives.

Efficient methods to detect or minimize spurious vectors are still very much sought after and recently, the author (from Owodunni and Nickels (2003)) presented an algorithm that is aimed at reducing the occurrence of spurious vectors caused by large local velocity gradients. It was demonstrated that by performing the window correlations using a reconstructed version of the original images where the reconstructed particle image size can be arbitrarily specified, large velocity gradient effects can be reduced while maintaining the particle image size required for highest accuracy and low occurrence of particle image overlaps.

To further help detect spurious vectors, a validation technique was also designed which involves testing how stable (i.e. how sensitive) the correlation-peak coordinates are to slight changes in correlation parameters such as changes in the particle image intensity, their size or the reversing of the order of the correlating images. A displacement vector whose peak coordinates in the table deviate excessively from the previous location when changes are made to these correlation-governing parameters are tagged as invalid vectors. Further details of this advanced image processing algorithm are presented in Appendix F.

In PTV, the efficiency with which particle images can be isolated from the background pixels is important to the accuracy and valid data yield. Particle tracking algorithms capable of efficiently discerning between particle images and high intensity background pixels, detecting overlapping particle images have been described by several authors including Carosone et al. (1995); Guezennec et al. (1994); Maas et al. (1993); Okamoto et al. (1995); Stitou and Riethmuller (2001). Many of these Particle Tracking algorithms employ particle image intensity threshold techniques to efficiently separate the particle images from the background.

1.6.3 Post-processing accuracy

In PTV, the interpolation of the randomly located vectors onto a regular grid will also introduce some error in the measurements. The accuracy of schemes for remapping the scattered velocity vectors onto a regular grid have been presented and reviewed by Agui and Jimenez (1987); Cohn and Koochesfahani (2000); Spedding and Rignot (1993); Stuer and Blaser (2000). Common grid interpolation techniques include local least squares fit and inverse distance methods such as the Adaptive Gaussian Window (AGW) technique to velocity data. The least squares fitting technique involves evaluating the velocity on a regular grid by fitting polynomials to the neighbouring vectors and then evaluating the velocity from the fitted polynomials. The inverse distance method works by weighting the contribution of the neighbouring vectors according to their distance from the point where the velocity is sought.

We do not discuss the accuracy of these techniques here but we note that the accuracy in the remapping of the randomly spaced vectors and the accuracy in determining the vorticity generally improve with increased ‘valid’ data density.

Methods to extract vorticity information from the velocity measurements and the accuracy of the voriticity calculations are also described by Cohn and Koochesfahani (2000); Fouras and Soria (1998); Lourenco and Krothpalli (1995). Finite difference schemes, or schemes which involve the direct fitting of a least-squares polynomial to the velocity vectors followed by differentiating the polynomial have become standard methods to calculate the vorticity. [1]

1.7 Spatial resolution

The spatial resolution of the velocity measurements is another very important performance indicator. We define the spatial resolution as the density of the number of independent velocity measurements in a region of interest (ROI). The resolution of PIV measurements must be high enough to reveal the flow structures. It is especially important in turbulence measurements that the spatial resolution of the velocity measurements be high so that the smallest scales can be revealed at higher Reynolds numbers without confining the measurements to a smaller region in the flow field.

The spatial resolution of the measurements is greatly dependent on the camera's pixel resolution, the seeding density and particle image size.

For the same region of interest in the flow, a higher camera resolution allows smaller seeding particles to be used to achieve the same particle image size. Smaller particles in turn allow higher seeding densities, which in turn means that more particles can be tracked or smaller interrogation windows can be defined to give more vectors per unit area of the flow field.

While advances in camera technology will continue to improve measurement resolution, the velocity data density can be optimised using efficient interrogation algorithms.

Potentially, the highest spatial resolution measurement that can be achieved is when each individual particle image is tracked. To achieve high resolutions while maintaining high valid-data yield, several forms of super-resolution interrogation algorithms have been designed and are now commonly implemented.

Optimum spatial resolution measurements can now be achieved using hybrid PIV-PTV techniques, which combine PIV with particle tracking so that particles in high seeding density flows can be tracked. See Guezennec et al. (1994);

---

[1] Expression for vorticity $\omega = \nabla \times \vec{V} = \left(\frac{\partial V_z}{\partial y} - \frac{\partial V_y}{\partial z}\right)\vec{i} + \left(\frac{\partial V_x}{\partial z} - \frac{\partial V_z}{\partial x}\right)\vec{j} + \left(\frac{\partial V_y}{\partial x} - \frac{\partial V_x}{\partial y}\right)\vec{k}$. We note that for 2D velocity measurements, only the third component in the expression for vorticity can be evaluated since we only measure the x and y velocity components V_x and V_y respectively.

Guezennec and Kiritis (1990); Keane and Adrian (1995). The hybrid PIV-PTV algorithm is a two-stage process where, at the first stage, an estimate of the displacement of each particle is obtained and at the second stage, each particle is tracked, using the estimated displacement from the first stage to aid in the search for the matching particle image pair. The initial particle displacement estimate is obtained by standard correlation analysis where the displacements of clusters of particles are determined. The displacement of each individual particle image in the cluster is then determined by conducting a search centred on the cluster displacement location in the second image, for each particle's matching pair.

Another approach to super-resolution measurements are iterative PIV techniques described by Hart (1999) and Scarano and Riethmuller (1999, 2000). These methods serve to improve the spatial resolution of PIV measurements beyond the basic interrogation window. The key concept in these methods is as follows: First, initial correlations are performed on interrogation windows large enough to contain sufficient particle images for high valid-data-yield. These windows are then subdivided into smaller interrogation windows on which a second set of correlations are performed. The smaller windows are correlated, but now with corresponding interrogation windows in the second frame offset by an amount equal to the initial displacement obtained from their parent window correlations. The results from the previous correlations of the larger windows therefore serve to predict the location of the best cluster pattern match for the particle clusters in the smaller windows. The division of the windows followed by their interrogation can be performed iteratively, until window sizes of the same order as the particle images are reached therefore achieving super-resolution.

We note that achieving a high valid data yield at the earlier iterative stages is important to realising a high valid data yield at the last stage since the correlation at a subsequent stage, is performed over an area in the second image, centred at the predicted particle cluster location obtained from the previous stage. Therefore, at each iteration stage, the resulting displacement vectors will typically be validated and where spurious vectors are identified, an average of the neighbouring 'valid' vectors is generally assigned in place of the spurious vector.

The vectors from the PIV stage of the hybrid PIV-PTV algorithm are also validated to improve the pairing success rate at the tracking stage.

Figure 1.11 shows the velocity vector field obtained using a 2-stage iterative algorithm where the first stage correlations were performed on a starting window size of 74×74 pixels to obtain an estimate of the displacement of the cluster pattern as shown in Figure 1.4. At the second stage, the interrogation window size was reduced to 37×37 pixels and a search for the best cluster pattern match was implemented around the predicted locations from the first stage.

An example of super-resolution measurements obtained using the hybrid PIV-PTV technique is also shown in Figure 1.12 where the particles from the same singly-exposed pair of images are tracked using the vectors shown in Figure 1.11 to guide the search for the each particle's true pair.

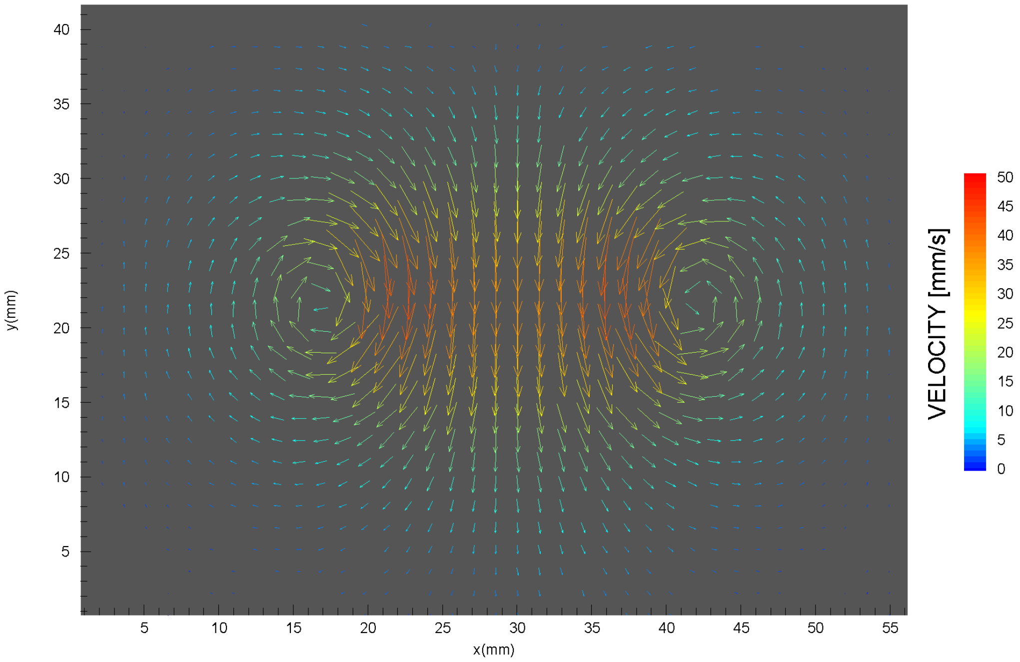

Figure 1.11: Velocity vector field of a laminar vortex ring flow obtained using a 2-stage iterative correlation. The first stage vectors are shown in Figure 1.4. Each vector represents the local velocity over a 1.5mm×1.5mm (i.e. 37×37 pixel interrogation) area of the recorded flow field. The achievable maximum number of valid vectors at the second stage is 1131.

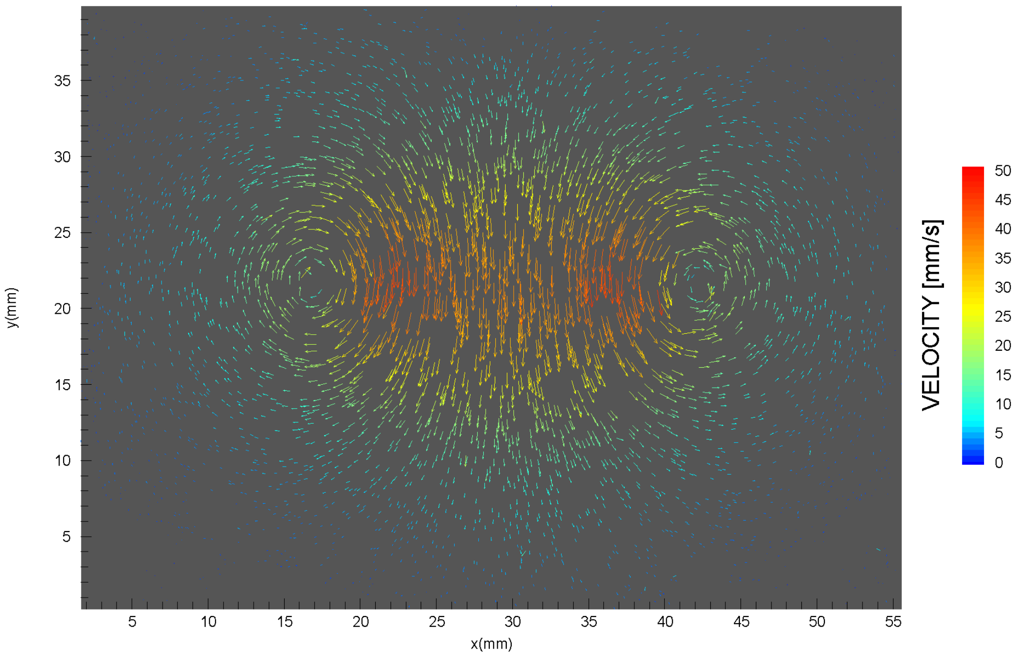

Figure 1.12: A sample PTV vector field for a laminar vortex ring flow obtained using the hybrid PIV-PTV approach. The particles were tracked using PIV data shown in Figure 1.11 to guide the search for the correct particle image pairs. Approximately 3700 particles were paired out of about 4500 particle images identified in the first image. Images analysed with the author's 2D PIV/PTV software package.

1.8 Extension of PIV/PTV to 3D measurements

The ever increasing camera resolution and the development of much improved algorithms have enabled valuable two-dimensional experimental data to be obtained using the classical arrangement. Many flows of interest are however three-dimensional in nature and a better understanding of their complex structures demands a complete specification of all three flow velocity components throughout the volume of interest.

Consequently, the extension of the classical planar PIV measurements to full three-dimensional measurements of flow velocity in a volume is now a major area of active research.

Some innovative hardware arrangements have been presented for measurements of all three components of velocity in a plane and throughout a volume. In this section, we discuss the operating principles of the most notable three-comoponent planar and volume PIV arrangements and we highlight their performance limitations and advantages.

1.8.1 Three-component, planar measurement systems

Although planar, three-velocity component measurements systems only provide velocity information in a plane of the flow field, these systems have proved useful. For completeness, we describe two widely used arrangements namely, the Stereoscopic PIV arrangement Prasad (2000b); Prasad and Adrian (1993) and the Dual-plane arrangement Raffel et al. (1995).

1.8.1.1 Dual plane PIV

In the dual-plane arrangement, planar measurements of all three velocity components are obtained using particle image information from two overlapping planes. One implementation of the dual plane technique is given by Schlicke and Greated (2002) where all three components of velocity were obtained by analysing three sequential images of the flow field taken at equal time intervals by alternating between two overlapping light sheet planar positions. The first and third images are recorded at the same illumination at time t and $t + 2\Delta t$

respectively, while the second image is recorded by an overlapping light sheet at time $t + \Delta t$. Just as in the classical 2D processing, the images are divided into interrogation windows over which the particle image cluster displacements will be calculated. Windows from the first image are then correlated with the corresponding windows from the second image to give the cluster in-plane displacement as revealed by the highest peak in the correlation table whose normalized height is P_{12}. Similarly, the windows from the second image are correlated with their corresponding windows in the third image to give an in-plane cluster displacement and a corresponding correlation-peak height P_{23} in the normalized correlation plane.

Since the height of the peak in the normalized correlation plane is a measure of the quality of the correlations, therefore if P_{12} is greater than P_{23}, then the particles are likely to be moving in the direction of the light sheet displacement. The out-of-plane cluster displacement is then estimated from an expression relating the difference between these peak heights and the light sheet thickness and intensity profile. By assuming that the laser beam is Gaussian in shape and the loss of correlation due to in-plane motion of particles out of the interrogation windows is the same for both sets of correlations, Schlicke and Greated (2002)) gave an expression for the-out-of-plane displacement D_z as:

$$D_z = \frac{\Delta z^2 \left(\ln P_{12} - \ln P_{23}\right)}{8\left(z_2 - z_1\right)}$$

where Δz is the width of the light sheet and z_1 and z_2 correspond to the light sheet position in the first and second images respectively.

Another dual-plane implementation was presented by Brucker (1996) where two sets of images are recorded at a predetermined time intervals using two overlapping light sheets of different colours to illuminate the flow at the same time. An expression based on the same parameters as those described above was given for the out-of-plane displacement.

Cenedese and Paglialunga (1989) also used two overlapping light-sheets of different colour to track the out-of-plane path of a particle, taking into account the change in the colour of the particle images along the path of the illuminated sheet.

1.8.1.2 Stereoscopic 2D-PIV

The 2D Stereoscopic PIV arrangement solves the problem of perspective error associated with single CCD camera systems and this technique has become popular for the measurement of the three velocity components in an illuminating laser sheet plane. In the stereoscopic approach, images of the particles in the illuminated plane of the flow are simultaneously recorded from two different angles using two cameras. Interrogation areas defined in real space are mapped onto each view. 2D correlations are subsequently performed on the images from each view to give the projected planar (i.e. x- and y-) velocity components for each view. All three components of the flow velocity vectors are then recovered by performing view reconstruction using corresponding window-segment displacements in the two views.

1.8.2 Three-component, volume measurement systems

Systems which allow measurements of all 3 velocity components at several points within a volume, thus, extending the spatial range of measurements to the third dimension, include the multi-camera 3D PTV arrangement, the Holographic PIV technique, the scanning PIV and the Defocusing PIV system.

1.8.2.1 Multi-camera 3D PTV

The multi-camera three-dimensional particle tracking arrangement is based on imaging the same region in the flow field from multiple views. This technique allows tracer particles to be tracked in a three-dimensional observation volume and has proved very useful in the Lagrangian study of fluid motion. See Guezennec et al. (1994); Maas et al. (1993); Malik et al. (1993); Soloff et al. (1997); Virant and Dracos (1997).

The space coordinates of the particles are determined by triangulation. Correspondence between particle images recorded in the multiple images are established and the 3D particle coordinates are determined.

The shared experience among researchers is that the particle concentration must be low to prevent several particle image overlaps when the particles are projected from 3D space onto camera image planes. This requirement limits the number of particles that could be imaged to a few thousands - typically in the region of about 1000 to 3000 vectors. For example in the experiments of Maas et al. (1993), about 1000 particles were successfully identified and tracked within an observation volume of $200\times160\times50\text{mm}^3$ with three 750×600 pixel, camera arrangement.

The spatial resolution offered by stereoscopic 3D PTV is however, typically too low to enable the macro and micro scales that exist in turbulence to be investigated at higher Reynolds number unless measurements are confined to a small region of the flow field. Further difficulties associated with the technique include optical access requirement for three or more cameras imaging the flow from different angles. The alignment of the cameras and the calibration procedure to obtain the correspondence between the different views could also be time consuming.

1.8.2.2 Holographic PIV

A more complicated technique for full 3D measurements is the Holographic PIV (HPIV) arrangement. A major advantage of the Holographic technique is the high spatial resolution of velocity measurements it offers.

Holographic PIV utilizes the holography technique to record two separate holograms of the 3D particle field flow onto the same holographic film. The interference pattern formed on the film by the scattered coherent light from the tracer particles and a reference light are recorded on a photographic film. From the 3D reconstruction of the flow field, slices (planes) are recorded and interrogated to give the planar velocities, which are subsequently combined to give the local three-component velocities throughout the volume of the flow.

Holographic PIV measurements have been reported to yield vectors in the order of about 10^6 on a regular grid and tens of thousands of particles have been reportedly tracked in volume sizes in the order of $50\times50\times50\text{mm}^3$, (see Barnhart et al. (1994); Meng and Hussain (1991, 1995a,b); Sheng et al. (2003); Zhang et al. (1997)). However, holographic PIV measurements is, at the moment, limited to single realizations due to the time needed to develop the holographic film and the subsequent reconstruction and interrogation of the image. Hence, Holographic PIV, for now, does not realistically lend itself to flow investigations that require several images of the flow to be recorded such as requirements for stable turbulence statistics. Furthermore the technique involves a highly complex and expensive set-up for recording, hologram reconstruction and data evaluation, which typically takes several hours to complete. Also, the feedback on the experimental set-up is slow since holograms need to be developed before problems with the experimental arrangements are detected.

A digital counterpart to holographic film recording is currently being developed to address the difficulties and limitations associated with film-based holography. However, early indications point to low spatial resolution. See Meng et al. (2004).

1.8.2.3 Scanning PIV

Another hardware arrangement for three-component velocity measurements throughout a volume is the Scanning Particle Image Velocimetry (SPIV) described by Brucker (1995, 1997). In this arrangement, the flow is sampled in depth by a rapid scanning light sheet and the three dimensional vector field is obtained by performing 2D spatial cross-correlation of particle images in the acquired 2D images and then using the continuity concept, to obtain the out-of-plane component of the velocity.

The scanning PIV technique requires very high frame rate cameras and fast light scanning to scan the whole flow field in a period of time shorter than the temporal scales of the flow. A very short illumination time interval Δt, which may be required to achieve small relative particle image displacements may also demand greater hardware speed or limit the depth of the volume that could be adequately measured. The range of applications of the scanning PIV technique is therefore dependent on available camera speed and illumination frequency. We also note that optical access and synchronization of the scanning system can be complex.

1.8.2.4 Defocusing DPIV

A relatively recent addition to three-component, whole volume velocity measurement techniques is the defocusing PIV. See Pereira and Gharib (2002); Pereira et al. (2000); Willert and Gharib (1992). This method uses out-of-focus imaging via a mask with three triangularly arranged apertures to extract the position of particle images in the third dimension i.e. along the camera optical axis. Here the distances between the resulting triangular sets of defocused particle images reveal the depth positions of the particles. An advantage of this arrangement is that the three apertures share a common optical axis so that the flow field can be imaged from a single direction. This low optical access requirement therefore offers the possibility of higher resolution imaging through relatively higher seeding compared to 3D stereoscopic PTV. However, the requirement that three defocused images of each particle must be correctly identified as belonging to the same particle significantly compromises the resolution of the method. The aggregate number of overlapping particle images also increases matching ambiguities.

After much modification to the original configuration presented in their 1992 paper, the current state of the system is now reported to give between 2,000 to 10,000 velocity vectors.

1.9 Project motivation, objective and approach

Many of the 3D systems described here involve elaborate and complicated hardware arrangement, which in turn is accompanied by high equipment cost and several experimental complexities. This, combined with their critically low performance in one or more important aspect, has hindered their widespread use. A system which is capable of providing three-component velocity information over a flow volume and that still measures up well across all aspects of performance is therefore still very much sought after. The objectives of this research are therefore as follows:

- To design and develop an economical and simple, high-performance three-dimensional PIV system.

- To demonstrate the system's potential to give highly accurate, high resolution three-dimensional flow velocity measurements.
- To identify and evaluate the experimental parameters and the features of the system that are critical to its performance.
- Design and specify high performance measurement procedure for the system.

To address this need for a simple high-performance alternative to the existing three-dimensional techniques we present in this dissertation, the development of a system which uses a single 3CCD colour camera and multi-coloured illumination to identify and track the locations of the particles suspended in the flow. The key concept here is to use colour information to identify the 'depth' of the particles in the flow (i.e. their distance from the camera). Because the colour of the illumination varies along the depth, the depth (i.e. z-coordinate) of each particle is revealed by the colour of light it scatters. The other two particle coordinates (i.e. the x- and y-coordinates) are obtained from the particles' positions on the image. Hence, the position of each particle in three-dimensional space can be identified. By comparing two flow images a short time apart, we can then calculate the all three velocity components of each particle within the volume of interest.

The proposed system may be viewed as an extension of Van Meel and Vermij (1961) streak photography system where tracers were illuminated with discrete multicoloured light along the camera's optical axis thus producing multicoloured streak images of the tracer particles. However this streak photography arrangement of Van Meel and Vermij (1961) gives low spatial resolution. Also, the accuracy (and dynamic range) is low due to the use of discrete bands of coloured light.

We confirm the proposed system's potential to yield accurate, high-resolution 3D measurements by applying it to an axisymmetric laminar vortex ring flow.

1.10 Overview of dissertation

In this chapter, the basic principles of the PIV technique were presented. We explained the measures of PIV performance and discussed how the performance is influenced by the experimental parameters, interrogation algorithms and hardware arrangements. We described some notable hardware arrangements that have allowed all three components of velocity at many points in a plane and a volume to be measured. We highlighted the limitations these innovative 3D systems. Lastly, we introduced a potential high-performance alternative approach to three-component-volume measurements where the spatial positions of the particles are determined using continuously varying coloured light and a single camera.

In Chapter 2 the construction of the colour 3D PIV system is described. The characteristics of the system components, how they fit together for image acquisitions and their influence on the system's performance are presented. The experimental arrangements for the laminar vortex ring and uniform motion experiments, which were used to evaluate the potential of the system, are also presented.

In Chapter 3 we give some general background information on how the images acquired with the system can be efficiently interrogated to give three-dimensional velocity information.

In Chapter 4 we explain how the colour flow images acquired in this project were analysed.

We present in Chapter 5 the measurements of the simulated uniform flow and laminar vortex ring flow made with the new system.

Lastly in Chapter 6 we summarize our research findings and make recommendations to further improve the performance of the colour PIV system.

Chapter 2

The 3D colour PIV-PTV system: Image acquisition and experimental arrangement

2.1 Introduction

In this chapter, we describe the hardware components that make up the colour 3D PIV-PTV system, their arrangement and how they fit together to enable the acquisition of particle images of fluid flow experiments. These are analysed to give velocity information about the flows under investigation. Two sets of flow cases were investigated in this work and the arrangement of the apparatus for each set of experiments is described. The first experiment was a simulated uniform flow experiment, performed to determine the accuracy of the 3D PIV-PTV system and the second experiment was the vortex ring experiment, which was carried out to further confirm the feasibility of the colour PIV system as a practical tool for three-dimensional velocity measurements. The author designed the system and the non-standard-issue apparatus were manufactured in the Cambridge University Engineering Department workshops.

We begin by first describing the system hardware arrangement. The operational characteristics of each component are then described followed by a description of how we acquire the colour images of the particle-laden flow field.

We proceed to explain how the particle image colours are resolved so that they can be related back to their positions in space. We explain how the characteristics of the flash and camera affect the quality of the acquired images and the choice of illumination region and depth. We also estimate the uncertainty introduced by the camera and flash in determining the particle position from the colour.

We then describe how the region of interest in the flow field is calibrated to obtain a database of colour versus position information that was used to determine the spatial coordinates of each particle image. Finally, we describe the experimental arrangement for the uniform flow and laminar vortex ring flow including the choice of seeding particles and working fluid.

2.2 The 3D colour PIV-PTV system arrangement

An illustrative drawing of the optical arrangement of the 3D PIV-PTV system for both measurements of uniform flow and vortex ring flow is shown in Figure 2.1.

The system's optical arrangement consists of a spark flash source (Pal flash) capable of producing two independently pulsed, short-duration, high-intensity, parallel beams of 'near-white' light whose spectral composition covers the visible spectral range of the electromagnetic spectrum (wavelengths between 400–700nm). An opaque disk with a 50mm square aperture is secured in front of the Pal flash to block out the unused part of the light beam. The light from the flash travels through a system of lenses and a filter arranged along the optical path of the beam. The light beam is focused onto the 200mm×25mm linear dichroic filter using a 50mm×50mm cylindrical plano-convex lens of 200mm focal length. The central wavelength of the filter varies linearly along its length thus only allowing wavelengths within the local band pass wavelength to be transmitted through thereby transforming the incoming light into an output beam of 'rainbow' coloured light whose local wavelength varies linearly along its depth. This multicoloured diverging beam is then collimated with an identical cylindrical plano-convex lens to produce a 50mm thick parallel beam of multicoloured light.

Arranged some distance beyond the second cylindrical lens is the flow to be investigated. The multicoloured beam illuminates the region of the interest within the seeded flow field. A 10 bit, 3-CCD multi-spectral digital camera, positioned approximately 1m away from the light beam and with its optical axis normal to the beam direction is used to record images of the micro sized tracer particles suspended in the flow. The camera is fitted with an f2.8/105mm Micro Nikkor lens to give a field of view of 57mm×43mm and a depth of field of about 15mm at this distance at an aperture setting of f-2.8. The position of each particle along the optical axis of the camera is revealed by the colour of the light it reflects. Images of the flow field are recorded at predetermined time intervals and the motion of the particle images between recordings is determined to obtain the velocity field information about the flow.

A schematic representation of the vortex ring experimental set-up is shown in Figure 2.1 and pictures of the vortex ring experimental arrangement are shown in Figure 2.2. The experimental arrangement for the uniform flow experiment is shown in Figure 2.3 where the vortex ring generating equipment is now replaced by the uniform flow block positioned on top of a small platform jack. The vortex ring generator equipment is mounted on a motorized 2-axis translation stage to allow for the precise positioning of the region of interest inside the beam. In the case of the uniform flow experiment, the uniform flow block and platform are mounted on the translation stages in order to achieve an accurate prescribed displacement of the block along the x- and z-directions, therefore, simulating uniform motion of the tracer particles 'frozen' inside the block. The whole experimental facility is housed in an enclosure to prevent unwanted illumination of the flow field by external sources of light e.g. ambient light.

The system's image-recording mode is the 'double frame–single exposure' mode where the flow field is illuminated twice in quick succession with each illumination registered on a separate frame of the camera output. The recorded images are transferred from the camera to an on-board frame grabber and stored on a computer hard disk. The digital images are subsequently interrogated with a custom software package, written by the author, to give three-dimensional velocity field information about the illuminated volume of the flow.

The system also doubles as a planar 2D system when the colour filter is removed and the thickness of the light beam is reduced.

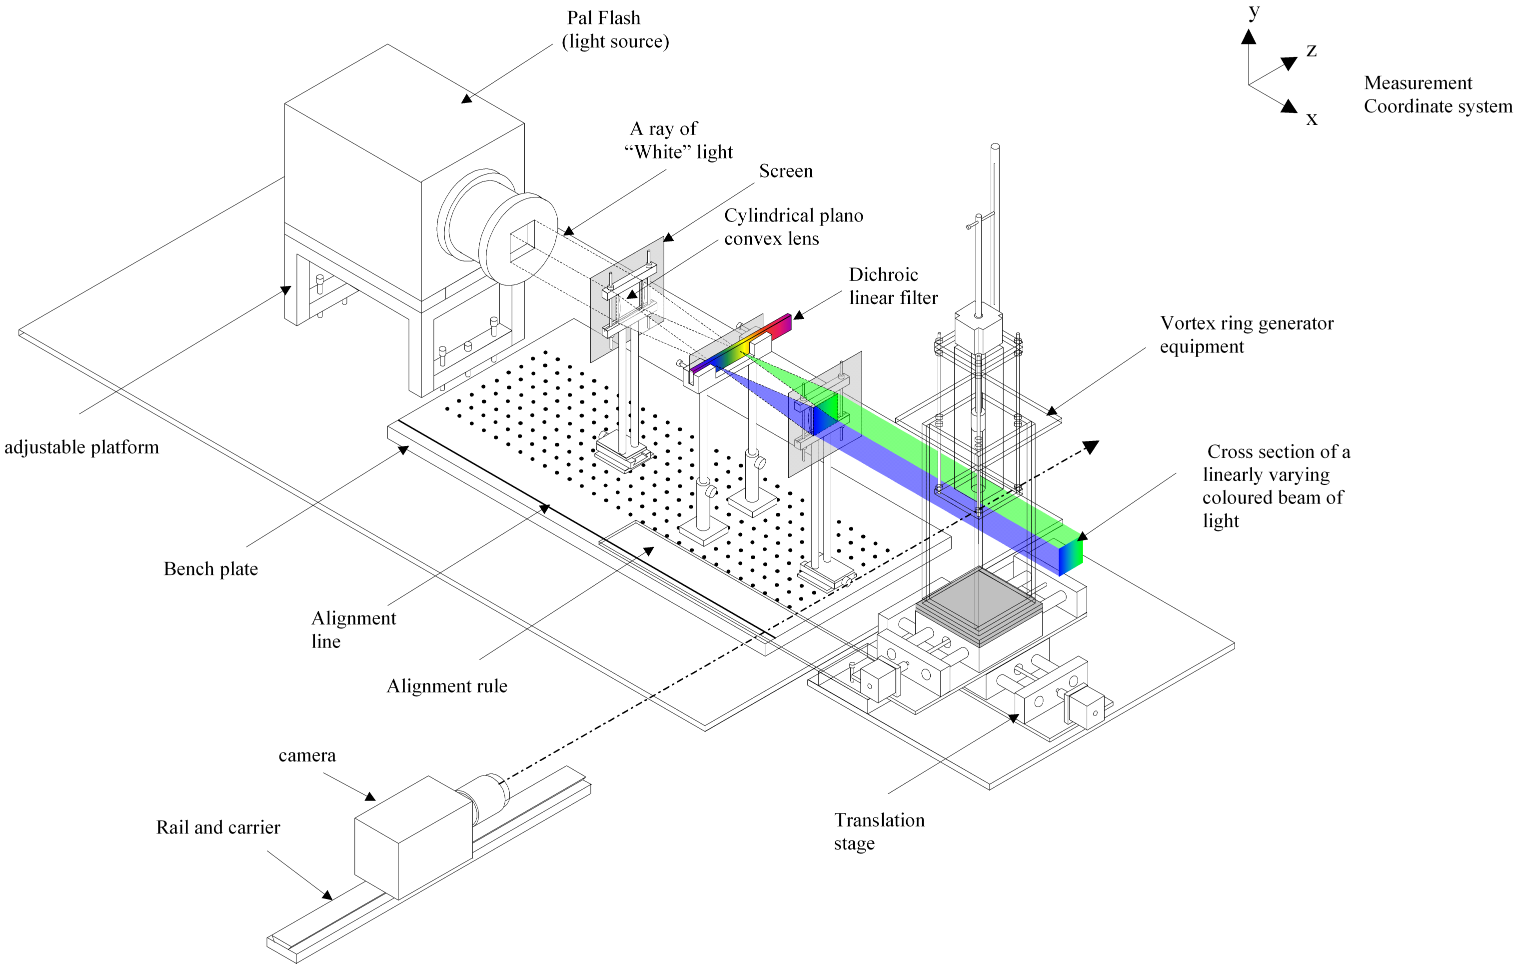

Figure 2.1: Schematic layout of the colour 3D PIV-PTV system.

2.3 System components

2.3.1 The light source

The source of illumination used in the experiments is the PalFlash 500 (from Pulse Photonics Ltd.). The Pal Flash 500 operates using open spark technology. Energy is stored by the high-energy discharge capacitors in the unit and then released as a very short duration, high intensity spark via a set of electrodes. The Pal Flash is capable of producing two independent flashes at very short time intervals apart. However, the recharge time of the capacitors is about 5 seconds.

The internal optics of the Pal Flash unit are shown in Figure 2.4. The unit consists of two pairs of electrodes arranged along the optical axis of the flash. Each pair of electrodes provides a single flash of between 100 nanoseconds and one micro second duration for an electrical output of 6 Joules of energy, which is converted to light. The top and bottom electrode of each pair are adjusted to give a gap size between 5 to 10mm.

The 35mm diameter lens fitted in front of the back electrodes focuses the back electrode spark onto the spark of the front electrodes and is in turn, focused onto a 110mm diameter condenser lens mounted in front of the front pair of electrodes. The focus of the 110mm lens is adjusted by releasing a pinch handle on the lens holder and sliding the lens along the adjustable slot to give an even collimated light beam output of both gaps. The Pal flash is also fitted with a focus light, which is installed behind the back electrode with another 35mm lens in front to focus the light from the lamp onto the back spark. The spark is composed of all the wavelengths in the visible spectrum, with the intensity of the wavelengths around the blue region significantly higher than other wavelengths.

In order to improve the spatial stability of the spark and also prevent random spontaneous trigger of the flash, argon gas was fed through tubes running from the back panel of the flash unit into the gap between each electrode pair. The gap size setting and gas flow rate required to achieve stability are dependent on the ambient conditions and is therefore a matter of trial and error and familiarity with the equipment. For this work, the gap size of the front electrodes was set to 10mm while that for the back electrodes was set to 6mm and the gas flow was

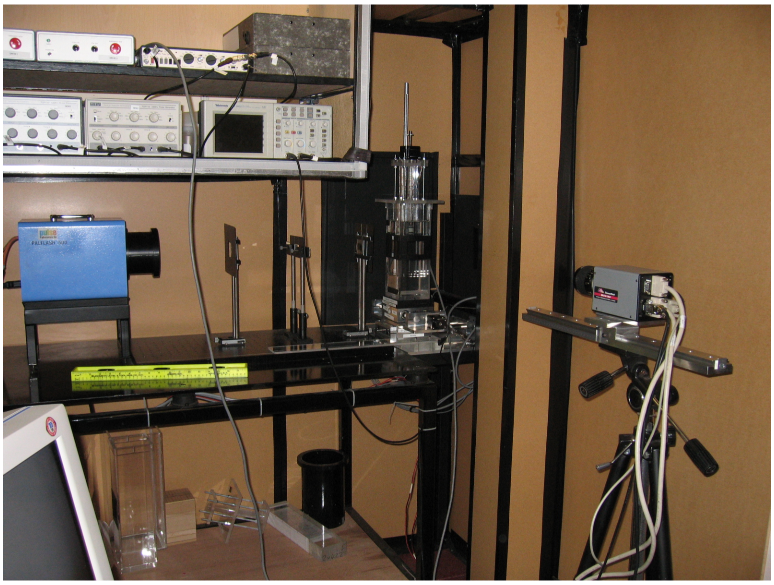

(a) Picture of the vortex ring experimental arrangement.

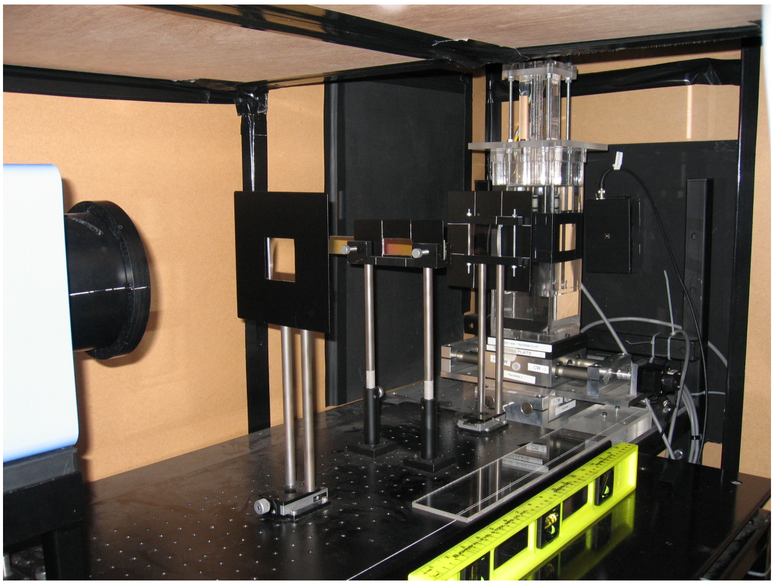

(b) A close-up view of the vortex ring generator and the lens-filter arrangement.

Figure 2.2: Picture of the vortex ring experimental arrangement. Pictured are the optics hardware, vortex ring generator, the pulse generators and the drives and traverses.

Figure 2.3: Picture of the experimental arrangment for the uniform flow simulation. Shown sitting on the traverse system is the platform jack on which sits the uniform flow block. The block contains seeding particles suspended-frozen inside it.

subsequently adjusted until the flash did not trigger spontaneously, but readily when intended. Triggering of the flashes was achieved with two 5-volt positive pulses via the BNC connectors located at the back of the unit. The Pal Flash unit was mounted on an adjustable platform for easy alignment of the flash with the lens-filter arrangement.

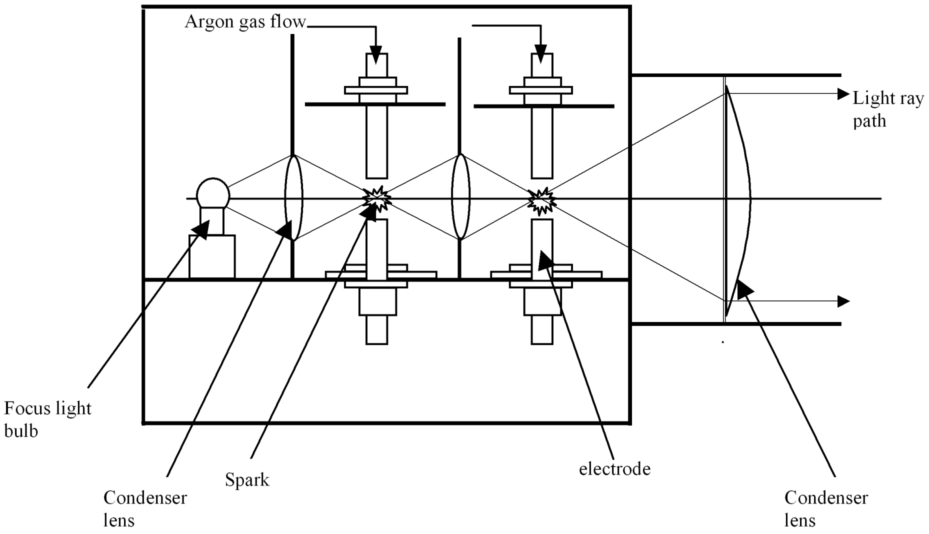

Figure 2.4: Optical layout of the Pal Flash 500.

2.3.2 The cylindrical lenses

A pair of 50×50mm, Valumax® BK7 plano-convex cylindrical lenses of focal length of 200mm (from Newport Corporation) mounted on bar-type lens holders were used to focus the light beam from the Pal Flash onto the filter and then re-collimate the resulting diverging multi-coloured light beam as shown in Figure 2.1. The cylindrical lenses are held in position by cylindrical lens holders mounted on miniature single-axis metric stages to allow precise adjustments of their distances from the filter. The long focal length of the lens was selected to allow the incoming light to converge on the face of the filter with a small angle of incidence as required for good filter performance.

2.3.3 The filter

A Schott Linear Variable Interference filter (Veril S 200 from H.V.Skan Ltd.) was used in the experiment to convert the white light beam to one whose colour varies linearly along the depth of the beam, see Figure 2.1. The filter is 200mm long, 25mm wide and 5mm thick. It is a band-pass interference filter whose central wavelength λ_m varies linearly over its length. The filter uses the interference effect to transmit wavelengths within a specific pass-band of the electromagnetic radiation, while wavelengths outside the pass-band are rejected through reflection. The variation of the centre wavelength along the length of the filter is shown in Figure 2.5. The spectral range of the filter is 400 – 700nm (i.e. blue to red) spanning a spectral length of 125mm. The maximum spectral transmittance τ_{max} and Half Width (HW) at central wavelengths of 450nm, 550nm and 650nm are:
450nm: $\tau_{max} = 40.7\%$ HW (nm) = 15
550nm: $\tau_{max} = 51.1\%$ HW (nm) = 14
650nm: $\tau_{max} = 55.3\%$ HW (nm) = 14

2.3.4 The camera

Mounted on a tripod and rail carrier assembly is a DuncanTech MS3100 3-CCD digital camera (from DuncanTech Inc.) used for the image capture in the experiments. This 10 bit, ½ inch interline transfer 3-CCD digital multi-spectral camera uses a colour separating prism and three Charged Coupled Device (CCD) imaging sensors to simultaneously acquire images in 3 spectral bands (RGB), within the 400-800 nm sensitivity of the sensors, through a common aperture. The pixel resolution of each CCD is 1392(H)×1040(V) pixels giving a total of 4.3 Million pixels of data for the camera. Each CCD pixel measures 4.65×4.65 microns. The camera's data transfer rate is 14.318 MHz max (i.e. 14.318 million pixel data per second) to give a free run mode frame rate of 7.5 frames/second.

A functional diagram of the Duncan Tech MS3100 camera is shown in Figure 2.6 below. Broadband light from the image target enters the camera through the lens. The prism optics divides the light based on wavelength such that a different spectral band exits the prism at each of the three exit faces. The

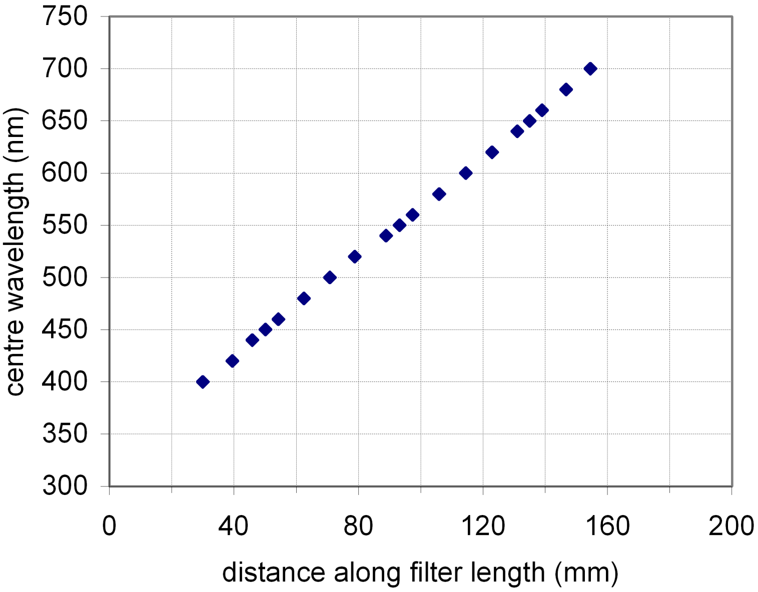

Figure 2.5: Variation of the filter's central wavelength along its length (data courtesy of Schott).

range of wavelengths included in each band is a function of the coatings on the faces of the prism. Optical trim filters that are placed between the exit plane of the prism and the array further narrow the spectral band of light that arrives at each of the three imaging sensors. The output signal of each CCD array is conditioned and digitised to a 10-bit digital value. The camera can output up to 32 bits of parallel pixel data that can be configured for either 8-bit or 10-bit resolution.

When configured for 8-bit operation, the lower two bits of data are dropped and the camera can output up to four 'sets' or 'taps' of image data for a total of 32 bits. In 10-bit mode, the camera can output up to three 'sets' or 'taps' of data for a total of 30 bits. For the experiments described in this dissertation, the camera was configured to operate in 10-bit mode to give 1024 grey level images (i.e. the pixel intensities can be any value between 0 and 1023 counts.)

The image data is output as digital pixel values from the CameraLink digital output connector on the rear panel of the camera. Camera configuration and control options such as its operating resolution, and the gain settings, and exposure times for each channel are set independently with the camera's RS-232 communications interface.

The camera's spectral response to the linearly varying coloured light beam is shown in Figure 2.7. The plot was obtained from image recordings of an air bubble, (approximate size 0.8mm), trapped in a block of water-clear resin, as it is traversed along the camera's axis, through the thickness of the 'rainbow' colour light beam.

2.3.5 The camera lens

The DuncanTech camera is fitted with a 105mm f2.8 micro Nikkor Lens to give angles of view of 3.8° and 2.9° along the x and y directions respectively as determined from Figure 2.8. The lens aperture is set at f-2.8. The camera is positioned approximately 1m from a reference plane in the illuminated volume to give a field of view of 57.3mm in the x-direction and 42.7mm in the y-direction at a reference pixel scale of 0.0410mm/pixel. The depth of field at this camera distance and f-number setting is estimated to be 15mm. We note that this value for the depth of field is based on the author's perception of the depth within which all particle images appear focused.

2.4 Image acquisition

The connection of the image acquisition hardware is shown in Figure 2.10 and the time sequence of events to capture a pair of singly exposed images is illustrated in Figure 2.11.

2.4.1 Camera's image capture mode and frame straddling

To achieve rapid successive illuminations of camera frames for high temporal image resolution and maintain control over the timings of the image capture, the camera's image acquisition mode was set to an externally controlled mode and then used with the Frame Straddling technique.

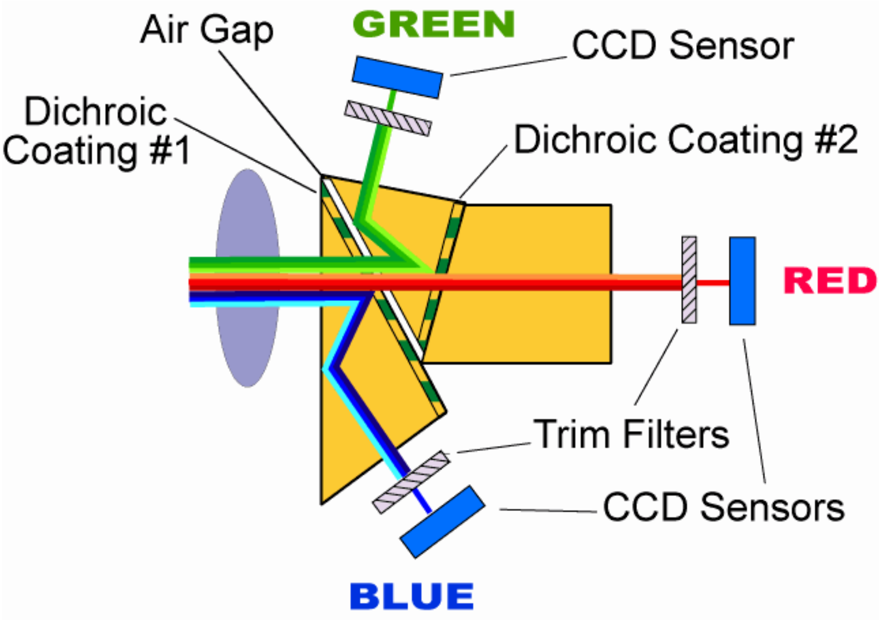

Figure 2.6: The camera's internal optics. Incoming light is split into 3 bands and each band is registered on a different CCD array. Image courtesy of DuncanTech Inc.

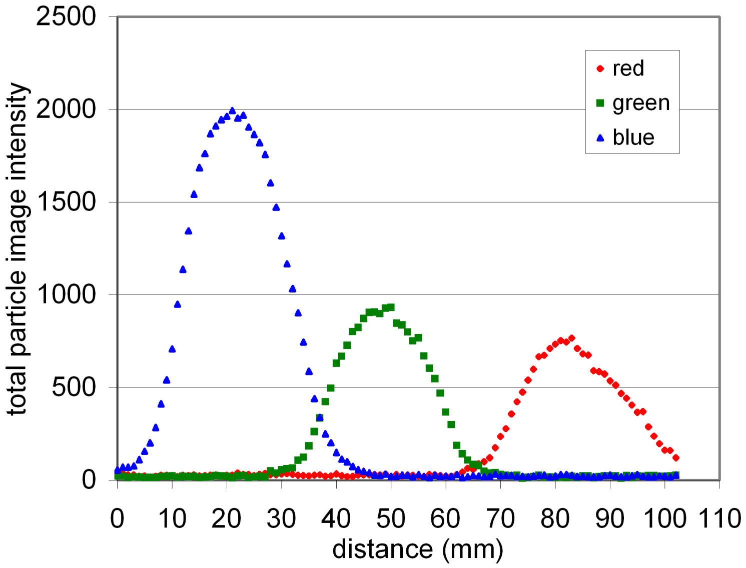

Figure 2.7: The camera's spectral response to the linearly varying colour light.

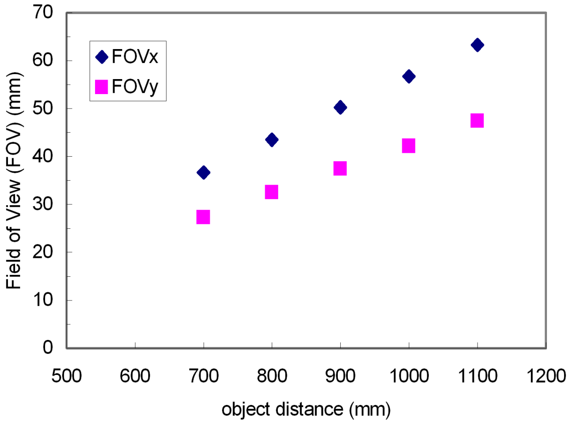

Figure 2.8: Variation of the camera's field-of-view with object distance for the 105mm micro Nikkor lens.

2.4.2 The camera's edge-control mode

The Edge-Control mode is one of the camera's externally triggered mode of image capture and was selected as the camera's mode of operation. This mode allows the precise control of the start of image acquisition and the frame exposure time. The Edge-Control mode is also well suited for use with the frame straddling technique since the readout time of a previous frame is overlapped with the exposure time of the next frame. Figure 2.9 illustrates the edge-control mode of operation for the camera. Precise control of the start of each image acquisition is achieved by sending a 5 volt TTL signal as an external trigger input to the camera. The active edge of the trigger signal initiates the readout of the currently exposed frame's data and starts the exposure of the next frame. The exposure time of the frames is equal to the time between two successive leading edges of the external trigger signals and must be at least equal to the camera's frame readout period of $\frac{1}{7.5}$ seconds (i.e. approximately 133ms) if incomplete image transfer is to be avoided.

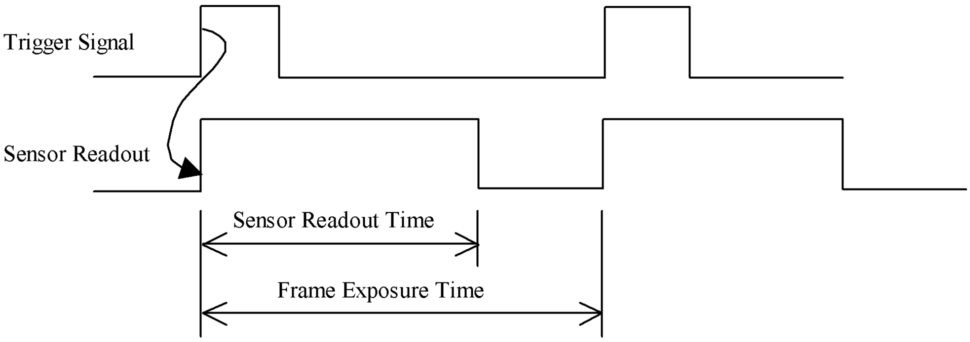

Figure 2.9: The camera's externally triggered Edge Control image capture mode.

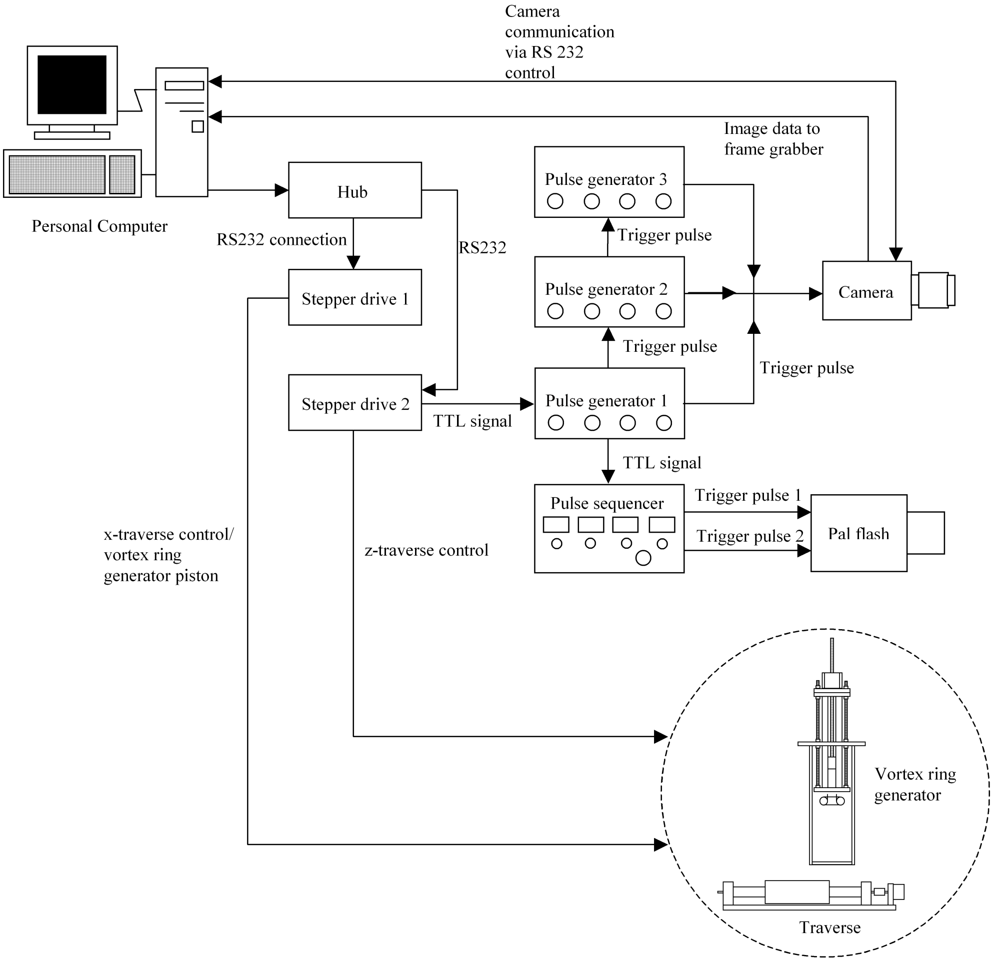

Figure 2.10: Connection diagram of the image acquisition system.

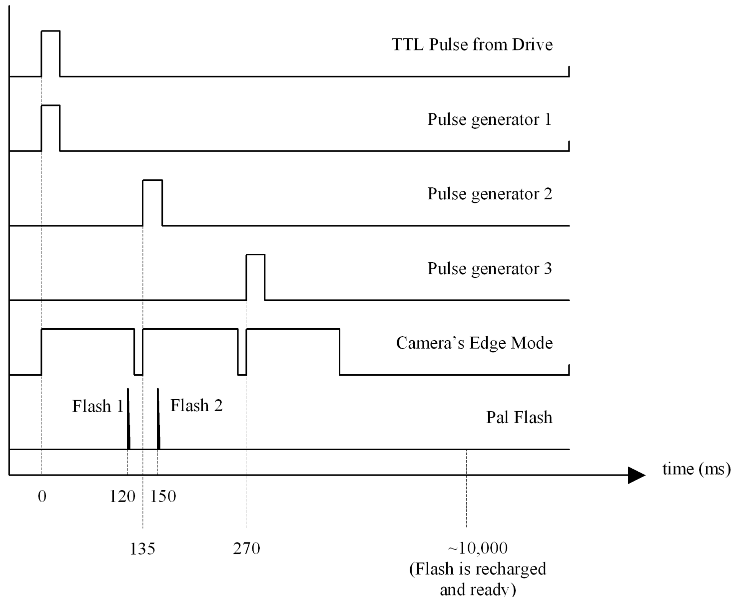

Figure 2.11: Timing diagram showing the sequence of events involved in the capture of two singly-exposed images as the camera's edge control mode is combined with the frame straddling technique.

2.4.3 Frame straddling

Figure 2.11 illustrates how the frame straddling technique is combined with the Edge mode of the camera. The frame straddling technique allows us to illuminate the flow field twice in quick succession and still be able to record each illumination on separate camera frames.

Frame straddling has proved useful in high speed flow measurements where the illumination time between the two images need to be a lot smaller than the available camera frame rate. The technique involves illuminating the flow field near the end of the exposure of one frame and illuminating the flow field again near the start of the exposure of the next camera frame to give an illumination time interval Δt of 30ms as illustrated in Figure 2.11.

For a continuous image capture, the cycle time is determined by the recharge time of the Pal Flash. If the cycle time is split between the frame exposure times, the influence of camera dark current noise on the recording will be overwhelming. To prevent long frame exposures, the exposure times for the pair of illuminated frames in each cycle was set to 135ms and a third trigger pulse was used to signal the end of the exposure of the second illuminated frame as shown in Figure 2.11. The use of 3 triggers to acquire two illuminated frames however means that 3 frames were captured per cycle i.e. two illuminated frames and one extra frame which is exposed for the remaining duration required to recharge the flash. These extra frames, which ‘see’ no flash illumination, were not saved to the hard drive.

2.4.4 Synchronization of image acquisition hardware

The flash timings are controlled by a Pulse Sequencer while three Pulse Generators control the exposure times of the camera frames. A Digital Frame grabber board installed in a Personal Computer captures and stores the image data on the computer’s hard disk.

Precise positioning and control of the 2-axis translation stage and the vortex ring generator piston are achieved with a multi-axis motion control system consisting of two programmable micro-stepping motor drivers and a Motion Control Network Hub.

The hub connects to the PC for programming and the drives are connected to the piston actuator and the traverse motors. The drive are powered by two 30 VDC power supplies to control the motion of the piston and the traverses.

The 3540i stepper drives and Hub444 are from Applied Motion Products® while the frame grabber is Active Silicon’s 64 bit-66MHz PCI Phoenix DIG48 PCI CameraLink Digital Frame grabber whose maximum image acquisition rate (pixel clock) is 60MHz. The computer is a 1.67 GHz AMD Athlon MP 2000+ PC with a 512Mb of RAM and a storage capacity of 56 Gigabytes. The Pulse Sequencer is the PAL SEQ 401 from Photonics Analysis Ltd while the three Pulse generators - two TGP 10MHz Pulse Generators and one TG105 5MHz Pulse generator - are from Thurlby Thandar Instruments.

The image capture sequence is programmed on the Applied Motion SiNet Hub Programmer interface installed on the PC. The required motion of the traverses and vortex ring piston plus a TTL trigger signal to initiate image recording are programmed and the program is subsequently downloaded to the drives via the Hub.

The acquisition of a pair of singly-exposed images is initiated when, at the prescribed time, the programmed TTL signal is sent from a designated drive to one of the Pulse generators.

The TTL signal is received by the first Pulse generator, which immediately and simultaneously sends a trigger pulse to the Pulse Sequencer, the camera and the second Pulse Generator. The second pulse generator, in turn sends a delayed pulse to the camera while simultaneously triggering the third pulse generator. The third Pulse generator, on receiving the pulse also sends another delayed trigger pulse to the camera. Altogether, the camera receives three sequential trigger pulses to capture a set of three image frames.

The trigger signal received by the Pulse Sequencer from the first Pulse Generator coincides with the start of the exposure of the first camera frame to be exposed to the flash illumination. The Pulse Sequencer can produce four 5-40 VDC pulses in sequence at predetermined delay times on receiving the input trigger pulse. The Pulse Sequencer triggers the flashes by sending a pair of pulses to the Flash unit.

The time between the pulses sent from the Pulse Sequencer to the Pal Flash unit is set so that the first flash is triggered towards the end of the exposure of the first camera frame and the second flash is triggered near the beginning of the exposure of the next camera frame to give two singly-exposed images separated by an illumination time interval Δt which is the difference between the delay times set for the two flash trigger pulses.

The raw 10-bit image data for each camera frame is output from the back panel of the camera via CameraLink connections to the frame grabber where it is reformatted in hardware as 48-bit colour data in RGB ordering before it is saved to disk. The size of each captured image on disk was 8,686,080 bytes in 48bit RGB format, so each set of 2 captured illuminated images required approximately 18 Megabytes of hard disk storage space.

2.5 Definition of particle image colour

When images of the region of interest in the seeded flow field are acquired as described, the particle images appear coloured according to their depth positions in the flow field. The particle image colours must therefore be defined so that the position of each particle can be decoded from its colour.

Figure 2.12 shows the 5×5 pixel area of a particle image as recorded on the blue and green CCD channels. It can be observed from the particle image intensity profiles that:

- The shapes of the image profile of the same particle registered on the CCDs are different.
- The coordinates of the maximum pixel intensities of the particle image on each CCD are different.

Because the particle image intensity profiles are different, the ratios of red, green and blue intensities for a particle image pixel will vary from pixel to pixel. This is still true when the peak intensities recorded all three CCDs are aligned. The variation in the colour ratios of the pixels and the fact that the peak intensities from all 3 CCDs will not necessarily be at the same image coordinates is reportedly[1] due to the camera's internal optical arrangement.

We therefore define the colour of a particle in terms of the ratio of each of the particle's image intensity recorded on the 3 colour channels to the sum of all three: The particle image intensities on each CCD channel are summed over a specified pixel area, centred at the maximum-intensity pixel location. The ratio of each intensity to the sum of all three, is then determined to give the particle's colour in RGB order, i.e. (red, green, blue) where:

$$\begin{aligned}
\text{red} &= \frac{I_{\text{red}}}{I_{\text{red}} + I_{\text{green}} + I_{\text{blue}}} \\
\text{green} &= \frac{I_{\text{green}}}{I_{\text{red}} + I_{\text{green}} + I_{\text{blue}}} \\
\text{blue} &= \frac{I_{\text{blue}}}{I_{\text{red}} + I_{\text{green}} + I_{\text{blue}}}
\end{aligned}$$

For example, if the total pixel intensity of a particle image is 2000, 3000 and 6000 counts on the red(R), green(G) and blue(B) CCDs respectively, then the

---

[1] From correspondence with the camera's manufacturer.

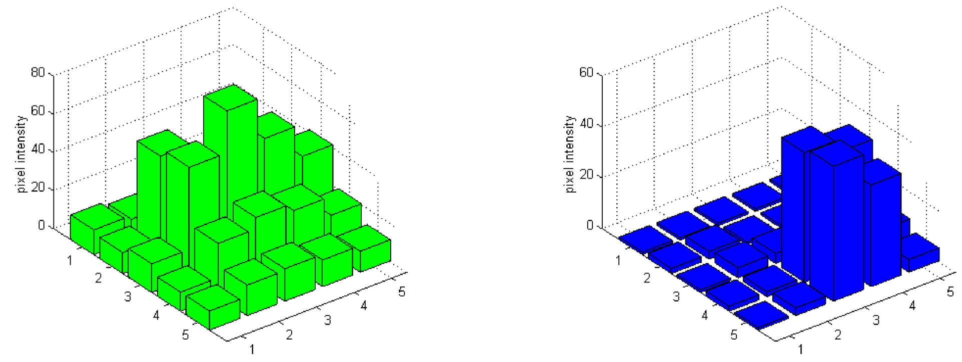

Figure 2.12: Profile of a particle image on the Green and Blue CCD's.

colour of the particle is determined as (0.18, 0.27, 0.55) in RGB order.

The size of the image area that is summed to give the particle image intensity is governed by accuracy considerations. The number of overlapping particle images must be kept to a minimum to avoid the inclusion of intensity contributions from other particle images. This means that for a given particle image size, the seeding density (i.e. the intended spatial resolution) and the depth of the illuminating light beam must be adjusted accordingly to accommodate the specified summation area which should be specified large enough for good convergence of the particle image colour. In this work, the particle image intensity sums were taken over the square area whose length spans about twice the estimated average particle image diameter.

2.6 Imaging system noise

In addition to the occurrence of particle image overlaps, the level of uncertainty in resolving a particle's z-coordinate as determined from its colour will also depend on the features of the imaging equipment. We note that the accuracy of the x- and y-coordinates of the particles is also dependent on how accurately the z-coordinates are resolved because the x-, y- and z-coordinates are related by equation (3.1).

The ability of the imaging system to produce highly repeatable imaging conditions is therefore crucial to obtaining consistent colour ratios and ultimately, consistent particle coordinates. The camera noise and the repeatability of the flashes are crucial system characteristics that affect the accuracy of the particle coordinates and in this section, we estimate the magnitude of the error introduced by the camera and Pal Flash in determining the z-position. We note that a good initial estimate of the measurement uncertainties is important as it allows one to design/arrange the experiments so that the flow length scales and particle displacements are greater than the measurement error.

2.6.1 Camera noise

Figure 2.13 shows the variation of the camera's dark current noise with gain, for all three channels. The camera noise variation shown was obtained by covering the lens to prevent the image frames from being exposed to any illumination throughout their exposure time of 135ms.

The gains for the three CCD channels are independently set on the camera's RS232- communications interface (DtControl) and can be varied from values of 2 to 36.

As can be observed in Figure 2.13, the camera noise worsens with increasing camera gain. A low camera gain setting reduces the camera noise but the particle image intensities are also reduced and there appears to be no optimum gain setting since the particle image colour combinations vary with particle depth and their intensities also vary with depth and particle size and the flashes. We note that flash-1 gives a significantly higher intensity than flash-2 so, it is expected that a greater level of uncertainty is associated with the resolved colours of the particle images in the second images of each pair of illuminated images[1].

In the experiments, the images were recorded at a camera gain setting of 30 for all three colour channels. We give a rough estimate of the uncertainty introduced by the camera in locating the z-position of the particles as follows:

For particles residing in the central Green-Blue (G-B) illumination region, the intensities of the particle image intensities when summed up over a 7×7 pixel area on the blue and green CCD area of the order of 5,000 intensity counts. To

---

[1] Flash-1 intensity is about 30-40% higher than flash-2.

simplify this analysis, we assume that the intensity recorded on the red CCD is negligible and we take the red CCD particle image intensity to be zero.

At a camera gain setting of 30, on both green and blue CCDs, the camera noise intensity contribution to the pixel intensities is about 10.8 counts. The contribution of the camera noise to the summed particle image intensities is therefore about 530 i.e. $(7 \times 7 \times 10.8)$.

A registered colour of (0, 5000, 5000) counts could therefore, in the absence of noise, vary from (0, 4470, 5000) to (0, 5000, 4470) which in terms of ratios give a variation from (0, 0.47, 0.53) to (0, 0.53, 0.47).

From Figure 2.18 the variation in colour from (0, 0.5, 0.5) to (0, 0.53, 0.47) represents a 300 micron change in position along the z-direction.

This simple analysis therefore, indicates that the particle z-position error due to the camera noise can be as much as 300 microns.

Since the z-position error should decrease with higher particle image intensity to noise ratio, we can expect the z-position of the particle images from flash-1 to be resolved to greater accuracy than those obtained with flash-2 because the intensity of flash-1 is greater than the intensity of flash-2.

To help reduce the effect of the camera noise in the images acquired in this project, the mean pixel noise intensity was subtracted from each particle image pixel that was involved in determining the particle image colour.

2.6.2 Flash jitter

The flash jitter also introduces some level of uncertainty to the resolved z-position of the particle images. For a stationary particle to reflect the same band of wavelengths for each triggered flash, the light rays from the Pal-Flash must follow the same path. Deviations in the light paths due to changes in the ignition point of the sparks therefore introduce a degree of uncertainty to the resolved colour of the particle images.

An estimate of the z-position uncertainty introduced by the flash jitter was from the recordings of several images of an 800 micron size gas bubble trapped inside a block of water-clear resin. The large size of the bubble was chosen to allow high intensity particle images to be recorded at a very low camera gain setting,

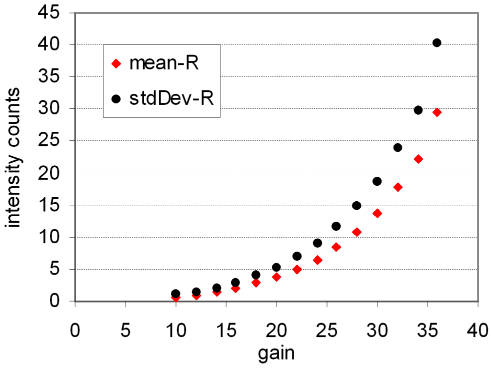

(a) red CCD pixel noise.

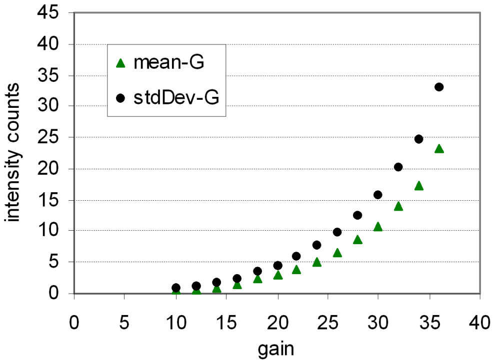

(b) green CCD pixel noise.

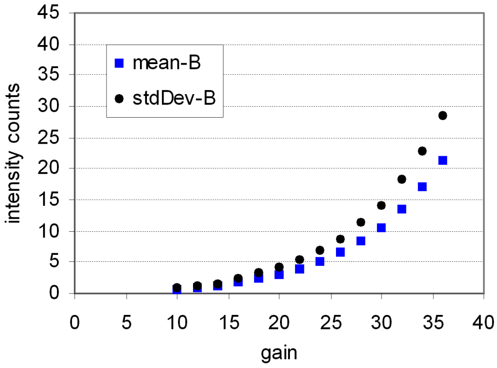

(c) blue CCD pixel noise.

Figure 2.13: Camera's mean pixel noise intensity for a 10-bit image recorded with the camera lens covered to prevent light entering through the lens.

hence making the contribution of camera noise to the recordings negligible.

Figure 2.14 shows the values obtained for the green colour component for a hundred image recordings of the gas bubble. Standard deviations in green colour ratio of 0.0084, 0.0114 and 0.0026 were determined for recordings performed with the front and back flashes and the focus light respectively.

Taking the variation of the green colour ratio with depth to be approximately 0.1 per mm in the Green-Blue central region of the light beam (see Figure 2.18 and Figure 2.19) then, these standard deviations translate to a z-position uncertainty in the order of about 80 microns for the front flash (i.e. flash-1), 110 microns for the back flash (i.e. flash-2), and 30 microns for the constant light source.

As expected, the uncertainty in the colour of the images recorded with the focus light is relatively lower than the estimated uncertainties for both flashes. This is because the focus light emanates from a fixed position. However, the estimated z-position uncertainty from the recordings made with the focus light is still of the same order of magnitude as the estimated uncertainty for the flashes to suggest that there is another major contributing source to the measurement uncertainty — perhaps the camera's readout-noise.

2.6.3 Measurement of the illumination time interval

While the illumination time interval, Δt, between flash-1 and flash-2 does not affect the accuracy to which the particle positions are determined, the accuracy to which the velocity is estimated is however dependent on the measured illumination time interval. A good approximation of the particle velocity therefore requires an accurate measurement of the illumination interval.

Measurements of the time interval between the flashes was taken directly from the Tektronix TDS 2002 Two Channel Digital Oscilloscope (pictured in Figure 2.2) which showed the response of the Photodiode placed in the beam path, beyond the flow arrangement as pictured in Figure 2.2(b). A plot of the readings of Δt taken from the oscilloscope reading for 100 pair of flashes is also shown in Figure 2.15. The oscilloscope gives the time interval between the responses of the photodiode to both flashes to two decimal places. The observed maximum deviation from the preset Δt value of 30ms was 0.02ms. This converts to a

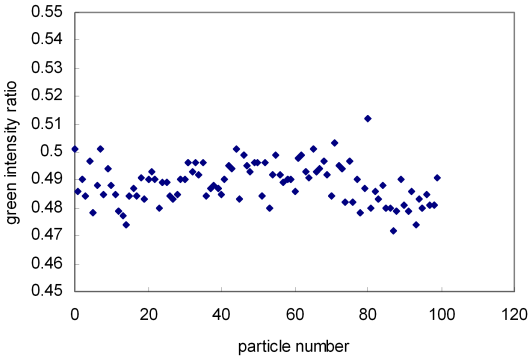

(a) flash-1 jitter.

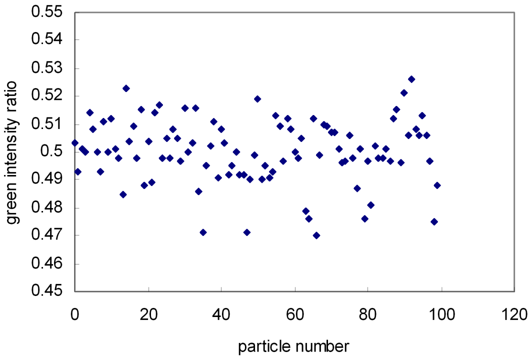

(b) Flash-2 jitter.

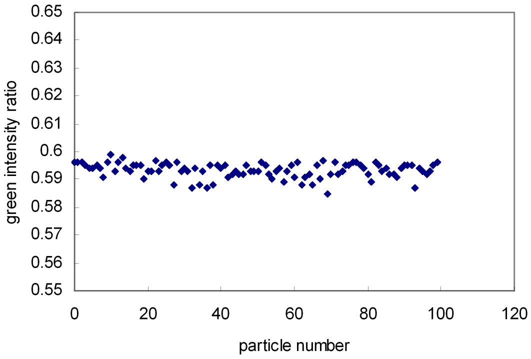

(c) Focus Light jitter.

Figure 2.14: Variation in the resolved green colour ratio for 100 image recordings of a 800 micron air bubble using both flashes and the focus light of the Pal-Flash.

negligible maximum uncertainty of 0.07% for the desired time interval of 30ms set on the pulse sequencer.

A large illumination time interval of 30ms was chosen to get a large particle displacement relative to the expected measurement error. This value for Δt gives a good enough compromise between maintaining a reasonable relative particle displacements in the interrogation blocks and low tracking error while ensuring that the particle displacements in most areas of the flow field ROI are larger than the expected measurement uncertainties.

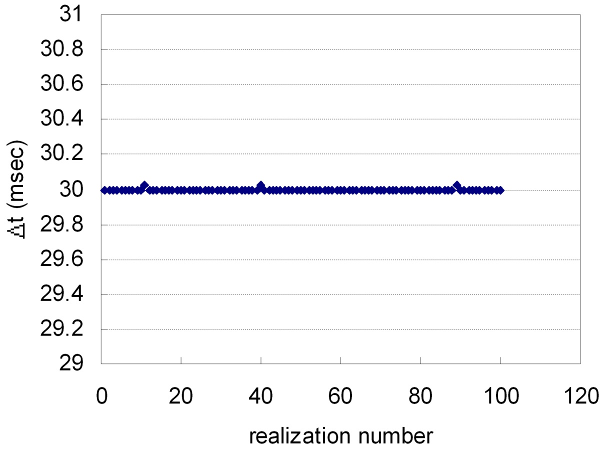

Figure 2.15: Δt as read from the oscilloscope, for 100 pairs of flashes.

2.7 Calibration of the multicoloured light

2.7.1 Choice of calibrated region and depth

From the preceding discussion on the camera and flash characteristics and the definition of the particle image colour, we note that for an unambiguous and accurate measurement of the particle image colours, the following conditions must be met:

- The particle images must be registered on at least two CCDs. In other words, the particles must reside at locations within the illuminated volume where at least 2 CCDs are adequately sensitive to the wavelength of the light at those locations.

- The particle image pixels must not be saturated as this will lead to erroneous colour determination.
- The particles must reside in locations where their image intensities will be high enough to minimize the effect of camera noise. Also, the colour graduation along the depth must be steep enough to keep the relative particle position errors low.

These requirements impose restrictions on the size of the region of the rainbow-coloured light beam suitable for interrogation the flow field. From the camera's spectral response shown in Figure 2.7, the region along the depth of the light beam which best satisfies these requirements is the Blue-Green overlapping region spanning a depth of about 12mm.

For the region spanning 0-30mm along the z-direction (as indicated in Figure 2.7), the pixel intensities on the blue channel CCD are too high compared with the other two channels. The blue component of the pixels will therefore become saturated long before any measurable intensity is recorded for the other two components. The camera noise also has a greater influence on the intensity recordings for the lower intensity CCD channels and as a result, the uncertainty in the colour of the particles residing in this region will be larger.

Similar arguments can also be made for the 45mm to 60mm region and the region beyond the 70mm mark.

We also note that the Red-Green (R-G) overlapping region cannot be used exclusively or in combination with the Green-Blue (G-B) region because the intensities in the R-G region are relatively much lower. The particle image intensity to noise ratio will be lower and the z-position uncertainty, will therefore be greater in the Red-Green overlapping region. Furthermore, the image intensities of the particles in the blue-green overlapping region will become saturated long before any measurable or satisfactory image intensity can be read out for those particles residing in the R-G region.

In the experiments, the central 12mm thick section within the Green-Blue overlapping region of the multicolour light beam, was used to illuminate the flow fields. This region of the multicolour beam was therefore calibrated to give a database of colour vs. position along the beam's depth.

2.7.2 Calibration procedure

The calibration arrangement is illustrated in Figure 2.16.

The camera's optical axis is aligned with the z-axis traverse motion so that the image of an object initially located at the centre of the digital image will remain at the same pixel coordinates as the object is traversed along the z-direction.

A white flat plate inclined at 45° to the incoming multicoloured light beam was mounted on the motorized 2-axis translation stage and then traversed along the depth of the beam in 1mm steps along the z-direction. At each z-location, the flashes illuminate the plate and the camera records its images. The colours at selected points on the face of the plate were subsequently extracted from the recorded images by relating the coordinates of the points back to the image coordinates using the expression given in equation (3.1a). The calibration points on the plates are the intersections between the plate and the user-specified calibration grid lines along the y- and z-directions as shown in Figure 2.17. The number of the calibration points is therefore determined by the step size movement of the plate and the user-specified calibration grid spacing δ_y and δ_z along the y- and z-directions respectively. The calibration plate's location along the camera's optical axis and the coordinates of the calibration points are measured from a reference plane as shown in Figure 2.16. The plate's z-distance is measured from the reference plane to the point of intersection between the plate and the line $x = FOV^x_{z=0}$. The reference plane origin is taken as the bottom left corner of the reference plane.

The images of the plates were taken at a camera gain setting of 9 for all 3 CCD channels. At this gain setting, the contribution of camera noise to the recorded intensities was very low. To reduce the error in the calibration due to flash jitter, three recordings of the plate were made at each z-location and the median of the colour extracted for each calibration point was chosen as the colour of light at that point's location in the multi-colour beam.

We note that the time required to complete the acquisition of a specified number of calibration plate images is determined by the flash recharge time while the time available to perform the calibration is determined by the gas cylinder capacity and the flow rate of the argon gas being fed to the flash. In this project, a calibration plate step movement of 1mm along the z-direction was specified so that the calibration could be completed before the argon gas was exhausted.

It also be noted that the number of calibration points that can be extracted from the plate images is limited by the PC memory available to hold the calibration data during the image analysis stage where the colour of each particle image matched to a position using the calibration point information.

In this project, about 45,000 calibration points were extracted at a grid spacing of $\delta_x = 1\text{mm}$, $\delta_y = 1\text{mm}$ and $\delta_z = 1\text{mm}$ within a $55 \times 30 \times 25\text{mm}^3$ calibration volume. For each flash, approximately 2.5 MB was needed to store the spatial coordinates of the calibration points, their corresponding image coordinates and their colour.

The steps involved in extracting the image coordinates of a calibration point and its colour from the calibration plate images are as follows:

1. The y- and z-calibration grid point coordinates (y_p, z_p) are specified.
2. The corresponding x-coordinate x_p is determined from the relationship:

   $$x_p = FOV^x_{z=0} - \frac{(z_p - z)}{\tan\alpha}$$

   where z is the plate's distance from the reference plane measured along the line $x = FOV^x_{z=0}$ as indicated in Figure 2.16.
3. The image coordinates (x', y') of the point are then determined using equation (3.1).
4. The colour of the calibration point was subsequently averaged over a small neighbouring image region. The sum of the intensities over a 7×7 pixel area centred at (x', y') was taken for each channel and the ratio of the three (R, G, B) intensity sums to give the colour for the point.

Figure 2.18 and 2.19 show the variation of colour with depth over a y-plane in the calibrated volume.

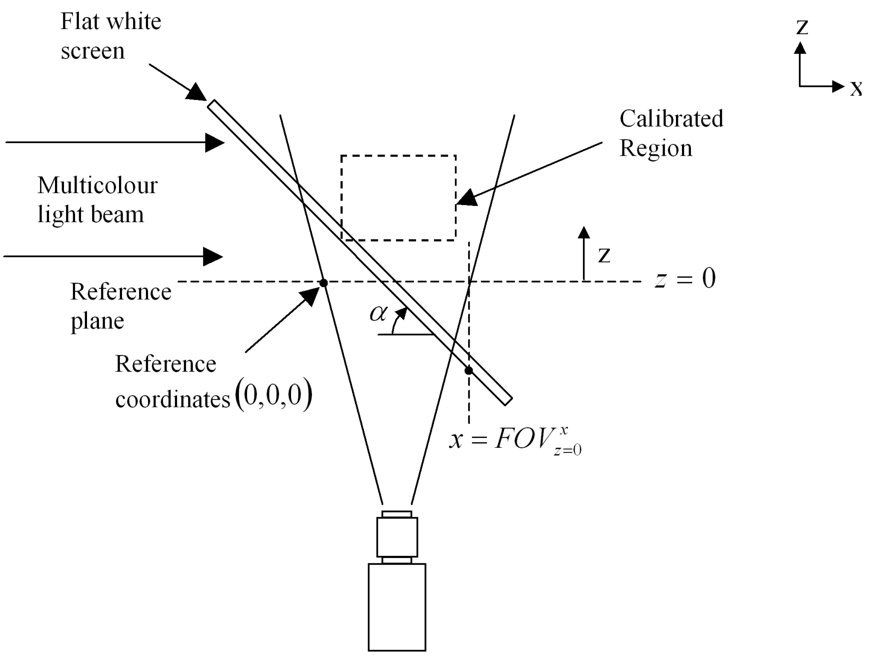

Figure 2.16: An illustration of how the white calibration plate is positioned relative to the camera and the multi-colour light beam.

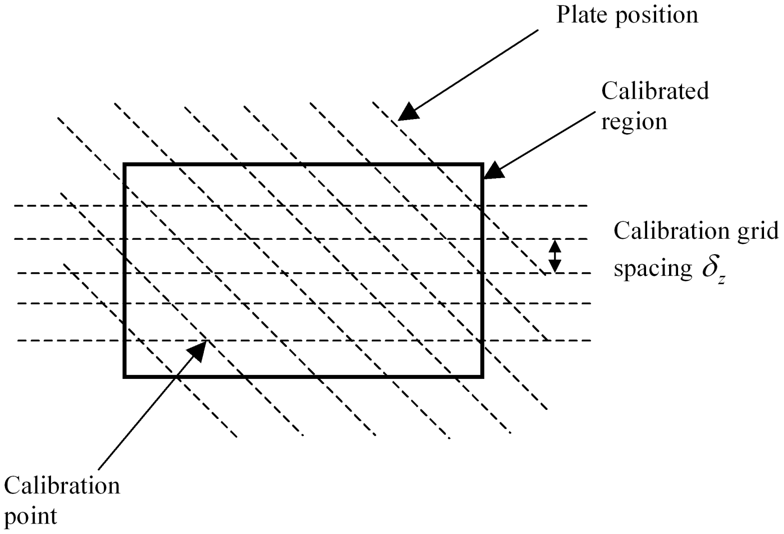

Figure 2.17: The calibration points in the y-plane, as the plate is traversed along the z-direction. The calibration points are the intersections between the user specified z-grid lines and the plate.

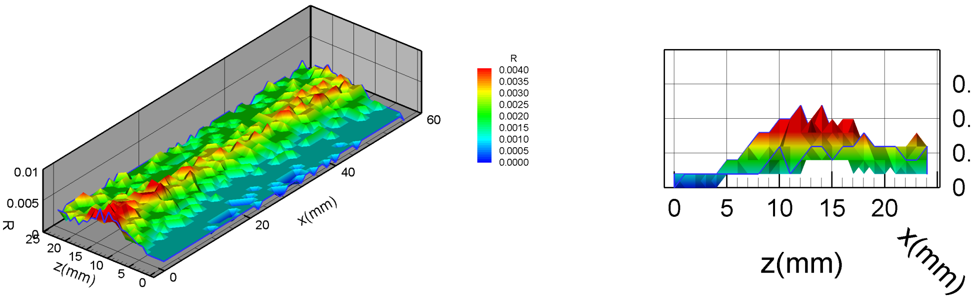

(a) red intensity ratio.

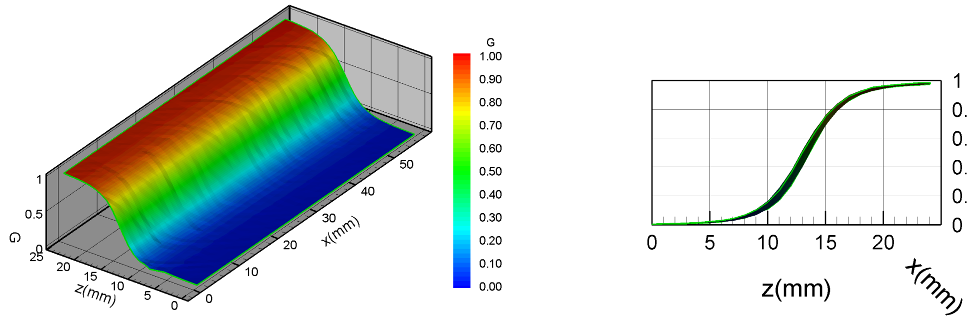

(b) green intensity ratio.

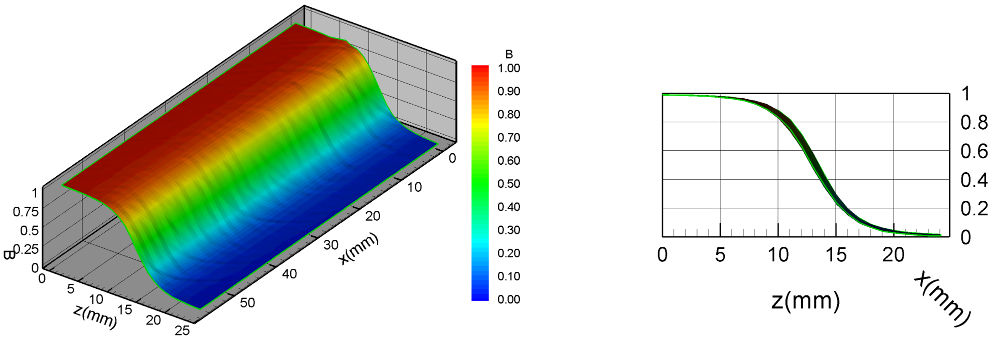

(c) blue intensity ratio.

Figure 2.18: Variation of colour over the plane y = 20mm within the calibrated region of the multicolour light beam from flash-1.

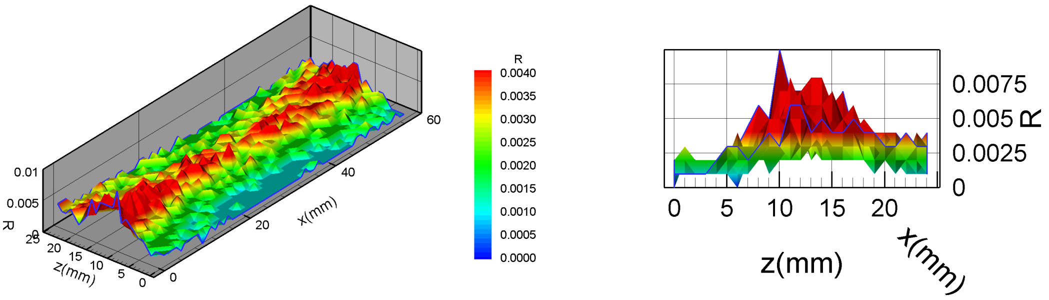

(a) red intensity ratio.

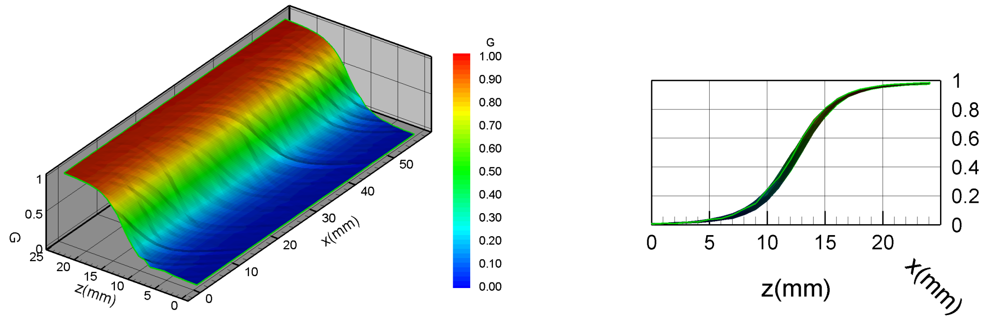

(b) green intensity ratio.

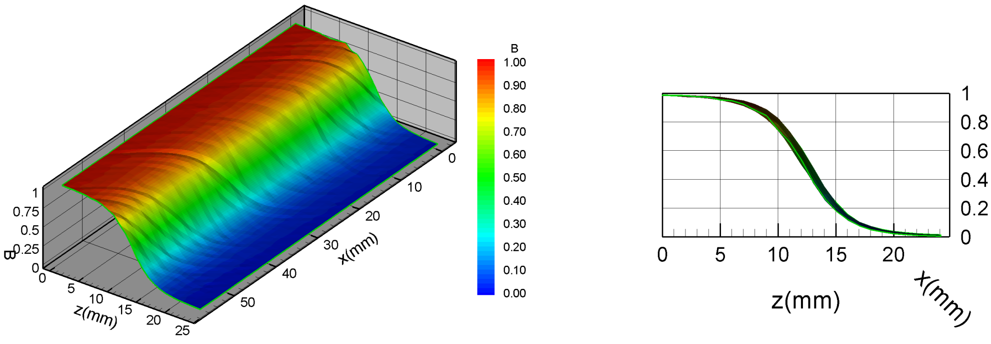

(c) blue intensity ratio.

Figure 2.19: Variation of colour over the plane y = 20mm within the calibrated region of the multicolour light beam from flash-2.

2.8 The tracer particles

The seeding particle choice for the experiments was Talisman-10 particles from Plascoat Systems Ltd., selected among the seeding particles listed in Appendix G. These particles are hard, irregularly shaped thermoplastic white powders commercially used for the protective coating of metals. They are available in a size distribution of 5 microns to 250 microns and their density is approximately 0.99g/cc. These particles exhibit excellent light reflection qualities and are also highly resistant to chemicals.

For the chosen illuminated ROI depth, particle sizes between mesh sizes of 210 and 250 microns were used to meet the requirement for adequate light scattering while maintaining low number of particle image overlaps and high spatial resolution measurements (see experimental notes in Appendix A.1).

At the experimental lens aperture setting of f-2.8 and camera distance of about 1m from the illuminated region, the particle image size[1] was estimated to be about 3 pixels.

We note that when particle images overlap, their position accuracy is reduced since the resolved colour for each overlapping particle image is likely to include intensity contributions from the other overlapping particle images. For good measurement accuracy, it is therefore important that the occurence of particle image overlaps is kept to a minimum.

Figure 2.20 and Figure 2.21 show the result of simulations performed to give an estimate of the number of particle image overlaps that can be expected, in the experiments, as a function of the particle concentration and size for the illuminated field size in our experiments. Particle images with flat intensity profile were used in the simulations to help determine the number of overlaps. Although the real particle image intensity profiles are closer to Gaussian than to a flat profile, the flat intensity profile used in the simulation can be interpreted as the diameter over which the real particle image intensities are to be summed to determine their colour.

In the simulations, the particle images that touch were also counted

---

[1] We define the particle image diameter as the number of pixels that covers 68.3% of the light intensity recorded along the central horizontal or vertical line of pixels.

as overlapping particles so Figure 2.20 and Figure 2.21 somewhat give an overestimate of the number and percentage of overlapping particles respectively.

Assuming that the intensities of the pixels within a 7×7 pixel area was used in the determination of the particle colours, then from Figure 2.20, we see that at the target seeding density corresponding to an average particle separation of 2.5mm to 3mm, we can expect about 100 cases of overlaps (which translates to 200 overlapping particles) out of about 1900 particles.

2.9 The working fluid

The fluid chosen for the vortex ring experiments was water mixed with Industrial Methylated Spirit 64 O.P. (from Fisher Scientific) to give a specific gravity of 0.99 at $25°C$ to match the density of the particles. The dynamic viscosity of the mixture, at $25°C$, using value for 90-10 %wt. water-ethanol mixture, is 0.00132 Pa.s. while that of pure water at $25°C$ is 0.00089 Pa.s.

2.10 Uniform flow experimental arrangement

2.10.1 Uniform flow – introduction

Figure 2.3 shows a picture of the experimental arrangement for the uniform flow simulation. The simulation of uniform flow was achieved by moving particles suspended in a block of resin by prescribed distances along the x- and z-directions using the 2-axis translation stage. The uniform flow simulation enabled the accuracy of the system to be estimated since the exact displacements of the blocks are known.

2.10.2 The uniform flow block

The uniform flow block, shown in Figure 2.22, was constructed by mixing quantities of Cenosphere particles with water-clear Polyurethane resin in a 50mm cube Perspex block container. The mixture was subsequently degassed in a vacuum chamber and then left out to harden, thereby freezing the particles at

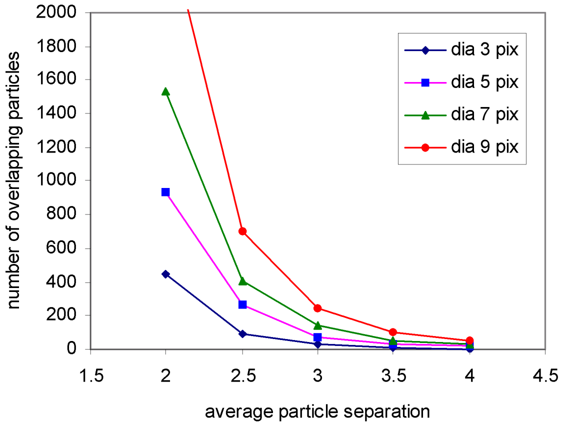

Figure 2.20: Prediction of the variation in the number of overlapping particles with the particle size and seeding density for a simulated volume of size $57 \times 43 \times 12\text{mm}^3$ projected onto an image of size 1392 × 1040 pixels.

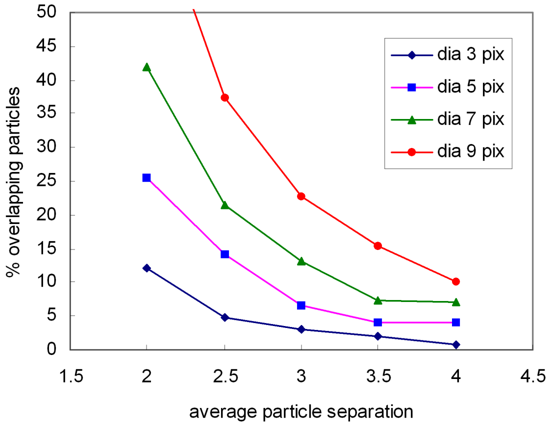

Figure 2.21: Number of particle overlaps as a percentage of the number of particle images in the simulated volume from Figure 2.20.

their random positions in the block. The block used in the experiment contained particles at an average particle-separation of about 3mm.

2.10.3 The translation stage

The motorized 2-axis translation stage shown in the pictures in Figure 2.2, Figure 2.3 and Figure 2.23 consists of two similar traverses arranged 90° on top of each other. Each traverse was constructed from a 130mm × 130mm × 46mm Quad Block (from IGUS) through which runs a pair of 300mm long, 20mm diameter stainless steel shafts. The ends of the shafts are held firm in the slots of the Shaft End Blocks, which are in turn fastened onto a 350mm × 130mm base plate. Fitted along the centre of the quad block is anti-backlash nut through which runs an anti-backlash lead screw. The screw lead is 1mm and its accuracy is 0.0006mm/mm. A computer controlled 1.8° (NEMA size 13) stepper motor coupled to the lead screw drives the quad block. The motor of each traverse is controlled by the Applied Motions® micro stepping Drives, thus, enabling the traverses to be moved with a precision equal to the lead screw accuracy. The maximum travel of each traverse is 130mm.

2.10.4 Uniform flow experimental settings

To simulate uniform flow, the 2-axis traverse system was moved back and forth, several times, over a 1mm distance along the x- and z-directions to give particle displacements of (1mm, 0mm, 1mm).

We illuminate the central 12mm thick region of the uniform flow block. At each traverse position, a pair of images of the particles suspended in the block was recorded using flash-1 and flash-2 illuminations to record the first and second image in the pair respectively.

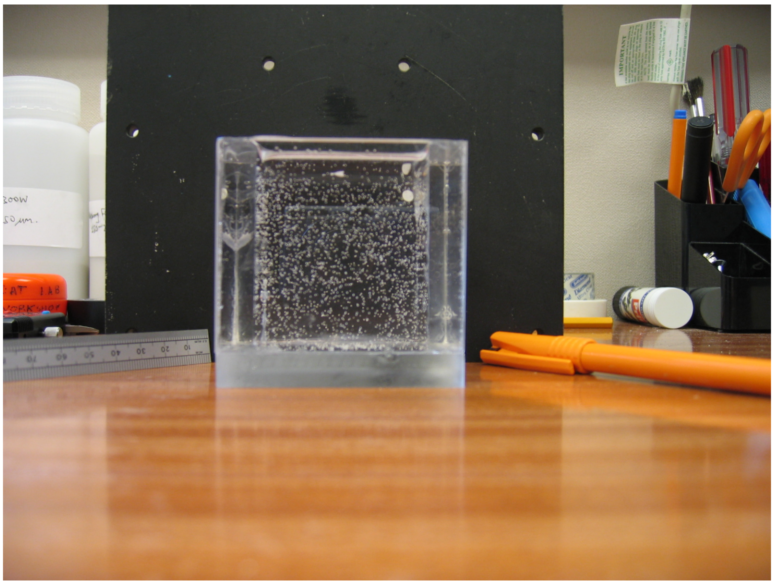

Figure 2.22: A uniform flow block with the tracer particles suspended in the resin mixture.

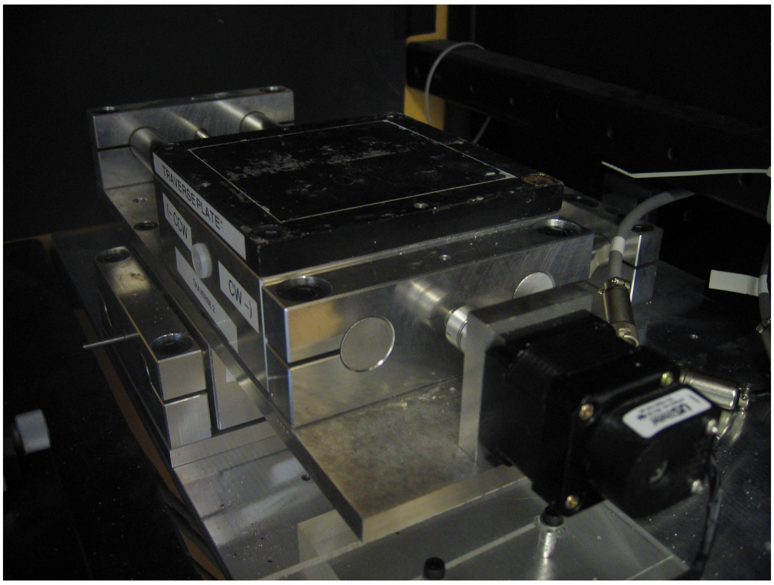

Figure 2.23: The 2-axis translation stage.

2.11 Vortex ring flow experimental arrangement

2.11.1 The vortex ring generator equipment

The vortex ring experimental arrangement is shown in Figure 2.24 and Figure 2.2 where it is mounted on the 2-axis translation stage. The vortex ring generator apparatus consists of a piston and cylinder arrangement in a clear 100mm × 100mm × 300mm deep Perspex tank filled with the working fluid. The cylinder is constructed from a 50mm × 50mm × 250mm long Perspex block through which a 20mm diameter hole is drilled. A 90mm square Perspex plate is fitted flush to the base of the block to further extend its flat base. The piston is attached to the end of the lead screw of a computer controlled hybrid linear actuator (NEMA size 23). The actuator's step angle is 1.8° hence giving a linear travel of 0.002in (0.0508mm) per step. Rotation of the lead screw about its axis is prevented by fitting a small screw through its radius and into the groove running along the length of a fixed pole. The piston motion profile is programmed via the Applied Motion SiNet Hub interface onto the micro-stepping drive, which controls the actuator. The maximum piston travel is approximately 200mm.

Vortex rings are produced by the movement of the piston through the cylinder. Fluid is sucked into the cylinder during the upward stroke of the piston and vortex rings are produced by ejecting fluid out through the cylinder orifice, during the downward piston stroke.

Located on opposite sides, near the edge of the orifice, is a pair of narrow holes through which dye can be injected into the flow to enable the visualization of a cross section of the vortex rings. The cylinder orifice diameter D_o is 20mm.

To reduce high screening effects caused by higher particle seeding densities, the internal dimension of the region of the tank where the vortex ring is observed is reduced to 60mm × 60mm by fitting a wall of perspex plates parallel to the walls of the tank and 30mm from the centre of the cylinder.

2.11.2 Laminar vortex ring flow experimental settings

Figure 2.25 shows the programmed piston motion. Low Reynolds number laminar vortex rings are produced by moving the piston at an average velocity $\bar{U}_p$ of

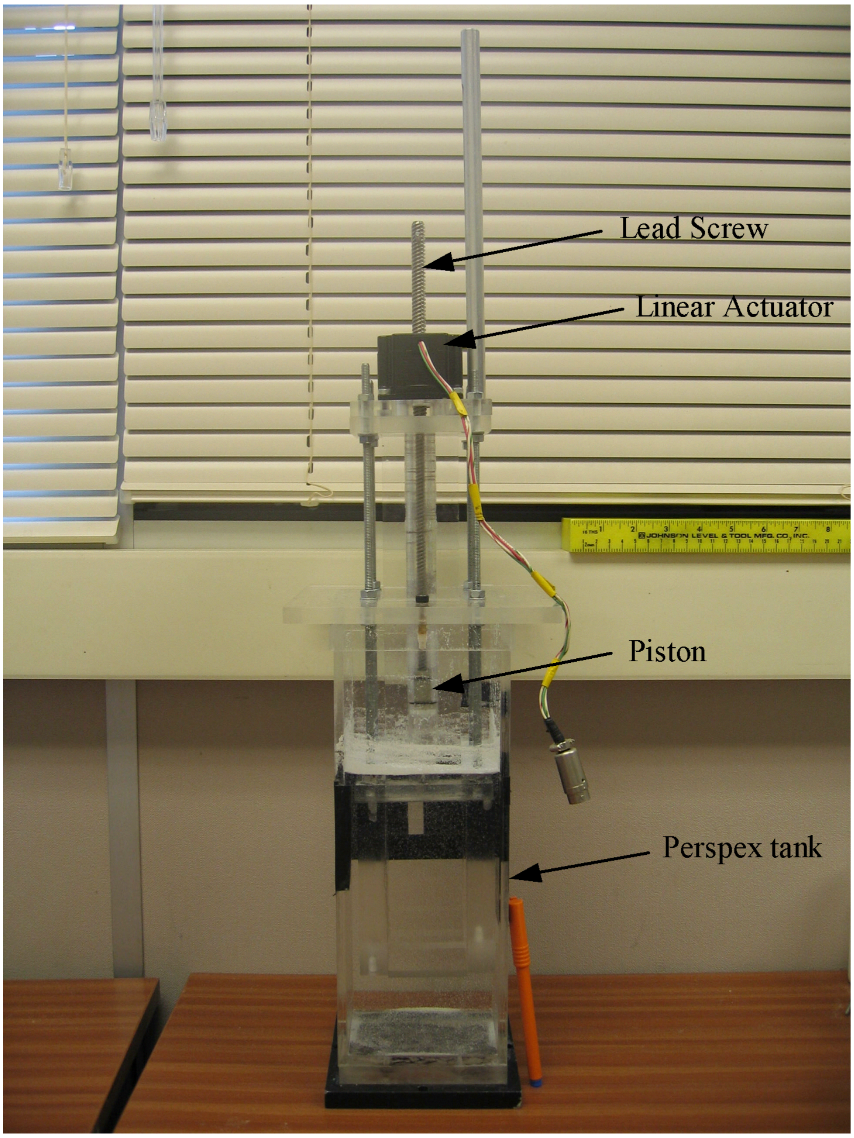

Figure 2.24: The vortex ring generator.

48.6mm/s to give a piston-stroke L_p equal to 30mm and a L_p/D_o ratio of 1.5. The corresponding Reynolds number ($Re = \frac{\rho \bar{U}_p D_o}{\mu}$) is approximately 720 (based on the fluid properties at $25°C$). The seeding particles used for the vortex ring experiments are the Talisman-10 particles at an estimated seeding density corresponding to an average particle separation of 2.5mm – 3mm. The time between two successive image illuminations, Δt, was set to 30ms on the pulse sequencer and confirmed with the photodiode response. For each downward piston motion of the vortex ring generator, the flash-1 images of the laminar vortex ring flow field were captured 1.52 seconds after the start of the piston motion.

Due to the restrictions imposed on the useful depth of the illumination, the whole 3D structure of the vortex ring could not be recorded in one image. However, because the flow is axisymmetric, we only captured the front half of the vortex ring in four 12mm thick overlapping sections, by moving the z-traverse in 5mm steps along the z-direction. At each traverse position, the experiment was repeated several times to give multiple realizations of the illuminated section of the flow field. For each section, twenty pairs of images were interrogated. The vectors from all 4 sections were subsequently combined and interpolated onto a regular rectangular 3D grid to reveal the 3D structure of the front half of the vortex ring.

A sample flow image acquired during the vortex ring experiment is shown in Figure 2.26.

We present in Chapter 5 the results of the investigations of the vortex ring experiments and the uniform flow experiments at these settings. But first, we explain in Chapter 3 and Chapter 4 how the acquired flow images are interrogated.

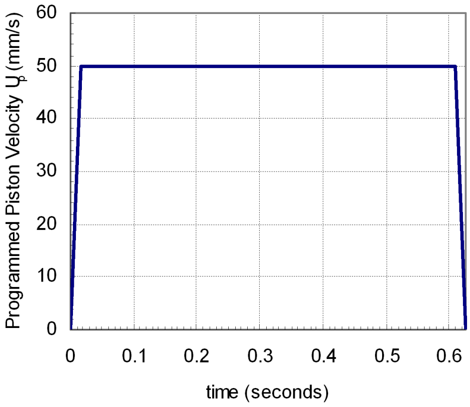

Figure 2.25: Programmed piston velocity profile for the downward stroke.

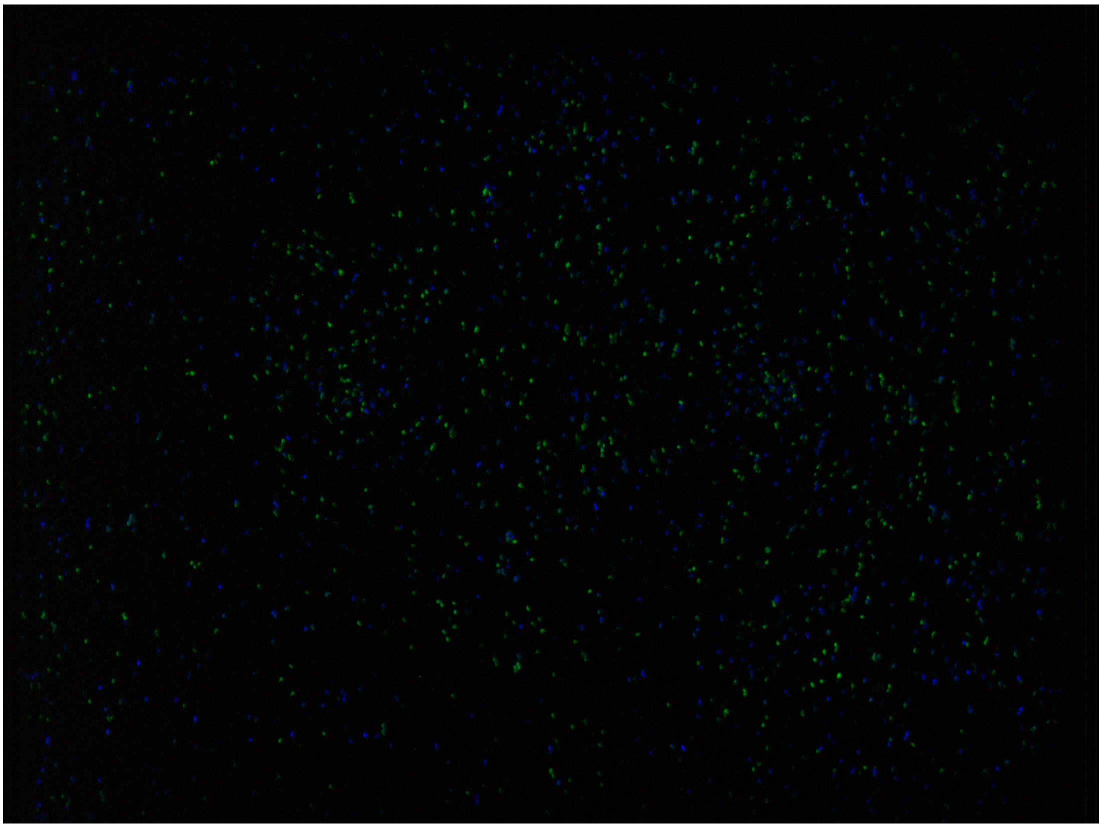

Figure 2.26: A sample flow image acquired during the vortex ring experiment.

Chapter 3

Fundamentals of 3-Component velocity measurements in a volume using a single camera

3.1 Introduction

In this chapter, we give some general background to interrogating flow images coded with information about the z-spatial coordinates of the particles.

Unless stated otherwise, all references to correlations refer to cross-correlations of a pair of singly-exposed particle images. Details of the simulations presented in this chapter can be found in Appendix D.

3.2 Interrogation of images

The first steps during the interrogation of images like the colour flow images acquired in this project is to extract the particle images and determine their 3D spatial coordinates in real world space. The displacement, hence velocities of the identified particles can subsequently be determined to give a velocity vector field to represent the flow field.

For the highest spatial resolution measurements, the particle displacements can either be obtained by iterative (multi-stage) correlation, or by pairing the particle images using the hybrid PIV-PTV scheme.

The advantage of achieving super-resolution by iterative PIV correlation is that the resulting velocity vectors are on a uniform grid. However, super-resolution via iterative correlation is likely to give less accurate measurements because the vectors will typically be located at the centre of the interrogation block regardless of the position of the particles within the interrogation blocks and the variation of the flow velocity gradients across the interrogation blocks. See section 1.6.1.2. Also, at higher iteration stages, the particle sizes become comparable to the interrogation area sizes. As a result, particle images are likely to encroach the boundaries of the blocks hence resulting in poor estimation of the particle displacements from the correlation of partial particle images instead of the whole particle. Furthermore, at higher iterative correlation stages, some interrogation areas will contain no particles. These empty interrogation areas will still need to be interpolated if all the velocity grid points are to be filled.

The author therefore contends that accurate high-resolution measurements are best achieved by tracking each particle either directly or via the hybrid approach where we identify each particle from the images and then determine their displacements as the difference between their image centroid location on the pair of recorded images. The randomly spaced vectors can then be interpolated onto a regular grid if one wishes.

3.2.1 Pairing of particles: Direct particle tracking vs. hybrid PIV-PTV

The choice between pairing particle images directly, using the nearest neighbour approach or via the hybrid PIV-PTV method will depend on the particle spacing/displacement ratio.

Figure 3.1 shows, for a range of spacing-displacement ratios, the likelihood that the nearest neighbour of a particle is that particle's true pair. We observe that when the spacing/displacement ratio of the particles is less than about 5, the probability of correctly matching particles drops sharply.

Achieving a high spacing/displacement ratio would require either sparse seeding or small particle image displacements. However, sparse seeding will result in low resolution measurements while small particle image displacement requirements could mean restricting the measurements to smaller flow regions so that larger particle image pixel displacements relative to the measurement uncertainties is maintained.

An alternative way to achieve high resolution and still maintain high valid-data-yield, accuracy and dynamic range is to use a hybrid PIV-PTV technique where initial displacement estimates of particle clusters are obtained by correlations before each particle is paired using its estimated cluster displacement to guide the search for its pair.

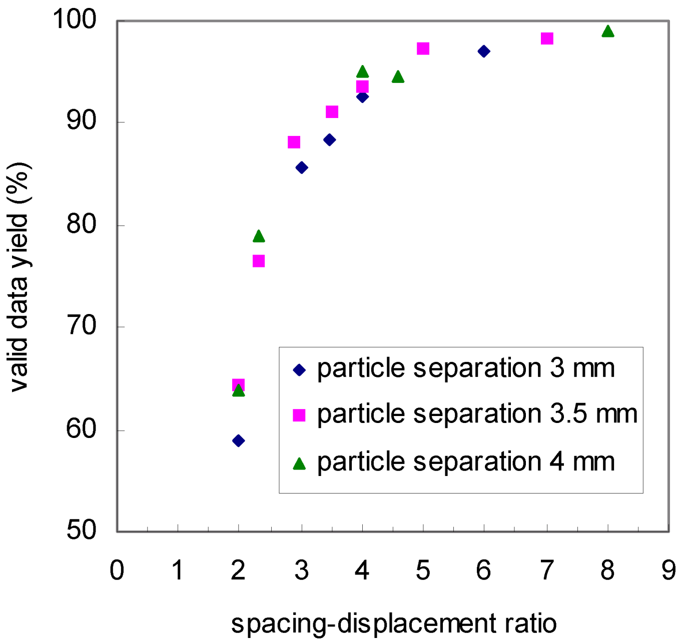

Figure 3.1: Plot showing the probability of correctly pairing particle images over a range of spacing-displacement ratios using the 'nearest-neighbour' approach. Results obtained from the simulations of uniform flow.

3.3 Hybrid PIV-PTV algorithm

The steps involved in the interrogation of the recorded pair of singly-exposed images using the hybrid PIV-PTV algorithm is as follows:

1. From the recorded images of the flow field, construct the particle image field on a 3D grid. See section 3.3.1.
2. Divide the grid into interrogation blocks[1].
3. Cross-correlate corresponding interrogation blocks either directly or via FFT.
4. Locate the peak in the block of correlation values.
5. Convert the peak coordinates to real world displacement vector
6. Repeat steps 2–5 until all the interrogation blocks are analysed.
7. Pair individual particles in one image frame to its corresponding image in another frame, by searching for a matching particle around the estimated location.

Steps 2 to 6 could be performed iteratively for super-resolution or to give a better displacement estimate for each particle in the cluster. We note that the correlation peak does not need to be located to sub-pixel accuracy if the desired result is to be achieved by tracking each particle.

We illustrate steps 2 to 6 in Figure 3.2 and Figure 3.3[2] which show the correlation process using the direct approach and FFT respectively.

---

[1] We define two terms that will be used frequently when the correlations on a 3D grid are discussed. We define a **Voxel** as a three-dimensional equivalent of a pixel and we define an **Interrogation Block** as a three-dimensional equivalent of a two-dimensional interrogation window, made up of voxels.

[2] For correlations via FFT, half of the size of the resulting correlation table in any direction corresponds to positive search distances while the other half corresponds to negative search distances. We also note that for FFT based correlation, the 3D array $g(i, j, k)$ is reversed to be $g(L-i, M-j, N-k)$ before it is transformed. See notes on discrete cross-correlation via FFT from Bracewell (1999). The 3D correlation table will also need to be rearranged, after performing the inverse FFT, before searching for the peak value in the 3D table.

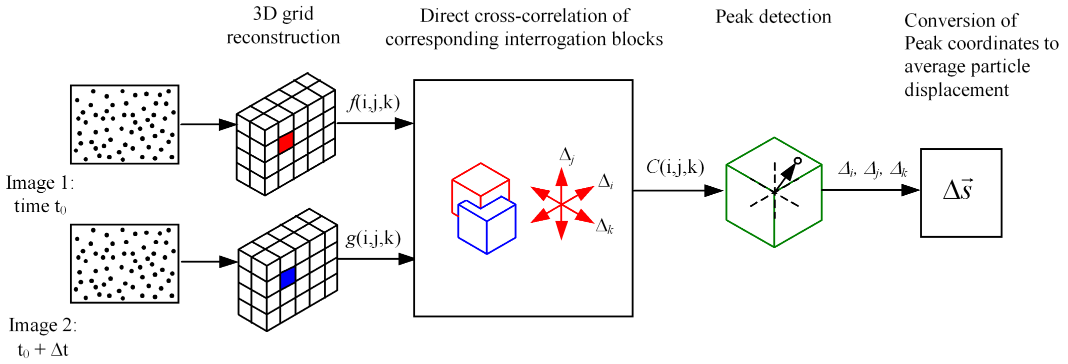

Figure 3.2: The steps involved in the interrogation of a pair of particle images by direct computation of the correlation function.

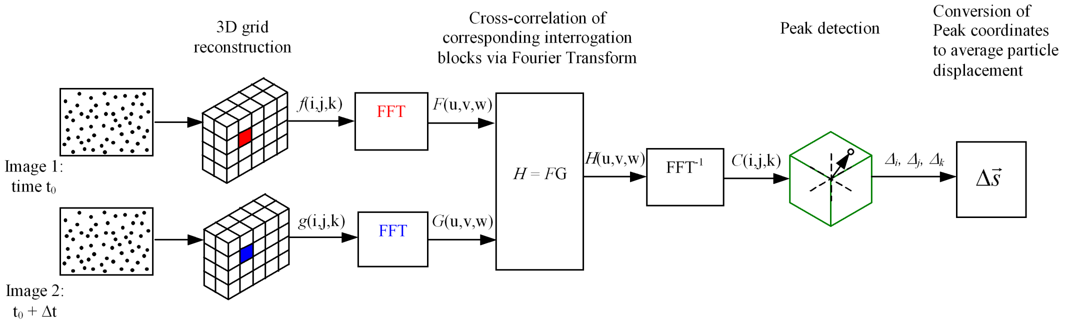

Figure 3.3: The steps involved in the interrogation of a pair of particle images using FFT-based correlation technique.

3.3.1 Reconstruction of the particle image field

We list the steps involved as follows:

1. Identify the particle images from the digital images.
2. Relate the particle image positions back to real world coordinates.
3. Specify the size and location of the Region of Interest (ROI) in space.
4. Define a 3D grid to represent the ROI.
5. Map the particle spatial coordinates onto the grid coordinates.
6. Using a suitable model, reconstruct each particle centred at its grid coordinates.

3.3.1.1 Identifying the particle images and determination of the particle image centroids

The first steps in the image analysis process involves identifying the particle images and relating their centroid positions in the image back to real world coordinates.

We present, in section 4.2, Chapter 4, a method for identifying the particle images based on determining a threshold intensity value to help separate the particle images from the background.

A common method of determining the centroid of the particle images is to weight the contribution of the particle image pixels using the expression given in Appendix E.2. We describe the procedure particular to our 3D colour PIV system in section 4.3 in Chapter 4.

3.3.1.2 Relationship between image space and real world space

During the conversion of the particle image coordinates to spatial coordinates, the camera lens perspective must be taken into account to avoid perspective error in the resulting x- and y-spatial coordinates of the particles.

For a camera fitted with a lens which has a non negligible angle of view, the field of view increases with the imaging distance as shown in Figure 3.4. Images on different planes along the optical axis will therefore scale by different amounts on the image plane. This change in image scale with imaging distance has several implications on the accuracy of the measurement:

First, the images of points at the same x- and y-coordinates but different z-coordinates in real space, will not be recorded at the same position on the image plane, unless the points lie on the optical axis. Therefore, when a point that is not on the optical axis, is displaced along the z-direction, its image coordinates change. Secondly, when two points on different z-planes are displaced by the same amount, their image displacements will not be the same since the imaging distances are different.

The induced change in a particle's coordinates on the image plane as a result of its movement along the z-direction or the difference between the observed particle image displacement of two similarly displaced particles residing in different z-planes, is known as the perspective error.

The induced change in particle image position due to lens perspective is illustrated in Figure 3.5 which shows the projection of the 3 points shown in Figure 3.4, onto the image plane. In this illustration, we see that although points 1 and 2 have the same spatial x- and y-coordinates, their image coordinates are different. This difference in particle image position is the perspective error that will be observed in the measurements if point 1 moves to point 2. The perspective error is further demonstrated as we see in the figure that point 1 and point 3 share the same image coordinates even though they are not located at the same position in space.

We give the relationship between real world coordinates (x_p, y_p, z_p) and image coordinates (x', y') as:

$$x_p = \frac{x'}{N_x}\left(FOV^x_{z=0} + 2z_p \tan\tfrac{1}{2}\theta_x\right) - z_p \tan\tfrac{1}{2}\theta_x \tag{3.1a}$$

and

$$y_p = \frac{y'}{N_y}\left(FOV^y_{z=0} + 2z_p \tan\tfrac{1}{2}\theta_y\right) - z_p \tan\tfrac{1}{2}\theta_y \tag{3.1b}$$

where the camera lens angle of view in the x- and y-direction are θ_x and θ_y

respectively. $FOV^x_{z=0}$ and $FOV^y_{z=0}$ are the field of view in the x- and y-direction respectively at a reference z-location. N_x and N_y are the image size (measured in pixels) along the x- and y-directions respectively.

From these expressions, we see that we can correct for the perspective error introduced by the lens into the x- and y-spatial coordinates of a particle if the particle's z-coordinate is known. We also note that any error associated with determining the particle's z-coordinate z_p, also affects x_p and y_p.

3.3.1.3 Modeling the particles on a 3D grid

When the spatial coordinates of the particles are determined, a 3D grid respresenting the flow ROI is defined on which each particle is represented by a suitable model.

The particles can be modelled as a sphere or a cube. The chosen model can also be defined to have a flat or a Gaussian intensity profile.

A flat intensity profile can be generated by applying the same intensity grey-level to the voxels that are within a particle radius from the grid location of the particle. We note that if a flat intensity profile is used for the particle reconstruction, the error in positioning the constructed particle on the grid can be as much as half a voxel size.

A Gaussian profile is set by weighting a specified maximum intensity value I_{max} against a Gaussian function, to give voxel intensities that decay the further away the voxel is from the particle coordinates. The intensity of each particle voxel is expressed as $I = I_{max}e^{-(\frac{l}{\sqrt{2}r})^2}$ where l is the distance of the voxel from the particle coordinates on the grid and r is the radius of the reconstructed particle.

If the preferred route to super-resolution is via Iterative Correlation, then it is recommended that the particles be modelled using a Gaussian profile so that they can be more accurately placed on the grid. The resulting peak in the 3D correlation table should also be resolved to sub-pixel accuracy after the last stage of the iterative correlation.

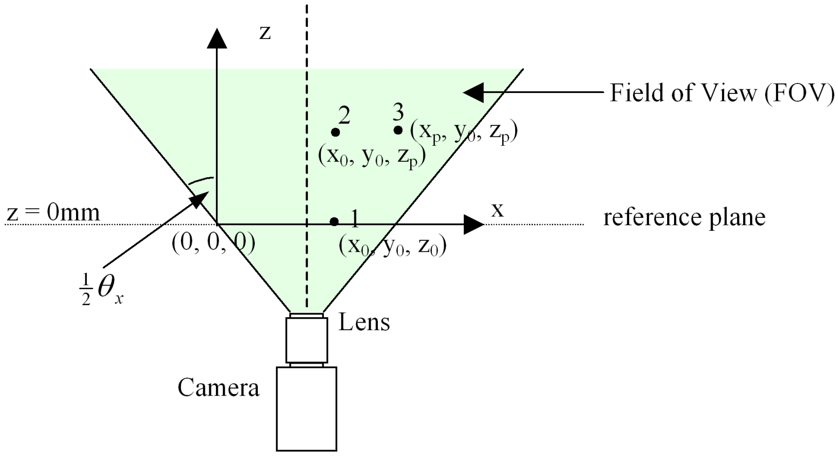

Figure 3.4: Depiction of the field of view covered by a camera lens having a non-zero angle-of-view.

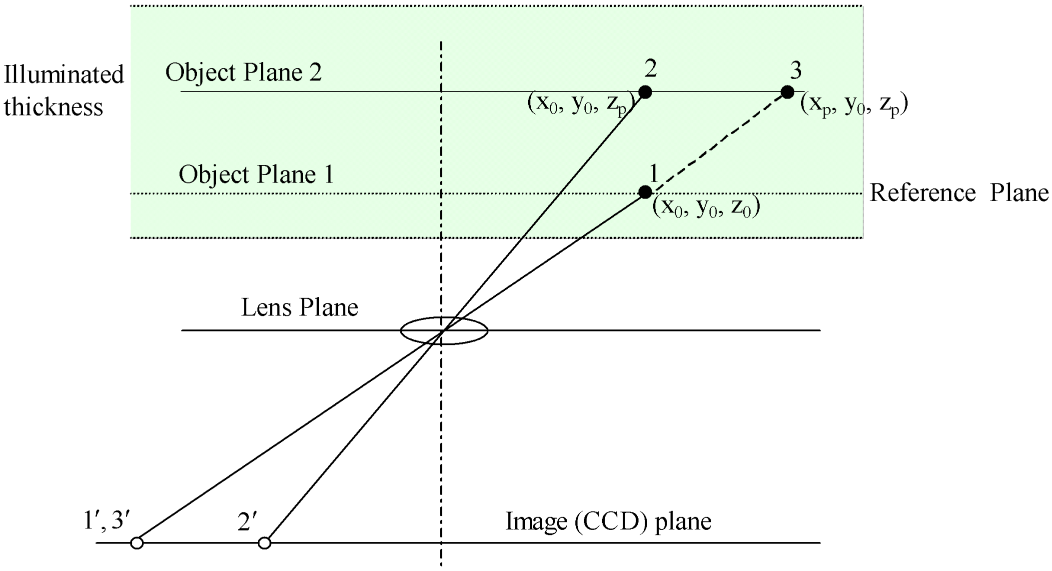

Figure 3.5: Projection of points in space on the image plane of a camera fitted with a lens having a non-zero angle-of-view.

3.3.2 Computation of the discrete 3D cross-correlation function

The discrete 3D cross-correlation function is expressed as:

$$C(\Delta_i, \Delta_j, \Delta_k) = \sum_{i=1}^{L} \sum_{j=1}^{M} \sum_{k=1}^{N} I_1\left(i, j, k\right) \cdot I_2\left(i + \Delta_i, j + \Delta_j, k + \Delta_k\right) \tag{3.2}$$

where the size (in voxels) of the interrogation blocks along the x, y and z-direction respectively are $L \times M \times N$.
I_1 and I_2 are the voxel intensities in interrogation block 1 and interrogation block 2 respectively.
$I_1(i, j, k)$ is the intensity of the voxel located at coordinates (i, j, k) in the first grid and $I_2(i, j, k)$ is the intensity of the voxel located at coordinates (i, j, k) in the second grid. Δ_i, Δ_j and Δ_k are the grid search coordinates along the x, y and z-directions respectively.

We note that the correlation function is not normalized. This is because the correlation is performedon a reconstructed grid where such effects as uneven particle image illuminationand background image noise do not exist. Furthermore, to address the effect of out-of-boundary particle motions, correlation techniques which involve every voxel in the first interrogation block being correlated with a voxel in the second grid will be implemented.

Savings in the computation time are also realized when the correlation function is not normalized.

3.4 Correlation speed

3.4.1 Direct correlation vs. FFT

The decision to compute the correlation function directly or via FFT is essentially based on computation time considerations and we give predictions for the costs of computing the cross-correlation function, for a pair of interrogation blocks, in Appendix C.

We define the correlation computational cost as the number of arithmetic operations required to evaluate the correlation function over the search distances. The larger the number of arithmetic operations, the longer the correlation time.

To correlate a pair of interrogation blocks whose lengths are N, the cost of evaluating the correlation function via radix-2 FFT is approximately $45N^3 \log_2 N$ while the cost of evaluating the correlation function directly is approximately $16N^3\Delta^3$, where Δ is the maximum correlation search distance in any direction.

For practical search distances Δ, the computational cost for Direct Cross-Correlation (DCC) will be greater than that of FFT based correlation; therefore, it would take longer to compute the correlation function directly.

The computational cost for DCC can however be drastically reduced by taking advantage of the sparse nature of the 3D grid. We note that when the particle images are reconstructed on a 3D grid, many of the grid voxels will not belong to any particle. The intensities of these voxels will be zero therefore they will not contribute to the value of the correlation function. When these zero-intensity voxels are omitted during the evaluation of the correlation function, the computational cost for DCC is reduced to approximately $16\gamma N^3\Delta^3$, where γ is the compression ratio and is defined as the number of non-zero intensity voxels divided by the total number of voxels in the grid.

We refer to the processing where only the particle voxels are used to evaluate the correlation function as Direct Cross-Correlation of a Compressed grid (DCCC).

A comparison between the predicted computing cost for DCCC and FFT reveals that correlations via DCCC is faster than FFT if $45 \log_2 N > 16\gamma\Delta^3$.

As an example, if an interrogation block of size 32^3 voxels contains 4 particles and each particle is reconstructed as a cube of size 3^3 cubic voxels, the compression ratio γ for the block will be:

$$\gamma = 4 \times \frac{3^3}{32^3} = 0.0033$$

Now assuming that the maximum search distance is 10 voxels, then the cost of computing the correlation function using the DCCC method will be:

$$16 \times 0.0033 \times 32^3 \times 10^3 \approx 1.7 \times 10^6 \text{ arithmetic operations.}$$

For correlation via FFT, the computational cost evaluates to:

$$45 \times 32^3 \times log_2 32 \approx 7.4 \times 10^6 \text{ arithmetic operations.}$$

Therefore, in this example, the DCCC approach is about 5 times faster and should be chosen if achieving high correlation speed is paramount.

Figure 3.6 gives a comparison of the computational cost between FFT and DCCC over a range of compression ratios, for three different block sizes.

For an even comparison, the search distance for DCCC was chosen to be half the block size to match the maximum displacement that can be correctly resolved by FFT. However, we note that the search distance for the DCCC approach can be specified independent of the interrogation block size.

From the derived expressions for the computational costs given in Appendix C, we note that the relative speeds between both approaches is a function of $\gamma\Delta^3$. Figure 3.7 gives the predicted relative speeds as a function of $\gamma\Delta^3$ for the same three block sizes of 32, 64 and 128 in Figure 3.6.

The range of correlation block sizes that would typically be specified during analysis will typically fall between 32 and 128 voxels. The prediction curves for interrogation block sizes of 32 and 128 voxels will therefore bound the values of the relative speeds between the two approaches when the interrogation block sizes fall within this range. From Figure 3.7, we can therefore deduce that DCCC is faster than FFT-based correlation when $(\gamma\Delta^3)^{-1}$ is greater than about 0.1.

We confirm the validity of the derived expressions for the computation costs, by comparing the predicted correlation times with those from actual analysis of simulated images. The comparisons which are presented in Table 3.1 and Figure 3.8 show that the predicted and actual relative speeds are in good agreement.

During PIV experiments, the values of the correlation parameters which determine the relative speeds between both approaches will first and foremost be chosen to give a high valid-data-yield. We therefore proceed to compare how the speeds of both FFT based correlation and DCCC are affected as changes are made to the correlation parameters in an effort to obtain high valid-data-yield measurements.

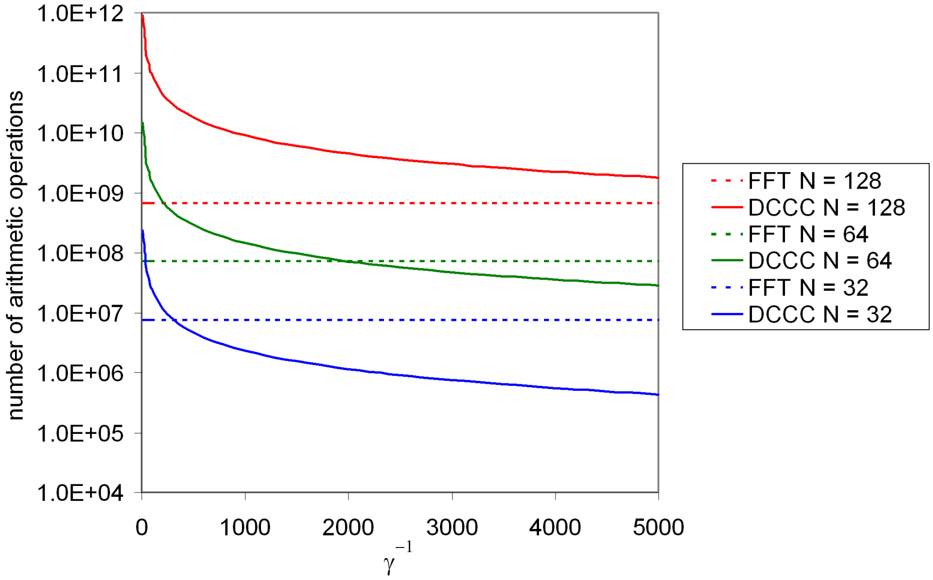

Figure 3.6: Prediction curves for the number of arithmetic operations needed to cross-correlate a pair of interrogation blocks directly (DCCC) and via FFT. Shown are predictions for three iterrogation block sizes over a range of compression ratios.

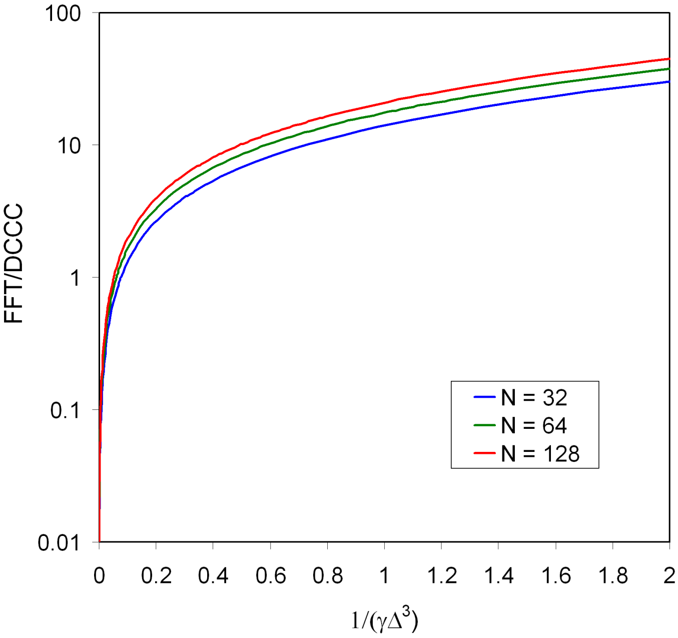

Figure 3.7: Semi-log plot comparing the costs of correlations between FFT based processing and DCCC.

$\frac{1}{\gamma\Delta^3}$	γ^{-1}	DCCC (sec)	FFT (sec)	FFT/DCCC	Predicted FFT/DCCC	Number of particles per block	Interrogation block length N	Search distance Δ
0.034	138	315	95	**0.3**	**0.44**	2	32	16
0.038	1241	208	125	**0.6**	**0.62**	8	64	32
0.075	2457	100	125	**1.25**	**1.24**	4	64	32
0.13	533	46.4	125	**2.69**	**2.05**	4	64	16
0.15	4912	51.2	126	**2.46**	**2.49**	2	64	32
0.155	634	45	94	**2.09**	**2.06**	2	32	16
0.254	1040	26.2	125	**4.77**	**4**	4	64	16
0.303	1241	27.5	124	**4.51**	**4.78**	8	64	16
0.6	2457	12.8	124	**9.69**	**9.48**	4	64	16
1.2	4912	5.6	125	**22.3**	**19.04**	2	64	16

Table 3.1: Comparison between actual relative correlation times and predictions.

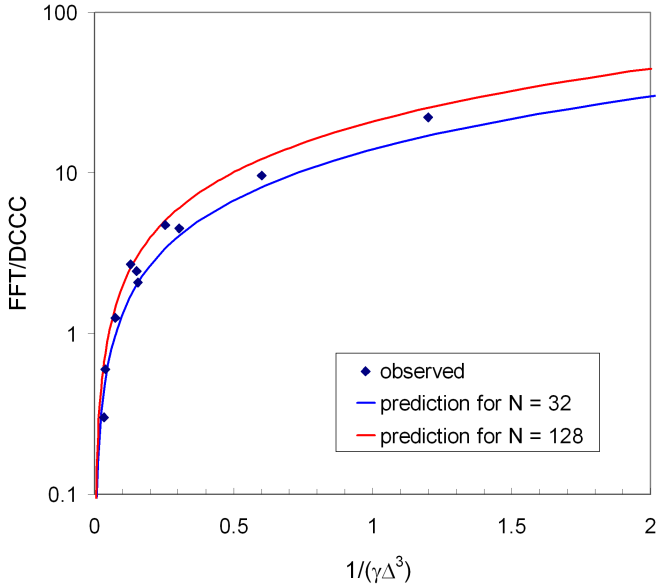

Figure 3.8: A plot of the observed relative correlation times in table 3.1. Also shown are prediction curves for interrogation block sizes of 32^3 and 128^3 voxels.

3.4.1.1 Effect of interrogation block size on correlation time

A major benefit offered by DCCC processing is that the size of the interrogation blocks can be specified to the exact required size. For radix-2 FFT based processing, the size of the correlation blocks must be a power of 2, therefore given an interrogation block size, the FFT block sizes will need to be increased to the next power of 2 size.

Although larger interrogation blocks reduce the number of blocks to interrogate in the grid, for DCCC processing, the overall cost of evaluating the correlation function for the entire grid is independent of the block size. However, for FFT based processing, the size of the correlation blocks affect the computational cost, and although larger correlation blocks mean less blocks to process in a grid, the overall computational cost increases. For example, increasing the FFT correlation block size from 64^3 to the next size of 128^3 voxels increases the correlation time by about 17%. See Appendix C.4.

3.4.1.2 Effect of particle displacement on correlation time

For FFT based correlation, the interrogation block size restricts the maximum resolvable particle displacement to half the block size. The chosen block size must therefore be at least twice the expected maximum particle displacement in order to resolve the highest displacements. When the FFT block size is increased, then the correlation time also increases.

For DCCC, the resolvable range of particle displacement is not strongly tied to the size of the interrogation block. The block sizes can be set independent of the particle displacements. In DCCC, larger particle displacements would require increased search distances and this will come at an increased cost to the computational time. The computational cost increases as the cube of the search distance.

We also note that regardless of the correlation method, larger search distances will generally produce, in the correlation table, more peaks due to randomly matched particles, thus increasing the likelihood that the highest peak will be a spurious vector peak. This likelihood increases with higher relative particle displacements that will typically accompany greater particle movements.

3.4.1.3 Effect of out-of-boundary motion of particles on correlation time

The movement of the particles across the block boundaries reduces the signal-to-noise ratio in the correlation table, hence lowers the valid-data-yield probability from the correlations. To maintain a high signal to noise ratio in FFT based processing, the sides of the second interrogation block in the pair can be extended so that all the particles in the first interrogation block are retained in the second interrogation block. The first correlation block is then “zero-padded” up to the same size as the second block before the pair is correlated. The increase in block size will however be accompanied by an increase in the correlation time as discussed earlier.

For DCCC processing, the out of boundary motion effect can be minimized by simply searching for cluster pattern matches on the second grid instead of explicitly defining the interrogation blocks on the second grid. No additional correlation times are incurred when the correlations are performed in this manner since the block sizes are not increased.

3.4.1.4 Effect of high velocity gradients on correlation time

Large particle displacements also give rise to larger relative particle displacements within the interrogation blocks. The relative particle displacements within the blocks will also increase if larger interrogation blocks are defined. But an increase in the relative particle displacements reduces the likelihood of obtaining valid vectors from the correlations.

To achieve a high valid-data-yield probability when the displacement gradient within the interrogation block is high, a bigger particle size can be specified during the reconstruction of the particles on the 3D grid. But larger particles reduce the compression ratio and for DCCC, the correlation cost increases due to the increased number of particle voxels to correlate. The correlation time for FFT based processing is however independent of the number of particles or the particle size, so for FFT, additional correlation time is not incurred when larger particles are constructed to overcome high relative displacements in the blocks.

3.4.1.5 Effect of choice of reconstructed particle model and size on correlation time

Between a cubic model and a sphere model, cubic models will typically take up more grid voxels, for a specified particle reconstruction size. Also, an increase in the model size will mean more grid voxels will be needed to construct the particle and more computer memory to store the voxel information.

The more voxels the model takes up, the less compressed the grid will become, and for DCCC, this will mean an increase in the correlation time. Increasing a cubic model size by increasing its length by a factor of 2 reduces the compression ratio by a factor of 8 (i.e. 2^3).

For FFT based correlation, the model type and size does not affect the correlation time.

3.4.1.6 Effect of grid resolution on correlation time

The grid resolution also affects the cost of computing the correlations. The coarser the grid, the less number of voxels representing the same volume in the flow field. Each interrogation block will therefore consists of a lesser number of voxels and the number of voxels needed to model the same particle size is less. The search distance, in terms of the number of voxels, is also reduced. A coarser grid size therefore reduces the number of arithmetic operations involved in computing the correlation function using both FFT and DCCC.

Changing the grid resolution however has a more pronounced effect on the DCCC processing time due to its cubic variation with both search distance and block size. This is illustrated in Figure 3.6 where we see that the relative cost of computing the correlation function changes with block size for a fixed compression ratio. For DCCC, evaluating the correlation function on a grid, which is 4 times coarser, improves the correlation speed by a factor of about 4000![1] So, grid correlations which take 10 seconds to complete would, on a grid which is 4 times

---

[1] For particles modelled as cubes, a grid which is 4 times coarser in all directions reduces the number of particle voxels by a factor of 64, assuming that the initial particle size on the grid is a multiple of the grid factor. The number of search coordinates in the correlation table is also reduced by a factor of 64 (i.e. by a factor of 4 in each of the 3 directions) to give a total reduction in computational cost by a factor of almost 4000.

finer in all directions, take about 11 hours!

3.5 Correlation valid-data-yield

Successful pairing of individual particle images relies on obtaining a good displacement estimate from the correlations so, the grid voxel scale should be chosen carefully. Defining a grid that is too coarse will result particles being placed poorly on the grid, which in turn introduces artificial particle displacement gradients. Furthermore, as the grid voxel size approaches the size of the reconstructed particles, each particle will be represented by a single voxel on the grid. The flow velocity gradients may now be too high to expect a high probability of obtaining a valid displacement vector from the correlations.

We show the effect of grid resolution on the correlation outcome in Figure 3.9 where the same pair of simulated particle images is analysed — using DCCC based Hybrid PIV-PTV technique — on two different grid resolutions.

The simulated velocity field is a ring of straight-line vortex rods, whose details are given in Appendix D.2.5. Figure 3.9(a) and Figure 3.9(b) both show the velocity vector field at the end of the PIV stage (left vector field) and the subsequent particle tracking vector field (right vector field). The results presented in Figure 3.9(b) were obtained from correlations on a grid whose voxels are half the size, in all directions, to that used to produce the vector field in Figure 3.9(a).

We see that for the coarser grid, there are so many spurious vectors that it is difficult to see the flow structure. While there are still some spurious vectors for the results obtained using the finer grid, the number of spurious vectors is much less and the flow structure can be observed.

3.5.1 Parameter settings for high correlation valid-data-yield

During the correlation process, when the out-of-boundary particle motion is compensated for by either defining a larger interrogation block or by floating the interrogation block in the second grid, then the factors which determine the correlation outcome are the number of particles that are contained in the blocks,

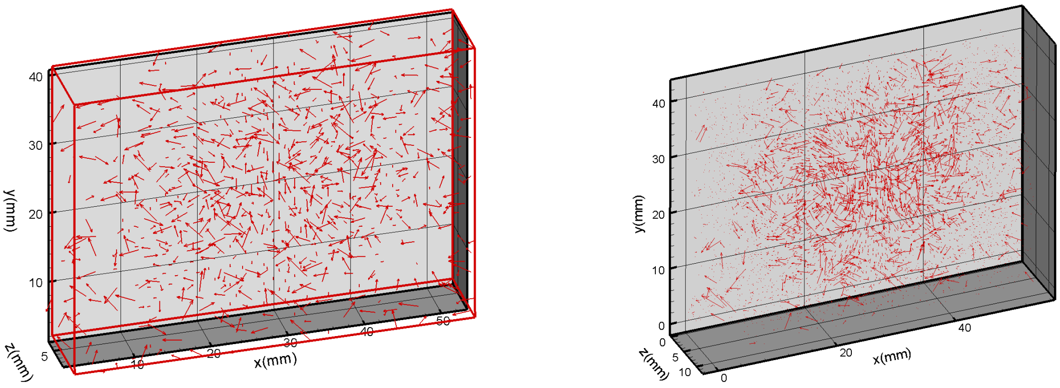

(a) Hybrid PIV-PTV analysis results obtained using a coarse grid.

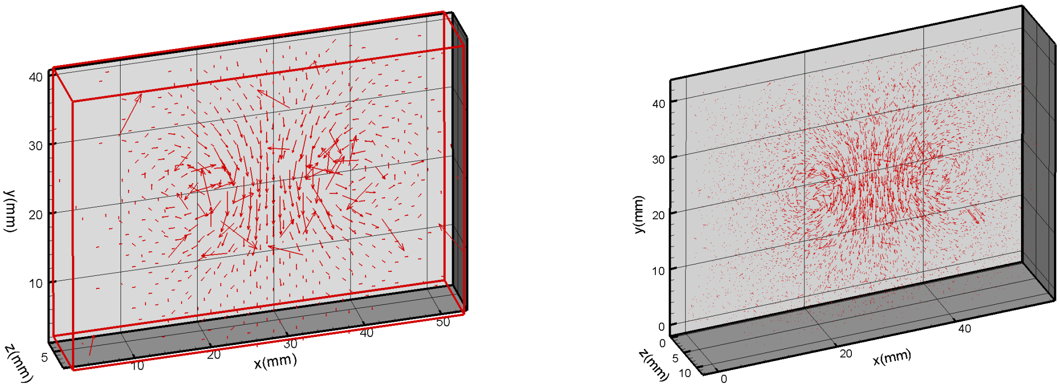

(b) Hybrid PIV-PTV analysis results obtained using a finer grid.

Figure 3.9: The unvalidated velocity vector field, for a pair of simulated particle images processed on a coarse and fine grid. Coarse grid voxel size: 400 microns. Finer grid voxel size: 200 microns. Specified reconstructed particle image size for both grids: 400 microns.

the reconstructed particle size and the relative particle displacement in the block.

How these correlation parameters combine to affect the correlation valid data yield can be deduced from the scatter plot in Figure 3.10 which gives the correlation valid-data-yield probability for several combinations of block seeding density, particle size and particle displacement gradients in a simulated shear flow.

The simulated shear is over the z-plane and the average block displacement gradient Δ_z in the correlation blocks is defined as $\frac{1}{2}\frac{\partial u}{\partial z}N\Delta t$ where N is the block size, Δt is the illumination time interval between the two images of the simulated flow field and u is the velocity along the x-direction[1].

We observe from Figure 3.10 that for a high valid-data-yield, the flow velocity gradient that can be tolerated depends both on the number of particles and the size of the particles in the blocks. For a given particle displacement gradient, the chances of obtaining a valid vector from the correlations are higher when the seeding density is higher. This is because higher seeding densities increase the chances of more true particle pairs correlating at similar rates while also reducing the chances of neighbouring particle clusters correlating better than the true clusters.

We note that for any meaningful correlation, there must be at least 2 particles in the interrogation blocks. From Figure 3.10, we deduce that when the number of particles in the blocks is more than 2, and the relative particle displacement is less than about 1 particle diameter, then we can expect a more than 95% chance of obtaining a valid vector. Hence, we recommended that to achieve a valid-data-yield greater than 95%:

- the block size should contain at least 3 particles
- and the relative particle displacements in the blocks should be no greater than about 1 reconstructed particle diameter.

It is important to note that while constructing bigger particles may give seemingly valid vectors, the resulting approximation of the cluster displacement is poorer

---

[1] We chose this definition for Δ_z to be similar to the definition of the average window displacement gradient definition given in the 2D simulation experiments reported by Keane and Adrian (1992).

when the relative particle displacement is of the same order as the flow length scales. More particles may then be incorrectly paired at the tracking stage of the Hybrid PIV-PTV process.

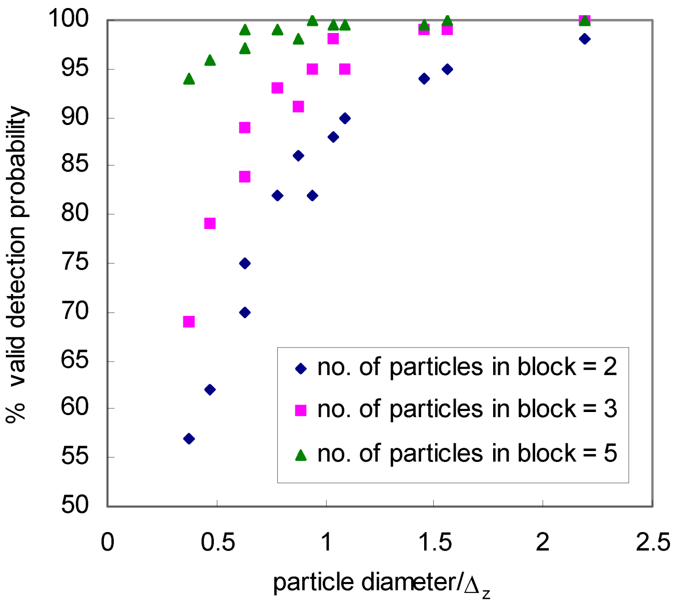

Figure 3.10: The effect of particle size, seeding density and flow velocity gradients on correlation valid data yield.

3.5.2 Vector validation and fixing

In real experiments, it may not be possible to set the correlation parameters to ensure that these conditions for high valid-data-yield are met for all the interrogation blocks. For example, the distribution of the particles in the flow field is likely to result in many situations where some of the interrogation blocks will contain less than the desired number of particles to ensure high valid -data-yield probability. It is therefore important that spurious vectors resulting from the correlation of such blocks are identified and removed to keep the valid-data-yield high.

In this section we present an effective novel vector validation and fixing algorithm, which can be incorporated into the PIV measurement chain.

3.5.2.1 Correlation peak stability validation

We observe that since the noise peaks in the correlation table are due to random correlations of the particles in the interrogation blocks, the positions of these noise peaks in the correlation table are also random. So, under the same flow conditions, a different set of correlating particles in the interrogation blocks should produce a different set of noise peaks in the correlation table. The only region in the correlation table where a peak will always be expected to be located regardless of the distribution of the particles are the coordinates that correspond to the average block displacements. Therefore, if one of the noise peaks in the correlation table is the highest peak, then we can expect the highest peak to be located at a different set of coordiantes if the distribution of the correlating set of particles is changed while still preserving the flow condition.

Such changes to the correlating set while still preserving the flow information can be achieved by simply reversing the order of the correlating images and correlating the interrogation blocks in the first grid over a region in the second grid whose size is dictated by the specified correlation search distance.

Assuming that the displacement gradient across the interrogation block is less than a particle diameter, the distance between any two peaks obtained from two different correlating sets of particles, should be within a particle diameter from one another if both peaks are formed as a result of the matching of true particle pairs. If the distance between the peaks is greater than about 1 particle diameter, then at least one of the peaks must be a noise peak and the interrogation block vector is tagged as spurious.

We make the following notes when Peak stability validation is achieved by reversing the order of the camera images:

- The correlations must imply an actual search for the best pattern match over the second grid, instead of simply ‘cutting out’ corresponding interrogation blocks of the same size from both grids. This is to allow for the reverse correlation to be different (not just in sign) from the forward correlation.

- When the order of the correlating images is reversed to obtain the second set of correlations, the signs of the coordinates of the resulting peak from the correlation must also be reversed before it is compared to the coordinate of the peak from the first set of correlations.

- The peak stability validation technique is most effective when the seeding density is high and the correlating search distance Δ is large.
- Peak-stability validation, as described here, adds to the overall interrogation time since one more set of correlations has to be performed to obtain vectors for the case where the order of the images is reversed.

3.5.2.2 Fixing spurious vectors

Preferably, the reconstructed particle sizes should be specified to be large enough to overcome the effects of local flow gradients and any artificial gradient that may have been induced by the uncertainty introduced during the determining the particle spatial coordinates and the imprecise positioning of the particles on the grid. But specifying an optimum particle size for high valid-data-yield can be difficult especially as the velocity gradients will vary from one region of the flow field to another. Too small a particle size will lead to spurious vectors due to larger relative particle displacements while a large size will slow down DCCC-based correlation.

We note that large velocity gradients will sometimes be confined to only a small area in the imaged flow field, hence, the number of interrogation blocks requiring larger particles will be a small fraction of the total number of interrogation blocks. In such cases, large savings in correlation time and high valid-data-yield can be achieved if the correlations are first performed using smaller particles followed by the re-correlation of the spurious blocks using a larger diameter particle set.

One way to incorporate this vector validation and fixing routine in the image analysis process will be to prepare larger particle sets alongside the original correlating set. The first set is used in the initial correlations while the 'larger particle-size' sets will be used during the subsequent re-correlations of spurious blocks. While as many sets of particles of increasing sizes could be prepared and spurious blocks can be re-correlated using increasingly larger particle sets until they are tagged as valid, it is recommended that just one other set be prepared.

The need to use excessively large particles just so that the block vectors can be tagged as valid indicates that the gradient is too high to even make the vector accurate enough to aid in the tracking of the particles.

The Peak-Stability vector validation and fixing process is shown in Figure 3.11. Here, the diameter of subsequent sets of reconstructed particles is increased over the previous set, by an amount equal to the initially specified particle diameter to be used for the original correlations.

We demonstrate the higher effectiveness of the peak stability validation method over the detectability-based validation method by applying both validation methods to detect the spurious vectors in Figure 3.9(b). Figure 3.12 shows the velocity vector field before and after detectability-based validation was applied while in Figure 3.13, we show the velocity vector field before and after Peak stability validation was applied to the data. We also show, in Figure 3.14, the results when some of the blocks identified as spurious during peak stability validation, were fixed.

We notice a much improved valid-data-yield with the Peak-Stability validation technique over the detectability-based validation method. For the Detectability validation method, we see that at the detectability value of 1.1, only some of the spurious vectors were removed. Some of the seemingly 'valid' vectors were also identified as spurious. Although it is expected that more spurious vectors will be identified as such at higher detectability value settings, more 'valid' vectors will also be removed thus, reducing the spatial resolution of the velocity data.

Next we present the image interrogation procedure that was designed specifically for the colour flow images acquired with the 3D colour PIV system described in Chapter 2.

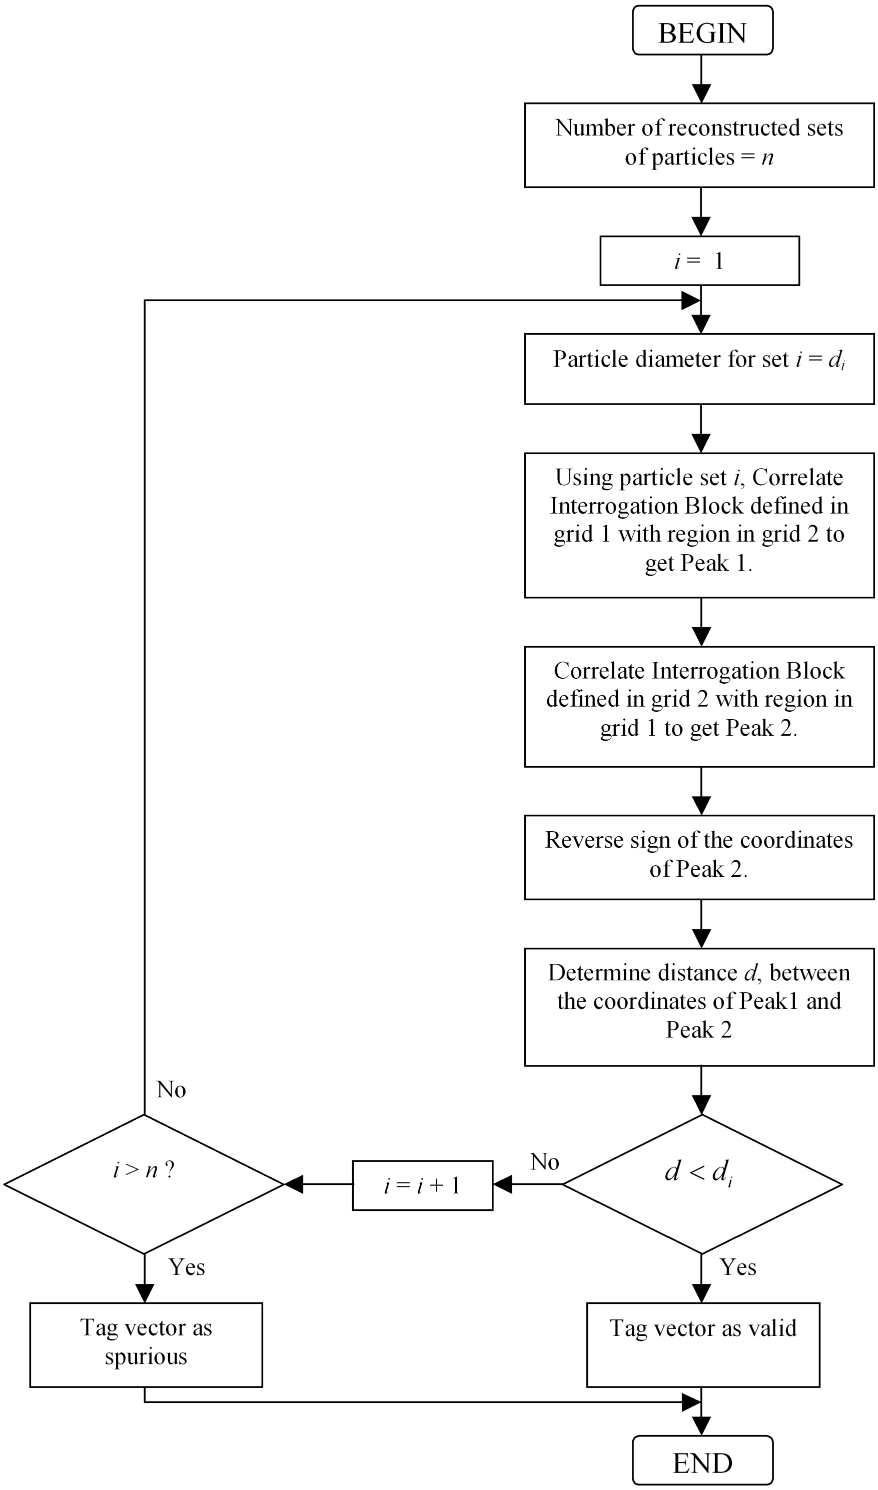

Figure 3.11: The Peak Stability validation and spurious vector fixing procedure.

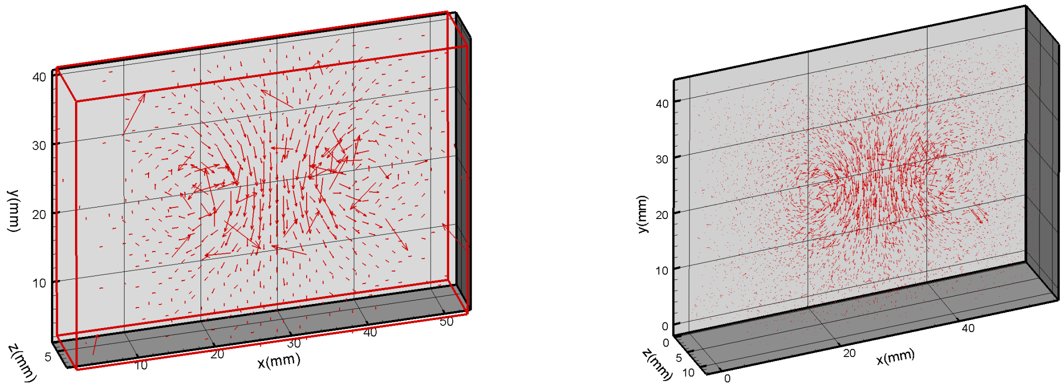

(a) Unvalidated PIV and PTV vector fields

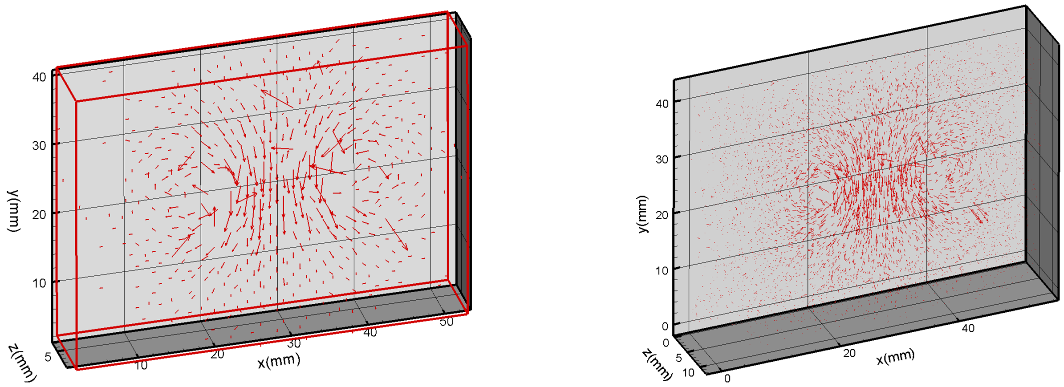

(b) Validated Vector fields using Detectability Validation value of 1.1.

Figure 3.12: Comparison between an unvalidated vector field and the same vector field after applying the ‘detectability’ validation.

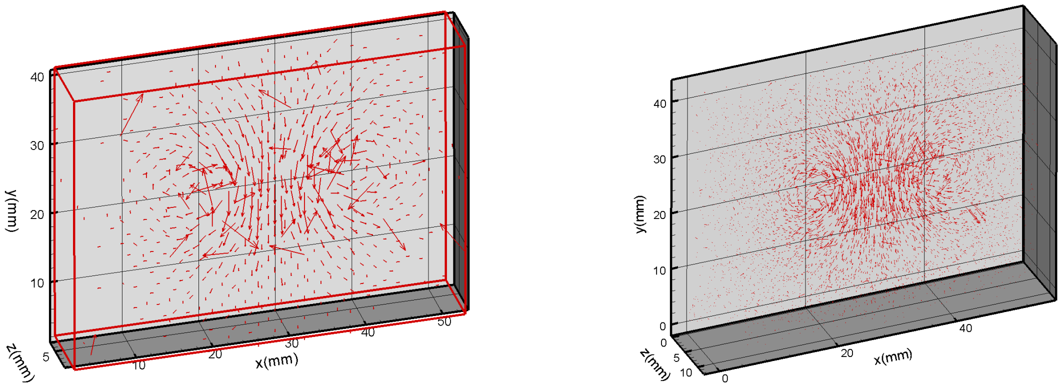

(a) Unvalidated PIV and PTV vector fields

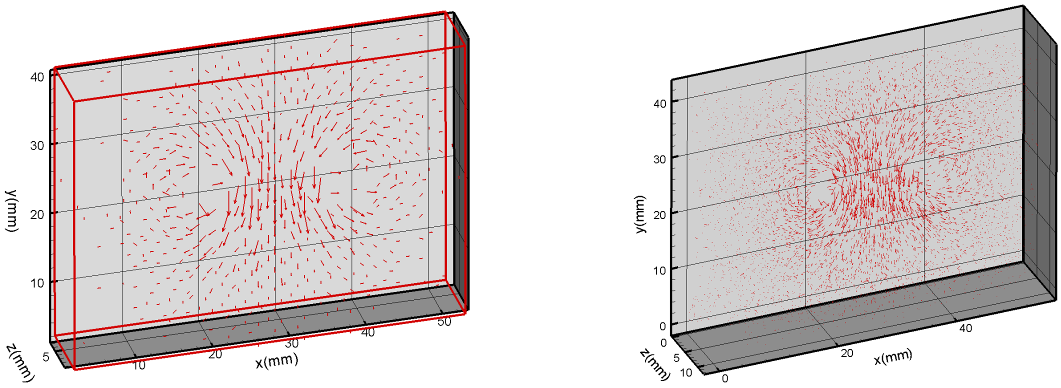

(b) Validated Vector fields using the Peak Stability validation

Figure 3.13: Comparison between an unvalidated velocity vector field and the same vector field after applying the peak stability validation method.

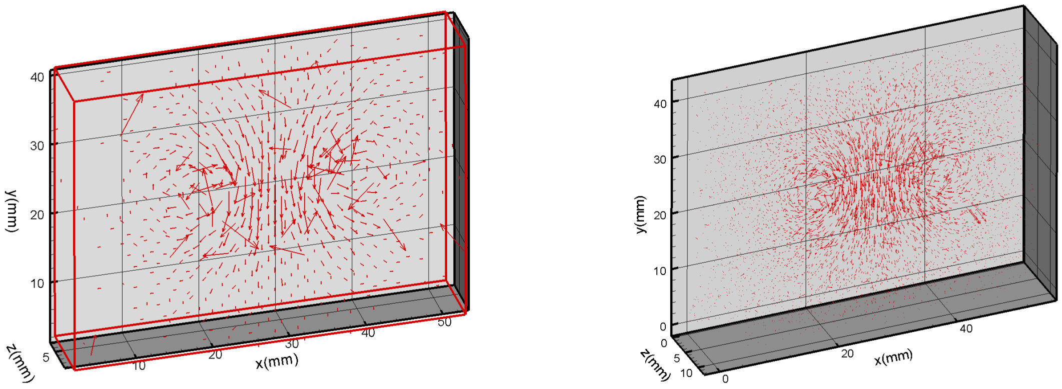

(a) Unvalidated PIV and PTV vector fields

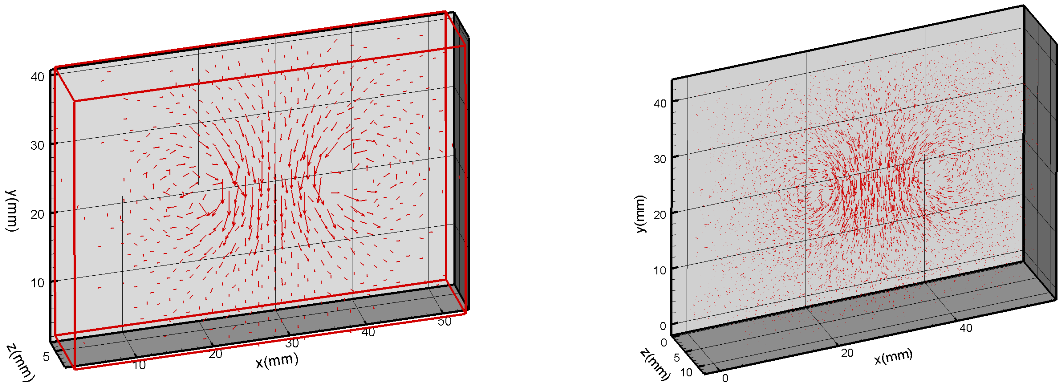

(b) Validated Vector fields using the Peak Stability validation plus Vector Fixing

Figure 3.14: Comparison between an unvalidated velocity vector field and the same vector field after applying the peak stability validation plus the vector fixing routine.

Chapter 4

The 3D colour PIV-PTV system: Image Analysis

4.1 Introduction

In this chapter, we describe how the digital colour images acquired with the experimental arrangement and procedure described in Chapter 2 were interrogated to give the 3D velocity information for the region represented by the images.

A flowchart of the image analysis procedure is presented in Figure 4.1. The process of analysing the images can be divided into the following stages:

- Particle Image identification
- Determination of the three-dimensional coordinates of the particle images
- Particle image pairing
- Data validation and fixing
- Post-processing

We explain these stages in the remainder of the chapter.

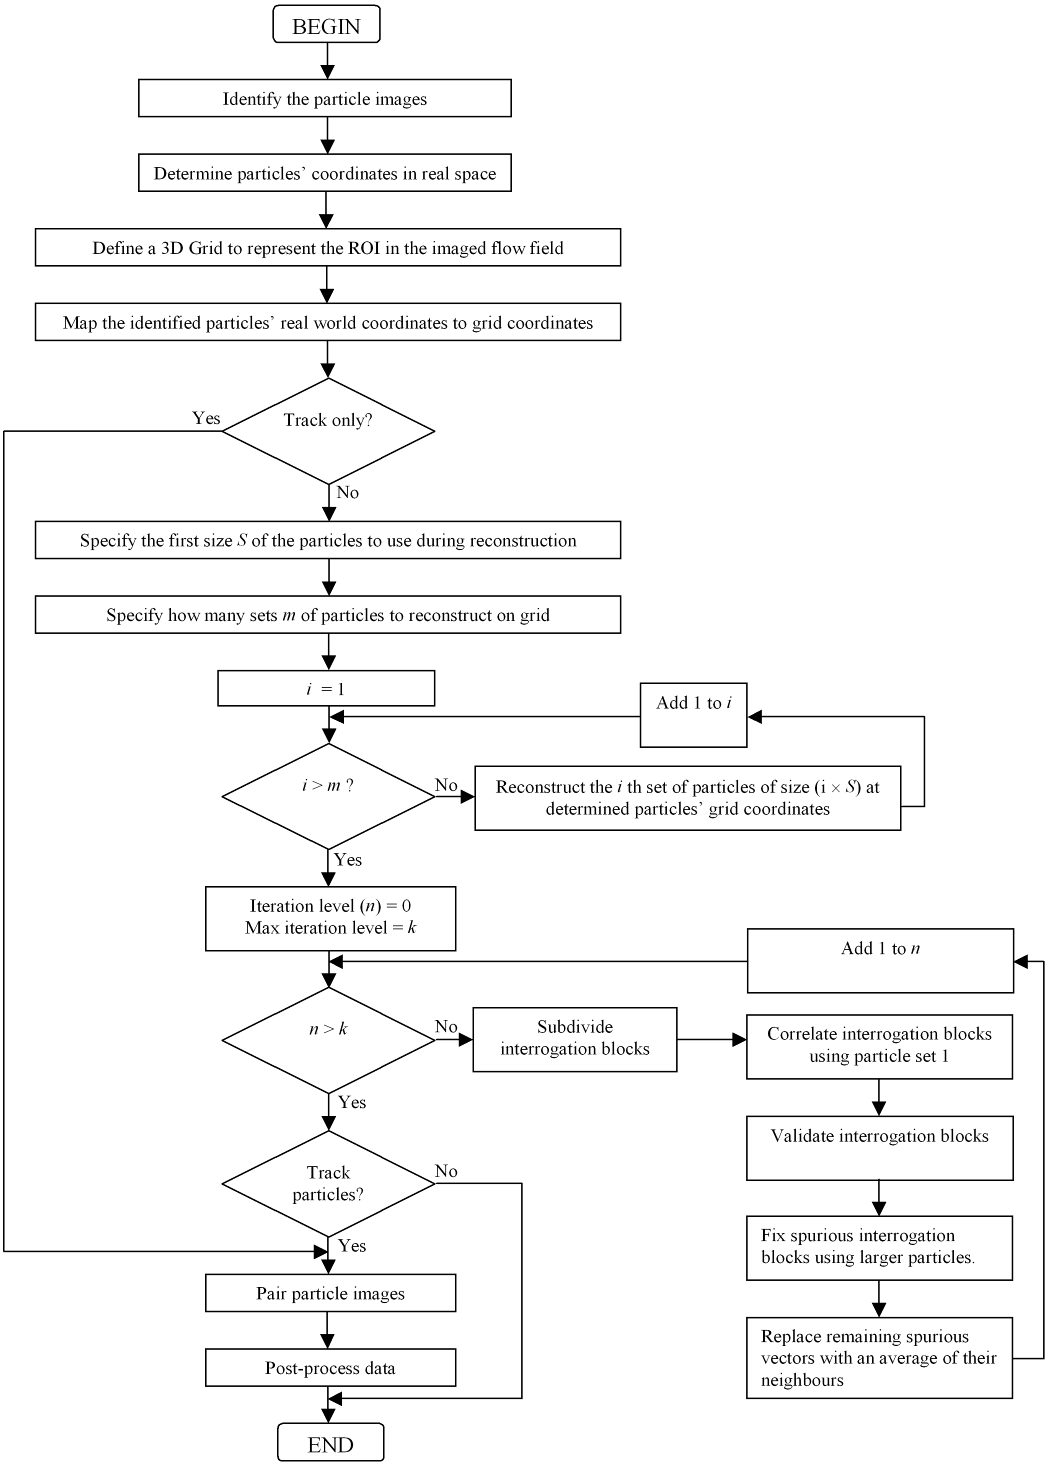

Figure 4.1: 3D image analysis process.

4.2 Particle image identification

The aim at this first stage is to correctly identify the images of the tracer particles from the digital images.

The task of identifying particles from the images has been a subject of extensive research due to the impact of the process on the valid data yield. The extraction of particle images from the digital images is a non-trivial process and care must be taken to ensure that as many of the particles are identified with a minimum number of erroneous identification of background pixels as belonging to particles. False identification of particles would result in a contamination of the particle data set, which could result in true particles being erroneously matched with false particles. Another possibility presented by such false particle identification is that true particles could be denied their true pair because they may have already been matched with false particles thus reducing the valid data yield and the number of correctly paired particle images.

While a lot of attention has been paid to particle detection methods, there appears to be no universal particle-identification algorithm that can effectively accomplish the task of detecting particle for an extensive range of experimental and imaging conditions that could be presented by experimental factors such as camera noise, reflective boundaries, non-uniform illumination, irregularly shaped particles and excessive particle overlaps. Particle detection algorithms are therefore usually customized to give optimum performance for the specific conditions with which one is presented. The task of identifying particle images can however be made a lot easier if proper steps are taken to improve the experimental conditions.

The high contrast digital images acquired in this work and the near-absence of boundary reflections makes the task of identifying the particle image pixels easier. For the detection of the colour particle images, a simple, yet robust particle identification algorithm was developed based on identifying the peak intensity pixel of the particle images. This fast algorithm was found to be highly efficient at detecting and splitting overlapping particle images and its effectiveness in identifying particle images is also independent of particle image size or particle image intensity profile.

The particle image identification procedure is as follows:

First, a greyscale image is constructed from the highest intensity among the three colour components for each pixel in the colour image.

The particle images in the greyscale image were then identified by locating their highest-intensity pixels. This was achieved by putting each pixel in the greyscale image through a set of elimination criteria to either accept it as a particle image peak intensity or reject it.

A pixel is tagged as the highest-intensity pixel of a particle image if it satisfies the following set of elimination criteria:

1. The pixel intensity must be higher than a threshold value. The threshold value is defined as the sum of the mean and twice the standard deviation of the pixel intensities in the image.
2. The pixel intensity must be higher than its nearest eight neighbours.
3. At least three of the pixel's nearest neighbours must be higher than the threshold intensity.
4. The average intensity of the nearest eight neighbours must be higher than the threshold value.

Each identified peak and its coordinates are taken to represent a particle image and its approximate location on the colour image.

In the occasional cases where a particle is represented by two or more adjacent peaks, all but one of the pixels in question are removed. Such cases are identified and fixed by scanning the resulting binary image, line-by-line and eliminating neighbouring peaks that are connected horizontally, vertically or diagonally to the current peak, as the scan progresses.

The process of constructing the grey image from the colour image and the subsequent extraction of the particle image peak-intensity pixels is depicted in Figure 4.2.

4.3 Determining the coordinates of the particles

For each identified particle, two sets of 3D coordinates are determined:

- We first determine the particle's real world coordinates.
- We then determine the particle's coordinates on a 3D grid on which the particle can be reconstructed.

4.3.1 Determining the particles' real-world coordinates

The steps involved in determining the spatial coordinates (x_p, y_p, z_p) of the particles are as follows:

1. Determine the particle image colour.
2. Relate the particle image colour to z-coordinates in space, z_p.
3. Determine particle image centroid (x', y').
4. From (x', y') and z_p, determine the x- and y-spatial coordinates (x_p, y_p) of the particle, using equation (3.1a) and (3.1b).

To determine the colour of a particle, the intensity of the image formed by the particle on each of the three CCDs is determined and the ratio of each intensity to the sum of all three intensities. This process is illustrated in Figure 4.2. See also section 2.5. The intensity of each particle image is determined by summing the pixel intensities over a specified area centred at the highest-intensity pixel of the particle image on each CCD. The pixel summation area was chosen to be just large enough to ensure satisfactory convergence of the colour ratios but still small enough to avoid contributions of pixels whose intensities are significantly dominated by the camera noise. A summation area of size 7×7 pixels was determined to be adequate for the particle images acquired in this work.

To further mitigate the influence of camera noise, we also subtract, from each pixel involved in the summation, the expected mean contribution to the intensity of the pixel, from the camera noise. We show the variation of the mean camera noise with gain setting in Figure 2.13. The pixels whose intensities are less than the camera noise intensity are omitted from the sums.

With the particle image colours resolved, the real world z-coordinate of each particle is subsequently obtained by comparing its image colour to those in the database of calibration points. The database of calibration points contain colour vs. position information for a set of calibration grid points defined within the illuminated volume in space (see section 2.7.2). To determine a particle's z-coordinate in space, we extract from the database, those calibration points whose image coordinates are within a specified radial distance (in pixels) from the image coordinates of the particle image. The extracted calibration points are arranged in ascending order based on their z-coordinates and then, a piecewise-linear curve is fitted through the points. The z-coordinates of the point on the curve that corresponds to the colour of the particle image is subsequently assigned as the particle image's real world z-coordinate.

The real world x- and y-coordinates of the particles (x_p, y_p), were subsequently determined by first locating their image centroids (x', y') on a designated CCD channel and then relating the particle image centroid position back to real world coordinates, using equation (3.1a) and (3.1b).

In this project, the particle image centroid was calculated using the green CCD particle images. We note from the definition of the particle image colour that the image of each particle must be recorded on at least two CCDs in order for the particle's colour to be resolved. The chosen CCD on which the particle image centroids are determined, must be one of the CCDs on which the particle images are properly recorded.

The particle image centroid was calculated using the pixels in the 7×7 pixel area centred at the highest-pixel intensity on the green particle image and applying the 2D centroidal approximation scheme given in Appendix E.2.

4.3.2 Converting real world coordinates to grid coordinates

To track the particles via the Hybrid PIV-PTV scheme, we chose an approach where the particles are first reconstructed on a 3D grid and subsequently correlated to give a first estimate of the particle displacements.

The size of the 3D grid (in voxels) is determined by dividing the size of the Region of interest (ROI) — in real world dimensions — by the voxel-scale specified by the user. The grid size (i.e. the number of grid voxels) along the i-direction,

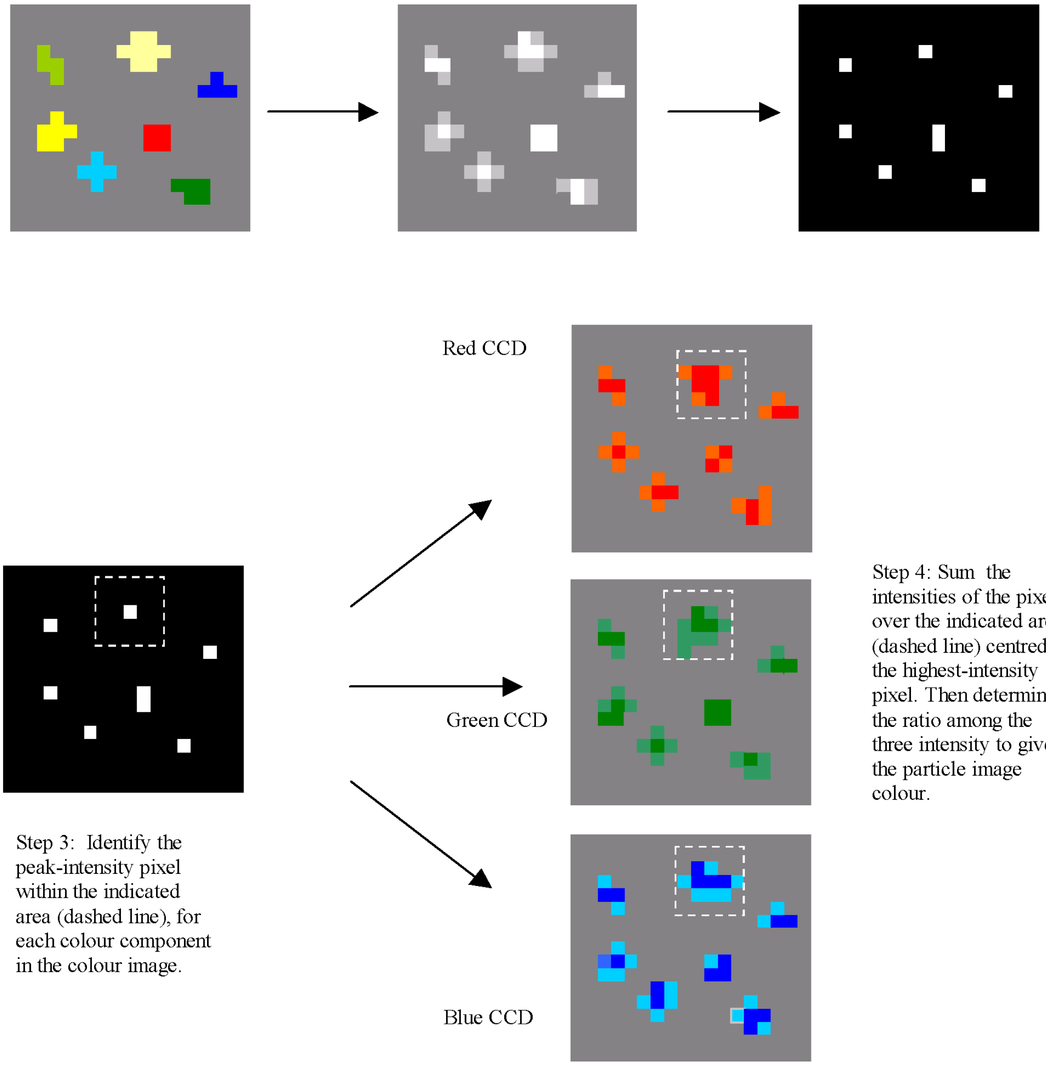

Step 1: Construct a grey image from colour image by extracting the highest-intensity colour component for each pixel.

Step2: Extract the highest-intensity pixel of each particle image from the grey image.

Step 3: Identify the peak-intensity pixel within the indicated area (dashed line), for each colour component in the colour image.

Step 4: Sum the intensities of the pixels over the indicated area (dashed line) centred at the highest-intensity pixel. Then determine the ratio among the three intensity to give the particle image colour.

Figure 4.2: The steps involved in the identification of particle images and the subsequent determination of their colours.

is expressed as:

$$\text{(grid size)}_i = \frac{\text{(ROI size)}_x}{\text{(Voxel Scale)}_x}$$

The grid size along the j and k direction are similarly defined.

The particle coordinates on the grid are subsequently obtained by dividing their real world coordinates relative to the ROI boundaries by the size of the region of interest i.e.:

$$\text{(particle grid coord)}_i = \frac{\text{(real world coord)}_x - \text{(ROI min coord)}_x}{\text{(ROI size)}_x}$$

The particle coordinates along the j- (y-) and k- (z-) grid directions are similarly expressed.

4.3.3 Representing the particle coordinates as a single value

Along with the particle coordinates in the real-world space and grid space, a unique index is also assigned to each particle to facilitate the referencing of its grid coordinates relative to the coordinates of other particles.

A lookup value is defined for each particle according to the expression:

$$\text{lookup}\,(i, j, k) = (k \times S_x \times S_y) + (j \times S_x) + i \tag{4.1}$$

where i, $j$and k are the particle grid coordinates (to the nearest integer voxel value) along the x, y and z-directions respectively and S_x and S_y are the number of voxels along the x- and y-directions respectively in the grid.

The particle coordinates and their lookup values are subsequently stored in order of their increasing lookup numbers as required by the fast algorithm that was designed to extract, from the list of particles, only those particles within any specified region in the grid during the correlation and the particle pairing process. We give details of this algorithm in section 4.4.2.3.

4.3.4 Removing invalid particles

We note that some of the particles identified in the images may need to be removed either because their colours are inaccurately resolved or the particle coordinates fall outside the Region of Interest.

We define invalid particles as:

1. Saturated particle images.
2. Particle images outside ROI.

These invalid particles are removed from the list of identified particle images.

4.4 Particle image pairing (particle tracking)

With the spatial coordinates both in the real world and on the grid now determined, the particle displacements can subsequently be determined by pairing particle images belonging to the same particle in two successive images.

Particle pairing is achieved either directly using the "nearest neighbour" approach or via the Hybrid PIV-PTV approach where an initial estimate of the displacement of clusters of particles are first obtained using correlations before implementing a search around the estimated displacement for the nearest neighbour for each individual particle within the cluster. We describe both approaches to particle tracking next.

4.4.1 Direct particle pairing

To pair the particles directly after their coordinates have been determined, a search is implemented on the second grid, around the coordinates of the particles in the first grid.

The particles from the first image are paired with those from the second image by progressively enlarging the search radius up to a prescribed limit (typically of the order of the average particle separation), centred at the initial particle coordinates. At each search radius, particle pairs that were identified were saved and removed from the list of particles available for pairing.

To speed up the pairing process, only the particles within the search radii were considered for pairing. These candidate particles were extracted from the list of particles using the fast particle extraction algorithm described in section 4.4.2.3.

We note that for direct particle tracking, a high pairing success rate is typically achieved only if the maximum particle displacement is much less than the average particle separation.

4.4.2 Particle pairing via grid correlations

When the particle spacing/displacement ratio is likely to result in more than one candidate pair being available within the search area defined on the second image, then the probability of mismatching the particles increases as the true pairs move away from the initial position and neighbouring particles move nearer to the initial position thus, increasing likelihood that the nearest candidate particle is not the correct pair.

In such situations, it is beneficial to obtain an initial estimate for the displacement of the particles so that the search for the nearest candidate can be performed nearer to the true pair.

Obtaining estimates for the particle displacements is accomplished via cross-correlation of clusters of particles on the 3D grid defined in section 4.3.2.

- First, the 3D model of each particle is constructed at its grid coordinates.
- Then the grids are subsequently divided into interrogation blocks and each block from the first grid is correlated with its corresponding pair in the second grid to give the average displacement of the particle cluster in the block.

Once the estimates for the displacements of the particle images are determined, a search for a matching pair is then conducted, on the second grid, around the predicted final position of each particle. The maximum size of the search region is set equal to the size of the particles used in their reconstruction since the final position of each particle can be expected to be within a particle diameter from the estimated position, if the relative particle displacement in the interrogation block is less than a particle diameter.

From the estimated final position of the particles, the search radius is progressively enlarged and as pairs are identified, they are stored and removed from the list of available pairing candidates.

The low spacing-displacement ratio realised in this work, demands that pairing of particles should be accomplished using this hybrid PIV-PTV technique.

4.4.2.1 Reconstructing the particles on the 3D grid

Before the correlations are performed, each particle is reconstructed on the grid. We discussed how the particles can be reconstructed in section 3.3.1 and we present the particle reconstruction algorithm in Figure 4.3.

For the images acquired in this work, we modelled the particles as cubes centred at the particles' coordinates on the 3D grid. A flat intensity profile was applied to the model with the intensity value specified as 1023.

Two sets of particles were constructed, with the particle size of the second set specified to be twice as large as the first set. The first set was used for the correlations while the second set was used to fix vectors that may have been spurious as a result of large particle displacement gradients. We describe the vector validation and fixing routine in section 4.5.

For each set of particles, the voxels in the 3D grid that represent particles were assigned lookup numbers, rearranged and stored in memory. Defining a lookup number for each particle voxel and the subsequent ordering of these voxels based on their lookup number is key to the fast extraction algorithm designed for the extraction of the particle voxels within the interrogation blocks. See section 4.4.2.3.

We note that during the particle reconstruction phase, we did not explicitly define the 3D grid by allocating enough memory to store all the grid voxels. Doing so would require a large memory storage area, most of which would be dedicated to storing 'zero intensity' voxels, which contain no particle information.

The amount of memory required to store all the grid voxels can exceed the amount of memory that is available on today's PCs. For example, given the $50 \times 30 \times 10\text{mm}^3$ illuminated volume of interest defined in this work, if the voxel scale is set equal to the image pixel scale of 0.041mm/pixel, then the 3D grid size will be approximately 1220 × 730 × 250 voxels. At 4 bytes per voxel for a declared variable of type 'long', on an Intel Pentium IV processor, this translates to a memory requirement of about 1.8 Gigabytes for a pair of 3D grids. This required amount of memory is beyond the capacity of today's PC's.

While defining a coarser grid by adjusting the grid voxel scale would reduce the memory requirement, correlations on coarse grids may give poor or many spurious block displacement estimates (see Figure 3.9). Furthermore, the correlation parameter settings may indicate that the DCCC approach will give faster correlations. However, efficient implementation of DCCC requires that the zero-intensity voxels be omitted before correlating the particle voxels and not during the evaluation of the correlation function.

In light of these considerations, we only allocated enough memory to store the particle voxel coordinates and their intensities.

For the vortex ring flow images acquired in this project, roughly 1100 particles were identified in the interrogated ROI per image. At the correlation settings given in section 4.7, approximately 1 MByte of memory was needed to store the intensity and grid coordinates of the particle voxels in a pair of correlation 3D grids.

Later, during the correlation process, we extracted from the stored list of particle voxels, those contained within the interrogation blocks and correlation search area.

4.4.2.2 3D correlation of particle cluster patterns

The correlations can be performed by computing the correlation function directly using the DCCC method or via FFT, however, the DCCC method was used in the processing of the images acquired in this work because it was faster. See section 4.7.

We describe both implementations of DCCC and correlations via FFT in this section. It should be noted that we did not determine the cluster displacement to sub-voxel accuracy because only a good estimate representing all the particles in the cluster is needed to help guide the search for the individual particle pairs.

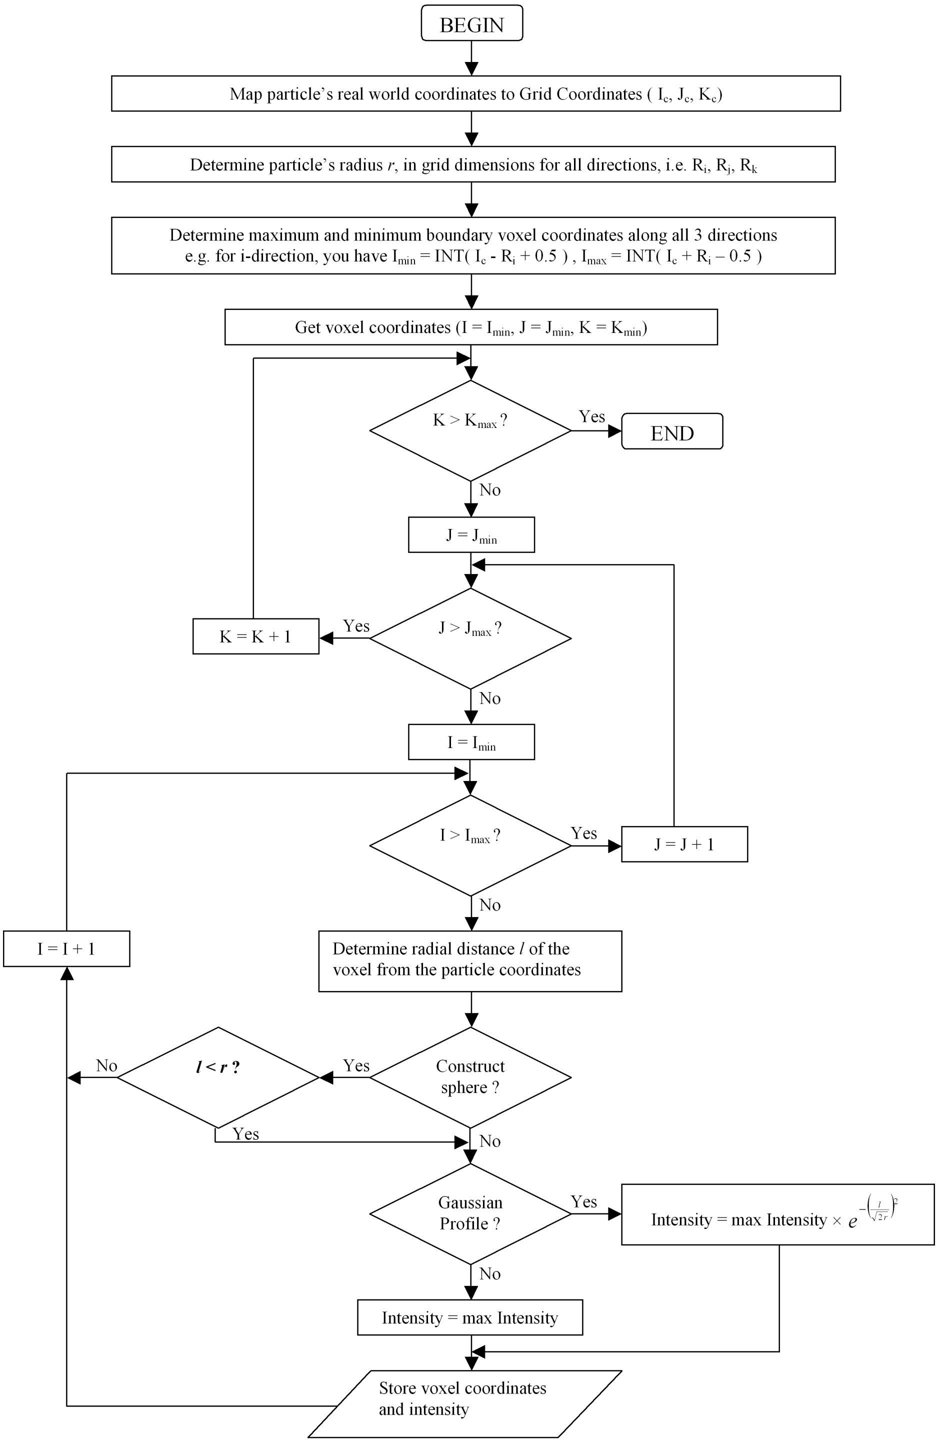

Figure 4.3: Algorithm for the construction of a particle on a 3D grid.

FFT-based correlation implementation

FFT based correlation is implemented as follows. The pair of grids representing the pair of images is divided into interrogation blocks and the blocks from the first grid are correlated with corresponding blocks in the second grid to give an average particle displacement vector.

The particles within each interrogation block are extracted from the list of particle voxels and then placed in explicitly declared 3D arrays. The pair of 3D array data is then correlated via FFT (see notes for Figure 3.3 in section 3.3). Each array is transformed and the results from both transformations are multiplied element-by-element. The inverse FFT of this product is then taken. The output of the inverse transform is then rearranged by sub-dividing the output into eight equal blocks and then swapping diagonally opposite blocks. The rearranged output is subsequently searched for the peak value to give the average particle-cluster displacement vector.

DCCC implementation

To implement DCCC, the first grid is divided into interrogation blocks but the second grid is not divided. Instead, for each block in the first grid, a search for the best cluster pattern match is carried out on the second grid. The regions involved in the correlation process are depicted in Figure 4.4. The particles (i.e. their voxels) in each interrogation block are extracted from the list of particle voxels and correlated over an explicitly constructed 3D search region in the second grid. The particle voxels within the boundaries of the correlating region in the second grid are extracted from the list of particle voxels and placed in an explicitly declared 3D array whose size is at least as large as the region involved in the correlations.

The expression for the correlation function C then takes the form:

$$C\left(\Delta_i, \Delta_j, \Delta_k\right) = \sum_{n=1}^{n_{max}} I_1[n] \cdot I_2\left(i + \Delta_i, j + \Delta_j, k + \Delta_k\right) \tag{4.2}$$

where $I_1[n]$ is the intensity of the nth extracted particle voxel located within the interrogation block and whose grid coordinates are i, j, k.

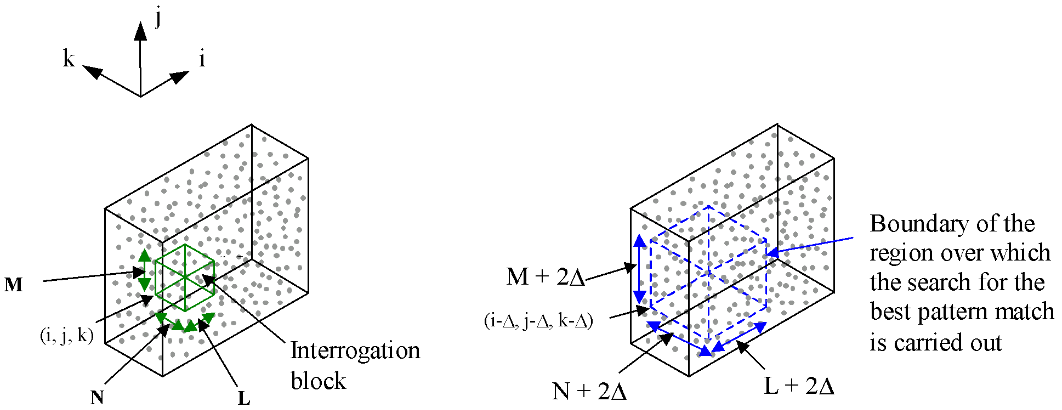

Grid 1

Grid 2

Figure 4.4: Depiction of the regions involved in the correlation of an interrogation block via DCCC.

4.4.2.3 Extraction of particle voxels

Before the interrogation blocks can be correlated, the particle voxels that reside in the blocks and search boundaries must first be extracted from the ordered list of particle voxels.

The speed with which the voxels within a specified region in the grid can be extracted from the ordered list of particle voxels is critical to the overall performance of the correlation algorithm. The time spent performing this task must be a very small fraction of the time spent evaluating the correlation function.

An extraction algorithm based on the ‘bisection method’ was designed to achieve the task of extracting the particle voxels from the list.

The particle voxels within the boundaries of any given region are extracted by scanning through the list of voxels and using the bisection method to jump to the next candidate voxel that falls inside the region when the current candidate voxel falls outside the boundaries of the region. The search for valid candidates stops when the list of candidate voxels is exhausted or when the coordinates of the current candidate voxel is greater[1] than the maximum boundary coordinates of the region.

---

[1] We define the greater coordinate as that whose lookup value is greater.

The voxel extraction process is presented in the flow diagram shown in Figure 4.5.

The bisection method was used to facilitate the identification of the smallest lookup value in the array list of particle voxels that is larger than the lookup value of a target boundary coordinate.

The implementation of the bisection method involves progressively halving the size of the array until the array entry whose lookup value is nearest to, but higher than the target boundary lookup value is left.

At each bisection stage, we compare the last element in the left half of the array and the first element in the right half of the array with the target boundary value. The left half of the array is discarded if its last entry is less than the target boundary lookup value ortherwise the right half is discarded. This process is repeated until the array size is reduced to one.

For an ordered list containing n particle voxel coordinate information, the correct entry in the list can be identified within $\log_2 n$ bisection stages. So, entry positions in a list of 1 million particle voxels can be located in less than 21 array divisions.

4.5 Vector validation and fixing

The successful pairing of the particle images using the hybrid PIV-PTV technique depends on how well the estimated cluster displacements represent the particle displacement.

Since the cluster displacements are used as a guide to identifying the right pair during the tracking stage, a spurious displacement vector from the correlations may lead to erroneous or no pairing of particle images. It is therefore important that appropriate steps are taken to ensure optimum chances of valid correlations.

The major sources of spurious vectors include:

- Insufficient number of particles in the interrogation block
- False particle identification.
- Large relative particle displacements within the interrogation block.

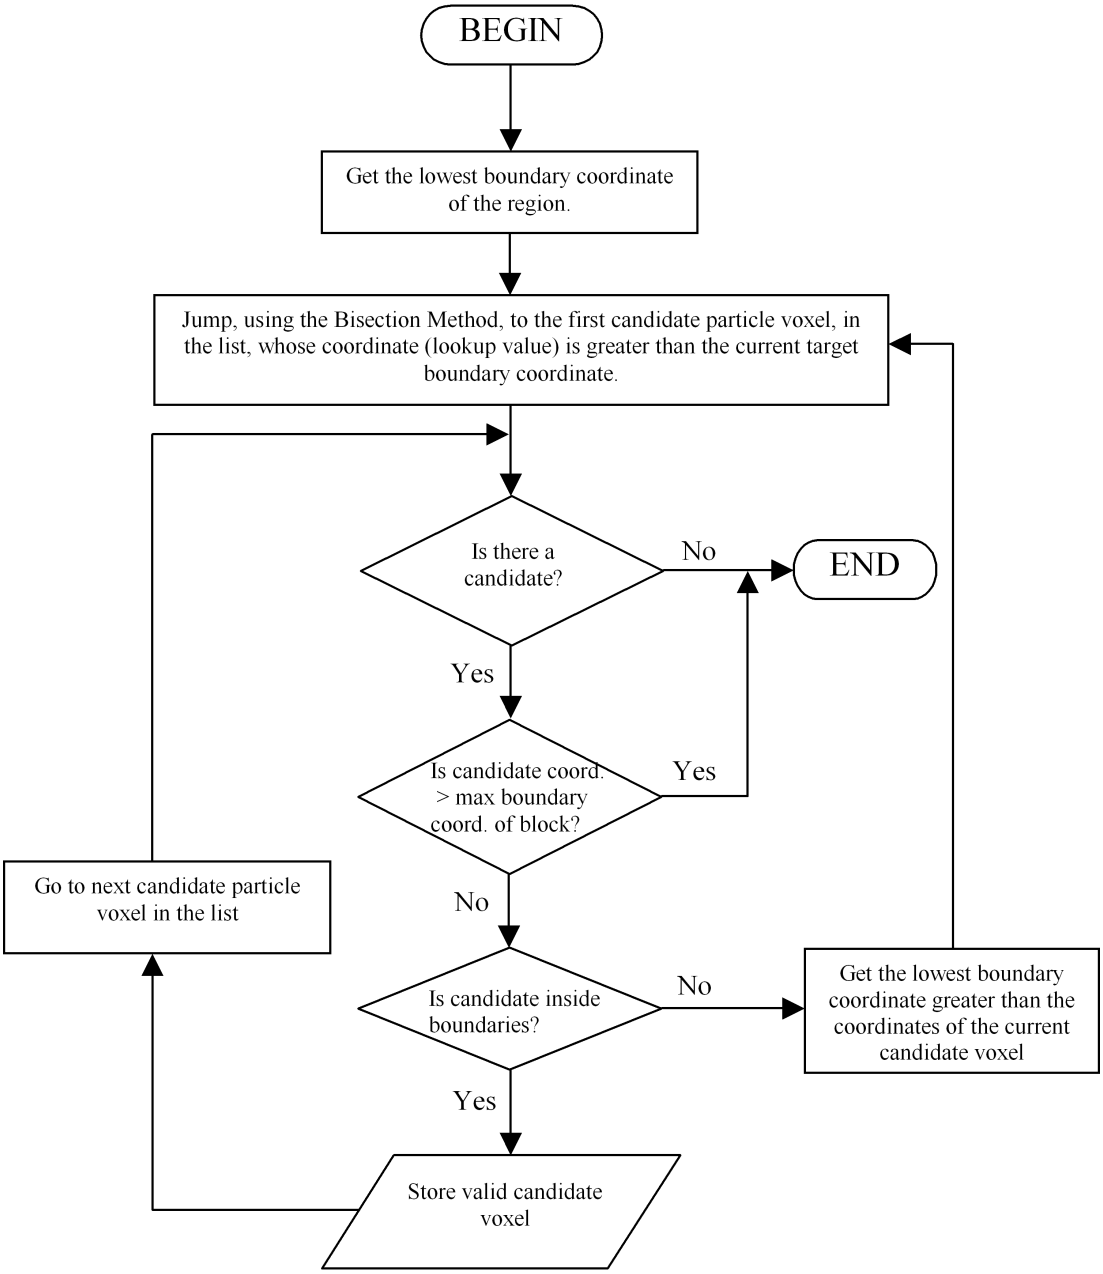

Figure 4.5: A flow chart of the algorithm that was implemented to extract particle voxels within a specified region in a 3D grid.

- Poor estimation of the particle positions on the grid. The main sources of this error are inaccurate colour determination and the resolution of the 3D grid relative to the size of the reconstructed particles.

The vectors from the correlations were validated and fixed using the Peak Stability validation and Fixing routine described in section 3.5.2. Vectors from correlation blocks with less than 3 particles were marked as spurious. For particles residing in interrogation blocks whose vectors remained spurious at the end of the fixing process, these particles were not paired.

4.6 Post-processing

For particles that have been paired, their velocities were determined by dividing the difference between their initial and final positions by the time interval between the illumination of the images.

The resulting velocity vectors in the velocity field were subsequently placed midway between the positions of the image pairs.

The randomly spaced velocities can also be remapped onto a regularly spaced grid and velocity derivatives like the vorticity can also be computed.

In our work, the remapping of these randomly spaced vectors onto a regular grid is achieved using the Adaptive Gausssian Window Averaging (AGW) (See Agui and Jimenez (1987)) where the velocity u at each regular grid point is determined according to the expression:

$$u\left(x\right)=\frac{\sum\limits_{i=1}^{N_v}\alpha_i u_i}{\sum\limits_{i=1}^{N_v}\alpha_i}$$

where: N_v is the number of randomly located velocity vectors within a specified interpolation volume.

u_i and x_i are the velocities and coordinates respectively of the randomly located vectors.

The weighting coefficient α_i is expressed as:

$$\alpha_i = \exp\left(\frac{-|x - x_i|^2}{H^2}\right)$$

where H which determines the width of the Gaussian function is expressed as $H = 1.24\Delta_v$, and Δ_v is the mean particle separation. The weighting coefficients α_i are also adjusted so that their sum is always equal to 1.

The vorticity, ω defined as

$$\omega = \omega_x\vec{i} + \omega_y\vec{j} + \omega_z\vec{k} = \nabla \times \vec{V} = \left(\frac{\partial V_z}{\partial y} - \frac{\partial V_y}{\partial z}\right)\vec{i} + \left(\frac{\partial V_x}{\partial z} - \frac{\partial V_z}{\partial x}\right)\vec{j} + \left(\frac{\partial V_y}{\partial x} - \frac{\partial V_x}{\partial y}\right)\vec{k}$$

is computed at each regular grid point (x, y, z) using the following 2nd order central Finite Difference approximation:

$$\begin{aligned}
\omega_x &= \frac{w\left(x, y + h_y, z\right) - w\left(x, y - h_y, z\right)}{2h_y} - \frac{v\left(x, y, z + h_z\right) - v\left(x, y, z - h_z\right)}{2h_z} \\
\omega_y &= \frac{u\left(x, y, z + h_z\right) - u\left(x, y, z - h_z\right)}{2h_z} - \frac{w\left(x + h_x, y, z\right) - w\left(x - h_x, y, z\right)}{2h_x} \\
\omega_z &= \frac{v\left(x + h_x, y, z\right) - v\left(x - h_x, y, z\right)}{2h_x} - \frac{u\left(x, y + h_y, z\right) - u\left(x, y - h_y, z\right)}{2h_y}
\end{aligned}$$

where h_x, h_y and h_z are the regular grid spacing along the x, y and z-directions respectively and u, v, w are the velocity components along the x, y and z-directions respectively.

4.7 Image interrogation settings

In this section we give the correlation parameter settings used during the interrogation of the uniform flow and laminar vortex ring flow images.

Interrogated ROIx : $5 - 55$mm
y : $5 - 35$mm
z : $9 - 19$mm
Correlation block size : $5 \times 5 \times 5\text{mm}^3$
Correlation Search distance Δ : 2mm
Grid voxel size : $0.2 \times 0.2 \times 0.2\text{mm}^3$
Reconstructed particle diameter : 0.6mm
Reconstructed particle shape : cube
Particle intensity profile : flat
Correlation method : DCCC-based hybrid PIV-PTV
Validation type : Peak stability validation and fixing routine.

Taking the average number of particles identified within the interrogated ROI to be 1100, then at these experimental settings, the predicted computational cost for FFT[1] and DCCC (see section 3.4) are as follows:

$$\begin{aligned}
\text{grid compression ratio } \gamma : &\quad \frac{1100 \times 3^3}{250 \times 150 \times 50} = 0.01584 \\
\text{DCCC } \left(16\gamma N^3 \Delta^3\right) : &\quad 16 \times 0.01584 \times 25^3 \times 10^3 = 3.96 \times 10^6 \\
\text{FFT } \left(45 N^3 \log_2 N\right) : &\quad 45 \times 32^3 \times log_2 32 = 7.4 \times 10^6
\end{aligned}$$

We present, in the next chapter, the results from the analysis of the uniform flow and vortex ring flow images.

---

[1] The radix-2 FFT was implemented, hence, for an FFT-based correlation, the correlation block sizes will be a power of 2.

Chapter 5

Experimental results

5.1 Introduction

In this chapter, we present the results of the simulated uniform flow experiments and the laminar vortex ring flow experiments. The images obtained from both experiments were analysed using the image analysis software written by the author and which incorporates the analysis procedure described in Chapter 4. The image analysis software was written in the 'C language'.

5.2 Uniform flow measurements

The uniform flow experiment was conducted to give an estimate of the accuracy that can be expected during the vortex ring flow experiment.

In this section we examine the effect of subtracting the mean camera noise from the particle image intensities on the accuracy of the postion measurements. We investigate the effect of restricting the ROI to only the region of the illuminated flow field where the variation of colour along the z-direction is greatest. We also confirm that higher intensity flashes give more accurate measurements.

The displacement of the particles in a $50 \times 30 \times 10\text{mm}^3$ volume of the uniform flow block were determined by interrogating pairs of images using the hybrid PIV-PTV algorithm. A sample particle displacement vector field is shown in

Figure 5.1. Shown are the vectors representing the displacements of the tracked particles within the interrogated ROI in the uniform flow block.

We determine the mean accuracy in locating the particles' positions in space by taking the difference between the 'real' uniform flow block displacement of (1, 0, 1)mm and the estimated particle displacements obtained from the image interrogations.

In Figure 5.2, we confirm that higher particle image intensity to noise ratio gives higher accuracy. We show the z-position accuracy for 2 sets of uniform flow images where the images in each pair in the first set are flash-1 images and the images in the second set are flash-2 images. We note that because flash-1 has a higher intensity than flash-2, the z-position error in the measurements is less for the pairs of uniform flow images obtained using flash-1. The x- and y-coordinates of the particles are also expected to be more accurate for the flash-1 images since the z-position is related to the x- and y-coordinates by equation 3.1. See also the discussion on the variation of particle image accuracy with particle image intensity to noise ratio in section A.2.1.

Figure 5.3 shows the mean z-position errors of the particles obtained when the first image in each pair is from flash-1 and the second image is from flash-2. We confirm, from the two sets of measurements shown in Figure 5.3 that z-position accuracy of the measurements is improved when the mean camera noise is subtracted from the intensities of the pixels involved in the determination of the particle colours.

In Figure 5.4, for the case where the mean camera noise is subtracted, we compare the z-position error when the interrogated ROI thickness is reduced from 10mm to the central 5mm thick section. We note that the z-position error is less for the 5mm thick ROI measurements, thus confirming that the accuracy of the measurements could be improved by restricting the correlation ROI to a smaller central region in the flow where the variation of colour with depth is greatest.

In Figure 5.5, we give the mean particle position error for the analysis where only the central 5mm thick region represented is interrogated and the mean camera noise is subtracted. The position error is obtained for six pairs of uniform flow images, where for each pair, the first image was recorded with flash-

1 illumination and the second image, from flash-2. The measurements show the z-position error is about 350 microns, which is an order of magnitude greater than the x- and y- position error, which we estimate to be about 30 microns.

We note that a significant part of the error along the x- and y- direction is due to the error associated with the resolved z-coordinates of the particles. Other contributions to the x-y coordinate error come from particle image overlaps, camera noise and the non-spherical particle shape.

By interrogating only the central 5mm thick region represented in the vortex ring flow images, and subtracting the mean camera noise from the images, we can therefore expect a mean position error of about 350 microns along the z-direction and a mean error of about 30 microns along the x- and y-directions. For our experiment, a 350 micron position error represents a 7% error relative to the depth of the measurement volume of 5mm.

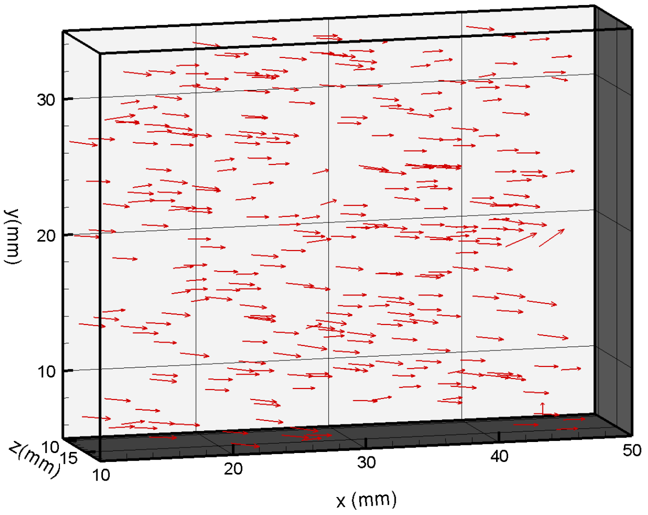

Figure 5.1: Sample uniform flow PTV vector field.

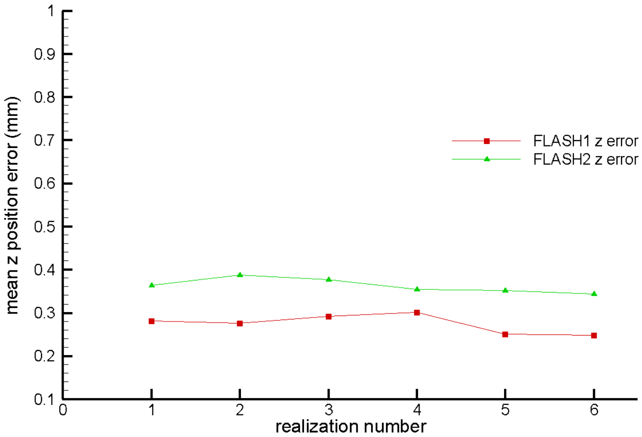

Figure 5.2: The z-position error from flash-1 uniform flow images and flash-2 uniform flow images.

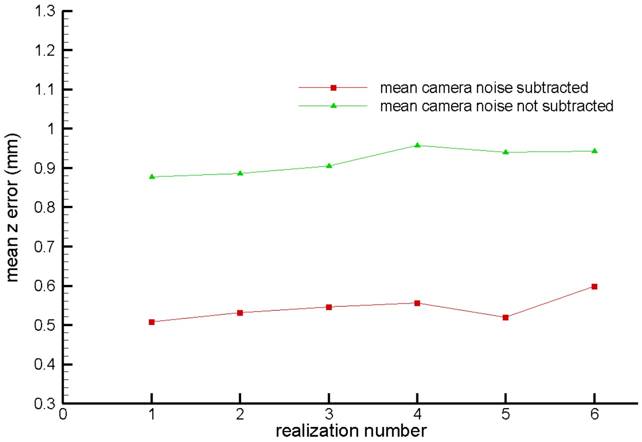

Figure 5.3: Effect of subtracting the mean camera noise on the z-position accuracy.

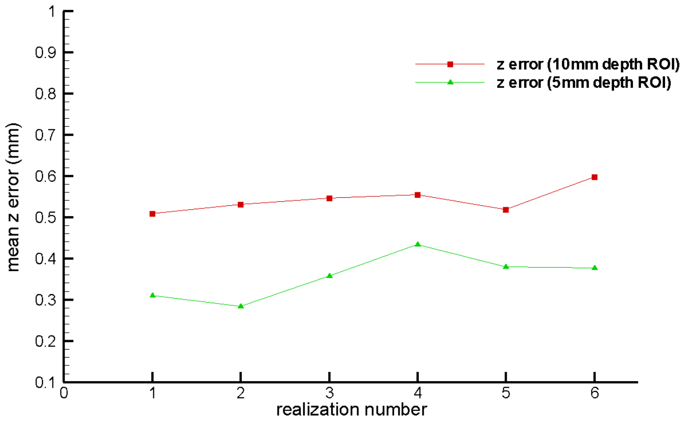

Figure 5.4: Comparison between the z-position error obtained from uniform flow measurements where the depth of the interrogation ROI are 10mm and 5mm.

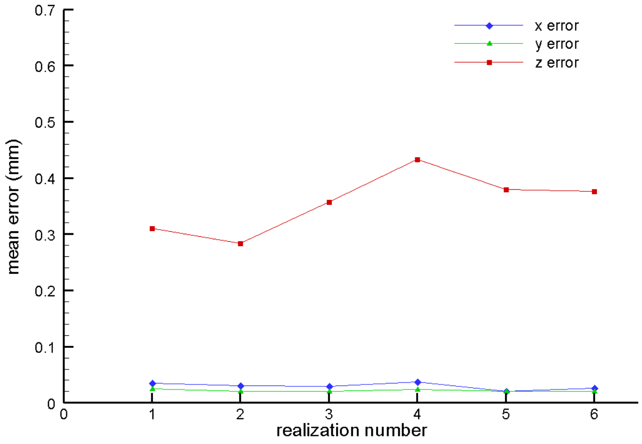

Figure 5.5: Mean position error of the particles in the central 5mm thick region of the uniform flow when the mean camera noise is subtracted.

5.3 Laminar vortex ring flow measurements

A sample instantaneous PIV velocity and vorticity field from the 2D measurements of the laminar vortex ring flow is shown in Figure 5.6 which is the same velocity field shown in Figure 1.11, but now with the vorticity field superimposed. It is this velocity and vorticity field that were used in the comparisons betweeen the 2D and 3D measurements which we present in this section.

We note that although the 2D measurements were obtained using the same equipment, the way in which the 2D images were acquired and subsequently analyzed are different from the 3D case:

First, the illuminating volume is much smaller and can be thought of as a light sheet. Aslo, the resulting image is a monochrome image. In other words, only one of the three CCD channels was used in the 2D image acquisition. Secondly the seeding density in the 2D mode is higher and the seeding particles are smaller. We note that a different 2D PIV system will give similar 2D images under similar imaging conditions.

For the image analysis stage, the 2D velocity results used in the comparisons were obtained using a completely different image analysis software that was also written by the author. We note that the final 2D velocity field results are not from any interpolation but from window correlations of the originally acquired monochrome digital flow images. Further details of the 2D image analysis algorithm and the experimental parameter settings can be found in Appendix F.

We show in Figure 5.7, a sample instantaneous velocity field obtained from the interrogation of a pair of laminar vortex ring flow colour images representing the central region of the ring. Shown are the vectors representing approximately 600 tracked particles out of a total of about 1150 particles identified inside the $50 \times 30 \times 10\text{mm}^3$ correlation ROI. Although the position of the velocity vectors are random, we can still see the structure of the vortex ring. We show the velocity vectors from another angle in Figure 5.8(b) for better observation of the vortex ring structure.

For the 3D measurements of the laminar vortex ring flow, it was important to achieve good z-position accuracy so that the regions where the velocities are low can be adequately resolved. Therefore, subsequent results presented here were obtained by restricting the depth of the interrogated ROI, for each image, to the central 5mm region, where the change in colour with z-distance is greatest.

As noted earlier, due to hardware restrictions, only the front half of the vortex ring was measured—in four 12mm thick overlapping sections. For each section, 20 realizations of the vortex ring flow were interrogated. The resulting PTV vector fields from all 20 realizations for all 4 sections were subsequently combined and then remapped to give an ‘averaged’ velocity field of the front half of the vortex ring flow.

The remapped velocity field and the vorticity isosurface for the front half of the vortex ring are shown in Figure 5.8. The PTV vectors are remapped on a grid with a 2mm × 2mm × 2mm grid point spacing. The vorticity is subsequently calculated on the regular grid, using the remapped velocity vector values.

We show in Figure 5.9, the smoothed vorticity isosurface and the velocity field over the x- and z-planes of symmetry. the expected vortical structure of the vortex ring can be observed from the velocity field over these orthogonal, symmetry planes. We show the smoothed vorticity isosurface from different viewing directions in Figure 5.10 to confirm that it takes the form of the expected ’half-doughnut’ shape. The vorticity field was smoothed by taking the average of the nearest 8 neighbours plus the central point.

To validate the 3D system, we compare the 3D velocity and vorticity measurements to the 2D measurements. The velocity and vorticity data used in the comparisons were extracted at points along symmetry lines that pass through the centre of the vortex ring whose coordinates are approximated as (30, 21.25, 13.25)mm.

The comparsions between 2D and 3D velocity and vorticity measurements are presented in Figure 5.11 to Figure 5.14.

From the figures, we see that there is good agreement between both 2D and 3D results. There is also good agreement between the v-velocity profiles taken along the two orthogonal planes of symmetry for the 3D measurement. The major error, as expected, was in the w-velocity (see Figure 5.12(b)), which led to relatively larger errors in the estimation of the x-vorticity component (see Figure 5.13(b)). The magnitude of the u- and w-velocities along the x- and z-

(horizontal) lines of symmetry are expected to be close to zero. However, we see that for the z-velocity component (see Figure 5.12(b)), there is much greater variation in the w-velocity magnitude. This is due to the relatively much larger error in estimating the z-positions of the particles.

Overall, the 3D velocity measurements agree well with 2D measurements with the biggest discrepancies between the two sets of measurements occurring around the core of the vortex ring where relatively fewer particles were tracked thus resulting in a coarse interpolation. We also note that for the 2D measurements, the gradient error (see section 1.6.1.2) will be greatest in this region hence, contributing to the discrepancies between the 2D and 3D measurements.

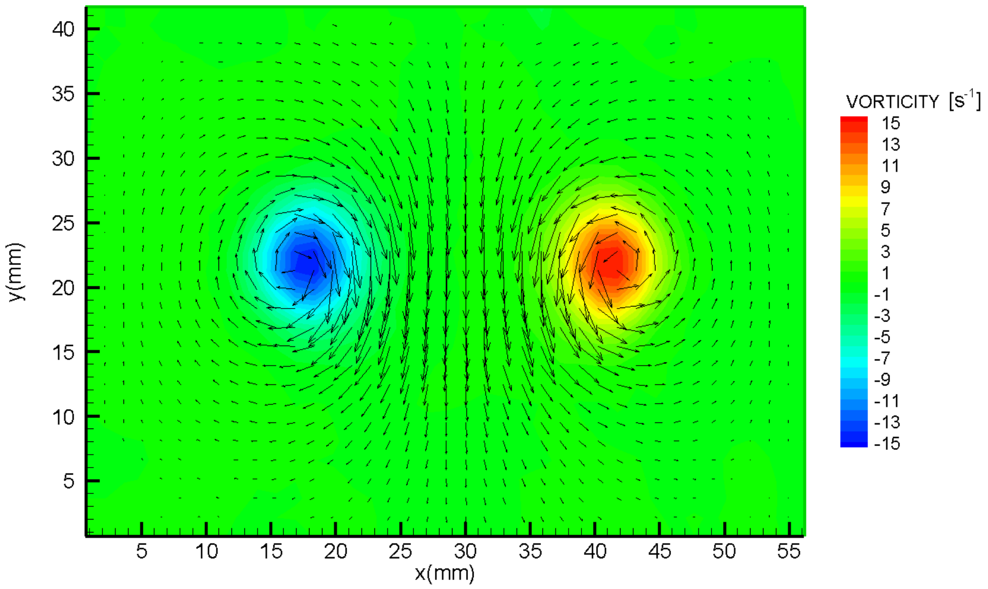

Figure 5.6: Instantaneous velocity and vorticity field from a single 2D PIV measurement of the laminar vortex ring flow.

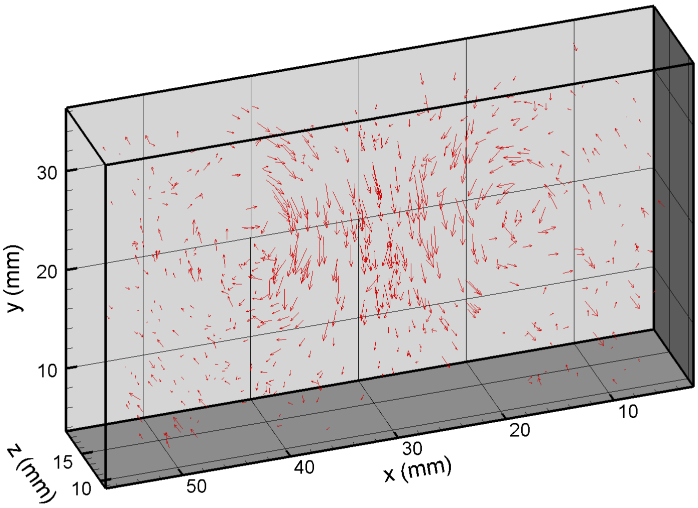

(a) A sample 3D instantaneous velocity vector field representing the central 10mm region of the laminar vortex ring flow.

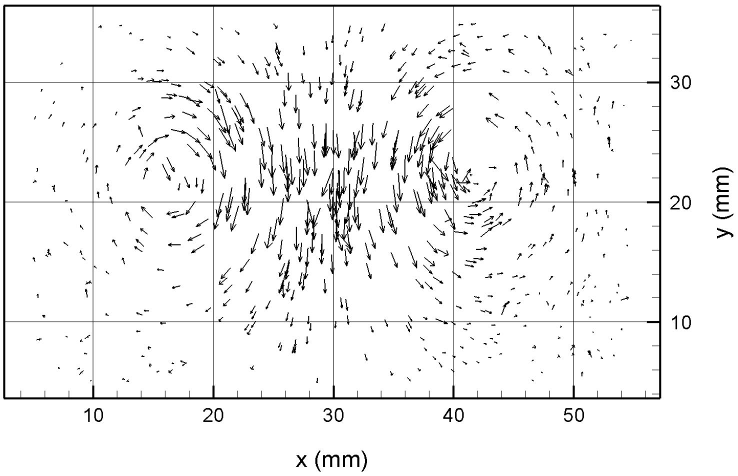

(b) The same 3D instantaneous velocity vector field viewed from a different angle.

Figure 5.7: Two views of the velocity field from one instantaneous 3D measurement of the central 10mm thick region of the laminar vortex ring flow field.

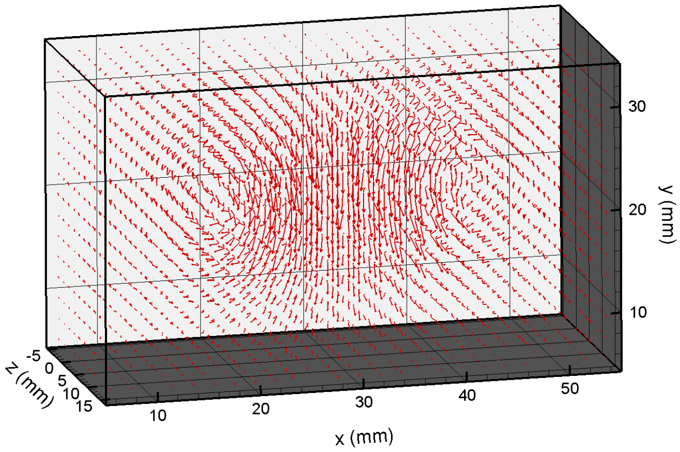

(a) The remapped velocity field for the front half of the vortex ring flow.

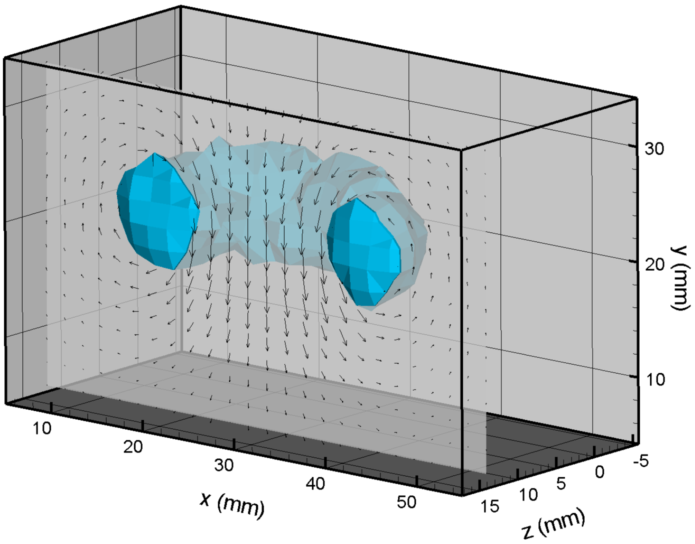

(b) The corresponding vorticity isosurface at 50% of the peak vorticity magnitude (computed at the grid points) shown with a section of the remapped velocity field.

Figure 5.8: Remapped velocity vector field and vorticity isosurface for the laminar vortex ring.

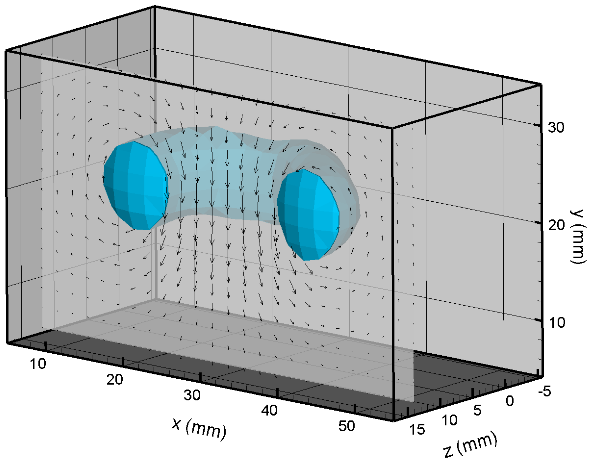

(a) The smoothed vorticity isosurface and the z-plane of symmetry section of the remapped velocity field.

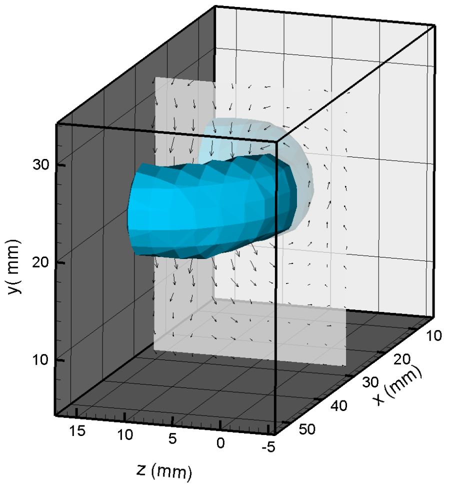

(b) The smoothed vorticity isosurface and the x-plane of symmetry section of the remapped velocity field.

Figure 5.9: The smoothed vorticity isosurface and the velocity vector field over the x- and z-planes of symmetry.

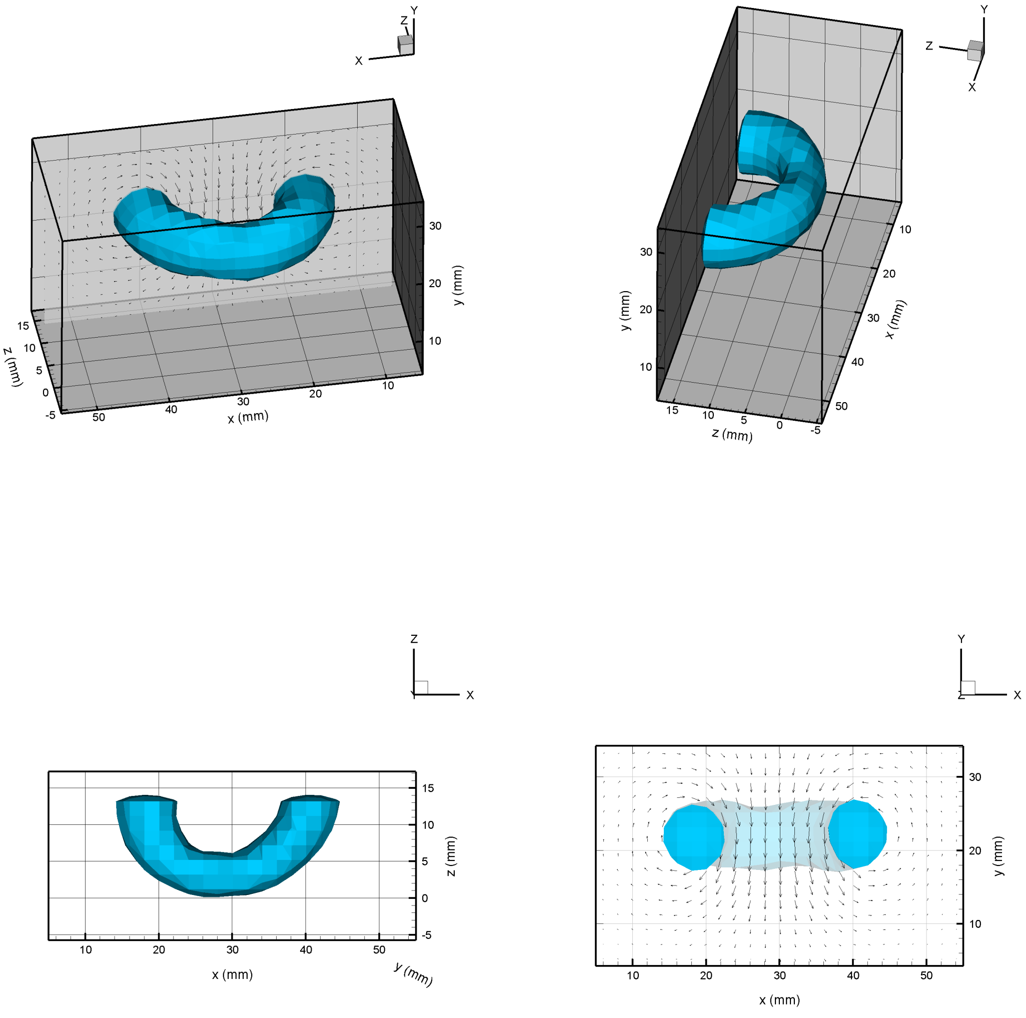

Figure 5.10: The smoothed vorticity isosurface viewed from different angles.

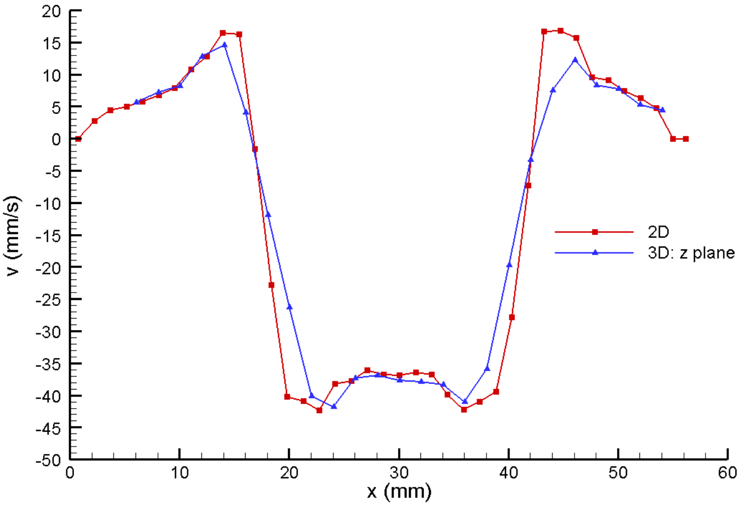

(a) y-velocity component profile along the z-line of symmetry from 2D and 3D measurements.

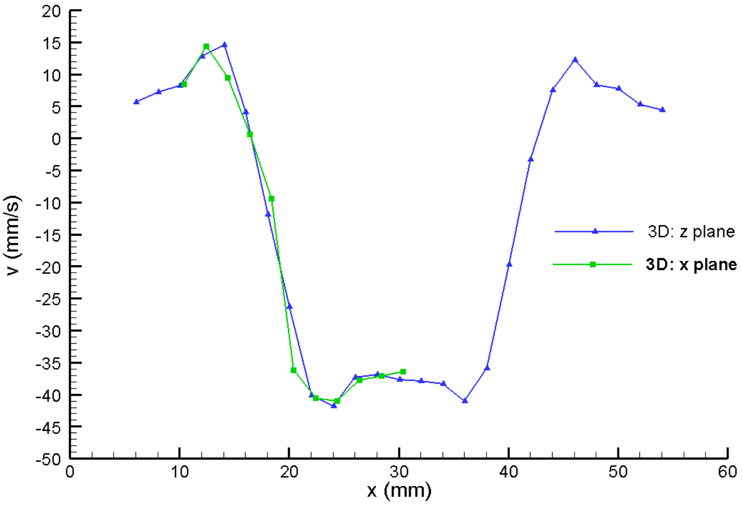

(b) y-velocity component profile along the x- and z-line of symmetry from the 3D measurements.

Figure 5.11: Profile of the y-velocity component along the x- and z-line of symmetry from the 2D and 3D measurements.

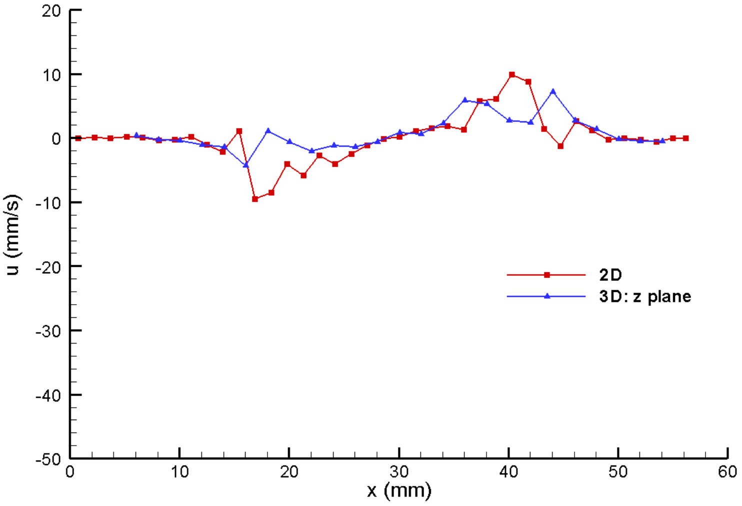

(a) x-velocity component profile along the z-line of symmetry from 2D and 3D measurements.

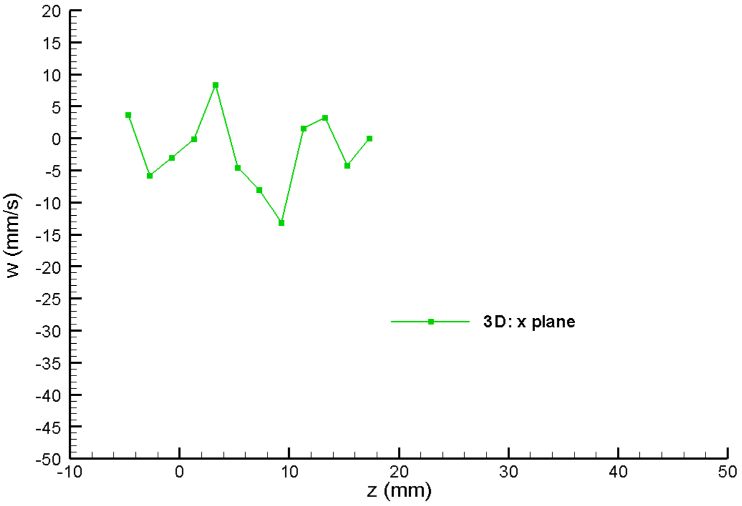

(b) z-velocity component profile along the x-line of symmetry from 3D measurements.

Figure 5.12: Profile of the x- and z-velocity component along the z- and x-line of symmetry from the 2D and 3D measurements.

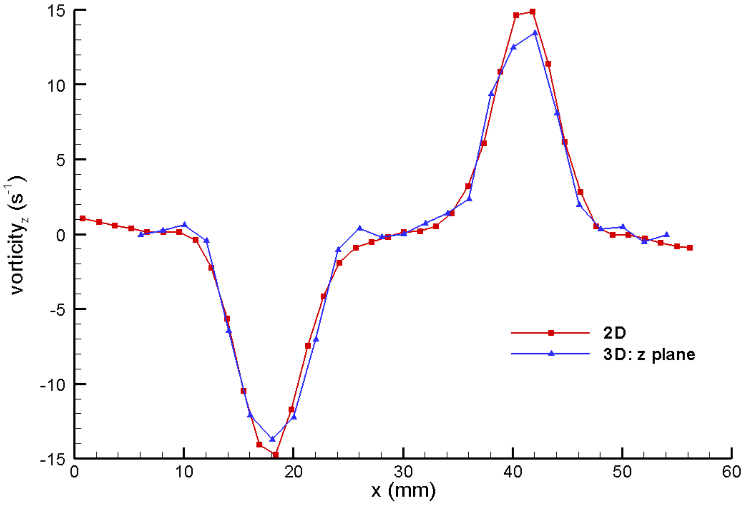

(a) Profile of the z-component of vorticity component along the z-line of symmetry from the 2D and 3D measurements.

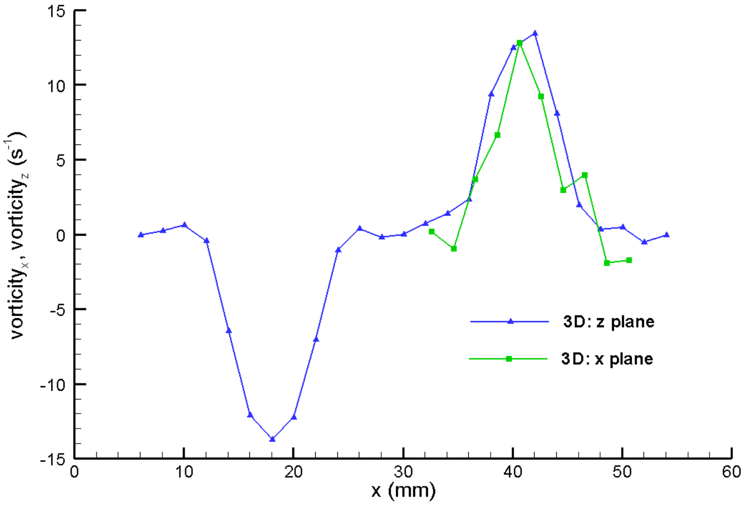

(b) Profiles of the vorticity components obtained along the z- and x-line of symmetry from 3D measurements.

Figure 5.13: Vorticity profiles from the 2D and 3D measurements.

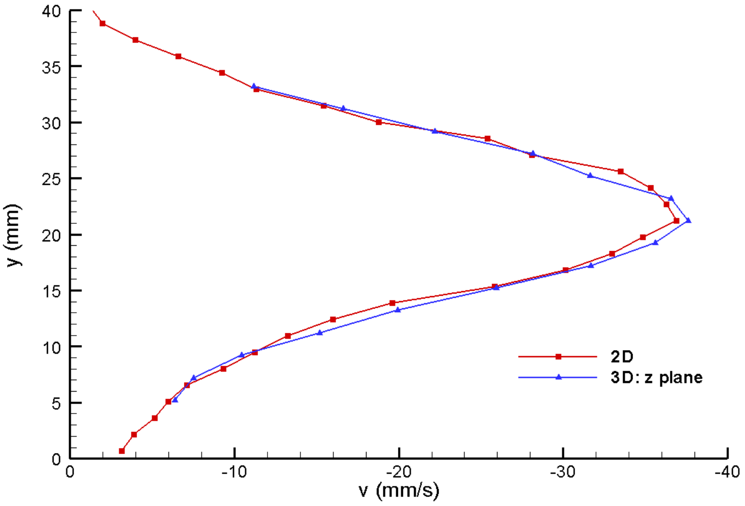

Figure 5.14: Profile of the y-velocity component along the vertical line of symmetry from the 2D and 3D measurements.

Chapter 6

Conclusions

6.1 Summary

In this project, we developed a single-camera, multispectral-illumination PIV system capable of giving accurate, high resolution, three-dimensional velocity measurements at many points in a volume of interest within a practical measurement time. We designed a high performance image acquisition and analysis procedure for the system and we demonstrated the system's ability to give true three-component volume measurements by applying it to investigate a laminar vortex ring flow.

We also identified and explained how the hardware characteristics and experimental parameters affect the speed, accuracy, resolution and depth of field of the measurements. From our investigation we summarize the major findings as follows:

- If the particle-spacing/displacement ratio is greater than 5, a high particle pairing success rate above 95% can be achieved by pairing the particles directly using the nearest neighbour assumption. When this spacing/displacement ratio cannot be achieved, then a high pairing success rate can be realised using the Hybrid PIV-PTV technique.
- To achieve a high probability of obtaining valid particle displacement vectors from the correlations, we recommend that the interrogation blocks should be large enough to contain at least 3 particles and the size of the

reconstructed particles on the grid should be greater than the particle displacement gradient across the block.

- The time required to correlate the particle images can be a significant part of the measurement speed. For a wide range of practical correlation parameter settings, Direct Cross-Correlation on a Compressed grid (DCCC) can be many times faster than FFT-based correlation.
- The particle image position is most accurate when the particle image diameter is between 1 and 1.5 pixels.
- Errors in the particles' z-coordinates will also affect the position accuracy in the x- and y- directions.
- High camera noise in the image reduces the accuracy of the particle image locations in all 3 directions, therefore, a high-intensity illumination together with a low-noise camera noise are recommended for higher accuracy.
- The accuracy to which the particle coordinates are determined is also affected by the flash jitter. A stable source of illumination is therefore necessary for highly accurate measurements.
- As the measurement depth increases, the screening effect and the number of overlapping particle images increase. Smaller particle sizes will therefore be required to reduce the number of overlapping particle images and the screening effect. Smaller particles on the other hand will require that higher intensity illumination be used to give high particle image intensity-to-camera noise ratio. Higher-intensity illumination will also be required to allow the camera lens to be stopped-down enough so that the particle images across the illumination depth are in focus.
- Also for good depth of field measurements, the illuminating source of light must possess a fairly even spectral composition over the visible wavelength range. The camera's spectral response to the illuminating beam should also be such that there is an extensive overlapping of the camera's colour bands so that the z-coordinates of the particles can be adequately resolved from their colours.

6.2 The performance of the 3D colour PIV system

The performance of any PIV arrangement is ultimately determined by the quality of the system's hardware as a system can only perform up to the allowable limits dictated by the characteristics of the hardware components. Therefore, any two arrangements based on the same principle of operation will perform differently across the various performance indicators, if the hardware components have different characteristics. The accuracy and resolution of the system reported in this dissertation is therefore specific to the hardware components used to acquire the images and does not reflect the full potential of the system arrangement.

We do not give direct comparisons between the colour PIV system and those 3D systems described in section 1.8.2, Chapter 1 because a complete performance specification is hardly quoted for many published experiments that have been performed using any of these 3D systems and secondly, our particular choice of experiment hardware (e.g. camera resolution and noise and seeding particle size) which affect the system performance is different from the type of hardware components used in experiments with these other systems. We can however, still make valid comparisons based on the general consensus on the performance of these other systems.

The spatial resolution achieved in the laminar vortex ring experiment is very much comparable to that which has been reported for the 3D multi-camera PTV system. We note that with higher quality imaging equipment (i.e. better camera and flash), we would have been able to use smaller particles at a higher seeding density to give higher spatial resolution beyond the range quoted for the Defocusing PIV system. With the off-the-shelf imaging equipment and the particle size selected for our work, we achieved measurements over a flow field volume of 57mm×43mm×12mm at an average particle separation between 2.5 and 3mm, with relatively small number of overlapping particle images. Approximately 1600 particles were identified in each of the digital images. Of this total, about 1100 particle images were determined to be inside the ROI of 50mm × 30mm × 10mm from which about 600 particles were tracked.

Also with the image analysis software developed for the system, we were able to track the particles at an average vector correlation rate of about 900 vectors/sec on an Intel Pentium IV processor PC, using the Hybrid DCCC based-correlation/tracking approach. Overall, the whole image analysis time for each image took about 30 seconds with most of the time spent at the pre-processing stage where each particle image needed to be identified and its z-coordinates in space determined. While the time spent at the preprocessing stage can still be improved, the image analysis time realized in this work is satisfactory enough to make the interrogation of several hundred images practical. This is a major advantage over Holographic PIV techniques which, only allow single realizations of the flow field, and often involves time-consuming image processing and analysis. The 3D colour PIV system developed in this project should therefore be well suited to investigating a much wider range of flows including turbulence, where several hundred images may need to be analysed.

Lastly, in our particular experiments, the z-position error was quoted to be in the order of 350 microns while the x- and y-position error was determined to be about 30 microns. These values are significantly higher than what has been quoted for systems such as the 3D multi camera PTV arrangement but this is not at all surprising, considering the limitations of the particular equipment that was used in this project.

6.3 Recommendations for a higher performance system

Using commercially available hardware, we have demonstrated the feasibility of achieving fast, accurate and high resolution flow velocity measurements using the colour PIV system arrangement. The performance of the system can also be expected to improve with advances in the technology of the underlying system components, in particular, the camera and flash. We therefore conclude this dissertation by specifying the flash and camera characteristics required for a much improved system performance:

Camera characteristics

- A cooled camera with a low read-out noise and dark current noise to keep the measurement uncertainties low and to reduce the need for a higher flash intensity.
- The camera's 3 colour bands should also overlap extensively over the visible spectral range, to allow the colour of the particles over a greater illumination depth to be resolved.
- Also, a high resolution, high frame-rate camera is desirable to enable smaller structures to be observed and to allow higher temporal resolution measurements to be taken respectively.

Flash characteristics

- A high intensity flash so that the camera gain setting can be set low to give a high intensity-to-noise ratio for the particle images. A higher-intensity flash will also allow a better depth of field to be achieved by stopping down the camera lens. Furthermore, a higher intensity flash will allow smaller particles to be used to improve spatial resolution and reduce overlapping particles as we gain more depth.
- A fairly even spectral distribution in the visible wavelength range to give good camera response to the light over larger depth of field.
- A highly stable illumination source with minimal jitter. The ignition point of the flash must be stable.
- Fast recharge time to allow for faster calibration and good temporal resolution measurements.

Appendix A

Experimental notes

A.1 Achieving a good depth of field

In this section, we highlight some important experimental issues that will need to be addressed when arranging a 3D experiment using the single camera.

For volume measurements, one would typically like to obtain measurements across a good depth of the field — preferably a depth size similar to the planar (i.e. x-y) dimensions. We also want the images of the particles along the depth to be in focus because the images of out-of focus particles will be much larger than those particles in focus thus presenting the possibility of more particle images overlapping unless the seeding density is reduced at a cost to the spatial resolution.

In addition to aiming for a large Depth of Field, we also want as much of the Area of View (i.e. CCD area) to represent only the Region of Interest (ROI) for highest spatial resolution.

For typical lenses, the Depth of Field (DOF) is a function of:

- the Field of View (or the imaging distance) and
- the lens aperture setting (i.e. the f-number setting).

Also, experiences with different lenses reveal that:

- at the same Field-of-View and lens f-number setting, a similar Depth-of-Field is realized for typical lenses.

- Also, at the highest lens aperture setting, the DOF is significantly smaller than the planar (x-y) dimensions of the Area of View.

While moving the camera further away from the target increases the Depth of Field, the Field of View also increases and the Area of the CCD occupied by the ROI becomes smaller as a result. Therefore it appears that the only real option to achieving a large depth of field while maintaining a good spatial resolution is to stop-down the lens i.e. make the lens aperture smaller.

Stopping-down the lens to achieve the desirable Depth-of-Field, however, means less of the light that is reflected by the particles will reach the CCD. Low particle image intensity however translates to a low particle image intensity-to-noise ratio hence an increase in the measurement uncertainties, (see Figure A.4).

Using larger particles will improve the amount of light that is reflected. However, larger particles typically translate to lower spatial resolution because the seeding density will need to be reduced to keep the number of overlapping particles and the screening effect low.

A less problematic option to achieving high particle image intensity-to-noise ratios is to combine higher intensity illumination with a low-noise camera. The cost of high-performance low-noise cameras and powerful illumination sources can however be very high. Investing in highly reflective particles in the first instance could go a long way to reducing the requirements for high illumination power and low-noise cameras.

A.2 Particle position accuracy

When the vectors are validated and the particles have been successfully paired, the particle displacements are taken as the difference between the initial and final positions of the particles in space. Since it is the particle image centroids that are mapped back onto real world coordinates, the precision to which a particle's displacement in real world space is resolved will therefore depend on:

- the accuracy to which the particle image centroids are determined and
- the accuracy to which the particle image centroids are related back to the real world coordinates.

In section 3.3.1.2, we gave a relationship between the position of a particle in space and its image coordinates. In this section, we investigate how the size of the particle image[1] and common sources of image noise affect the accuracy to which the particle image centroid is determined.

A.2.1 Particle image centroid accuracy

Major factors that affect the accuracy to which the particle image centroids are determined include camera noise, non-spherical particles, boundary reflections, reflections from other particles and particle image overlaps.

In this section, we present how the aforementioned sources of image noise affect the accuracy of the particle image displacements. We model these contributing sources to the image noise as normally distributed noise intensities added to the Uniform Flow simulations of 8-bit particle images. The intensity profile of the simulated particle images is Gaussian and details of the particle image simulations are given in Appendix D.

To show how increasing image noise influences the measurement uncertainties, we used two separate sets of normally distributed intensities. The first set of noise intensities added to the image had a distribution with a mean and standard deviation of 5 intensity counts. The other set was an intensity distribution with a mean and standard deviation of 10 counts.

From the results of the simulations that are presented in Figure A.1 to A.4 we make the following deductions:

- The displacement mean bias error increases with noise. See Figure A.1
- The random error also increases with noise. See Figure A.2
- Zero image noise gives zero rms error but there is a bias error associated with small particle image diameters. The bias error is minimized when the particle images is larger than 1 pixel. See Figure A.3.
- The rms error is minimized when the particle image diameter is between 1 and 1.5 pixels. See Figure A.4.

---

[1] See Appendix D for the definition of particle image diameter used in this thesis.

- To summarize, a particle image diameter between 1 to 1.5 pixels gives the highest accuracy. Also the higher the particle image intensity to noise ratio, the more accurately the particle image centroids can be determined.

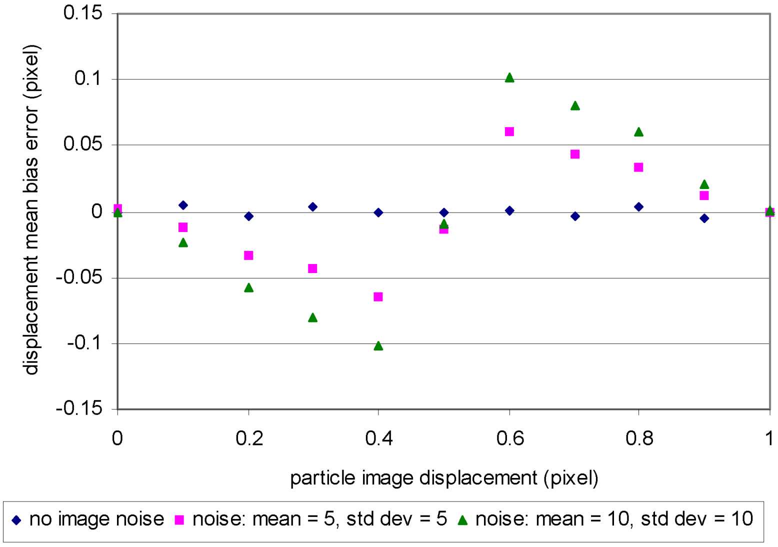

Figure A.1: Effect of noise on particle image displacement mean bias error for a simulated particle image diameter of 2 pixels.

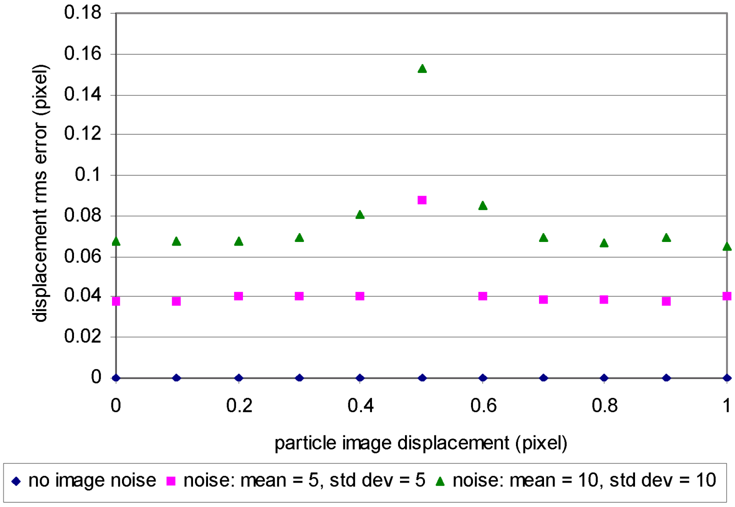

Figure A.2: Effect of noise on particle image displacement rms error for a simulated particle image diameter of 2 pixels.

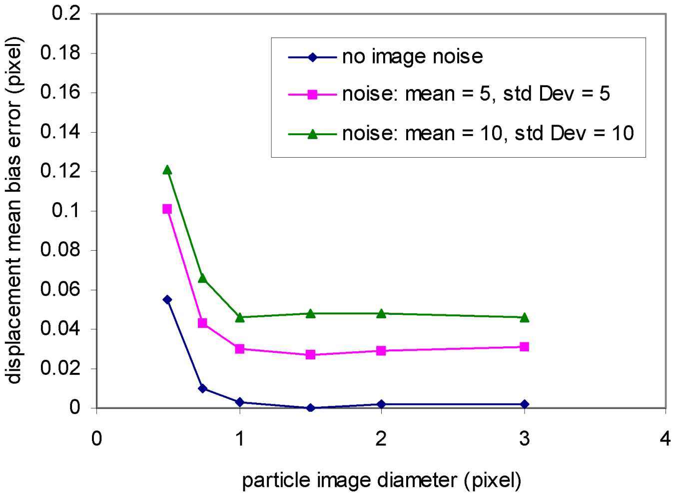

Figure A.3: Variation of particle image displacement bias error with particle image size.

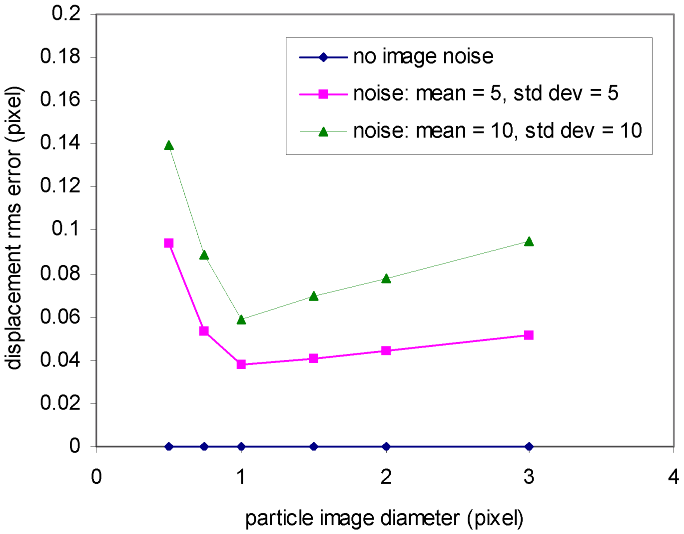

Figure A.4: Variation of particle image displacement rms error with particle image size.

Appendix B

Radix-2 FFT computational cost

The computational costs for the radix-2 FFT algorithm are as follows:

B.1 Computational cost: 1D FFT

The computational cost of computing the Discrete Fast Fourier Transform of a 1D array of data of length N is given as:
Complex number multiplies[1]: $\frac{N}{2}\log_2 N$
Complex number adds[2]: $N\log_2 N$
OR
Real number multiplies: $2N\log_2 N$
Real number adds : $3N\log_2 N$

B.2 Computational cost: 2D FFT

The FFT of a 2D array data is computed by first transforming each row, and subsequently replacing each row with its transform then followed by the transforming of each column, and subsequently replacing each column with its transform. The Fast Fourier transform of a 2D array of size $N \times N$ therefore requires a total of $2N$ one-dimensional transforms i.e.

---

[1] 1 complex number 'addition' operation = 2 real adds

[2] 1 complex number 'multiplication' operation = 4 real multiply + 2 real adds

Real number multiplies: $4N^2 \log_2 N$
Real number adds: $6N^2 \log_2 N$

B.3 Computational cost: 3D FFT

For a 3D array of size N^3, the 3D FFT can similarly be computed by first performing 1D transforms on the arrays along the x-direction and subsequently replacing each x-array by its transform. The procedure is then repeated along the y direction and z-direction in succession. The Fast Fourier transform of a 3D array of size $N \times N \times N$ therefore requires a total of $3N^2$ one-dimensional transforms.
The computational cost in performing 3D FFT is as follows:
Real number multiplies: $6N^3 \log_2 N$
Real number adds: $9N^3 \log_2 N$

Appendix C

Operation count for the discrete 3D cross-correlations

C.1 Definition of the 3D discrete cross-correlation function

For the correlation of a pair of 3D interrogation blocks, the discrete cross-correlation function is given by the expression:

$$C(\Delta_i, \Delta_j, \Delta_k) = \sum_{i=1}^{L} \sum_{j=1}^{M} \sum_{k=1}^{N} (I_1(i, j, k)) \cdot (I_2(i + \Delta_i, j + \Delta_j, k + \Delta_k))$$

where the size of the interrogation blocks are $L \times M \times N$.
I_1 is the voxel intensity in interrogation block 1.
I_2 is the voxel intensity in interrogation block 2.
$I_1(i, j, k)$ is the intensity of the voxel located at coordinates (i, j, k).
Δ_i, Δ_j and Δ_k are the correlation search coordinate distances along the x-, y- and z-directions. Δ is the maximum correlation search distance. For example, a ± 32 voxel correlation search distance means $\Delta = 32$. The highest value of the correlation function is subsequently determined by searching the resulting 3D correlation table.

Notes: The correlation function can be evaluated directly by computing the expression as it is or via Fast Fourier Transform (FFT).

The computational cost of processing the correlation function can be expressed in terms of number of arithmetic operations involved in the computations.

The cost of computing this function can be defined in terms of the number of individual pixel comparisons plus the number of array element interrogations needed to determine the peak correlation from the resulting correlation table.

It is assumed that for direct computation of the discrete cross-correlation function, the second interrogation block is allowed to float thereby resulting in a complete overlap in the two interrogation blocks at every interrogation search distance.

Here, for simplicity, we assume that the interrogation block lengths $L = M = N$.

For the iterative correlation algorithm computation described here, the interrogation block lengths are halved at each subsequent iteration stage to give an eight-fold increase in the number of interrogation blocks over the previous stage. For direct computation of the correlation function, beyond the first stage, the correlation search distance Δ is reduced to a value α, which is of the same order as the particle image diameter and is kept constant for subsequent stages to ensure that its value does not go below 1 after many iterations (stages).

The computational costs involved in computing the correlation function directly and via FFT are approximated as follows:

C.2 3D direct correlation

C.2.1 Single Stage Correlation

To compute $C(\Delta_i, \Delta_j, \Delta_k)$ directly for all search coordinates $(\Delta_i, \Delta_j, \Delta_k)$ within the maximum boundary equal to Δ requires:

$$\begin{aligned}
&N^3\left(2\Delta+1\right)^3 \text{Multiplications}\\
&\left(N^3-1\right)\left(2\Delta+1\right)^3 \text{Additions}\\
&\left(2\Delta+1\right)^3 \text{Peak Search}
\end{aligned}$$

C.2.2 Iterative (multi-stage) correlation

The computational cost involved in computing $C(\Delta_i, \Delta_j, \Delta_k)$ directly for all $(\Delta i, \Delta_j, \Delta_k)$ during an iterative correlation process requires:

$$\begin{aligned}
\text{Multiplications:} \quad & N^3 \left(2\Delta+1\right)^3 + \sum_{i=1}^{i=n} \left[\left(2^{3i}\right) \left(\frac{N}{2^i}\right)^3 \left(2\alpha+1\right)^3 \right] \\
\text{simplifying} \quad & N^3 \left(2\Delta+1\right)^3 + nN^3 \left(2\alpha+1\right)^3 \\
\text{Additions:} \quad & \left(N^3-1\right) \left(2\Delta+1\right)^3 + \sum_{i=1}^{i=n} \left[\left(2^{3i}\right) \left(\left(\frac{N}{2^i}\right)^3 - 1 \right) \left(2\alpha+1\right)^3 \right] \\
\text{simplifying} \quad & \left(N^3-1\right) \left(2\Delta+1\right)^3 + \left(2\alpha+1\right)^3 \left[nN^3 - \frac{8}{7}\left(8^n-1\right) \right] \\
\text{Peak Search:} \quad & \left(2\Delta+1\right)^3 + n\left(2\alpha+1\right)^3
\end{aligned}$$

where n is the number of iterative stages beyond the first stage.

Assuming $N^3 \gg 1$ and $\Delta \gg 1$, and for practical combinations of N and n, $\frac{8(8^n-1)}{7nN^3} \ll 1$, then these expressions can be approximated as:

$$\begin{aligned}
&\approx \quad 8\Delta^3 N^3 \left[1 + n\left(\frac{2\alpha+1}{2\Delta}\right)^3 \right] \text{multiplications} \\
&\approx \quad 8\Delta^3 N^3 \left[1 + n\left(\frac{2\alpha+1}{2\Delta}\right)^3 \right] \text{additions} \\
&\approx \quad 8\Delta^3 \left[1 + n\left(\frac{2\alpha+1}{2\Delta}\right)^3 \right] \text{peak searches}
\end{aligned}$$

This means that the additional time taken to perform n additional stages of iterations during a multi-stage direct 3D cross-correlation is approximately $n\left(\frac{2\alpha+1}{2\Delta}\right)^3$ times that taken to perform the first stage.

C.2.3 Compression based iterative correlation

For compression based processing where the image is compressed so that the ratio of the non-zero intensity voxels are a fraction γ of the number of voxels in the correlation grid, then, the computational effort is approximately:

$$\begin{aligned}
\text{Multiplications:}\quad & \gamma N^3\left(2\Delta+1\right)^3+n\gamma N^3\left(2\alpha+1\right)^3\\
\text{Additions:}\quad & \left(\gamma N^3-1\right)\left(2\Delta+1\right)^3+\left(2\alpha+1\right)^3\left[n\gamma N^3-\frac{8}{7}\left(8^n-1\right)\right]\\
\text{Peak search:}\quad & \left(2\Delta+1\right)^3+n\left(2\alpha+1\right)^3\\
& \text{where } 0\leq n\leq \tfrac{1}{3}\log_2 n_p,\ \ n=0,1,2,3,\ldots\\
& n_p \text{ is the number of particles in the window.}
\end{aligned}$$

For single stage correlation, the computational cost (Addition + multiplication) is approximately:

$$\left(2\Delta+1\right)^3\left(2\gamma N^3-1\right)$$

which can be further approximated as

$$16\gamma N^3\Delta^3$$

assuming $\Delta \gg 1$ and $\gamma N^3 \gg 1$.

C.3 3D correlation using FFT

C.3.1 Single stage correlation

The evaluation of the 3D correlation function for all (Δ_i, Δ_j)requires:
2 FFTs, N^3 complex multiplies and 1 inverse FFT. See Figure 3.3.
The computational effort required to determine the 3D correlation function, based

on the FFT computational costs given in Appendix B.3, then becomes:

Real Multiplications:	$18N^3 \log_2 N + 4N^3$
Real Additions:	$27N^3 \log_2 N + 2N^3$
Total (add+multiply):	$45N^3 \log_2 N + 6N^3$
Peak Search:	N^3

C.3.2 Iterative (multi-stage) correlation

The computation of the cross-correlation function for several iterative stages requires the following number of arithmetic operations:

$$\begin{aligned}
\text{Add + multiply:}\quad & \sum_{i=0}^{i=n} \left(2^{3i}\right) \left[45 \left(\frac{N}{2^i}\right)^3 \log_2 \left(\frac{N}{2^i}\right) + 6 \left(\frac{N}{2^i}\right)^3\right] \\
& \sum_{i=0}^{i=n} \left[45N^3 \log_2 \left(\frac{N}{2^i}\right) + 6N^3\right] \\
& \sum_{i=0}^{i=n} \left[45N^3 \left(\log_2 N - i\right) + 6N^3\right] \\
& 45\left(n+1\right) N^3 \log_2 N - 45N^3 \frac{n\left(n+1\right)}{2} + 6\left(n+1\right) N^3 \\
& \left(45N^3 \log_2 N + 6N^3\right) \left[1 + n - \frac{45n\left(n+1\right)}{2\left(45 \log_2 N + 6\right)}\right] \\
\text{Peak search:}\quad & \left(n+1\right) N^3 \\
& \text{where } 0 \le n \le \log_2 N,\ \ n = 0, 1, 2, 3, \ldots
\end{aligned}$$

Assuming N is large so that $45 \log_2 N \gg 6$, we get:

$$\begin{aligned}
\text{Multiply+add:}\quad & \left(45N^3 \log_2 N\right) \left[1 + n - \frac{n\left(n+1\right)}{2 \log_2 N}\right] \\
\text{Peak search:}\quad & \left(n+1\right) N^3
\end{aligned}$$

C.4 Effect of block size on total correlation cost

For a given grid size and fixed maximum search distance, the computational cost associated with computing the correlation function directly is independent of the interrogation block size. This is because, the total number of blocks to correlate varies linearly with the reduction in size of the block.

However the computational cost associated with correlating a given grid using FFT changes with the interrogation block size. When an initial block length N is halved i number of times, then the number of blocks to correlate increases by a factor of 2^{3i} and the new computational cost becomes:

$$\begin{aligned}
\text{additions + multiplications:} \quad & \left(2^{3i}\right)\left[45\left(\frac{N}{2^i}\right)^3 \log_2\left(\frac{N}{2^i}\right) + 6\left(\frac{N}{2^i}\right)^3\right] \\
= \quad & \left(45N^3 \log_2 N + 6N^3\right)\left[1 - i\frac{45}{(45\log_2 N + 6)}\right] \\
\approx \quad & \left(45N^3 \log_2 N\right)\left[1 - \frac{i}{(\log_2 N)}\right]
\end{aligned}$$

where i is the number of times the initial block length N is halved in all directions. The total computational cost involved in the peak searches, however, is the same as for the original block size. i.e. $2^{3i}\left(\frac{N}{2^i}\right)^3 = N^3$

So when the interrogation block is reduced in size, although there are more blocks to correlate, the total computational cost is less.

We can also deduce that during iterative correlation, when the block length is halved in all directions after each iterative stage, then the computational time at the next stage is reduced by a factor of approximately $\frac{1}{\log_2 N}$ over the previous stage or approximately $\frac{i}{\log_2 N}$ over the first stage.

We note that for the converse case where the initial block of length N is doubled i number of times in all directions, the new computational cost is approximately

$$\left(45N^3 \log_2 N\right)\left[1 + \frac{i}{(\log_2 N)}\right]$$

Appendix D

Artificial particle image generation and simulation details

D.1 Particle image generation

In the simulations reported in this study, particle images with flat and Gaussian intensity profiles were generated.

To generate a particle image with a flat intensity profile, we assigned the same intensity value to all pixels, within a specified radius from the particle image location.

To generate a particle image with a Gaussian intensity profile, we followed the same method used by Huang et al (1997). We describe the particle image intensity profile as:

$$I\left(x,y\right)=I_0\exp\left[-\frac{\left(x-x_0\right)^2+\left(y-y_0\right)^2}{\sigma_p^2}\right]$$

where the centre of the particle image is located at (x_0, y_0) with a peak intensity of I_0 and $\sigma_p = \frac{d_p}{\sqrt{2}}$, where d_p is the particle image diameter which we define as the size of the image area over which 39.5% of the scattered light is registered.

We also give an alternative definition of the particle image diameter as the number of pixels that covers 68.3% of the light intensity recorded over the central horizontal or vertical line of pixels.

The intensity of the particle image pixels is computed by integrating the above expression over the area of the pixel using the expression:

$$I\left(x,y\right)=I_0\int_{y-0.5}^{y+0.5}\int_{x-0.5}^{x+0.5}\exp\left[-\frac{\left(x-x_0\right)^2+\left(y-y_0\right)^2}{\sigma_p^2}\right]dxdy$$

For 3D simulations, the particles were colour-coded according to their z-coordinates in space. The colour of each particle, in RGB ratio format, was obtained by matching the particle's z-position in space to a colour in a prepared database of colour vs. position information. Each particle's colour was subsequently multiplied by the intensities $I\left(x,y\right)$ of the particle image pixels to give a colour image in RGB format.

D.2 Simulation details

D.2.1 Simulation information for table 3.1

- Simulation performed on an Intel 1.8Ghz Pentium IV processor, 1 Gbyte RAM.
- Simulated Grid size: 512 × 512 × 128 voxels
- Particle model: cubic

D.2.2 Simulation details for Figure 3.1

- Simulated image size: 1404 × 1065 pixels
- Voxel scale in all directions: 0.0413mm/pixel
- Simulated volume: $58 \times 44 \times 20\text{mm}^3$.

Notes: Each point represents the results from 10 simulated particle images.

D.2.3 Simulation details for Figure A.1 and Figure A.4

- Simulated image size: 1392 × 1040

- Approximate number of particles in each noisy image: 1600
- Approximate number of particles in each noise-free image: 150
- The particle image centroid was determined using the expressions in Appendix E.2.

Notes: To reduce the number of overlapping particles, we first divided the image into smaller windows of equal area based on the user-specified average particle spacing. Then 1 particle was randomly positioned in each window to give an even particle image distribution. In Figure A.3 and Figure A.4, the errors are the average errors of the 9 simulated sub-pixel particle image displacements.

The centroid of each particle image was determined using the pixels with a pixel-square area equal to the same area as that which bound all the particle image pixels in a noise-free image. For example, if a 5×5 pixel area was used in determining the centroid of the particle images in a noise-free image, then a 5×5 pixel area, centred at the highest intensity pixel was also used to determine the particle image centroids in the noisy images.

D.2.4 Simulation details for Figure 3.10

- Grid size: $512 \times 512 \times 64$ voxels
- Interrogation block size: $32 \times 32 \times 32$ voxels
- Particle model: sphere
- Particle sizes used: 3, 5 and 7 voxels.

Notes: A valid vector is identified as that from correlations where the highest peak in the correlation table is due to the pairing of particle images belonging to the same particle.

The simulated flow is a shear flow whose velocity profile varies linearly along the z-direction but is constant along the x and y-directions. i.e. the shear is over the y-plane.

The simulated flow velocity is zero along the line z = 32 voxels.

D.2.5 Simulation details for Figure 3.9

Simulation of the motion of particles in a velocity field of a series of joined up straight line vortex rods arranged to form a ring.

Simulated volume size: 57mm × 43mm × 12mm
Average particle separation: 1.5mm
Correlated Region of interest (ROI): 54mm × 42mm × 12mm
Interrogation block size: $3 \times 3 \times 3\text{mm}^3$
Search distance Δ in all directions: 2mm
Image size: 1392 × 1040
Particle model: cubic
Reconstructed particle size: 0.4mm
Correlation grid scale: 0.2mm/voxel
Vortex ring centre coordinates (in mm): (29,22,6)
Vortex ring diameter: 20mm
r_c: 3mm
ω_o: 1 s^{-1}
Δt: 1 sec
Number of vortex rods: 100
Number of time divisions: 100
NOTES:

- For particle displacements smaller than a voxel size, the approximate cluster displacement may be zero. So in the PIV vector field, the vectors for the blocks where the approximate displacements are zero, will not show in the plots. However, this does not mean that the blocks are spurious. We see from the PTV results that some of the particles in these interrogation blocks, were tracked as should be the case if the blocks are not spurious.

- The induced velocity $\vec{V}$, at any point in the flow field by a vortex rod is given by the expression:

  $$\vec{V} = \frac{\Gamma}{4\pi} \oint \frac{\vec{dl} \times \vec{r}}{r^3}$$

  where r is the distance of the point from the vortex rod segment of length

$\left|\vec{dl}\right|$ and Γ is the circulation which is expressed as:

$$\Gamma = \pi {r_c}^2 \omega_o \left(1 - e^{-\frac{r^2}{{r_c}^2}}\right)$$

where r_c is the vortex rod core radius and ω_o is the vorticity at $r = 0$.

Appendix E

Sub-pixel estimators

We present sub-pixel estimators for the peaks in the correlation tables and the particle image centroid.

E.1 1D sub-pixel interpolators

Consider the correlation peak/highest-intensity particle image pixel I_i and its nearest two neighbours along the i-direction as shown below,

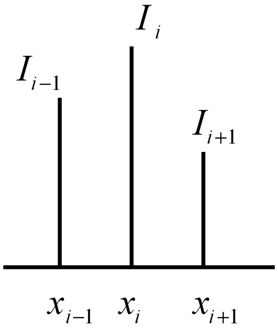

Figure E.1: A Correlation peak and its neighbours.

Where
$x_{i-1} = x_i - 1$
$x_{i+1} = x_i + 1$
The peak coordinate x_{peak} along the i-direction is estimated to sub-pixel accuracy

using one of the following sub-pixel estimators ε:
Defining $x_{peak} = x_i + \varepsilon$ then

E.1.1 Gaussian estimator

$$\varepsilon_{gaussian} = \frac{1}{2}\left(\frac{\ln I_{i-1} - \ln I_{i+1}}{\ln I_{i-1} - 2\ln I_i + \ln I_{i+1}}\right)$$

E.1.2 Parabolic estimator

$$\varepsilon_{parabolic} = \frac{1}{2}\left(\frac{I_{i-1} - I_{i+1}}{I_{i-1} - 2I_i + I_{i+1}}\right)$$

E.1.3 1-D centroid estimator

$$\varepsilon_{centroid} = \frac{I_{i+1} - I_{i-1}}{I_{i-1} + I_i + I_{i+1}}$$

The sub-pixel estimation along the j-direction is given by similar expressions, with the i-index replaced by j.

E.2 2D sub-pixel approximation of a particle image centroid or correlation Peak

The centroid ($centroid_x, centroid_y$) of a particle image or a correlation peak can also be determined using the expression:

$$centroid_x = \frac{\sum\limits_{j=n-N}^{j=n+N}\sum\limits_{i=m-M}^{i=m+M} x_{i,j} I_{i,j}}{\sum\limits_{j=n-N}^{j=n+N}\sum\limits_{i=m-M}^{i=m+M} I_{i,j}}$$

$$centroid_y = \frac{\sum\limits_{j=n-N}^{j=n+N}\sum\limits_{i=m-M}^{i=m+M} y_{i,j} I_{i,j}}{\sum\limits_{j=n-N}^{j=n+N}\sum\limits_{i=m-M}^{i=m+M} I_{i,j}}$$

where N, M = 1,2,3,.....
$I_{i,j}$ is the intensity value or correlation value at coordinates $x_{i,j}$ and $y_{i,j}$ in the

image or correlation table.
m and n are the coordinates of the peak intensity of the particle image or the peak in the correlation table.

Appendix F

A high performance 2D PIV image interrogation algorithm

A high performance image algorithm should possess the following features:

- high resolution
- robust to large gradients
- high accuracy and valid-data-yield
- high speed correlation

Here, we describe how a high performance algorithm incorporating the ideas presented here was used to obtain the velocity vector plots shown in Figure 1.11 and Figure 1.12. We show the process in Figure F.1.

Achieving high resolution

Use iterative correlation or Hybrid PIV-PTV technique. The Hybrid technique should give higher resolution and better accuracy. For iterative correlation, if the initial interrogation windows are defined to contain approximately 10 particles to meet the recommendations of Keane and Adrian (1992), then the windows should only be sub-divided once to give the highest resolution. In order words, iterative correlation should not be more than a 2-stage process. Further window subdivision will result in many windows containing no particle images. Furthermore, the accuracy of the measurements will most likely be reduced as the particle image sizes become comparable to the window sizes hence leading to these smaller windows not containing whole particle images.

We note that the correlations of the windows containing no particles will obviously be meaningless. These empty windows, if they can be identified as such, will still need to be interpolated if they are not to be left blank. So achieving super-resolution using the Hybrid technique and then subsequently interpolating onto a regular grid if required may be a better alternative.

Achieving fast image interrogation speed

Since identifying particle images is a first step in particle tracking, we can exploit the fact that the number of particle images and their approximate position are known to achieve fast processing.

We can reconstruct the particle images on a compressed grid (i.e. one whose background pixel intensities are zero) and correlate using DCCC instead of FFT when $8\gamma N^2\Delta^2 < 30N^2 log_2 N$.[1]

We could further increase the correlation speed by reconstructing the correlation particle images on a grid which is coarser than the original particle image pixel resolution. If the grid is made twice as coarse in both directions as the original image dimensions, then for DCCC processing, the computational time is reduced by a factor of 16.

We note that at the final stage of iterative correlation, we revert back to the original image to achieve high accuracy. We note that at this last stage,

---

[1] Similar to the 3D case (see Appendix C), we give a prediction for the 2D correlation computational cost as follows:

$$\begin{aligned} \text{DCCC:} \qquad & \approx 8\gamma\Delta^2N^2\left[1+n\left(\frac{2\alpha+1}{2\Delta}\right)^2\right] \\ \text{FFT:} \qquad & \approx \left(30N^2\log_2 N\right)\left[1+n-\frac{n\,(n+1)}{2\log_2 N}\right] \end{aligned}$$

where the interrogation window size is $N \times N$ and n is the number of iterative stages beyond the first stage. We also assume that during iterative correlation, the length of the interrogation window at any stage is half of the interrogation window length at the previous stage. Furthermore, the search distance α, is kept constant for all stages beyond the first stage.

the correlation search distance is about a particle diameter, so the time spent correlating the original image (via the direct computation of the Cross-correlation function) is still about the same order as that spent in the earlier correlations of the compressed grid via DCCC.

The total correlation time for the particle images whose results are shown in Figure 1.11 and Figure 1.12 is about 3.7 seconds for all 3 stages of correlation, hence equating to a vector processing rate of approximately 300 vectors/sec for PIV processing and approximately 1000 vectors/sec for Hybrid PIV-PTV processing, on a 2GHz Intel Pentium 4 processor PC with 1 GByte of RAM. For the first two correlation stages where the reconstructed image was used (i.e. DCCC), the correlation time was about 0.7 seconds of the total correlation time of 3.7 seconds.

Achieving high valid-data-yield

Because the particles are reconstructed using DCCC, we can set the experimental parameters to give the particle image size required for highest accuracy and not worry too much about high velocity gradients since the particles can be reconstructed to any required size on the grid.

Furthermore, because we can obtain the number of particles contained in each window, we can also define the interrogation window sizes that will ensure that we have the recommended number of particles in each window for high valid data yield. We can also validate the vectors obtained from the correlations by examining the number of particles contained within each window. For example, if a window contains far less than the recommended number of particle images necessary for high valid-data-yield, then the window may be tagged as spurious. Also if the window contains no particles, then the correlation of the original image will not be due to particle image correlations therefore that window should be tagged as empty.

A highly efficient data validation algorithm such as the Peak Stability validation and fixing routine described in section 3.5.2 in Chapter 3 will help to improve the valid-data-yield.

Maintaining high accuracy

If the conditions for DCCC to be faster than FFT based correlation hold, we can set the experimental particle image to the size required to give highest accuracy. The optimum particle image size is between 1 to 1.5 pixels. We do not need to have the original particle image size to be larger than this value just to overcome the high velocity gradients since we can always reconstruct particles to any size on the reconstructed correlation grid.

Experimental and correlation details for Figure 1.11 and Figure 1.12

Correlation type: DCCC (second interrogation window floats)
Machine: 2 GHz Intel Pentium IV PC with 1GB RAM
Actual particle seeding size: 37 to 53 microns.
Image Size: 1392 × 1040 pixels
Image Pixel Scale: 0.0407 mm/pix
Light sheet thickness: $\approx$ 0.8 mm
Correlation Grid Pixel Scale: 0.0814 mm/pix
Reconstructed Particle image size: 0.25 mm
Reconstructed Particle model: Sphere
Reconstructed Particle intensity Profile: Flat
Δt: 30ms
Interrogation window Size: 3 mm × 3 mm
Search distance Δ: 1.5 mm
Number of Iterative Correlation Levels (stages): 2. So, in Figure F.1, the value for k will be 1.
Validation method: Peak Stability validation and counting the number of particles in the windows to see if they are empty.

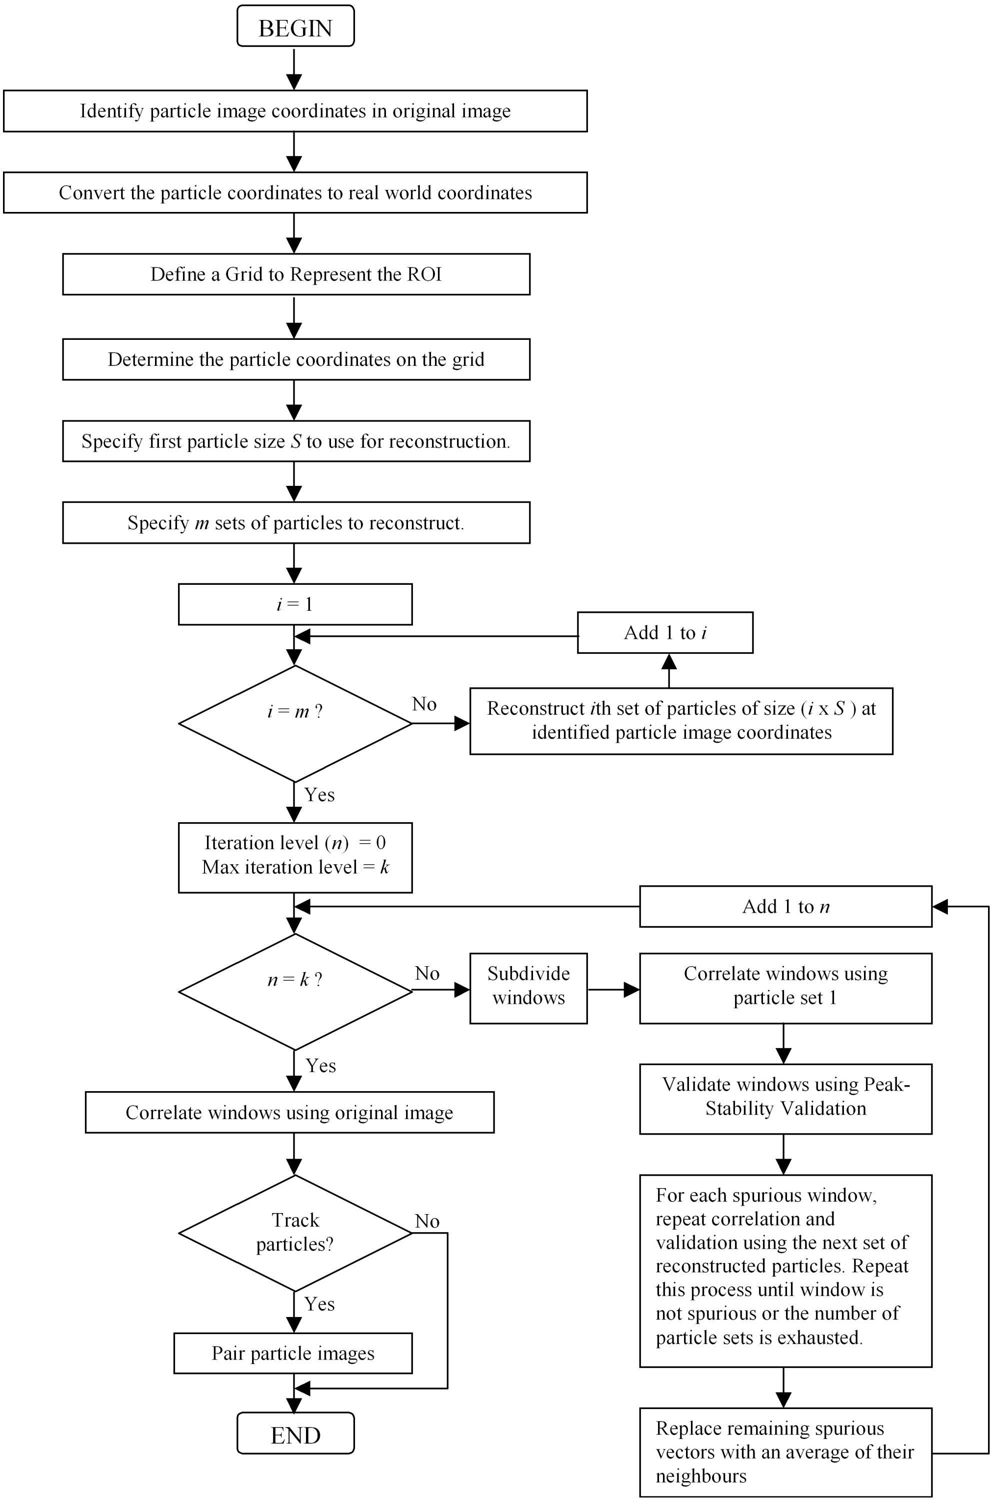

Figure F.1: The high performance 2D image analysis algorithm.

Appendix G

Seeding particles and their characteristics

Particle	Specific gravity	Size range	Reflectivity	Colour	Shape	Hardness	Chemical resistance
Cenospheres	0.65–0.95	5–300μm	High	Off-white	Spherical	High	Inert
E-Spheres	0.65–0.95	5–500μm	High	Off-white	Spherical	High	Inert
Pliolite	1.03–1.05	$\sim$ 2mm	High	White	Granular	Friable	Low
Polystyrene	1.05	5–600μm	Low	Clear	Spherical	Medium	Low
SCHCS	0.7–1.3	50–120μm	High	Silver	Spherical	High	Inert
Talisman	0.99	250μm (max)	High	White	Irregular	Medium	High

Table G.1: Seeding particles for liquid flows

SCHCS: Silver-Coated Hollow Ceramic Spheres

Suppliers

Cenospheres: Trelleborg Fillite, U.K.

E-Spheres: ABC chemicals Co. and Envirospheres PTY Ltd, Australia

Silver-coated Ceramic spheres: Potters Industries Inc, U.S.A.

Pliolite: Eliokem, France (U.K).

Talisman: Plascoat Systems Ltd., U.K.

Polystyrene: Micro-beads AS, Norway.

References

Adrian, R. (1984). Scattering particle characteristics and their effect on pulsed laser measurements of fluid flow: speckle velocimetry vs. particle image velocimetry. *Applied Optics*, 23:1690–1691. 3

Adrian, R. (1986). Image shifting technique to resolve directional ambiguity in double-pulsed velocimetry. *Applied Optics*, 25:3855–3858. 13

Adrian, R. (1991). Particle-imaging techniques for experimental fluid mechanics. *Annual Review of Fluid Mechanics*, 23:261–304. 12

Adrian, R., Meinhart, C., and Tomkins, C. (2000). Vortex organization in the outer region of the turbulent boundary layer. *Journal of Fluid Mechanics*, 422:1–54. 18

Agui, C. and Jimenez, J. (1987). On the performance of particle tracking. *Journal of Fluid Mechanics*, 185:447–468. 32, 130

Barker, D. and Fourney, M. (1977). Measuring fluid velocities with speckle patterns. *Optics Letters*, 1:135–137. 5

Barnhart, D., Adrian, R., Meinhart, C., and Papen, G. (1994). Phase-conjugate holographic system for high resolution particle image velocimetry. *Applied Optics*, 33:7159–7169. 40

Bracewell, R. (1999). *Fourier Transform and its applications*, volume 3rd Edition. McGraw-Hill Publishing company. 88

Brucker, C. (1995). Digital-particle-image-velocimetry in a scanning light-sheet: 3D starting flow around a short cylinder. *Experiments in Fluids*, 19:255–263. 41

Brucker, C. (1996). 3-D PIV via spatial correlation in a color-coded light-sheet. *Experiments in Fluids*, 21:312–314. 38

Brucker, C. (1997). 3D scanning PIV applied to an air flow in a motored engine using digital high-speed video. *Measurement Science and Technology*, 8:1480 – 1492. 41

Carosone, F., Cenedese, A., and Querzoli, G. (1995). Recognition of partially overlapped particle images using the Kohonen Neural network. *Experiments in Fluids*, 19:225–232. 32

Cenedese, A., Gregorio, F., Pococco, A., and Querzoli, G. (2000). *Effects of Images Compression on PIV and PTV Analysis*. Kluwer Academic Publishers. 18

Cenedese, A. and Paglialunga, A. (1989). A new technique for the determination of the third velocity component with PIV. *Experiments in Fluids*, 8:228–230. 38

Christensen, K. (2004). The influence of peak-locking errors on turbulence statistics computed from PIV ensembles. *Experiments in Fluids*, 36:484–497. 18

Cohn, R. and Koochesfahani, M. (2000). The accuracy of remapping irregularly spaced velocity data onto a regular grid and the computation of vorticity. *Experiments in Fluids [Suppl.]*, pages S61 – S69. 32

Dudderar, T. and Simpkins, P. (1977). Laser speckle photography in a fluid medium. *Nature*, 270:45–47. 5

Fouras, A. and Soria, J. (1998). Accuracy of out-of-plane vorticity measurements derived from in-plane velocity field data. *Experiments in Fluids*, 25:409–430. 32

Goss, L., Post, M., Trump, D., and Sarka, B. (1991). Two-color particle image velocimetry. *Journal of Laser Applications*, 3:36–42. 13

Guezennec, Y., Brodkey, R., Trigui, N., and Kent, J. (1994). Algorithms for fully automated three-dimensional particle tracking velocimetry. *Experiments in Fluids*, 17:209–219. 32, 33, 39

Guezennec, Y. and Kiritis, N. (1990). Statistical investigation of errors in particle image velocimetry. *Experiments in Fluids*, 10:138–146. 34

Hart, D. (1998a). The elimination of correlation errors in PIV processing. *9th International Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal*. 30, 31

Hart, D. (1998b). High-speed PIV analysis using compressed image correlation. *Journal of Fluids Engineering*, 120:463–470. 18

Hart, D. (1999). Super-resolution PIV by recursive local-correlation. *Journal of Visualization, The visualization Society of Japan*, 10. 30, 34

Huang, H., Dabiri, D., and Gharib, M. (1997). On errors of digital particle image velocimetry. *Measurement Science and Technology*, 8:1427–1440. 20, 28, 30

Huang, H., Fiedler, H., and Wang, J. (1993). Limitation and improvement of PIV. Part I: Limitation of conventional techniques due to deformation of particle image patterns. *Experiments in Fluids*, 15:168–174. 28

Huang, H., Fiedler, H., and Wang, J. (1993b). Limitation and improvement of PIV. Part II: particle image distortion, a novel technique. *Experiments in Fluids*, 15:263–273. 28

Keane, R. and Adrian, R. (1990). Optimization of particle image velocimeters. Part I: Double pulsed systems. *Measurement Science and Technology*, 1:1202–1215. 12, 29

Keane, R. and Adrian, R. (1991). Optimization of particle image velocimeters. Part II: Multiple pulsed systems. *Measurement Science and Technology*, 2:963–974. 12

Keane, R. and Adrian, R. (1992). Theory of cross-correlation of PIV images. *Applied Scientific Research*, 49:191–215. 22, 29, 30, 104, 176

Keane, R. and Adrian, R. (1995). Super-resolution imaging velocimetry. *Measurement Science and Technology*, 6:54–768. 34

Landreth, C., Adrian, R., and Yao, C. (1988). Double pulsed particle image velocimetry with directional resolution for complex flows. *Experiments in Fluids*, 6:119–128. 13

Lourenco, L. and Krothpalli, A. (1995). On the accuracy of velocity and vorticity measurements with PIV. *Experiments in Fluids*, 18:421 – 428. 28, 32

Maas, H., Gruen, A., and Papantoniou, D. (1993). Particle tracking velocimetry in three-dimensional flows. *Experiments in Fluids*, 15:133–146. 32, 39, 40

Malik, N., Dracos, T., and Papantoniou, D. (1993). Particle tracking in three-dimensional flows. Part II: particle tracking. *Experiments in Fluids*, 15:279–294. 39

Meinhart, C., Prasad, A., and Adrian, R. (1993). A parallel digital processor system for particle image velocimetry. *Measurement Science and Technology*, 4:619–626. 17

Melling, A. (1997). Tracer particles and seeding for particle image velocimetry. *Measurement Science and Technology*, 8:1406–1416. 26

Meng, H. and Hussain, F. (1991). Holographic particle velocimetry: a 3D measurement technique for vortex interactions, coherent structures and turbulence. *Fluid Dynamics Research*, 8:3352. 40

Meng, H. and Hussain, F. (1995a). In-line recording and off-axis viewing technique for holographic particle velocimetry. *Applied Optics*, 34:1827–1840. 40

Meng, H. and Hussain, F. (1995b). Instantaneous flow field in an unstable vortex ring measured by holographic particle image velocimetry. *Physics of Fluids*, 7:9–11. 40

Meng, H., Pan, G., Pu, Y., and Woodward, S. (2004). Holographic particle image velocimetry: from film to digital recording. *Measurement Science and Technology*, 15:673–685. 41

Merzkirch, W. (1987). *Flow Visualization*. Academic Press, second edition edition. 26

Meynart, R. (1980). Equal velocity fringes in a Rayleigh-Benard flow by a speckle method. *Applied Optics*, 19:1385–1386. 5

Meynart, R. (1983a). Instantaneous velocity field measurements in unsteady gas flow by speckle velocimetry. *Applied Optics*, 22:535–540. 5

Meynart, R. (1983b). Speckle velocimetry study of vortex pairing in a low-Re unexcited jet. *Physics of Fluids*, 26:2074–2079. 5

Nogueira, J., Lecuona, A., and Rodriguez, P. (1997). Data validation, false vectors correction and derived magnitudes calculation on PIV data. *Measurement Science and Technology*, 8:1493–1501. 28

Nogueira, J., Lecuona, A., and Rodriguez, P. (2001). Identification of a new source of peak locking, analysis and its removal in conventional and super-resolution PIV techniques. *Experiments in Fluids*, 30:309–316. 28

Okamoto, K., Hassan, Y., and Schmidl, W. (1995). New tracking algorithm for particle image velocimetry. *Experiments in Fluids*, 19:342–347. 32

Owodunni, O. and Nickels, T. (2003). Improving the performance of DPIV in flows with large gradients. *56th American Physical Society, DFD Meeting, New Jersey, USA*, Paper KC.003. 31

Pereira, F. and Gharib, M. (2002). Defocusing digital particle image velocimetry and the three-dimensional characterization of two-phase flows. *Measurement Science and Technology*, 13:683–694. 42

Pereira, F., Gharib, M., Dabiri, D., and Modarress, D. (2000). Defocusing digital particle image velocimetry: a 3-component 3-Dimensional DPIV measurement technique. Application to bubbly flows. *Experiments in Fluids [Suppl.]*, pages S78–S84. 42

Prasad, A. (2000a). particle image velocimetry. *Current Science, Review Article*, 79:51–60. 6, 28

Prasad, A. (2000b). Stereoscopic particle image velocimetry. *Experiments in Fluids*, 29:103–116. 37

Prasad, A. and Adrian, R. (1993). Stereoscopic particle image velocimetry applied to liquid flows. *Experiments in Fluids*, 15:49– 60. 37

Prasad, A., Adrian, R., Landreth, C., and Offutt, P. (1992). Effect of resoulution on the speed and accuracy of particle image velocimetry interrogation. *Experiments in Fluids*, 13:105–116. 20, 28

Raffel, M., Gharig, M., Ronneberger, J., and Kompenhans, J. (1995). Feasibility study of three-dimensional PIV by correlating images of particles within parallel light sheet planes. *Experiments in Fluids*, 19:69–77. 28, 37

Raffel, M., Willert, C., and Kompenhans, J. (1998). *particle Image Velocimetry: A Practical guide*. Springer. 26, 27

Scarano, F. and Riethmuller, M. (1999). Iterative multigrid approach in PIV image processing with discrete window offset. *Experiments in Fluids*, 26:513–523. 34

Scarano, F. and Riethmuller, M. (2000). Advances in iterative multigrid PIV image processing. *Experiments in Fluids*, pages S51–S60. 34

Schlicke, T. and Greated, C. (2002). Application of dual-plane particle image velocimetry to the measurement of breaking waves. *Measurement Science and Technology*, 13:1990–1995. 37, 38

Sheng, J., Malkiel, E., , and Katz, J. (2003). Single beam two-views holographic particle image velocimetry. *Applied Optics*, 42:235–250. 40

Soloff, S. M., Adrian, R. J., and Liu, Z. (1997). Distortion compensation for generalized stereoscopic particle image velocimetry. *Measurement Science and Technology*, 8:1441–1454. 39

Spedding, G. and Rignot, E. (1993). Performance analysis and application of grid interpolation techniques for fluid flows. *Experiments in Fluids*, 15:417–430. 32

Stitou, A. and Riethmuller, M. (2001). Extension of PIV to super resolution using PTV. *Measurement Science and Technology*, 12:1398–1403. 32

Stuer, H. and Blaser, S. (2000). Interpolation of scattered 3D PTV data to a regular grid. *Flow, Turbulence and Combustion*, 64:215–232. 32

Stuer, H., Maas, H., Virant, M., and Becker, J. (1999). A volumetric 3D measurement tool for velocity field diagnostics in microgravity experiments. *Measurement Science and Technology*, 10:904–913. 20

Van Meel, D. and Vermij, H. (1961). A method for flow visualization and measurement of velocity vectors in three-dimensional flow patterns in water models by using colour photography. *Applied Scientific Research*, 10:109–117. 43

Virant, M. and Dracos, T. (1997). 3D PTV and its application on Lagrangian motion. *Measurement Science and Technology*, 8:1539–1552. 39

Westerweel, J. (1993). Digital particle image velocimetry- theory and application. *PhD. Thesis. Delft University Press.* 16

Westerweel, J. (1994). Efficient detection of spurious vectors in particle image velocimetry data. *Experiments in Fluids*, 16:236–247. 30, 31

Willert, C. and Gharib, M. (1991). Digital particle image velocimetry. *Experiments in Fluids*, 10:181–193. 5, 6, 16

Willert, C. and Gharib, M. (1992). Three-dimensional particle imaging with a single camera. *Experiments in Fluids*, 12:353–358. 42

Zhang, J., Tao, B., and Katz, J. (1997). Turbulent flow measurement in a square duct with hybrid holographic PIV. *Experiments in Fluids*, 23:373–381. 40