Multiple CCD Array Digital Particle Image Velocimetry

Timothy Peter Dewhirst

A thesis submitted in fulfilment of the requirements

for the degree of Doctor of Philosophy

to the

University of Edinburgh

1998

Abstract

Digital image processing is a technologically limited subject, and has seen tremendous growth over the last ten years with the exponential increase in computing power commonly available. This trend can be seen in Particle Image Velocimetry (PIV) as the migration from traditional wet film cameras to fully digital systems. The use of digital cameras in PIV measurements has been common place for a number of years. However, the majority of the reported systems suffered from low spatial and temporal resolution due to the architecture of commonly available cameras. The aim of this thesis was to construct a system which was entirely digital, with high spatial resolution and flexible timing abilities, removing the need to use image shifting. In order to create such a system, the state of the art in digital imaging technology was reviewed. The conclusion of this study was that in order to achieve the aims set, a multiple CCD (Charge Coupled Device) system would be required.

The prototyping and construction of both two- and four-CCD array cameras is described. The software required to analyse images from low spatial resolution sources (relative to film) is also described, including the precautions that must be taken to ensure unbiased and accurate velocity measurements. Finally, the multiple CCD-array systems were applied to a number of flows to prove the concept. These include the study of breaking waves, wave motion over submerged pipes, and extension of the PIV technique to enable the measurement of accelerations.

Declaration

I declare that this thesis was composed by myself and that the work contained therein is my own, except where explicitly stated otherwise in the text.

(*Tim Dewhirst*)

March 1998

Acknowledgements

I would like to thank my family - Alison, my parents, Claire and Andy, for their support and encouragement.

I would also like to thank my supervisor Clive Greated for his continued support, and also being an excellent source of money for equipment and trips to other universities. Frank Morris has played an essential part in the running of experiments, and all of our lasers would have died without him. All of the other technicians in both the mechanical and electronics workshop have shown me great patience and have always provided a solution.

Michael Jakobsen's work with the acceleration measurements was invaluable, and greatly appreciated. His encyclopedic knowledge of electronics made much of the work possible and also helped me later on in constructing further timing devices.

Finally, my thanks to my fellow Ph.D.'ers, who have lightened many an evening, and who have allowed a good deal of drinking without having to sit on my own. This includes John "flood victim" Pullen, Jon "small but mighty" Entwistle, Zbigniew "small body, large brain" Sroczynski, little Tommy Haydon, Paul "Chunky" Rowland, Al "teletubby" Kelsey, Marieta "Maz" Soliva, Al "R2D2" Arnott and many others.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 History of Particle Image Velocimetry

Particle Image Velocimetry (PIV) has become a well established technique for the measurement of instantaneous full field velocity information in fluid flows, and as such it is almost exclusively used for the recording of this information being unsurpassed in flexibility of use.

A very great deal of work has been carried out into refining and improving the accuracy of the technique, and as a result PIV is now a commonly used and widely accepted measurement technique which allows a unique insight into the instantaneous velocity distribution within fluid flows [50, 56]. The growing acceptance of PIV has also resulted in the manufacture of equipment for use in PIV becoming commercialised, and this has allowed many of the technological barriers imposed by market forces to be removed. The range of flows which can now be handled by PIV has increased, again increasing the popularity of this tremendously useful technique.

PIV has always been a technologically limited field of study, and this trend looks set to continue into the foreseeable future. This thesis will outline a solution to a particular problem in this field, namely the rapid acquisition of a sequence of flow field images.

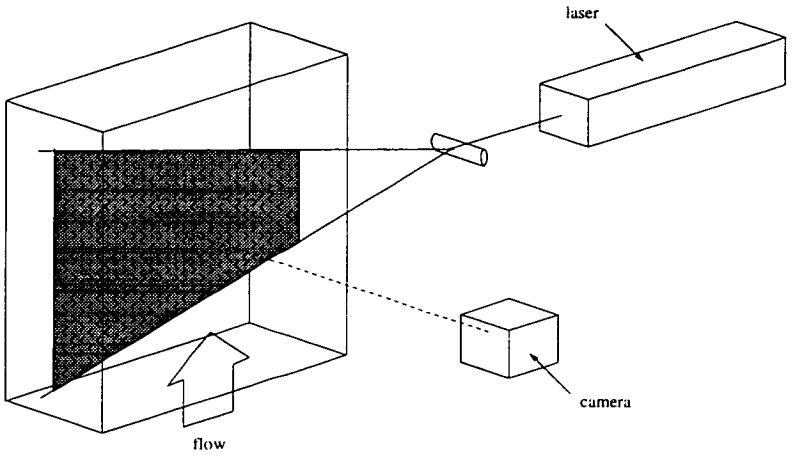

Figure 1.1: PIV setup

PIV is neither a new technique nor complicated in concept. The concept is a simple one, and it is only when we are looking at the accuracy of the method that complexity creeps in.

In order to record a fluid flow the flow has to be *seeded* with appropriate particles which will follow the fluid motion faithfully. The flow is then illuminated, usually with a stroboscopic sheet of high intensity light, and the light scattered from the particles is imaged by a camera placed normal to the illuminating sheet of light (see figure 1.1). The multiple images of each tracer particle are then matched, and so the local fluid velocity can be found. The only barriers which stand between the simplistic concept and the reality of reliable, high quality vector maps are technological and hardware based. The sufficient illumination of the flow, seeding of the flow, and capture of the particle images are of prime concern in the real experimental setup.

1.1.1 Flow Illumination

There are a number of ways of illuminating a fluid flow. Typically, either with a pulsed laser and a cylindrical lens system, or a continuous wave (CW) laser and

a scanning beam system are used [57, 41, 8, 5] , but lower powered sources have also been used [55, 17] including flash lamps and ordinary fluorescent lamps.

The system shown in figure 1.1 consists of a pulsed laser (for example a Nd:YAG) which is then passed through a series of lenses to alter the beam diameter, with finally a cylindrical lens expanding the beam in one direction only creating a light sheet [50]. The use of pulse lasers has the advantage that the pulse duration is in the nano-second region, and so they are ideally suited to high velocity ($\rightarrow$ trans-sonic) air flows, freezing the motion of the tracer particles very effectively. This can be more formally stated as a requirement that the particle image should move less than one of its own diameters during the duration of the illuminating pulse. This requirement is one of the main factors which limits the use of the scanning beam system to lower velocity flows [18].

The maximum measurable velocities achievable depend also on the minimum pulse separation, and this in turn depends on a number of factors. If a single cavity, dual-Q switched laser is used, the minimum pulse separation is not only strongly dependent on the optical resonator properties of the laser, but also the flash lamp capabilities. While this would appear to be a neat solution requiring only a single laser, in general this is not considered to be a good solution. The energies of the pulses are difficult to match and keep matched, and the energy per pulse is greatly reduced. The ideal solution is to use two lasers which are both Q-switched only once to allow the independent generation of the illuminating pulses, as well as allowing virtually any inter pulse separation [60].

However, pulsed lasers have the disadvantage that they have a noisy, non stable beam profile in general (although again this depends of the design of the optical resonator), and the use of a cylindrical lens will give rise to a Gaussian illumination intensity profile across the sheet. The noisy beam profile of the pulsed

laser is not of great concern in conventional PIV, but becomes important when extending the technique to measure, for example, three-dimensional velocities by the dual-plane technique (see chapter 7) [51]. The Gaussian intensity profile can also lead to problems with the analysis procedure, and in exceptional cases the captured image will need to be normalised before processing to remove these effects [62].

The scanning beam system was developed as a solution to the above problems at the University of Edinburgh by Gray *et. al.* [18, 20] for use in water wave studies and has now become a widely used system in PIV. This illumination system uses a constant wave (CW) laser to provide the illumination, and the result is a pseudo-static light sheet. The main advantages of this system are that it effectively uses all of the available laser power, and thereby allows a large area ($\geq 1m$ wide) to be illuminated, while simultaneously removing the Gaussian intensity variation across the illumination plane seen with the pulsed laser system. A typical scanning system is shown in figure 1.2. The laser beam is directed through a telescope arrangement allowing the beam to be reduced in diameter and collimated, and then directed onto the surface of a high precision rotating polygon mirror. The mirror lies at the focus of a parabolic mirror, and as the beam is swept over the surface of the parabola it is re-collimated into a pseudo-static sheet of light. The time taken to complete one scan (i.e. one full sweep of the parabola) is the strobe time, or the pulse separation to use the same terminology as with a pulsed laser system.

A standard PIV setup to investigate water waves is shown in figure 1.3, showing a scanning beam system and placement of the camera relative to the experimental apparatus. This is the configuration used in the majority of the work carried out in this thesis, as well as by many others [37, 41, 18].

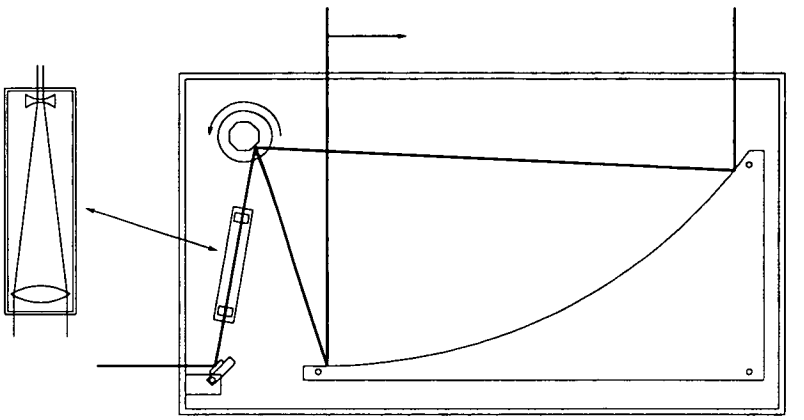

Figure 1.2: Scanning beam system

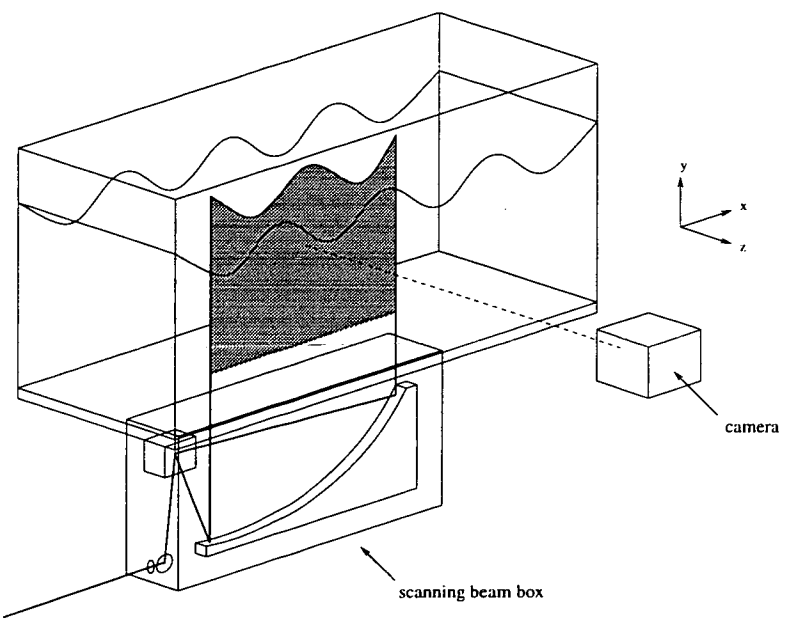

Figure 1.3: Water-wave study using PIV

fluid	seeding material	size
liquid	conifer pollen	$\approx 70\mu m$
	glass spheres	$10 - 100\mu m$
	silver coated hollow glass spheres	$10 - 100\mu m$
	gas bubbles (oxygen, hydrogen)	$50 \rightarrow \mu m$
gas	oil	$1 - 10\mu m$
	smoke	$1 - 50\mu m$

Table 1.1: Types of seeding material

1.1.2 Seeding

In section 1.1 it was stated that the seeding material will follow the fluid flow under investigation, and in an ideal case this would be true. However, the type of seeding that is used depends very strongly on the flow that is under investigation. The main specifications for the seeding material are that it should be small relative to the length scales present in the flow, not introduce any two-phase effects into the flow, follow the flow faithfully, and preferably have good scattering characteristics. A number of seeding materials have been tried out since PIV development began, and some of the more popular materials are listed in table 1.1 [25, 18]

The current preference at the University of Edinburgh is either for pollen or hollow glass spheres, both of which are relatively cheap and easy to use. However, the glass spheres are the easier to use as they do not go "off" as pollen does when left in water over long periods of time, and so the constant cleaning of the tanks which was required when using pollen has been reduced.

Because the seeding material has to be uniform and homogeneous to allow all areas of the flow to be recorded without loss of information, no information can be obtained from the flow unless either a long exposure or multiple images of each particle image are captured. In PIV, two or more exposures are used, and

once this is done the image from the camera can be analysed to give local mean velocities across the area of the flow that was illuminated.

1.1.3 Data Extraction

PIV stems from work in the sixties carried out into *Scattered Light Speckle (SLS) Metrology* and solid body mechanics, in which a rough (relative to the wavelength of light) surface is illuminated by a coherent source. The light which is scattered from the surface then forms a speckle pattern in space which can be recorded photographically [56].

If the surface or body is moved after the speckle pattern has been recorded, and then re-illuminated and re-recorded, the resulting photographic image can be considered to be a non-uniform diffraction grating. Probing the photographic recording with a coherent source will produce a set of *Young's fringes*, from which the magnitude and orientation of the displacement can be determined easily. The area probed in the interrogation is commonly referred to as the *interrogation area.* If the negative is then divided into a two dimensional grid of interrogation areas, a complete but discrete map of displacements can be built up.

In effect, these early applications of SLS Metrology were the limiting case of PIV in which individual particle images cannot be discriminated, but rather the seeding density in the fluid is high enough to produce a speckle pattern. This is often the case if the flow is seeded with, for example, milk [1] in the case of a liquid flow. Once it was realised that a truly specular pattern need not be produced, PIV as it is now known was created. The migration to lower seeding densities also meant that any effects introduced by the high concentration of seeding material, for example two-phase effects, possible in water due to the large relative size of the

[1] milk is an emulsion (from Latin *e-mulgere* "to milk") of, mainly, water and fat. It is the small fat globules which act as the seeding material.

seeding material, were removed along with any difficulties in actually obtaining such high seeding densities.

1.2 Analysis of PIV images

1.2.1 Optical Analysis

The analysis of PIV negatives has also been an area of considerable interest, with some remarkable schemes being produced to allow the rapid interrogation of the whole flow field [57, 26, 56] . As PIV developed from speckle based techniques in which individual particle images are not discernible, the analysis of the photographic images to extract the velocity or displacement information was carried out by means of a statistical correlation. The correlation of two functions can be defined as:

$$f(\underline{s}) \otimes g(\underline{s}) = \int_{-\infty}^{\infty} f(\underline{s}')g^*(\underline{s}' - \underline{s})d\underline{s}' \tag{1.1}$$

with $g^*(\underline{s})$ denoting the complex conjugate of $g(\underline{s})$ and $\otimes$ representing a correlation. If we are dealing with a digital image, the correlation can be calculated in two ways. The correlation can be performed directly by a summation. Equation 1.1 would then be recast as:

$$f(\underline{s}) \otimes g(\underline{s}) = \sum_{\underline{s}'=-\infty}^{\infty} f(\underline{s}')g^*(\underline{s}' - \underline{s}) \tag{1.2}$$

or, for an NxM image in which no phase information is stored:

$$f(x,y) \otimes g(x,y) = \frac{\sum_{i=0}^{N}\sum_{j=0}^{M} f(i,j)g(i-x, j-y)}{\sum_{i=0}^{N} f(i,j) \sum_{j=0}^{M} g(i,j)} \tag{1.3}$$

The correlation can also be recast in terms of the *Wiener-Khintchine theorem* which allows the correlation to be calculated using Fourier transforms and gives

$$\mathcal{F}(f(\underline{x}) \otimes g(\underline{x})) = F(\underline{u})G^*(\underline{u}) \tag{1.4}$$

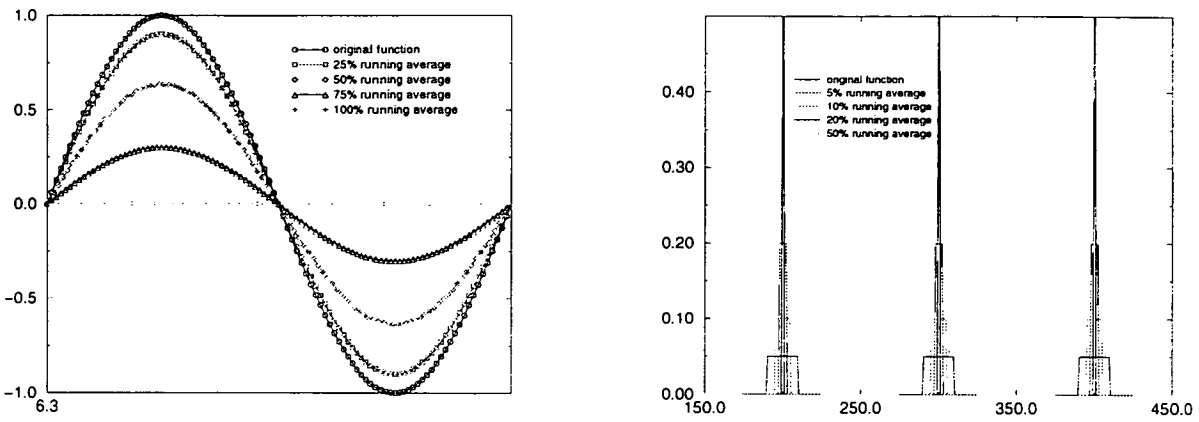

a) b)

Figure 1.4: Periodic signals with different levels of running average applied

with F(u) denoting the Fourier transform, $\mathcal{F}$, of f(x).

If the entire PIV negative was analysed in a single correlation, the result would be simply the mean velocity or displacement present in the recording. However, it is more often than not the spatial structures which are of interest, and so if the image is analysed in smaller sub-sections the result will be a map of the spatial velocity variations across the image.

In order to extract the spatial velocity information, the PIV image has to be divided into sub regions, with the simplest being the generation of a regular square grid of locations at which an interrogation is performed. The definition of the pitch of the grid is a trade off between the size of the interrogation area and the spatial frequencies within the flow that we want to try to resolve. The size of the interrogation area will directly influence the magnitude of the velocities that can be measured, and also the maximum spatial frequency that can be resolved.

The effect of the size of the interrogation area can be seen in figure 1.4. The interrogation area is essentially applying a running average or a low pass filter to the flow field, with the maximum spatial frequency that can be resolved being

preserved until the interrogation area or length of running average is the same as the period of the signal. However, the result is the original signal convolved with the spatial averaging function, and so the signal amplitude will be reduced, and depending on the nature of the function broadened. This would have been the case if we had chosen a periodic array of delta functions to illustrate this point (figure 1.4b). The delta functions would have been convolved with the spatial averaging function, but the area contained within the resulting function would remain constant, effectively reducing the amplitude of the periodic signal.

The size of the interrogation area is, however, usually defined by the particle seeding density present in the flow. The number of particle images which must be present within an interrogation area was a subject of study by, most notably, Keane & Adrian [27, 28]. The conclusion of their work is the almost gospel requirement that to give a valid vector probability of greater than 95% each interrogation area should contain roughly 10 particle image pairs and have a spatial velocity gradient of not more than 5%.

Now that the interrogation area has a minimum size to insure the detection of valid velocity vectors the grid size can be calculated. Following the Nyquist theorem, in order to resolve a signal containing a maximum frequency f, the signal must be sampled on a regular square grid at twice the maximum frequency, $2f$. If we are looking at an isotropic homogeneous turbulent flow, then the spatial resolution of our imaging system will not be able to resolve to the highest wave numbers, as these are dictated by the viscosity of the fluid, and will typically be higher than the interrogation area would be able to resolve. Therefore the highest spatial frequencies we could hope to resolve would be achieved by overlapping the interrogation areas. But how much can the interrogation areas be overlapped?

If we have an interrogation area of NxN pixels, then the highest spatial fre-

quency that can be measured will be:

$$f \equiv \frac{1}{N} \tag{1.5}$$

as all frequencies higher than this will effectively have a running average of over 100% of their period applied to them, and this will result in a constant factor (assuming they are locally periodic). Therefore, to sample this frequency correctly, the sample rate should be twice the highest frequency, i.e. $2f$. This then translates into a grid spacing of:

$$grid = \frac{1}{2f} \equiv \frac{N}{2} \tag{1.6}$$

This can be seen to be the 50% overlap so commonly employed in PIV analysis. Using a grid spacing of 25% or less of the interrogation area will result in the data being over-sampled, and so no further data will be extracted.

If we now return to the process of analysing the PIV image, the Wiener-Khintchine theorem allows a simple solution to the problem of extracting the local velocity data.

The PIV photograph is divided into a grid of interrogation areas, and at each interrogation area the local velocity needs to be determined. The most popular method was to probe the negative with a laser to give a set of Young's fringes [15]. This is the result of the non-uniform diffraction grating previously mentioned. Using a lens allows the far field result of the diffraction to appear in the focal plane of the transforming lens [13]. However, to calculate the reverse transform, the Young's fringes have be first of all multiplied by their complex conjugate. This is easily achieved simply by placing a screen in the focal plane of the transforming lens, as the intensity distribution of a coherent light source is its complex amplitude multiplied by its own complex conjugate. However this intensity pattern cannot be inverse transformed by another lens as it is incoherent.

The reverse transform can be carried out in a number of ways, but the most common are:

- Optically Addressed Spatial Light Modulator (OASLM)
- Capture the image using a digital camera

If the OASLM is used the OASLM performs the incoherent to coherent transform, and the remaining Fourier Transform is carried out by another lens. A *4-f* optical processing system has been set up with the output plane of the system containing the correlation. Ideally this should be the best solution as the optical bandwidth of the system is limited only by the quality of the optics and the OASLM. However, the development and relative scarcity of OASLM has meant that the second method for performing the reverse transform has become the more dominant. Work using an OASLM was carried out by Jakobsen [26] amongst others however with good initial results.

If the Young's fringes are captured by a CCD array, the resulting data can be transferred to a computer and the remaining transform carried out digitally by means of the computationally efficient, but still slow compared to a lens, *Fast Fourier Transform* (FFT) routine. The FFT is a routine that has become common place in many aspects of data processing, due mainly to the fact that it improved performance over the direct calculation of a Fourier Transform by an order of magnitude or two. The centroid of the displacement peaks in the correlation function are then found and the displacement can be recorded as a set of floating point numbers ready for further analysis or post-processing.

1.2.2 Digital Analysis

The optical processing methods were the dominant analysis techniques until recently, when the continuing exponential drop in computer prices and increase in

computing power made it feasible and affordable to carry out the whole procedure digitally. Typically, the time taken to analyse one point or interrogation area with a half-optical and half-digital system would be roughly two seconds. The majority of the time taken for analysing an interrogation area optically was used in physically moving the laser probe to another area on the PIV image. This process was usually carried out by a stepper motor driven system, and the result was an analysis which would take *hours* to process a single PIV image [18].

It is now common place for the entire analysis process to be digital, and as such the bottleneck of having a physical movement for each interrogation area is removed. The current analysis speeds range from 30-120+ points/second on a standard Pentium based desktop computer, and up to 1000 points/second on a high end DEC Alpha for an equivalent size of interrogation area. As all of the data is now digital, the use of dedicated hardware can allow the analysis process to be speeded up further. This dedicated hardware is most commonly in the form of an expansion card for the analysis computer, allowing the Fourier transforms to be calculated at a hardware, not software level. These expansion cards, normally called array processors or Digital Signal Processing (DSP) boards, allow relatively large interrogation areas to be transformed in milliseconds, rather than tenths of seconds.

The ubiquity of Intel based PC's in research labs has also lead to a number of commercial analysis software packages being developed and sold. Examples of commercial PIV analysis software would be AEA's VisiFlow, OFS' VidPiv or TSI's Insight. However, as the subject of PIV image analysis is still developing, most educational institutes have their own version of analysis software which can be freely modified.

Systems also exist which are entirely hardware based, being essentially com-

puters which have a very limited command base. The analysis program is typically stored as *firmware* [2], and this can allow the processor to take the output directly from a digital camera, and analyse the entire image at the frame rate of the camera. As such, this system is a near *real-time* processing system, although there are still bottlenecks present, most notably the capture and subsequent transfer of the image to the processing system. This system is ideal if very large amounts or long time sequences of PIV data are required. It can also be used to allow the near instantaneous interactive optimisation of experimental parameters, such as illumination levels, camera placement, seeding concentrations, etc.. However, such systems do have their drawbacks, the most notable of which is the lack of flexibility in analysis routines used. Such a system has been manufactured by Dantec Inc..

1.2.3 Autocorrelation

For all of the above discussion it has been assumed that there are multiple exposures of the scattering particles or speckle pattern per photographic image or *frame*. This is in essence the most simple description of PIV, as two or more exposures of each seeding particle are required to make a displacement or velocity measurement from the image. If it is the case that there are multiple exposures present, the analysis is termed an *autocorrelation*, as to extract the displacement information the interrogation area is correlated with itself to give a local measure of the displacement.

This method was, until recently, the most dominant method, and suffered from some well documented shortcomings [18, 62], namely:

- directional ambiguity of 180° (figure 1.5a)

[2] firmware is typically a program which is stored semi-permanently in an EEPROM (Electronically Erasable/Programmable Read Only Memory), and so although the instructions can be changed, this would usually be a job for the vendor of the equipment

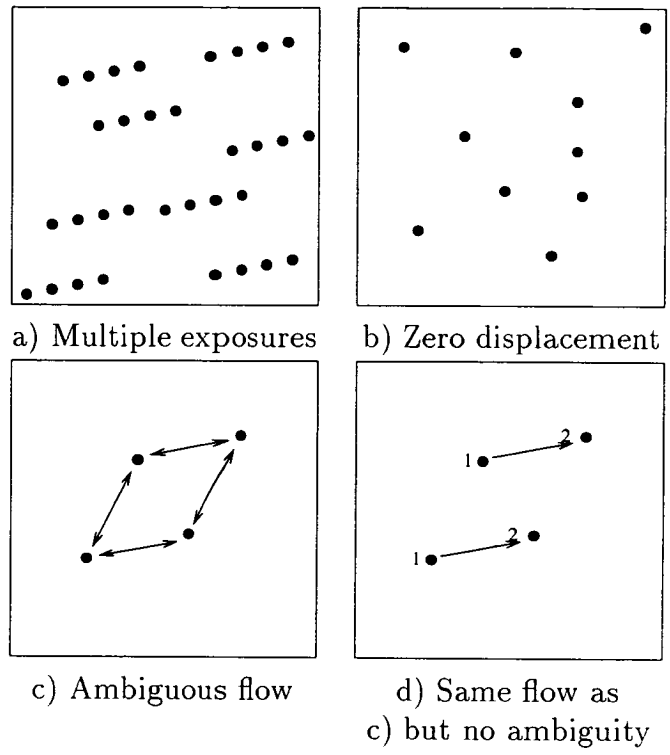

Figure 1.5: Possible interrogation areas

- low dynamic range or an inability to measure displacements as they approach zero (figure 1.5b)
- higher probability of random correlations (figure 1.5c)

Figure 1.6 shows a typical autocorrelation plane, which has a number of distinct features to it, after the notation of Keane and Adrian [27, 28, 34, 62]. The main feature to notice is that as this is an autocorrelation, there is a self correlation peak (R_P), which is always the tallest peak in the correlation plane (assuming that a valid correlation has been made). The peaks that are of interest, however, are the signal peaks (R_{D+} & R_{D-}) which are placed symmetrically around the self correlation peak. The number of signal peaks depends on the number of particle images per interrogation area. For example, if for each particle, four images of the particle are recorded on one frame of photographic film, there will be three

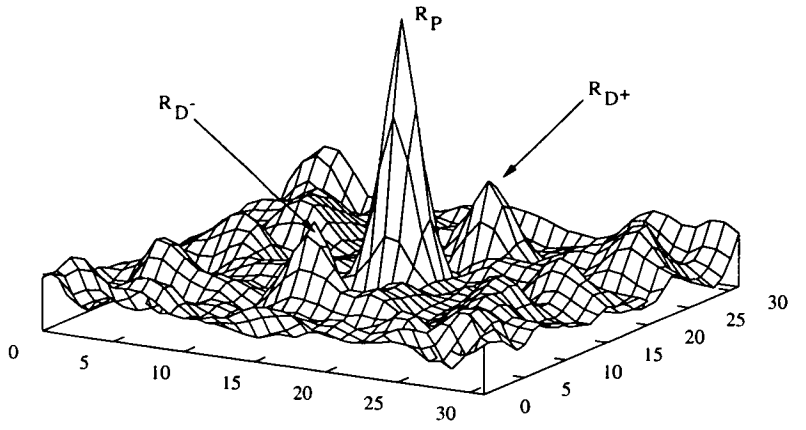

Figure 1.6: Autocorrelation plane

signal peaks on each side of the central correlation peak, assuming that the signal peaks can be resolved above the background noise present in the correlation plane. This noise is the result of random correlations and noise in the image. In general, however, for N images of the same particle there will be N-1 signal peaks.

Because each exposure of a particle image is, under ideal conditions, identical to all of the other exposures, the resulting train of particle images has a direction ambiguity inherent to it. The order of the particle exposures cannot be determined from inspecting the PIV image. Therefore, when carrying out the interrogation of the PIV image, the correlation plane of each interrogation area will reflect this ambiguity. This can be seen in the reflective symmetry of the signal peaks around the self correlation peak. This ambiguity can be ignored for some flow situations i.e. air flows in which the predominant velocity is in one direction and hence an *a priori* knowledge of the flow is available, but for any flow featuring a reversal a solution needs to be found.

A problem also arises as the velocity in the interrogation region approaches zero. If multiple exposures of each particle are present in the PIV image, then as the local velocity in the interrogation area approaches zero, so the sequential

images of the seeding particles will start to overlap each other, until they can no longer be resolved. This is very similar to the defining criteria for being able to resolve two star images with a telescope. As the particle images start to overlap, so the signal peaks will approach the self correlation peak, altering the centre of mass of the signal peaks. Thus, the velocity will become biased towards zero. As the particle images move closer still, the signal peaks will become indistinguishable from the central self correlation peak. As the local velocity is now zero, the interrogation area will contain a randomly placed set of seeding particles. The correlation of this interrogation area will reflect this fact, and will have only a self correlation peak and a noise background present. The highest random noise peak will now be chosen as the (obviously wrong) velocity measurement, resulting in a spurious vector or *outlier*.

Several techniques have been developed to remove some of these problems, most notably the technique of image shifting using rotating mirrors [47, 37], panning the camera [37] or electro-optic [30] devices. The concept of image shifting is a simple one, and the aim of all image shifting systems is to impose a velocity on top of the velocity distribution present within the flow. This is achieved by moving the image relative to the recording medium between the stroboscopic illumination of the flow. The resulting displacements present in the PIV image are then all biased to a value which is nearly uniform, and so all velocities extracted from the image will have a known direction (hence removal of directional ambiguity) and a non-zero value (extension of the dynamic range to include small and zero velocities).

However, image shifting also has its drawbacks. The movement of the image of the flow will result in any stationary boundaries present within the flow being smeared across the image in the direction of the shift. This is an obvious

drawback if it is the boundary layer which is being measured. The apparatus used in performing the image shift can also introduce distortions [37] which can only be removed by performing ensemble averaging of a large number of PIV vectors maps. Further, the majority of image shifting apparatus is external to the recording camera, and so introduces an additional level of complexity into the experimental proceedings.

1.2.4 Crosscorrelation

Another approach is to either label or store the particle images in such a way that the directional ambiguity is removed.

This is an area of research which has been widely investigated, with the most notable methods of approach being

- the alteration of the Point Spread Function of the optical system between exposures to allow post photographic developing filtering to distinguish between illumination pulses [31]
- the use of a gated image intensifier to "tag" the pulses so the first could be easily identified [17, 16]
- the use of a triple correlation function where the strobe illuminating spacing is not equal, and so a distinct triplet of images is used [12]
- the use of optics to split a photographic film plane in two, recording the first exposure one one half, and either one or both exposures on the second half [66]
- storing the sequential particle images on different frames or fields of a digital camera [9, 33, 35, 46, 52, 64]

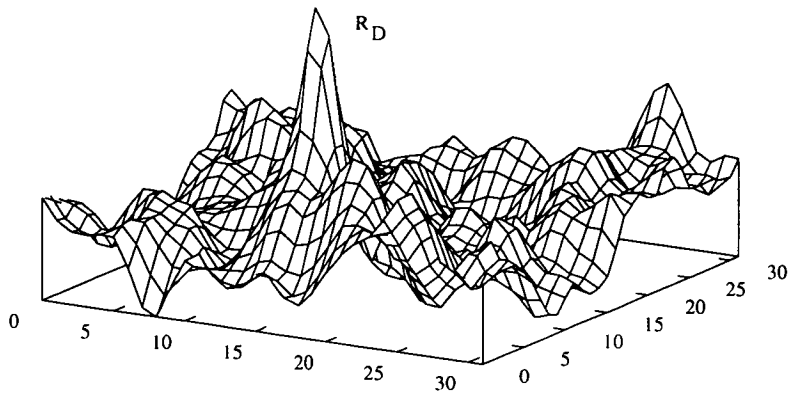

Figure 1.7: Crosscorrelation plane

The last area, storing the sequential images of the flow field on separate frames, is the approach adopted by the author.

If the sequential images of the flow field are stored in different locations, then the interrogation areas which correspond to the same spatial location in the object plane are extracted and correlated together to give a local displacement measurement. This is termed a *crosscorrelation.* Figure 1.7 shows a typical crosscorrelation plane, differing from the autocorrelation plane in that there is no self correlation peak (R_P) and there is only one signal peak (R_D). It can therefore be seen that there is no directional ambiguity, zero velocity flows are detected properly and there is a lower chance of random correlations (figure 1.5d).

1.2.5 Sub Pixel Resolution

As the analysis of the PIV images at some stage involves recording the PIV data with a CCD camera, the question of sub-pixel accuracy becomes important. In all cases, a correlation peak will be produced, and in order to extract the maximum amount of information from the interrogation area, we need to find the position of the correlation peak to sub-pixel accuracy.

When the predominant method of image analysis was half optical/half digi-

tal, the Fourier plane was digitised by a CCD array of typically either 256x256 or 512x512 pixels. The resulting image was then transfered to a computer and the Fourier transform was performed digitally. The resulting correlation peaks covered a large number of pixels, having a diameter of perhaps 10+ pixels. In these cases the central pixel location of the correlation peak is probably enough information to give an accurate measurement, as allowing for a maximum displacement of 64 pixels in a 256x256 interrogation area, the uncertainty will be $\pm\frac{1}{64}$. However, further information can be extracted by performing a centroiding or centre-of-mass routine on the peak, which is defined as being the ratio of the first order to the zeroth order moment of the peak.

However, with the ever increasing use of digital storage media, the interrogation areas have decreased in size to between 16x16 $\rightarrow$ 64x64 pixels. This, in conjunction with studies carried out into the sampling criteria for the PIV images [62, 64, 39] has lead to peak diameters in which only the pixels immediately surrounding the highest point in the correlation plane will contain any useful information. If this is the case, the routines used to calculate the sub-pixel displacements are referred to as *three point estimators.* The result of the estimator lies between $-0.5..0.5$, and the most commonly used estimators are listed in table 1.2 [62].

The more information that can be gained about the particle images the better tailored the peak-centre finding routine can be. The currently accepted best method for finding sub-pixel displacements is to apply a peak fitting routine to the correlation peak. Initially, a parabolic curve was fitted to the correlation peak, and this resulted in a great improvement over the traditional centroiding routine. However, as the particle images are essentially very well represented by a Gaussian intensity profile, the best results should therefore be obtained by

Method	Estimator
Centroid	$x = \frac{-R(i-1,j)+R(i+1,j)}{R(i-1,j)+R(i,j)+R(i+1,j)}$ $y = \frac{-R(i,j-1)+R(i,j+1)}{R(i,j-1)+R(i,j)+R(i,j+1)}$
Parabolic	$x = \frac{R(i-1,j)-R(i+1,j)}{2(R(i-1,j)+R(i+1,j)-2R(i,j))}$ $y = \frac{R(i,j-1)-R(i,j+1)}{2(R(i,j-1)+R(i,j+1)-2R(i,j))}$
Gaussian	$x = \frac{lnR(i-1,j)-lnR(i+1,j)}{2(lnR(i-1,j)+lnR(i+1,j)-2lnR(i,j))}$ $y = \frac{lnR(i,j-1)-lnR(i,j+1)}{2(lnR(i,j-1)+lnR(i,j+1)-2lnR(i,j))}$

Table 1.2: Three commonly used three point estimators

fitting a Gaussian curve to the correlation peak (this is a valid approach, as the correlation of a Gaussian with another Gaussian results in a Gaussian). This is the case, and in the majority of analysis a three point Gaussian fit estimator is used to provide a measure of the sub-pixel displacement.

1.3 State-of-the-Art PIV

It is the transition from wet film based cameras to digital CCD array cameras which is the most interesting in the authors view. The movement to a fully digital PIV process has allowed advancement in many areas, including the possibility of near real-time flow mapping and the removal of the need for image shifting for example, which allows PIV to be used in a wider range of situations with greater ease. The benefits also include the ability to see the images online, allowing for real time optimisation of laser powers, sheet geometries, etc. which would not be possible with a conventional film based system.

Advances in computing hardware are increasing the bandwidth of data that can be transfered from an imaging device to the host computer for storage and analysis. This is allowing the flexibility to record time series of images, whilst

image compression techniques can be employed to increase the bandwidth further. This compression can also be used to increase the speed of analysis while also allowing non-regular shaped interrogation areas to be used [19]. The compression of the images in a lossless form also allows the storage of a larger volume of data, which, although storage media has also increased exponentially (a factor of over 100 in under 10 years) is always useful.

The market for digital cameras and computer hardware for acquiring images from these cameras was until recently quite small. This made the variety of equipment that was available limited, whilst the lack of any real standards meant that the market was quite confused. The small market also meant that the prices for such equipment was often prohibitively high. For example, a 2048x2048 pixel CCD camera cost in excess of £20,000 three years ago, while now a camera with similar spatial resolution, based on CMOS technology can be purchased for £3,000. The maximum size of arrays is also increasing, and soon the question of spatial resolution will no longer be a problem.

The basic PIV technique has also been extended to include the capability to record time sequences of information, allowing the evolution of the flow to be measured, the use of holography and specialised drum scanning systems to allow volumetric, three component velocity measurements to be made and as such PIV has found applications in fields as diverse as flow through artificial heart values [22], mould filling [59] to large scale engineering experiments.

PIV is essentially a technologically limited subject, and as such many of the limitations which are common today will have simple and elegant solutions in the near future. This can be clearly seen in the domestic computer market, which is changing at a monthly pace. The ever increasing computer power that is being developed for high bandwidth applications involving multi-media will result in

PIV measurements having ever increasing resolution, both temporal and spatial.

Chapter 2

Digital Cameras

Digital cameras have become widespread in PIV experiments in the last few years. The first experiments that were carried out using digital cameras used what hardware was commonly available at the time, and even with the limited spatial and temporal resolution available some good results were obtained. One of the most striking examples of this would be the early work of C. Willert [65].

The digital camera has many features which appeal to the experimentalist. The removal of the need to process films, the possibility of acquiring a rapid sequence of images, the real time viewing of PIV images (and therefore optimisation of the illumination system), and more are all reasons for using a digital camera. Of course, digital cameras also have their drawbacks:

- Cost
- Limited spatial resolution
- Inflexible
- Requires dedicated hardware to store images

However, despite these limitations, the digital camera has become very popular in PIV experiments, and continues to be so.

Photographic film has a spatial resolution of up to 400 lines/mm, and typically for 100-400 ASA film we can expect to find a resolution of 100-300 lines/mm. Two typical film formats commonly used for PIV at the University of Edinburgh are 35mm (36x24mm) and 120 (120x100mm), which correspond to CCD arrays of 7,000x5,000 pixels and 24,000x20,000 pixels for a spatial resolution of 200 lines/mm. CCD arrays with these resolutions are typically considered to be beyond the attainable at the moment, with the largest arrays commonly in use being in the 1,000x1,000 to 2,000x2,000 pixel range.

There are companies that make very high resolution CCD arrays, including DALSA Inc., Kodak-Eastman and Philips. Apart from Philips all other companies making CCD arrays design each array as a single, monolithic architecture. This is an impractical way to design regular, repeating structures like CCD arrays for a number of reasons:

- CCD arrays are prone to defective pixels during manufacture, the larger the array, the more difficult it is to keep the central imaging area free of blemishes.
- New models require a total redesign, with all of the associated testing of the new design and re-tooling the production line.
- Such large designs are prone to crosstalk and smearing of the image.
- Cost. Low production run, very complex arrays cost a lot of money.

This is changing. Philips have developed a system whereby blocks of 1,000x1,000 pixels are literally stuck together on the same wafer to allow the simple construction of very large CCD arrays. The limiting factor for the size of the array is the dimension of the wafer, however the largest test array produced is 9,000x7,000

pixels, with a commercial version of 3,000x2,000 pixels in 35mm format available to buy. It is not inconceivable that in the not too distant future the current limitations of CCD arrays will be removed to allow single chip, high spatial and temporal resolution systems to be manufactured.

2.1 Spatial Resolution

The Nyquist sampling criteria states that in order to reconstruct exactly a continuous two dimensional signal, a square grid of twice the maximum spatial frequency present in the signal must be sampled. However, the aim of PIV is not to perfectly reconstruct the image, but rather to extract the local information about a group or *ensemble* of particle images, using the terminology of Westerweel [62].

It was first noticed in a series of experiments by Prasad *et. al.* [39] that in decreasing the resolution used in the second Fourier transform of an analysis by Young's fringes effectively no loss of accuracy was noticed in the velocity measurements obtained. This lead to the recommendation that in order to minimise errors in Digital PIV (DPIV) the optimal size of the particle images should be of the order of two pixels diameter. This result has since been confirmed by both mathematical models [62] and experimental results [64].

The conclusions which can be drawn from the above results indicate that the absolute spatial resolution of photographic film need not be mimicked in CCD's if they are to be used for PIV. This is an important result, because while CCD arrays with very high pixel counts (i.e. greater than 4,000,000) are available, for use in practical situations arrays of 1000x1000 pixels are currently the norm. This is due to a combination of price (£20,000 for a 2000x2000 CCD camera which will probably be dated in five years) and shuttering capabilities (most very large arrays are full-frame in construction - see section 2.2.2).

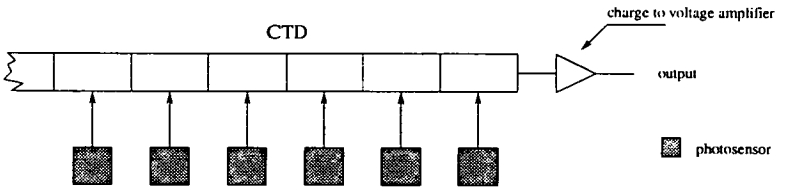

Figure 2.1: Simple CTD

2.2 Charge Coupled Devices

There are a number of different types of digital cameras available to buy, ranging from low resolution, interlaced designs aimed at the domestic TV and video markets, up to very large resolution arrays aimed at the professional stills photographer.

An excellent introduction to the operating theory behind Charge Transfer Devices (CTD) and CCD's can be found in *Handbook on Semiconductors* [11], with the main features of CTD's being presented here.

Charge Coupled Device (CCD) is a generic term for devices which produce charge as a result of an event, with the accumulation of charge then providing a measure of the event under investigation. These events can be anything which would produce a charge in the semiconductor sensing area, and so in general they are part of the electromagnetic spectrum, including X-rays and visible light. The event has to have enough energy to cause a band gap transition in the semiconductor, giving rise to an electron-hole pair. The hole is then removed through a drain and the electrons are stored in a potential well. Because the event has to cause a transition, the spectral response of the CCD will depend on the doping of the semiconductor material. To allow sufficient light to be gathered, CCD arrays also have an *integration time* over which the individual pixels integrate the effects of the events. This is exactly analogous to the exposure

time of a photographic camera.

If we consider the case of light, a typical row of a CCD device would look similar to figure 2.1, with a number of photosensitive cells being coupled to a CTD. The CTD works in the following fashion; the aim of the CTD is to transfer the charge accumulated in each photo cell to the single output stage without the individual packets of charge interfering with each other. This is achieved by having a basic unit of three *potential gates* or valves which can have one of three states, open, half-open or closed. The states of the gates are moved into the next cyclic state with the application of a clock pulse to the system (i.e. open→half-open, half-open→closed, closed→open) , as shown in figure 2.2. This has the effect that each trapped charge packet is moved along the CTD, towards the output amplifier while remaining separate from its neighbouring charge packets. The result, when combined with a charge to voltage amplifier, is an analogue voltage signal at the output of the CTD which is directly proportional to the charge stored at each pixel location.

The most common devices are for use in visible light, and typically consist of either a linear or a two-dimensional array of photo-diodes which convert the incident photons into a charge which is stored locally in a potential well until each well is read out sequentially. This is exactly analogous to the description above.

The way in which the two-dimensional arrays are constructed is important in determining which can be used for DPIV. There are a number of different ways in which the array can be constructed; a *full-frame interline transfer* device (figure 2.3), a *full-frame* and *frame transfer* device (figure 2.4) or a *line transfer* device (figure 2.5) being the most common. The output from these devices can be termed *progressive scan*, *full-frame*, *area scan* or *non-interlaced*, amongst others, but all

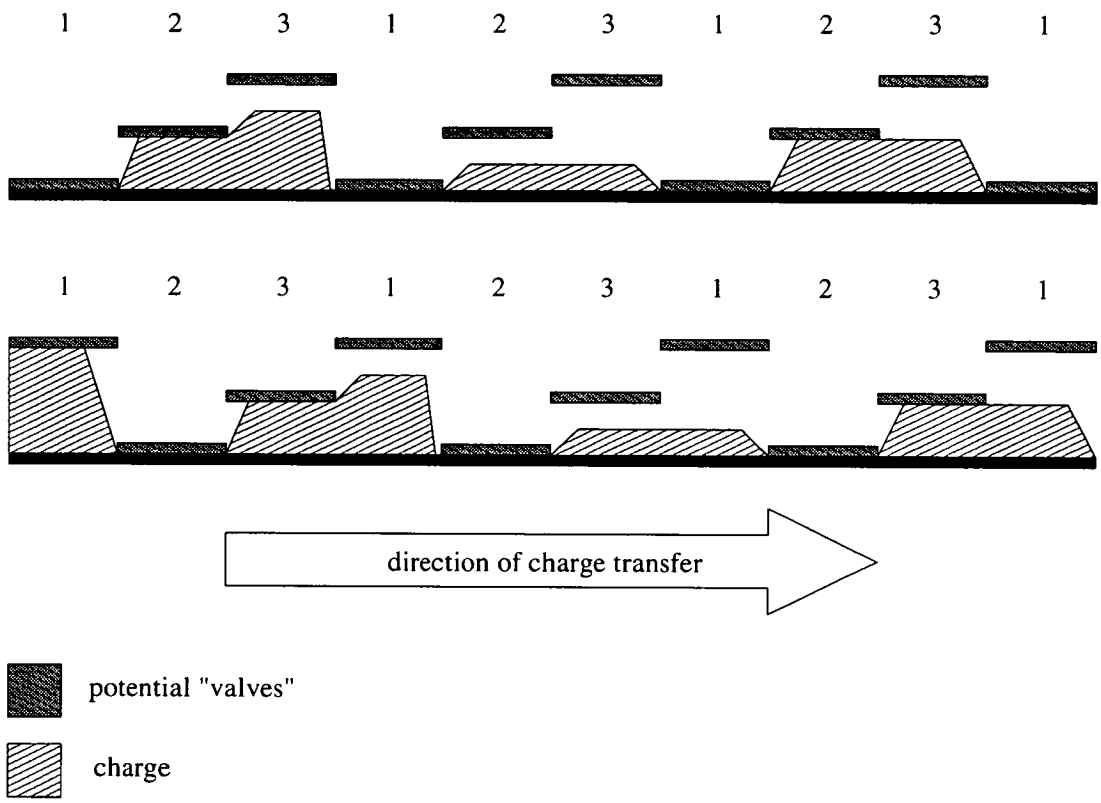

Figure 2.2: Schematic of a CTD

of the names refer to the fact that the entire CCD array is allowed to collect charge over the same time period, and the array is then read out sequentially in its entirety. This is important in allowing the full resolution of the array to be used, as we will see in Section 2.5.

The most important characteristics of any CCD array for use in DPIV have to be the fill factor, the shuttering capabilities and the sensitivity.

The fill factor of a pixel is defined to be the percentage or ratio of light sensitive area to other circuitry, such as charge transfer channels, interline storage areas, etc. within the active area. As all CCD arrays are defined by their pixel size or pitch in both directions, then the fill factor is important in determining how much light is lost per pixel, and also how much information is lost in the case of imaging small objects, for example particle images.

The shuttering of the device is important as if there is no electronic shuttering available, as with full-frame, line or frame transfer devices, the image formed on the CCD array will be prone to the effects of smear. This is the effect created when the pixels are still integrating the incident light whilst being read out, giving rise to a smearing of the image data in the direction of the readout.

Finally, the sensitivity of the device is very important, as otherwise the particle images from the seeding material in the flow may not register at all with the CCD array. The sensitivity of a particular CCD array depends on many factors, but mainly on the construction technique. The more silicon substrate the incident photons have to travel through before reaching the semiconductor junction, the less sensitive the device will be. Special CCD arrays are available which have been thinned either mechanically or chemically, but these are uncommon and expensive, with their main use being in astronomy. The architecture of CCD array which can be thinned is also limited to full-frame type arrays.

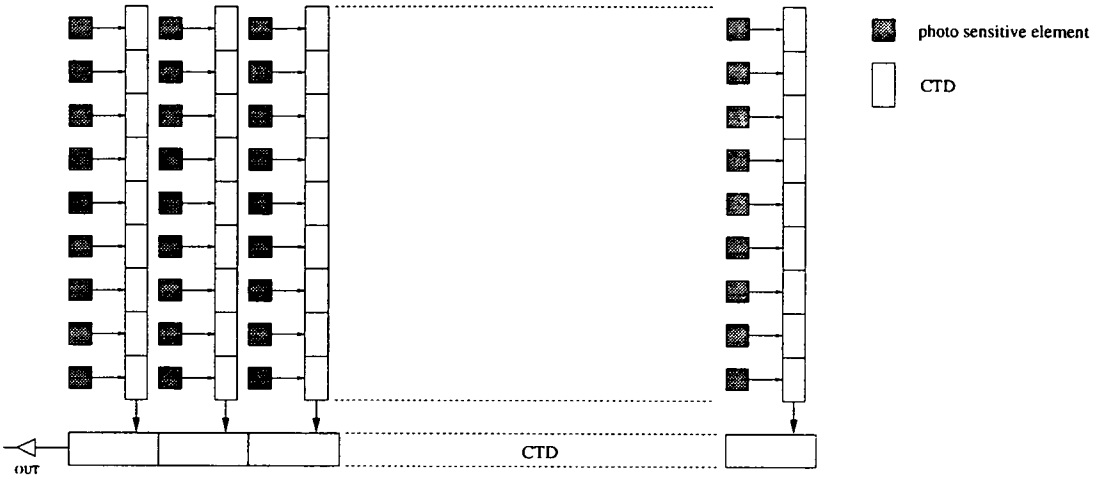

Figure 2.3: Interline transfer device

2.2.1 Full-Frame Interline Transfer

If the architecture of a CCD array is termed interline transfer, each pixel will have a photo sensitive area and an associated charge storage area. The charge generated during the integration of the array can then be transferred from the photo-cell to the storage area before being read out. As the name of the array type suggests, the location of the storage area is between the rows of photo sensitive material. The storage area also forms a CTD, and so for each row of pixels there is a CTD, the outputs of which form one final output CTD. A schematic layout of a full-frame interline transfer device is shown in figure 2.3.

The advantage of such an architecture is that the time taken to transfer the charge from the photo cell to the storage CTD is very short, and so in effect a very quick electronic shutter ($\mathcal{O}(\mu s)$) has been created. However, having the photo cell and associated CTD next to each other reduces the amount of space available for photo sensitive material and hence the fill factor of the array. This can have only one result, that of reducing the sensitivity of the device. A typical fill factor available from an interline transfer device would be that of a Pulnix

TM-9700, which have a pixel size of $11.6\mu m \times 13.3\mu m$ and a fill factor of 20%.

2.2.2 Frame and Frame Transfer

The architecture of a full-frame or frame transfer device is different from an interline transfer device, as for each pixel there are no separate photo-cells and storage areas. The photo sensitive area is also the CTD, and after integration the entire charge pattern is transferred down the photo sensitive area, line by line. The full-frame and frame transfer devices can be though of as an array in which each column is composed of a CTD.

The light gathering areas of full-frame and frame transfer devices are identical, but at the output stage there is a slight difference.

For full-frame devices, each vertical CTD column is clocked downwards a pixel at a time to a horizontal CTD, at one end of which is the output charge to voltage amplifier. The readout speed of this type of device is therefore limited by the rate at which the output CTD can operate, typically no higher than 20MHz. Therefore, for a 1000x1000 pixel array, each line will require $50\mu s$ readout time. This will give a frame rate of 20 frames per second, however the array will be integrating light for the entire duration of the transfer.

For frame transfer devices, there is an additional area of storage. This storage area is identical in size to the photo sensitive area, but is opaque to the incident light (see figure 2.4). Because of the identical nature of the photo sensitive and storage areas, the CTD columns can be transfered into an identical CTD column without having to wait for each row to be read out. Therefore, the entire array can be transferred into storage in the same time as it takes a full-frame device to read out one row of data.

These systems have the advantage that because the entire area of the CTD is photosensitive, little of the incident illumination is lost, providing a sensitive

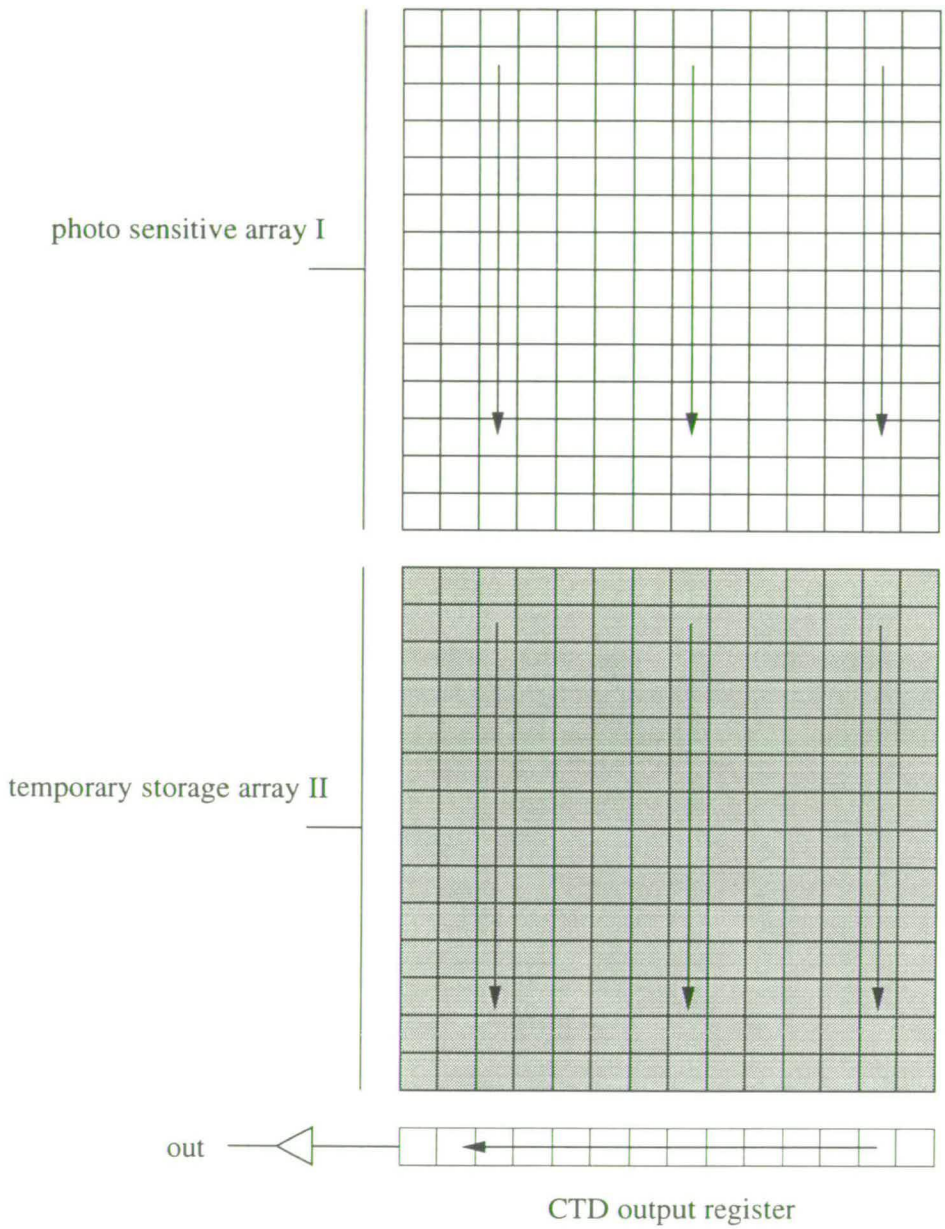

Figure 2.4: Frame transfer device

device. However, the method in which the data is transferred from the photo-sensitive area gives rise to slower transfer rates, leading to possible smearing of the image. In practice full-frame devices (for example the Kodak MegaPlus 4.2) often have a mechanical or electro-optic shutter placed in front of the array to try to reduce the amount of smear, but as most electro-optic shutters have a response time of $\mathcal{O}(40\mu s)$, some smearing is still possible. These cameras are an ideal replacement for film based systems, and are most frequently used as a direct replacement for film in i.e. professional photography.

This reliance on external shutters to stop incoming light from illuminating the CCD array of course makes these devices unsuitable for use in cross-correlation devices, as the time taken for either electro-optic or mechanical shutters to operate would restrict their use to only very slow flows.

However, one manufacturer of CCD based systems for PIV [68, 67] have taken advantage of this feature, and by altering the timing of the image transfer into the storage array has produced a camera which has built in image shifting without the associated distortions produced by the more traditional rotating mirror or panning camera systems.

2.2.3 Line Transfer

Line transfer (figure 2.5) devices again have photosensitive CTD's. However, once the image has been integrated, it is read out line by line at the clock rate into the output register, and then read out from the output register at the same rate to the output amplifier. Because this architecture employs photosensitive CTD's, line transfer devices should have the same benefits and pitfalls as frame transfer devices, namely high fill factors, and smearing of the image unless an external shutter is used.

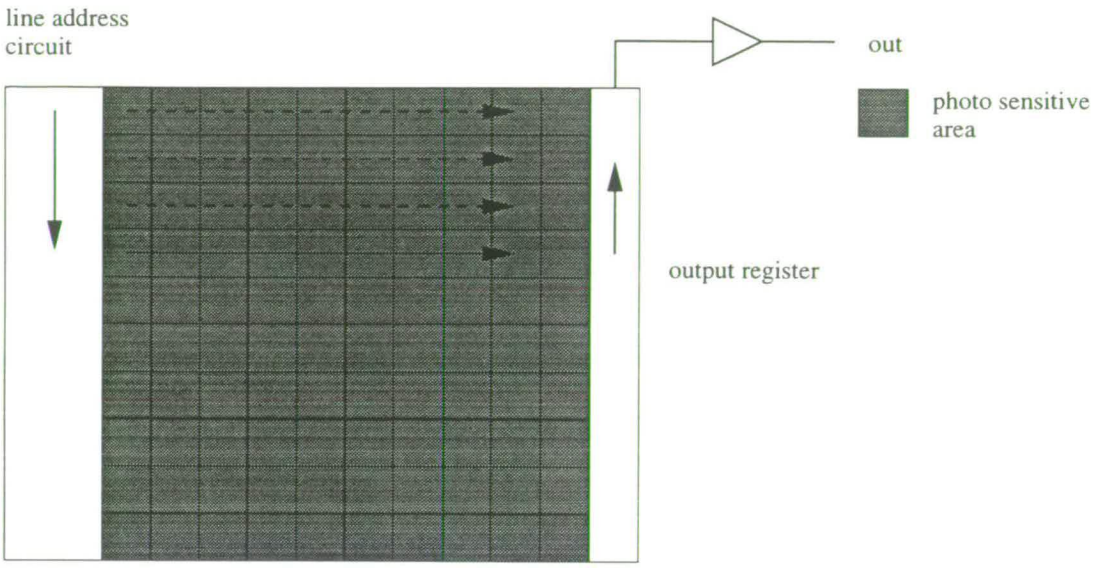

Figure 2.5: Line transfer device

2.3 Multiple Output CCD arrays

In order to increase the rate at which the image data can be read from a CCD array to a host computer, a number of CCD arrays now exist which have multiple output channels. The array is split into subsections, with each subsection having its own CTD output register and output charge to voltage converter (see figure 2.6). The subsections can then be read out simultaneously allowing a speed up in the data transfer rate of the captured image without increasing the frequency of the clock signal.

The division of the array can be in any format produced by the manufacturers. In the case of the Kodak ES1.0, the odd and even lines are read out to different output registers. However, this is not the only format in which an array can be produced, and some have been produced by Philips which have divisions similar to figure 2.6. Arrays with four subsections have also been produced, most notably by Thomson-CSF.

However, there is a drawback to this solution. Because each section has to have

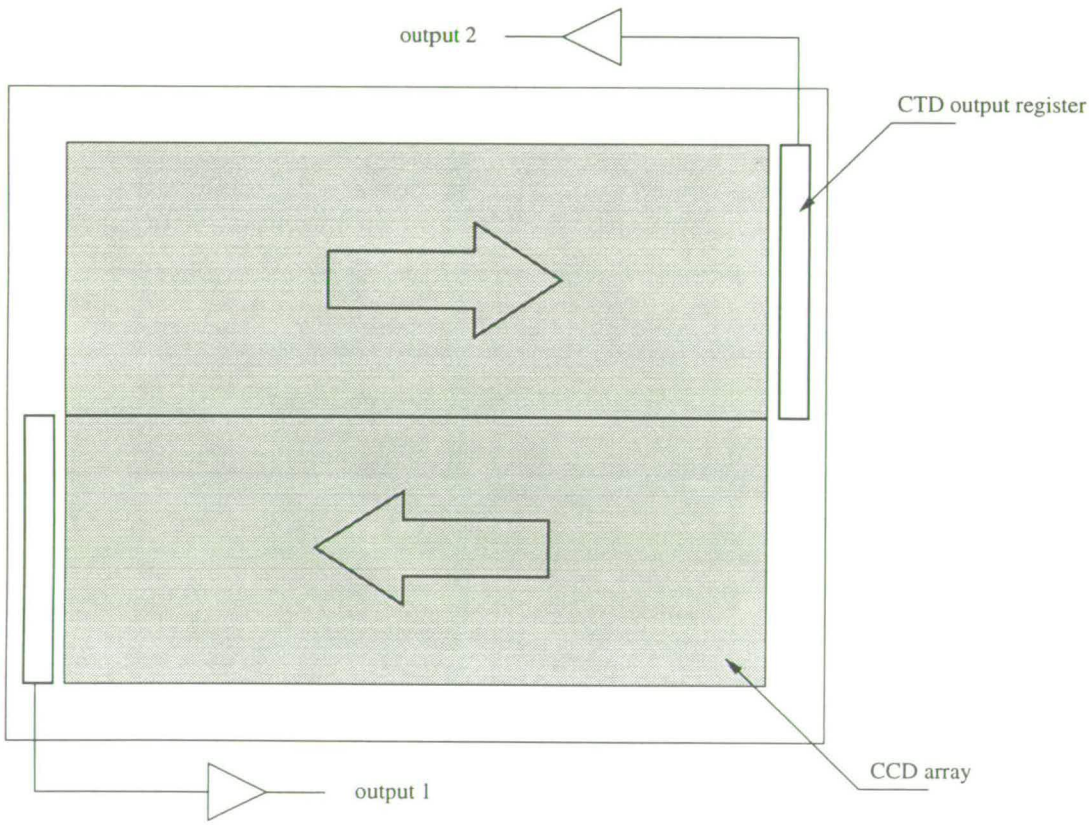

Figure 2.6: Architecture of a dual output CCD array

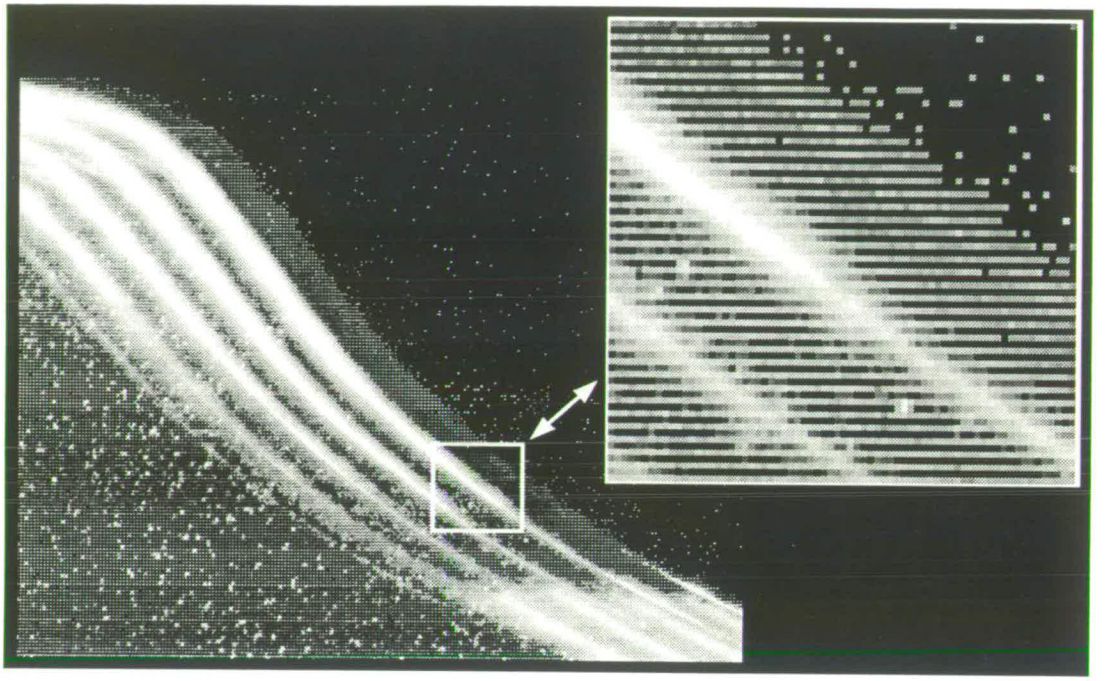

Figure 2.7: The effect of dual output CCD arrays

its own charge to voltage converter at the output of the section, the converters have to be very closely matched in order to allow a uniform output to be produced. Typically the amounts of charge present within each charge packet are small, and so any non-linearities within the output charge to voltage converter will appear in the output image as low level noise. An example of this can be seen in figure 2.7. This is an image that was taken under low illumination levels, resulting in only five bits of information, or 64 discrete grey levels, being registered by the digitiser.

The effect produced by the mismatched output converters produces an effect similar to interlacing (section 2.5), that is if the image were analysed, the correlation plane would contain strong periodic components present due to the alternating illumination levels. It can be seen from the figure that interlacing is not to blame for this effect as the free surface (of which there are multiple expo-

sures due to the experiment being run with illumination provided by a scanning beam illumination system) is complete. If these effects are allowed to persist in the image, then a high frequency component is added to the correlation which can be seen as very prominent noise. This effect can be removed by application of a smoothing filter in the spatial frequency plane, as will be described in Section 4.4.5.

However, the ES1.0 cameras which produced this result also have the ability to have the gain and offset of each output amplifier set and then stored to an EEPROM [1], allowing the output to be matched. As the EEPROM will continuously store these settings, the calibration should have to be done only once to remove this effect.

2.4 CMOS Devices

Both CCD and CMOS (Complimentary Metal Oxide Semiconductor) devices are made from silicon, and both detect light in exactly the same way. They can both have either a photo-gate (similar to a transistor in form) or a photo-diode to convert light into charge, and because both devices are made from the same material, they have the same or very similar spectral responses and quantum efficiency.

However, while CCD devices use a complex and closely packed architecture to allow the transfer of charge packets with clocking signals, CMOS devices have for each column of photo cells an associated output amplifier (passive pixel) or an amplifier for each individual photo cell (active pixel). Both technologies have their advantages and disadvantages, but are both available technologies.

[1] An EEPROM (Electronically Erasable/Programmable Read Only Memory) is a device which can store information indefinitely, is non-volatile (that is, retains the information when no power is present) and can be re-programmed by altering the voltage levels present at the IC.

The main advantage of CMOS technology over CCD technology is that because the entire photo sensitive area is made with CMOS technology, all of the associated clocking circuits, D/A converters, and any additional logic required can be incorporated onto a single integrated circuit, allowing a complete camera to be manufactured on one chip.

CCD technology on the other hand requires very precise clocking signals with complex forms (see figure 2.2), and this has the disadvantage that CCD arrays require typically a specialised clocking chip and also a number of non-standard voltages to be present. This leads to added complexity in the construction of the CCD camera and high power consumption. The design of CCD sensor boards (the board which directly interfaces to the CCD chip) is considered a black art. The absence of a requirement for complex clocking signals also allows CMOS devices to be read out to a computer more quickly, and the charge packet corruption that can occur in CTD's due to high frequency clocking signals is absent.

CMOS technology does have its own drawbacks, including low signal to noise ratios if passive pixel technology is used and low photo sensitive area if active pixel technology is used. Active pixel technology also has to overcome the non-uniform amplification of each of the amplifiers associated with each pixel.

Overall, CMOS technology is cheaper to produce, more flexible and consumes less power in comparison to CCD technology. However, for the range of digital cameras available on the market today, CMOS is still considered to be less sensitive than CCD, and CCD based chips have a superior image quality produced by over 25 years of development.

Because of the flexibility of CMOS, it is possible to make arrays which have an individual output amplifier for each pixel. The chip is then not organsied into rows and column, but rather is arranged as a memory chip, with the assertion of

address lines giving the value of the addressed pixel. This allows a great deal of flexibility in the output of the image. This allows the image to be sub sampled, translated, have different exposure times for different areas of the array, etc. with all of these effects being realised during the readout of the image. The lack of fixed geometry in the readout also means that the interfacing to a computer can be considerably simpler, because the signal does not have to conform to one of the video standards. This does however have its own drawbacks if the output of the array needs to be checked independently of the interfacing computer.

The photo-diodes and associated output circuitry can also be designed to have a logarithmic output. This can have the advantage that the sensor will have a dynamic range several orders of magnitude larger than a conventional CCD array, however most CCD arrays are chosen for their linear response to light, and the ease with which this can be checked.

2.5 Interlaced CCD arrays

Although CCD arrays have been used in astronomy since the sixties, the majority of CCD arrays have been designed to replace Vidicon-type [2] imaging tubes. The majority of these tubes and their driving electronics were designed to be used in television broadcast, and so their output formats are fixed to one of the video standards.

A conventional TV image or *frame* is *interlaced*, that is to say the image is split into odd- and even-lines, the odd- and even- *fields*. Each frame is read out either 25 (CCIR/PAL) or 30 (EIA/RS-170) times a second, and the arrays will typically have resolutions of 768x494 (CCIR/PAL) or 640x480 (EIA/RS-170).

[2]Vidicon-type devices are similar in appearance to a cathode ray tube. The image of interest is focused onto a surface which emits electrons due to the photo-electric effect. The magnitude of this emission is then converted into an output by scanning of the reverse side of the image "target" by an electron beam

Because of the demand for devices which output their image in these formats (i.e. for use in camcorder type video cameras), the majority of available CCD cameras have an architecture which limits their output to one of the common video standards.

For the CCD architectures shown above, interlacing is achieved by simply reading out either the odd- or even-fields in turn, and in practice most of the CCD arrays commercially available were until recently fixed by hardware to produce an interlaced output image.

The most commonly available CCD arrays are of the interline transfer variety, but do not have a storage site for each pixel. Instead, each adjacent odd-line pixel and even-line pixel share a storage site/CTD. Thus, to allow the device to operate at full resolution, either odd or even lines have to be read out in turn.

Interlacing has the undesired effect that while one field is being read out, the other is being exposed, and so for a fast moving target, the image will appear altered, as in figure 2.8 [3] where it can be clearly seen that between the odd and even fields being exposed, the free surface has moved giving rise to this very distinctive effect. The free surface can still be recognised, but if this image was an area of randomly distributed seeding material, then the particle images would not be round anymore. The particle images would consist instead of highly distorted images which would give rise to strong correlation effects giving an erroneous velocity measurement.

In some cases the flow is slow enough ($\sim \mathcal{O}(cms^{-1})$) to allow an interlaced camera to be used, but in most it is not, and so one field has to be discarded to allow the velocities to be extracted. In some cases this can be overcome by careful timing of the illumination pulses of the flow in a technique known as

[3]Thanks to Tim Barnes for this image

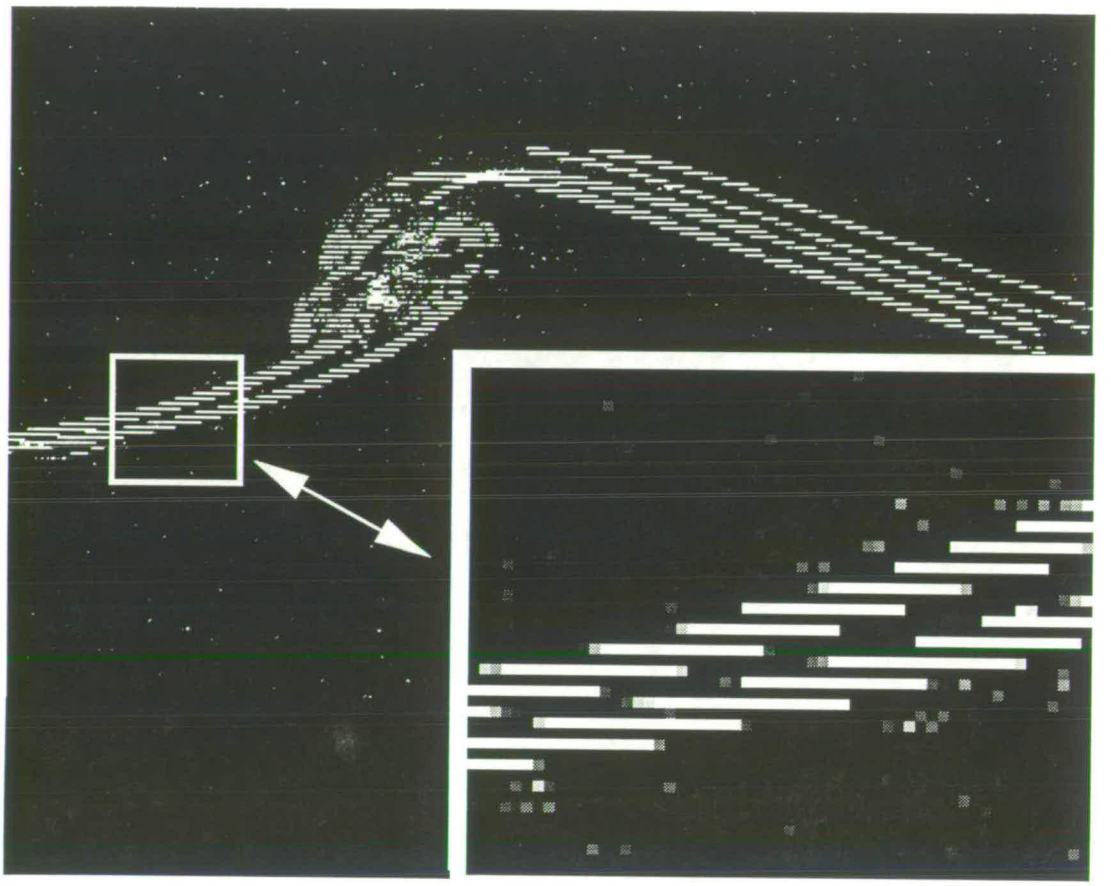

Figure 2.8: Plunging breaker captured with interlaced CCD camera

frame straddling (see section 1.2.4), allowing each illumination pulse to fall just on either side of the inter frame blanking period.

However, either of these methods result in at least a 50% reduction in vertical resolution as the optical setup has to be configured to allow a particle image to cover at least two field lines. The effect will then be similar in the vertical direction to the effect of a low fill-factor array.

2.6 Microlenses

In all CCD arrays there is going to be a considerable percentage of the illuminated area taken up by non-photosensitive material in the form of storage, CTD's,

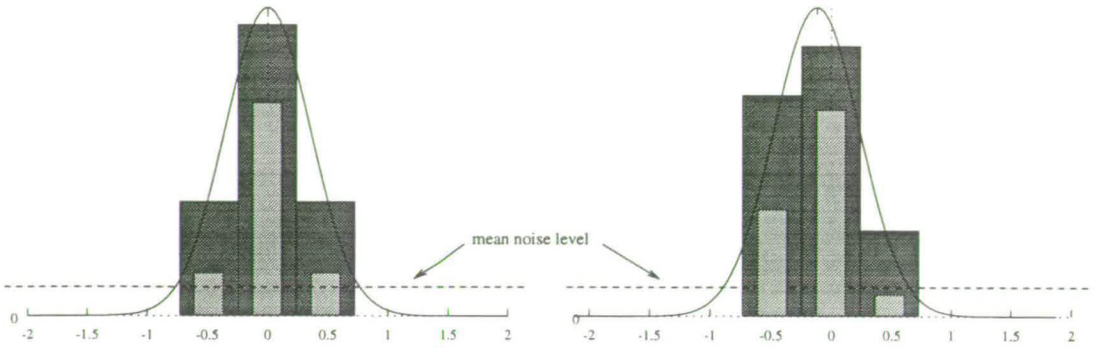

Figure 2.9: The effect of fill factors on information gained and light sensitivity - Full curves represent Gaussian particle image intensity profile, dark rectangles represent information gathered by full fill factor pixels, light rectangles represent non-full fill factor pixels.

control circuitry, etc.. The ratio of the photosensitive to non-photosensitive material is in general referred to as the *fill factor* of the array. If an array had a fill factor of 100% then no light would be lost, but in general CCD arrays have fill factors which are less than 100%. This has two major effects which are of direct consequence to digital PIV:

- lower light sensitivity
- lost information

The first point has been discussed already, and it would appear that interline transfer devices are possibly the worst in terms of low fill factor. The second point however applies to the situation when imaging a flow in which we are operating at the lowest possible resolution (i.e. a particle image covers two pixels in each direction [62, 64, 39]).

The fill factor of the array is low, and so the available light is lost between the pixels of the array. As the particle images are small, the amount of light that is available to the pixels adjacent to the pixel which contains the particle image centre will also be very low. If the particle image does now not lie directly

centred on a pixel, the probability of losing the information below the level of the background noise increases. If this situation happens, then the net result is in effect the same as having particle images which are too small, and so the accuracy of the sub-pixel displacement becomes questionable. This will result in displacements becoming biased towards integer values.

These drawbacks can be remedied to some extent by the use of microlenses. As the name suggests, these are small lenses placed on top of every pixel in the CCD array which act more as light guides than lenses, allowing the light that would otherwise be lost to be collected by the photosensitive area on the CCD array. This has the effect of increasing the fill factor of the CCD array, and is in general considered to be a good thing. However, it is often assumed that if a particular model of CCD array has microlenses, it will have a fill factor of 100% and this is more often than not not the case. For example, the cameras used in this thesis work are Pulnix TM-9701's, which have a fill factor of 60% with microlenses and 20% without.

2.7 Frame rate

Conventional interlaced CCD devices typically have a frame rate of either 25 or 30 frames/second, and for some low speed flows, this is fast enough to allow a time history of the flow to be obtained. However, if we want to move to a CCD based system which has a higher spatial resolution, we have to make a reduction in the frame rate. CCD devices in general are limited to clocking frequencies of up to 20MHz, due to the complex signals required to transfer the charge packets. If the frequency of the clocking signal is increased, the signal will be readout more quickly, but the noise levels in the image will increase. High spatial resolution, non-interlaced devices are therefore limited to very low

output rate (typically, output rates for a 2,000x2,000 pixel array would be $\mathcal{O}(5)$ frames/second maximum). While this is fine if we are using the high resolution CCD camera as a direct replacement for a film based camera, it is of little use in the rapid acquisition of a sequence of images of the flow under investigation.

On the other side of the data read out limit lie the low spatial resolution arrays which have very high read out rates. These typically have 512x512 pixels or (more commonly) less , while they can achieve frame rates of between 500 and 4500 frames per second. These arrays are an offshoot of military developed technology, and as such are still not widely available market items, with the price of the hardware, certainly not less than 50,000 USD, reflecting this availability. These arrays do have their exponents however, most notably Brücker [6, 4, 5] whose work in the field of scanning plane DPIV has seen tremendous success.

2.8 Output Formats

The format in which the video signal is output is also an issue when making assumptions about sub-pixel accuracy in PIV data.

The majority of commercially available CCD cameras output their data in analogue format i.e. the data is divided into lines, and each line is output along with the necessary timing information. This type of video signal is commonly known as *composite video*, and a typical line and field sequence can be seen in figure 2.10. When the data is output in this format, the signal can be recorded onto video tape or digitised by a frame grabbing board and computer. However, *jitter* can be present in the digitisation, generally introduced when the frame grabber board is trying to locate one of the many syncronising pulses present in the composite video output. This can lead to a slight alteration of the pixel values, due to the output voltage being sampled at a slightly incorrect time.

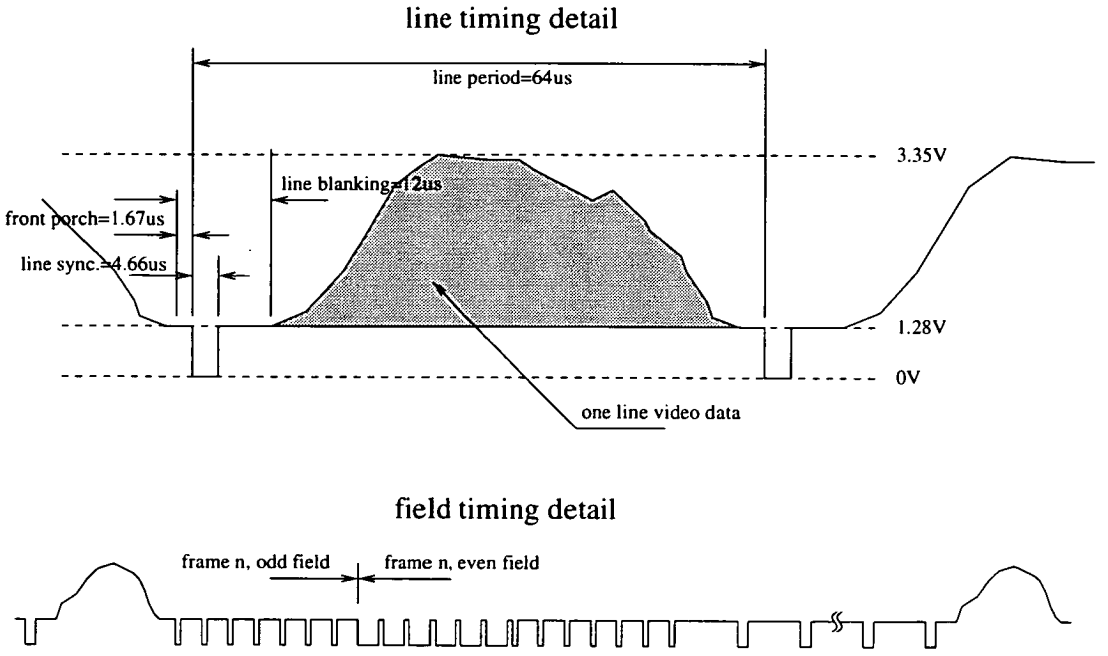

Figure 2.10: Typical line and field timing diagrams for a standard composite video output

If this happens, then the accuracy of the sub pixel measurement can be affected, and this is a possible source of uncertainty in the PIV measurement. However, it is becoming more common for scientific CCD cameras to have a digital output either alongside or instead of an analogue output.

The standard digital interconnect between camera and computer for such high speed signals is the RS-422 standard, consisting of an 8-bit wide parallel data path, with each bit being transfered as a balanced differential signal along a twisted pair. This allows circuitry at the receiving end to remove any noise which has accumulated along the cable path, and in general allows a very good data throughput. Extra information is also output along side the data, most commonly in the form of a pixel clock and a frame valid signal. These enable the removal of the jitter common with composite video signals.

The maximum data rate along an 8-bit wide RS-422 link should be 80MBytes/second

for a cable run of at most 10 feet, but to ensure all data is received in practice the data rate is much slower. In order to accomidate the output from multiple channel CCD arrays, the data path has to be widened to allow the extra information to be transferred in parallel. While this is physically possible, the grouping together of bundles of eight twisted pairs produces a very large, expensive and inflexible umbilicle between the CCD camera and host computer.

Advancements in available technology are altering this situation. The data can be transferred from camera to computer either by a high speed serial link over wire or fiber optic. A camera, almost uniquely suited to digital PIV is made by PCO, and comprises a high resolution cooled CCD array (over 1200x1000 pixels, true 12 bit). The data is transferred from camera to computer by a fiber optic link to a dedicated frame grabber card. Due to the clocking limitations the frame rate is still low, but the electronics have been designed to allow full control of the frame transfer event and shuttering. This allows the acquisition of two full frame images with virtually any interframe separation.

The wire based serial interface standard, commonly known as FireWire or IEEE-1394, is a serial link which operates at 100Mbits/s with future incarnations running at 200Mbits/s, 400Mbits/s and 1.2Gbits/s.

> The IEEE-1394 High Performance Serial Bus is a versatile, high-speed, and low-cost method of interconnecting a variety of personal computer peripherals and consumer electronics devices. The IEEE-1394 bus began life in 1986 as Apple Computer's alternative to the tangle of cables required to connect printers, modems, external fixed-disk drives, scanners, and other peripherals to PCs. The proposed standard (P1394) derived from Apple's original FireWire design, was accepted as an industry standard at the December 12, 1995 meeting

> of the Institute of Electrical and Electronics Engineers (IEEE) Standards Board.
>
> Isochronous data transmission lets even the lowest-speed implementation support two simultaneous channels of full-motion (30-frame-per-second),"broadcast quality" video and CD-grade stereo audio.
>
> *Taken from http://www.dtvgroup.com*

This interface standard also removes the need to use frame grabbers, as all information sent from the camera will be in the standard FireWire format, allowing the easy integration of a number of cameras which can be asynchronous to each other. Again, this is a standard that was developed for multi-media computing, but which will also advance and simplify DPIV experiments.

2.9 Spectral Response

As previously mentioned in this chapter, the charge that is produced in the photo cell of a CCD is due to a band gap transition at the junction of two dissimilar types of semiconductor (i.e. n-type and p-type doped). Because this band gap transition is, in the majority of cases, the result of an incident photon, the output of the device will depend on the wavelength of the incident photon. The spectral response of the device depends therefore on the material used, the geometry and the thickness of the photo cells.

This response is measured as a quantum efficiency, i.e. the ratio of the number of collected electrons and the number of incident photons per pixel. A number of spectral response curves are shown in figures 2.11 and 2.12.

The spectral response curves can be thought of in two ways. Firstly, if the curves of different devices were normalised, the relative sensitivity to the wave-

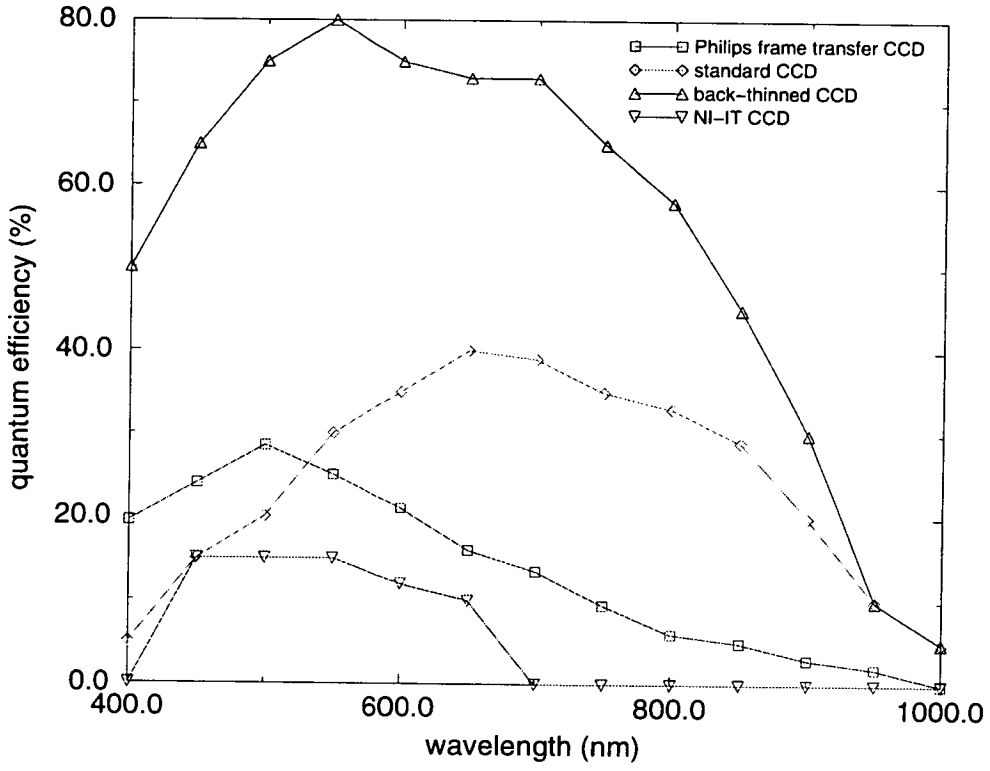

Figure 2.11: Spectral response curves for a number of different CCD architectures (NI-IT = non-interlaced (full frame) interline transfer)

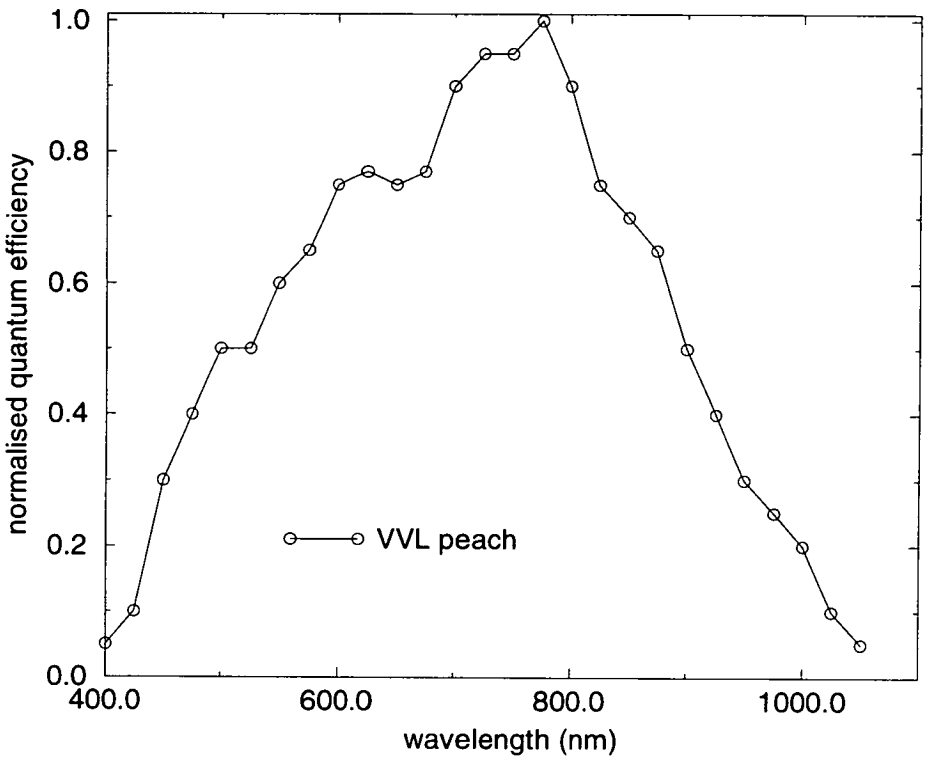

Figure 2.12: Spectral response curve for VVL peach CMOS array

length of illumination can be easily compared. Secondly, if the maximum quantum efficiencies or amplitudes of the normalised spectral response functions are taken the overall sensitivities can be compared without a dependence on wavelength. An example of this would be in the data sheets for the VVL peach camera. The normalised spectral response curve is shown in figure 2.12, and the maximum quantum efficiency is quoted as being 2 lux minimum illumination with a lens of $f1.2$.

The graphs clearly shows the difference architecture has on the sensitivity of the CCD array. The specially back-thinned CCD arrays have fill factors of anything up to 100%, and the smaller amount of material the incident photon has to travel through results in a much greater efficiency. Their spectral response is also shifted towards the blue-green end of the spectrum. This is mainly due to the fact that these arrays are illuminated from the back side, and therefore the incident photons do not have to pass through a layer of metal and oxide before reaching the semiconductor junction. The standard CCD and the VVL CMOS based sensors have similar spectral responses, but different sensitivities, due perhaps to a differing fill factor. Only a relative spectral response curve was available from VVL, however the minimum illumination level for the VVL array is 2 lux at $f1.2$, while for a Pulnix TM-9700 with 20% fill factor the minimum illumination level is 1 lux at $f1.4$. More interestingly, however, the full-frame interline transfer devices perform badly compared to the rest of the arrays. This is undoubtedly due to the low fill factors present in these arrays due to the large amount of storage required.

The other point of interest is to note the wavelengths at which the most commonly used PIV lasers operate. The argon-ion CW laser has a number of emission frequencies, between 488-514nm. The neodymium:yttrium-aluminium-

garnet (Nd:YAG) pulsed laser produces laser light at 1064nm, but is almost always frequency doubled by non-linear KDP crystals to 532nm. Both of these lasers produce the majority of their light in a region which is not the spectrally most sensitive in CCD's, CCD's being more sensitive to red light ($\approx$ 800nm). This is the reverse of photographic film, which is more sensitive to green light.

2.10 Noise

Noise is a problem for all electronic devices. This is due mainly to thermal effects, although in CCD arrays noise can also be generated by having a clock frequency which is too high. The generation of noise can be a problem in PIV due to the low incident light levels. If the scattered light is too low, or the noise level high enough, information is lost, and the accuracy of the PIV recording is reduced.

The thermal noise, commonly referred to as *dark current* produced by a CCD array is constant for a given temperature and exposure, and so can be removed at the output by removing a suitable bias from the output voltage. The generation of this thermal noise is, however, a random event, and so should be described by Poisson statistics with the number of events μ being equal to the number of incidents per unit time Λ multiplied by the time interval T. Therefore, the mean noise level will have a fluctuation, the standard deviation of which will be equal to the square root of the mean value.[4]

The rate of generation of the thermal noise doubles for every increase in

[4]For Poisson statistics:

$$\mu = \Lambda T \tag{2.1}$$

$$P(X = r) = e^{-\mu}\frac{\mu^r}{r!} \tag{2.2}$$

$$E(X) = \mu \tag{2.3}$$

$$V(X) = \mu \tag{2.4}$$

$$\sigma = \sqrt{\mu} \tag{2.5}$$

ambient temperature of $6 - 7°C$, and so cooling the CCD array can help to remove or reduce this effect.

Noise can also be generated by having a clock frequency which is too high. The rapid changing of the multi-level potential gates in a CTD introduces noise into the charge packets being transferred. The only solution for this is to reduce the clock rate, and this is reflected in the maximum read-out rates of the majority of CCD arrays being 20MHz.

Chapter 3

Construction of Multi-CCD Array Systems

3.1 Preliminary Considerations

The main requirements of this work were to build a camera which could be used with any type of flow, fast or slow, and which was capable of storing the sequential exposures of the flow in physically separate locations. The system would also have to be used with different types of illumination system i.e. scanning beam, pulsed laser or flash lamp.

The majority of CCD arrays which were available at the beginning of this project were not flexible enough to allow this design brief to be met. Most of the CCD cameras were aimed at either the security market (standard video resolutions) or the "scientific" market (high resolution, but long exposure times i.e. microscopy, astronomy). Therefore, a system had to be designed and built which would overcome these limitations.

Ideally, the CCD arrays should be asynchronously re-settable in order to allow the time at which the image is captured to be set by the experiment, and not vice versa. Unless the arrays have special features, then multiple arrays will be required to allow the separate storage of each exposure. The information from the CCD arrays also needs to be downloaded and stored in real time (that is, as

fast as the camera can produce data).

In the first instance, the idea of using a frame-transfer device seemed attractive, as the first image of the flow could be transferred to the temporary on-chip storage area while the array was being exposed again. However, the predominant illumination scheme at the University of Edinburgh is the scanning beam system. If a pulsed laser was used instead, the frame-transfer device would be ideal. The pulses could be triggered to fall either side of the transfer period. However, this would be syncronising the experiment to the camera, which we have already defined as being an undesirable characteristic. The synchronisation of the frame-transfer device to the scanning beam system is a considerably more complicated matter. To realise this synchronisation would certainly involve the modification of the driving electronics of the CCD camera. To do this, a knowledge of the electronics in the camera are required, but in the majority of cases camera manufacturers are not keen to release the operational details of their products. The modification could be made in conjunction with the company, but the results would be uncertain and the cost involved high. This option was therefore ruled out on the grounds of complexity, time involved and cost.

With these points in mind, the use of readily available, domestic video oriented arrays became the simplest option for the construction of a multiple CCD array system. The main construction requirements of this system were:

- Cost
- Simplicity of electronics required
- Simplicity of programming for acquisition of images
- Good spatial resolution
- Small overall size

- To have same optical path length as ordinary 35mm SLR camera

The desire to have the optical path length closely matched to that of a standard 35mm SLR camera is to allow the use of readily available, high quality lenses. The optics in many 35mm format lenses are of extremely good quality, and the economics of mass production make even a flat image plane, $f1.4$ 55mm lens cheap (£400) compared to, for example, a Hassleblad lens. The 35mm format also has a reasonable amount of room between the lens mounting and film plane ($\approx 42mm$), allowing plenty of room for optical components.

3.2 Prototype System

The initial ideas for the construction of the system, and more importantly how to overcome the problems of shuttering without the complexity of modification of the video camera circuitry, came from work carried out by Racca *et. al.* [42]. Their setup involved using a combination of beam splitting optics and Ferroelectric Liquid Crystal (FLC) cell "light valves" to expose each of the arrays in turn. The image captured was then stored locally in enough RAM to hold one frame of video data. This work was in turn inspired by the work of Bretthauer *et. al.* [3] which used eight CCD arrays arranged around an eight sided mirrored pyramid to allow the arrays to share the same view. This work was subsequently developed into a system capable of imaging up to a million images per second [46]. This camera, the Ultra High Speed Video camera UHSV-288, has been applied to very fast flows of the order of 0.3Ma, or between $70 - 180ms^{-1}$.

The next question was how exactly to configure the beam splitters, FLC's and linear polarisers so as to minimise the construction difficulties, optical distortions introduced and losses due to many optical components. The possible arrangements are shown in figure 3.1.

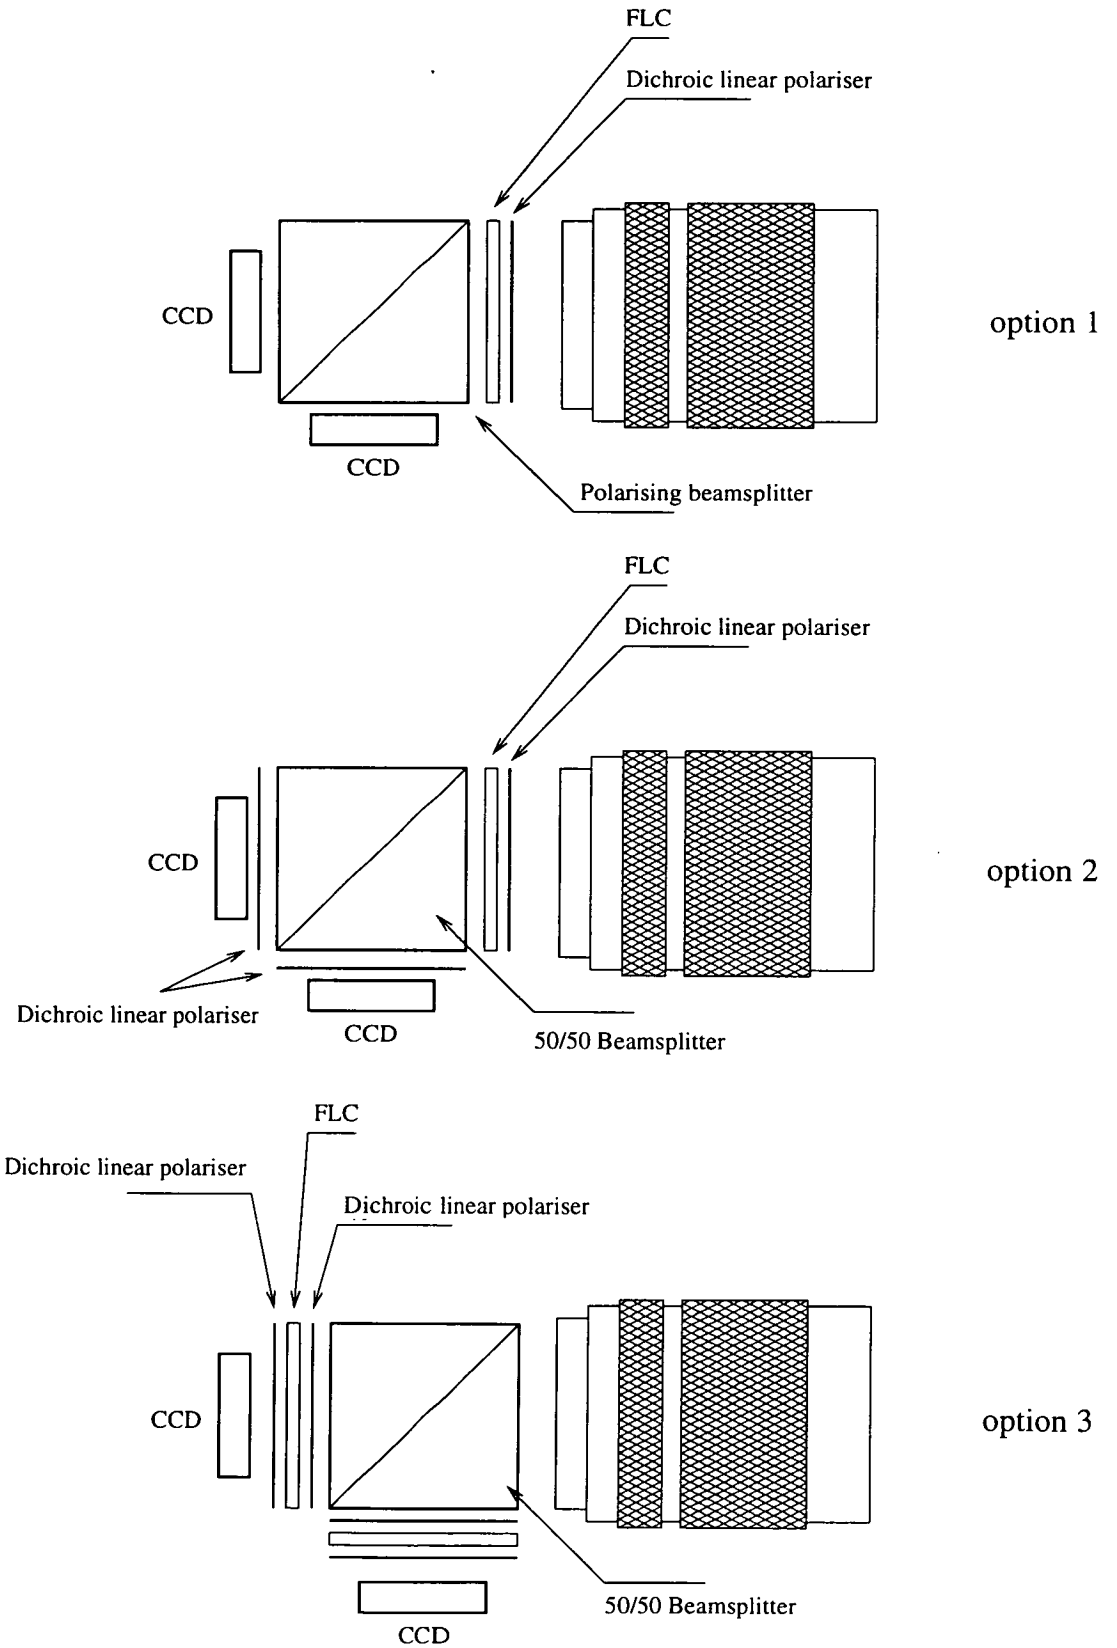

Figure 3.1: Possible arrangements of optical components

Optical Components	
	transmission
dichroic linear polariser	(randomly polarised) 40-45% (linearly polarised) 95%
unpolarised beam splitter	45%
polarised beam splitter	(for lin. polarised light) 95%
FLC	90%

Table 3.1: Transmission of optical components in multi-CCD array system

The losses in such a system can be estimated from the properties of the individual components. Approximate values for these components are shown in table 3.1. For option 1, the worst case transmission should be around 34%, as once the incident light has been linearly polarised, the FLC and polarising beam splitter will direct the light to the appropriate array with minimal losses. However, for options 2 and 3 the situation is worse. In both cases, the final transmission is about 15.5%. With such a large amount of incident light lost, and also a greater number of optical components (four as opposed to three), option 1 is the only really feasible solution. Option 3 does give the most flexibility with respect to timing of the array exposures, but the magnitude of the losses prohibit this arrangement.

3.2.1 CCD Arrays

As this first system was going to use FLC's to provide the shuttering, the electronic shuttering on the CCD arrays was of no interest. Moreover, if the cameras could be synchronised, then the images from both cameras could be captured by a colour frame grabber at the same time. This is possible as colour *video* frame grabbers have a separate channel for each of the primary colours, red, green and blue (RGB). The frame grabber assumes that as the three channels are coming from the same image, then the signals will be synchronous to each other. Thus,

array size	312(H) x 287(V)
pixel size	19.6μm(H) x 16.0μm(V)
array type	addressed CMOS
fill factor	unknown
minimum illumination	2 lux @ $f1.2$
other	external synchronously possible
	low power consumption ($< 150mW$)
	standard CCIR/PAL output signal
	all driving logic in single chip

Table 3.2: Specification of the VVL Peach

if we synchronise two independent cameras and feed their outputs into the frame grabber, the computer will be able to store both images at the same time. This has the obvious benefit that to store the images from up to three cameras, we only require one frame grabbing board.

The arrays chosen were VVL Peach cameras [1] (figure 3.2, table 3.2), which are designed for the security and video conferencing markets, as they are cheap and small. These cameras comprise a complete digital camera on a single integrated circuit, with the external circuitry being provided on a round printed circuit board 3cm in diameter.

The sensor on the VVL Peach cameras are of CMOS construction, and their output is a standard CCIR video signal. There is access to many of the shuttering and control features, in particular the auto-gain control (AGC) can be suppressed and the auto-exposure, which can be forced to be the entire field length. Both of these features are essential if the external shuttering is to work correctly. There is also access to the pixel clock and reset functions, which allows the arrays to be synchronised together. Lastly, these cameras were chosen for prototyping as they are extremely cheap. The cost per camera is less than £50 [2].

[1] VVL have since changed their name to Vision plc.

[2] These cameras have also now found success as the active component in a Fisher-Price child's video camera. Due to mass production, they can be bought in bulk in the US for $10 each

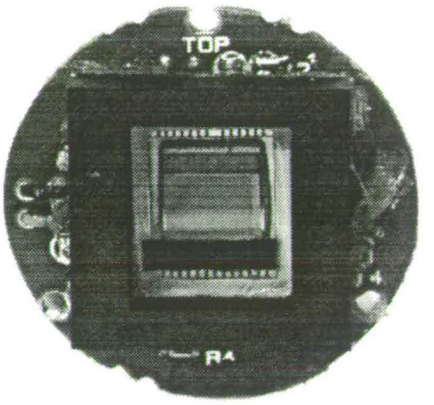

Figure 3.2: VVL Peach CCD array

These arrays have a low pixel count (312x287 pixels), but also have a special, non-documented feature. Although their output appears to be 25 frames/second, the array is not split into odd and even fields, but rather the whole array is actually running as a non-interlaced device outputting data at 50 frames/second. The resulting spatial resolution (in terms of vectors produced, 50% overlap of interrogation areas) would give 17x16 vectors at 32x32 interrogation area, and 37x33 with a 16x16 interrogation area. This is not large enough to do a great deal with, but for prototyping purposes it was adequate.

With a combination of careful timing and frame straddling the capture of four individual, but full resolution images of the flow could be achieved. The schematic of the system is shown in fig. 3.3, and a picture of the assembled prototype unit is shown in fig. 3.4, showing clearly the synchronising electronics, CCD arrays, polarising beam splitting cube and incident polariser and FLC cell.

The FLC cells used were supplied by James Gourlay [14] of the Applied Optics group within the Physics Department of the University of Edinburgh, and so high quality devices were readily available for experimentation. The FLC's have a switching time of typically $15 - 40\mu s$ although the ultimate speed of switching is

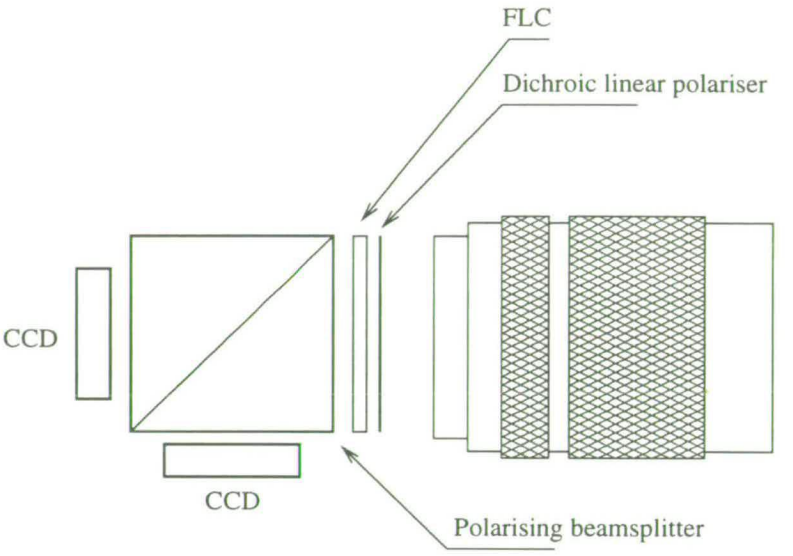

Figure 3.3: Prototype two-CCD array camera schematic

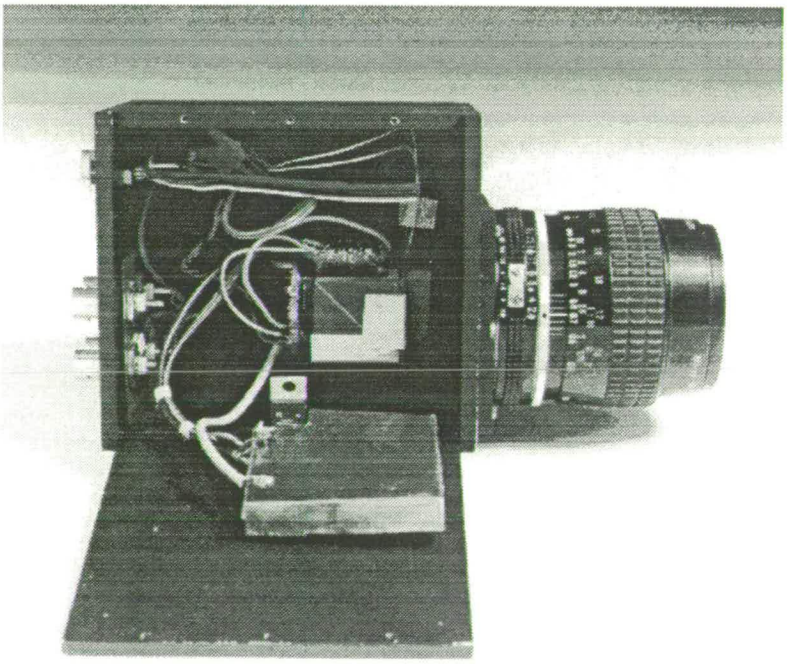

Figure 3.4: Completed prototype two-CCD array camera

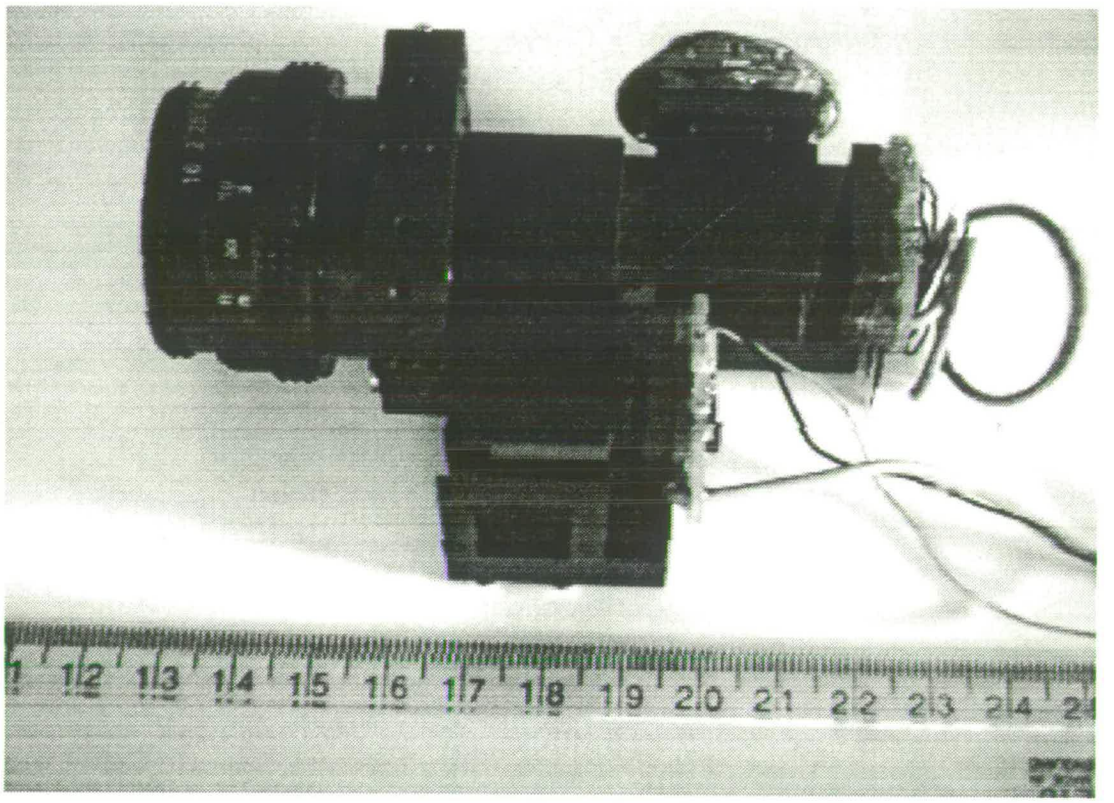

Figure 3.5: CCD arrays and beam splitting cubes on mounting

dependent upon the thickness of the liquid crystal layer. The thinner the layer the faster the switching time, but the degree of rotation of polarisation and the uniformity of the effect are both reduced drastically. The fastest devices available both in house and commercially are around the $40\mu s$ mark, which allowed a high contrast ratio ($> 100 : 1$) with a uniform coverage over the area of the array.

3.2.2 Construction

To mount all of the optical components together, it was decided to place them on an aluminium base plate, which could then be attached to the camera housing (see Appendix B, figure B.1 for the plans). The decision to only mount the components on one side also allowed ease of access for placement of linear polarisers and FLC's.

The CCD arrays were mounted by making a non-reflective aluminium collar

which fitted around the CCD chip. This collar could then be bolted to the base plate, ensuring that the arrays were in fairly good alignment to start with. This collar can be seen in figure 3.2. The collar and camera needed to be joined together in a rigid way, without the bonding material affecting either the CCD chip of printed circuit board. The solution settled upon finally was non-sumping silicone sealant, providing a secure bond which could also be removed if the need arose. The base plate, beam splitting cubes and three of the possible four cameras are shown in figure 3.5, along with a ruler and c-mount lens for scale.

3.2.3 Limitations

Although this prototype system does work it has two major drawbacks, spatial resolution and light efficiency.

The spatial resolution is low, but is not critical for a prototype system. The use of higher resolution arrays would be possible if the optical shuttering proved successful.

Light efficiency is the result of a number of factors. Firstly, the CCD arrays chosen were based on CMOS technology, allowing all of the controlling logic circuitry to be placed on the same single chip. While this is good from a manufacturing point of view it does lead to a lower quantum efficiency, in this case, resulting in the array being less sensitive than a regular CCD array. Future CMOS devices will undoubtedly continue to improve in sensitivity until they match conventional CCD technology.

Secondly, to use the system as intended the light has first to be linearly polarised. In most of the experimental setups at the University of Edinburgh it cannot be assumed that the scattered light is anything other than randomly polarised, and so a dichroic polariser or similar has to be inserted immediately behind the camera lens. This immediately leads to a reduction in scattered light

intensity of greater than 50%. The light that finally arrives at the arrays is $< 40\%$ of the original incident illumination.

While this transmitted intensity appears to be relatively high, the low sensitivity of the arrays used meant that the image of the flow became noise dominated. The use of more sensitive arrays would have been possible, but at this stage the availability of more flexible cameras increased. The flexibility of these cameras allowed the removal of many of the optics involved in the construction of the camera, and so reduced the light lost and possible distortions introduced by them.

3.3 Second Attempt

3.3.1 CCD Arrays

After the development of the first system, a number of digital cameras with some very useful features became available on the market. These cameras had all of the features required, namely they were asynchronously re-settable, were full frame inter-line transfer devices, had on board frame stores and were almost completely user controllable.

The devices chosen are completely asynchronously re-settable to an uncertainty of $5\mu s$, although typically the timing of the laser pulses is known to a higher accuracy than this. If the system is to be used in high speed air flows (i.e. requiring a pulse separation of $< 5\mu s$) , then this issue would have to be addressed more carefully, but some preliminary experiments carried out showed no great difficulties when used in conjunction with pulsed laser systems.

Two devices were constructed, with both systems based on the Pulnix TM-9701 CCD array cameras (figures 3.6 & 3.7). The specification of these cameras can be found in table 3.3.

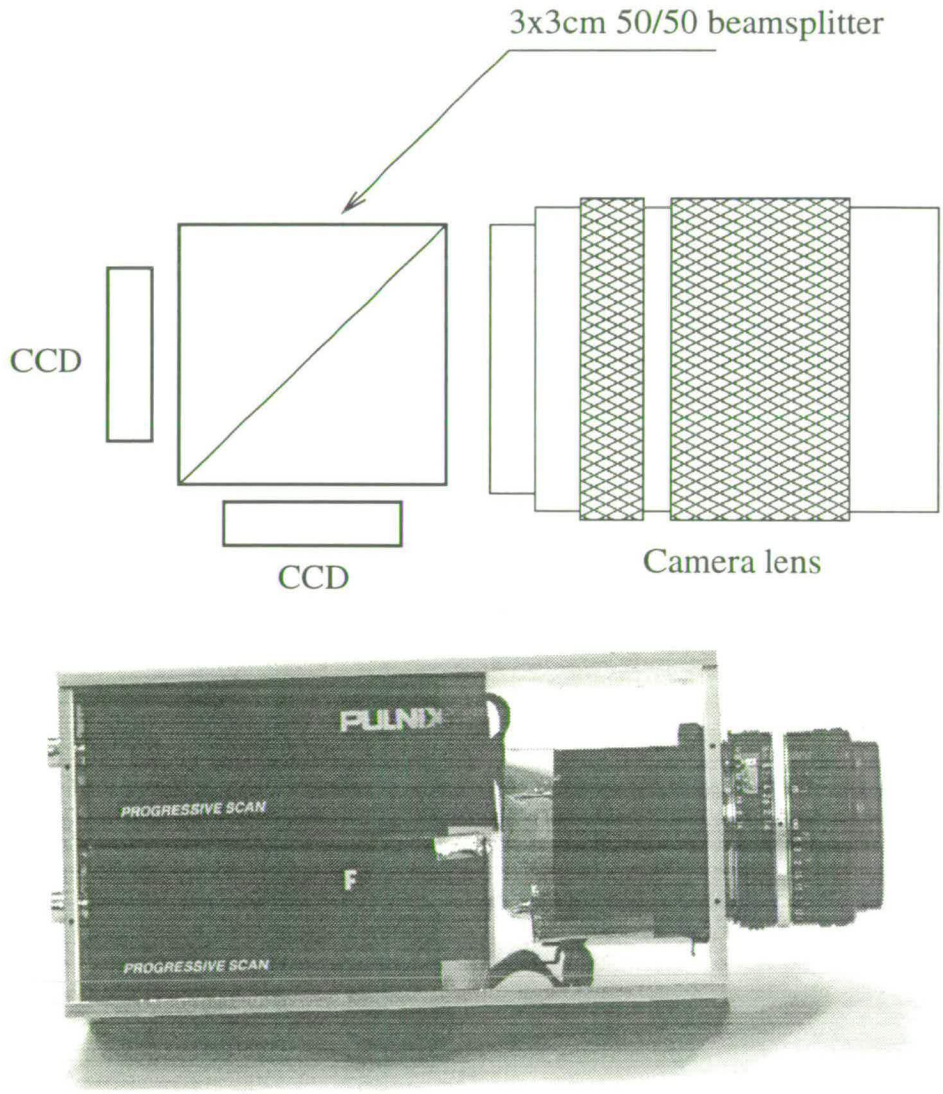

Figure 3.6: 2-CCD array system

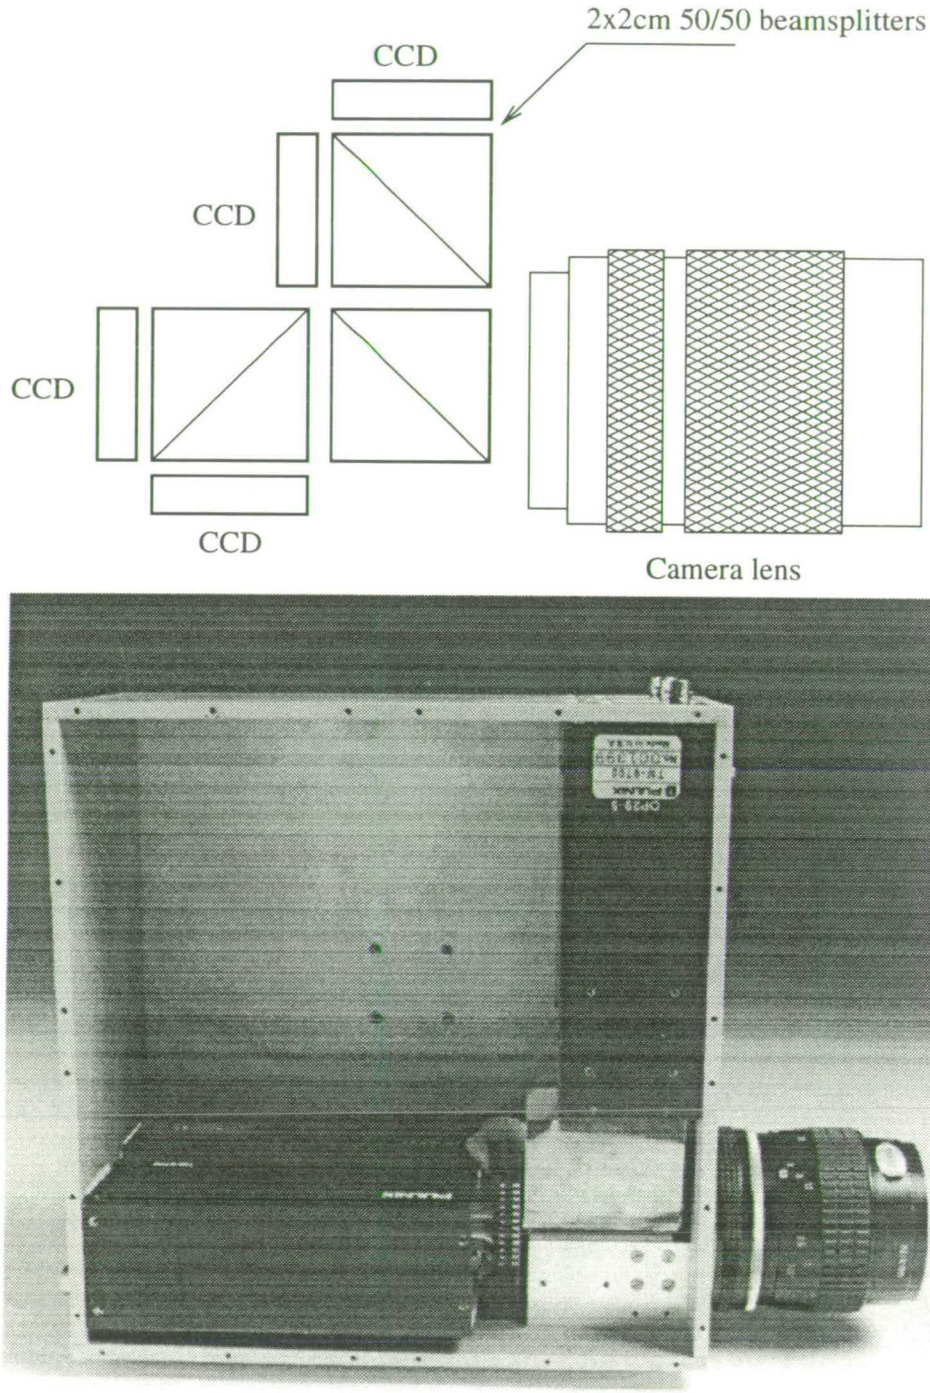

Figure 3.7: 4-CCD array system

array size	768(H) x 484(V)
pixel size	11.6μm(H) x 13.3μm(V)
array type	full frame interline transfer
fill factor	(9700) 20%, (9701) 60%
minimum illumination	1 lux @ f1.4
other	asynchronously re-settable (5μs uncertainty)
	digital output
	built in frame store
	flexible shuttering

Table 3.3: Specification of the Pulnix TM-970x

3.3.2 Construction

The first system constructed was a two-CCD array device, to be used as a general purpose cross-correlation camera, and also to provide a proof of concept. The optical arrangement of this new setup is considerably simpler that the prototype, consisting of a single large cubic beam splitter. The beam splitter is a standard BK-7 non-polarising 50/50 beam splitter, with the only extra specification being non-reflective coating. As previously mentioned, the lens is arranged such that the optical path length is the same as a standard 35mm single lens reflex (SLR) camera.

The construction of this first series of two-CCD array devices was again split into prototype and final stages. The plans of both cameras are in appendix B.

Prototype

The first design comprised of a central holder for the beam splitting cube, with each CCD array being held facing towards the appropriate face of the beam splitter by a floating steel plate. The object of this plate was to allow the CCD arrays to be moved in space with all six degrees of freedom. This would therefore allow the arrays to image the same portion of the image plane. The plates were held in place by thin layers of foam (which acted as springs) and sharpened M2

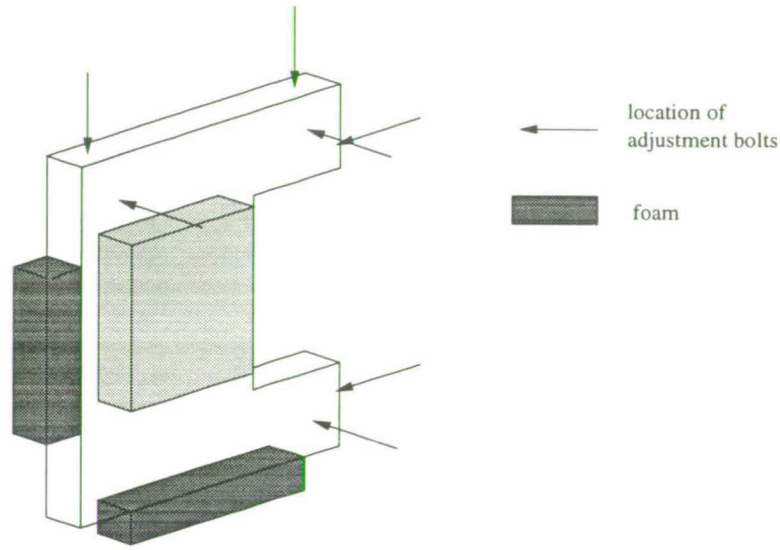

Figure 3.8: Mounting plates for CCD arrays

bolts (see figure 3.8).

Assuming the minimum accurate rotation of a bolt is 10°, and knowing the pitch (that is, the linear distance the bolt will move during one complete rotation) of M2 bolts to be 0.4mm, the smallest accurately repeatable linear movement will be approximately $11\mu m$. This is the same order of magnitude as one pixel, and so should give a translational accuracy of ± 1 pixel. However, it is perhaps more important to remove any rotational misalignment between the arrays. If the arrays have a linear translation between them, then the largest the misalignment can be is the same over the entire array. However, if there is a rotation with the centre of rotation at one edge of the array (to consider the worst case scenario), then the misalignment will be the angle of rotation multiplied by the size of the array. When the size of the interrogation areas used to analyse the image are 32x32 pixels, a misalignment of 10 pixels can represent a serious loss of information, however this would correspond to a rotational misalignment of less than 1° for an array of length 768 pixels.

This first design worked well and two complete cameras were made, one of which is used by Strathclyde University, the other used by Dalian University of Technology, China. Because of the complicated design which allowed the position of each CCD array to be adjusted completely, the arrays were almost entirely enclosed in metal. This had the effect of screening the arrays from each other very effectively.

Final Design

The second design was much simpler, after if was discovered that the adjustment of the arrays could be done by hand with a level of accuracy similar to that achieved by the multiple adjuster bolts. The benefit of having this simpler design also meant less time for assembly of the camera, simplicity in machining the necessary components, and the whole design was more rugged. The plans for the mounting components can be found in Appendix B, figures B.4, B.5, B.6. The only optical component again was a 3x3x3cm unpolarised BK-7 beam splitting cube. The CCD arrays were attached to the cube mounting with M2 bolts, and spacers of 1mm thick rubber sheeting. This allowed a slight adjustment of focal plane of each array.

All adjustments of the positions of the arrays had to be carried out by hand, but alignment of the arrays to within 5 pixels did not prove to be a problem, with the current camera in use at the University of Edinburgh being aligned to within 2 pixels in both the horizontal and vertical directions.

3.3.3 Four-CCD array camera

The four-CCD array camera was designed to the same set of criteria as the two-CCD array cameras described above. The only major difference was in order to keep the optical path length compatible with 35mm format SLR lenses, 2x2x2cm

beam splitter cubes had to be used. The design and construction of this four-CCD array camera was carried out by Michael Jakobsen, who was employed as a research associate at the time. The construction of the mounting for the beam splitting cubes was similar to that of the final two-CCD array system, with all components machined from solid aluminium. The arrays themselves were attached to this mounting again by M2 bolts with a spacer of 1mm rubber sheet. All of the arrays were aligned by hand, and apart from the further complexity of having four instead of two arrays to adjust, the arrays were again aligned without considerable difficulty to within 5 pixels. The CCD arrays used in the four-CCD array system were Pulnix TM-9700's, later being exchanged for TM-9701's.

3.4 Array Alignment

In order for the camera to be of any use, the arrays have to lie in the same area of the focal plane. The arrays are aligned by hand to give a result which is accurate to within 10 pixels. The resulting misalignment is then removed by computer during the analysis stage, and this process is covered in a later chapter. However, the arrays also have to lie within the depth of focus of the lens to ensure that all of the particle images are in focus during the experiment. This is achieved by mounting the arrays on rubber spacers which allow small adjustments to be made to their axial position. The camera is then used to image a real flow image with the lens at its minimum f-number, and therefore minimum depth of focus. The mounting bolts are adjusted for each array until all of the particles are in focus both across each individual array and for all arrays. The result is arrays which lie in the same focal plane with an uncertainty of $\pm 1.5 \mu m$ for an f-number=1.4, $\lambda = 532nm$. Two views of the arrays can be seen in figures 3.9 & 3.10.

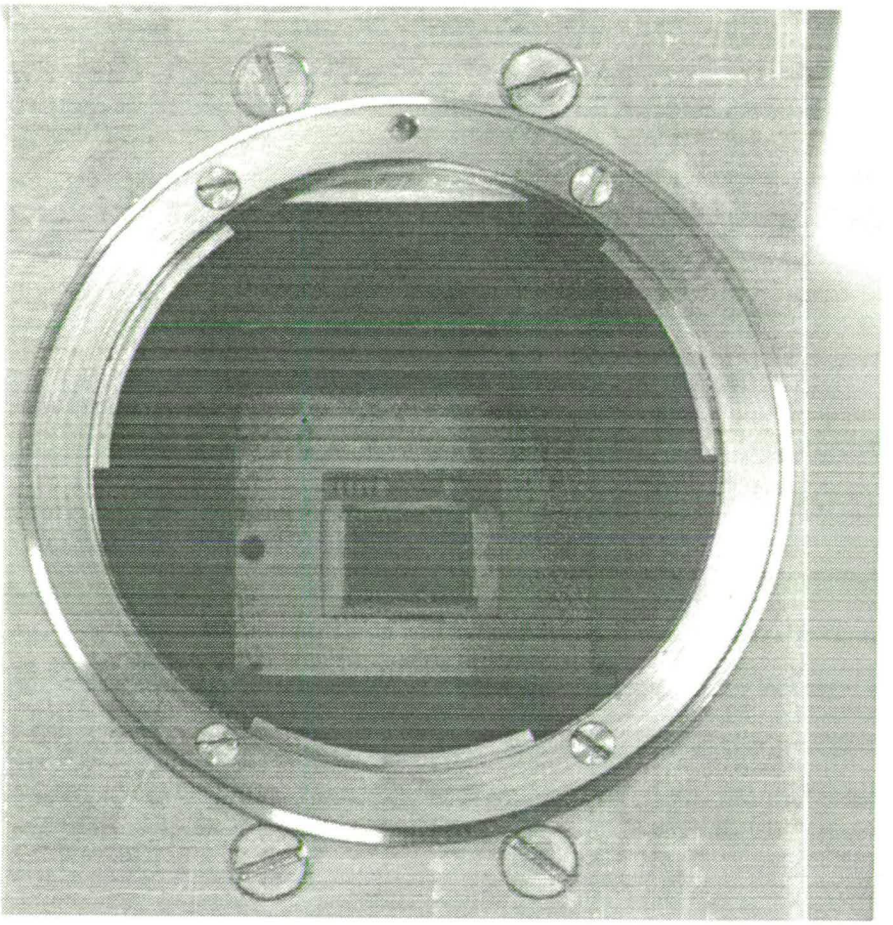

Figure 3.9: Arrays behind Nikon F-mount

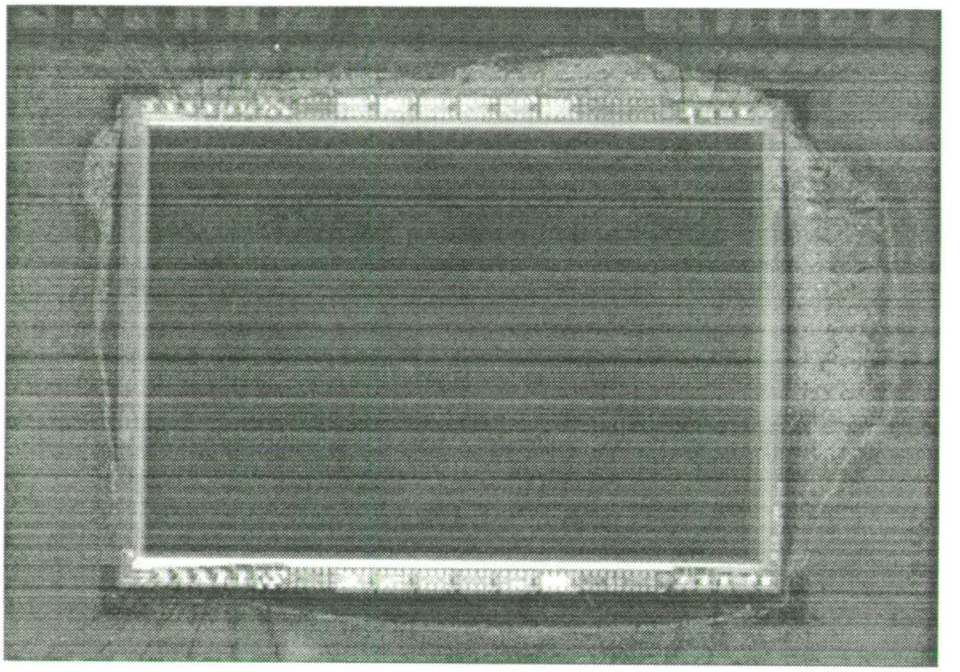

Figure 3.10: Both CCD arrays aligned on top of each other

3.5 Lenses

The 35mm film format has dimensions of roughly 36x24mm, giving a diagonal dimension of $\sim 43mm$. Because of the circular symmetry of all lenses, we can therefore assume that there will be a circular region in the image plane for which the monochromatic first-order aberrations present in the lens (Eq. 3.2) [21, 24] will be minimised. This will typically not be less than the longest dimension of a film frame in order to allow no noticeable effects to arise from any aberrations present in the lens. Chromatic aberrations need not be considered as we are using monochromatic light to illuminate the flows under investigation.

If we consider any aberations present in our optical system will produce a deviation from the ideal parabolic wavefront, the pupil function $P(x, y)$ of our lens will become complex [24]:

$$Q(x, y) = P(x, y)e^{i\kappa W(x,y;\eta)} \tag{3.1}$$

with $W(x, y; \eta)$ being the wavefront aberation function with units of distance (m) and $\kappa = \frac{2\pi}{\lambda}$.

The full first order aberations can now be written as:

$$W(x,y;\eta) = \frac{1}{8}S_1\left(\frac{r}{a}\right)^4 + \frac{1}{2}S_2\frac{yr^3}{a^4}\eta + \frac{1}{4}S_3\left(\frac{r^2+2y^2}{a^2}\right)\eta^2 + \frac{1}{2}S_4\left(\frac{r}{a}\right)^2 + \frac{1}{4}S_5\frac{ry^3}{a^4}\eta^3 \tag{3.2}$$

The terms in equation 3.2 are

- $W(x, y; \eta)$ is the wavefront aberration function
- x, y coordinates in lens plane, r: radial position in lens plane i.e. $\sqrt{x^2 + y^2}$, a: radius of lens
- η: object distance off axis
- S_1: Spherical aberration
- S_2: Coma
- S_3: Astigmatism
- S_4: Defocus or Field Curvature
- S_5: Distortion

The most important aberrations to consider are those which cause a defocusing of the image, i.e. spherical aberration, astigmatism and defocus, as they will alter the shape of the particle images, and thereby give rise to an effective loss of intensity for the out of focus images.

As the optical arrangements of both the two- and four-CCD arrays systems have been designed to maintain a constant optical path length, the image plane should appear as though it were a normal 35mm SLR. However, we are using

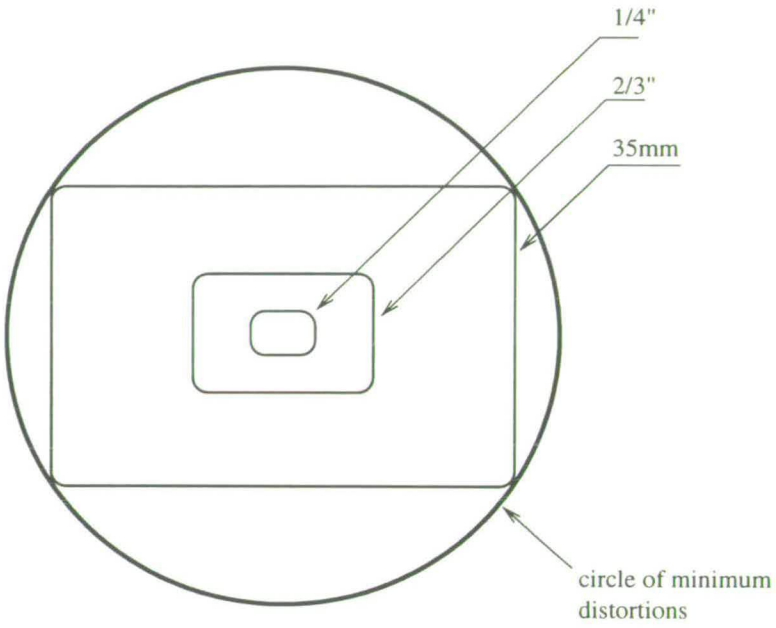

Figure 3.11: Relative coverage of CCD formats

CCD arrays which typically come in formats such as $\frac{1}{4}''$, $\frac{1}{3}''$, $\frac{1}{2}''$, $\frac{2}{3}''$ as they were designed originally for use with C-mount systems (typical back focal distance of 15-16mm). This therefore means we are using only the central 5-19% of the acceptable imaging area (see figure 3.11).

If we accept that we are using only a very small portion of the usable imaging area produced by the lens, then we should be operating within the paraxial region of the lens, and so spherical aberration should be reduced to a minimum. This should also be true of coma and astigmatism as the off-axis distance will be reduced to a minimum, and hence the aberrations also.

Aberrations which produce no defocus, but instead an alteration of the geometry of the image i.e. distortion, can be accounted for in software. As all processing of the PIV images is done digitally, any distortions can be removed during or prior to analysis. However, the prevention of these distortions is a better approach to take, and in general the statements made about the central area

of the image plane being relatively aberration free should apply equally well to geometric distortions.

The lenses which were used in the experiment carried out were all from the Nikon Micro-Nikkor series, which all have the property of a very flat field image plane. The lenses used were all either 55mm or 50mm focal length, and had either f2.8 or f1.4 maximum apertures depending on how much light was available to illuminate the flow.

3.6 Internal Reflections

In all of the above constructions, the light behind the lens is split by at least one beam splitting cube. This can lead to problems with off-axis imaging. If we are using a standard 35mm format SLR lens, with the optical paths matched, then we are using less than the central 20% of the available image. However, the field of view is considerably wider than this, and so the possibility of internal reflections becomes important (see figure 3.12). This is most often seen when using the four-CCD array system in conjunction with a scanning beam illumination system. Because of the large area of flow that can be illuminated (up to 1m wide), the possibility of light being scattered from another part of the flow being internally reflected and finding its way to the CCD array is high. This image then appears as a ghost image superimposed on top of the image of the region of interest. This can be solved easily, by either placing a mask around the edge of the region of interest, stopping the scattered light from further away reaching the camera, or by narrowing the illuminating sheet of light. However, if this problem is not addressed during the experiment, the results of the correlation process will be reduced or even completely wrong due to the two areas of flow being superimposed.

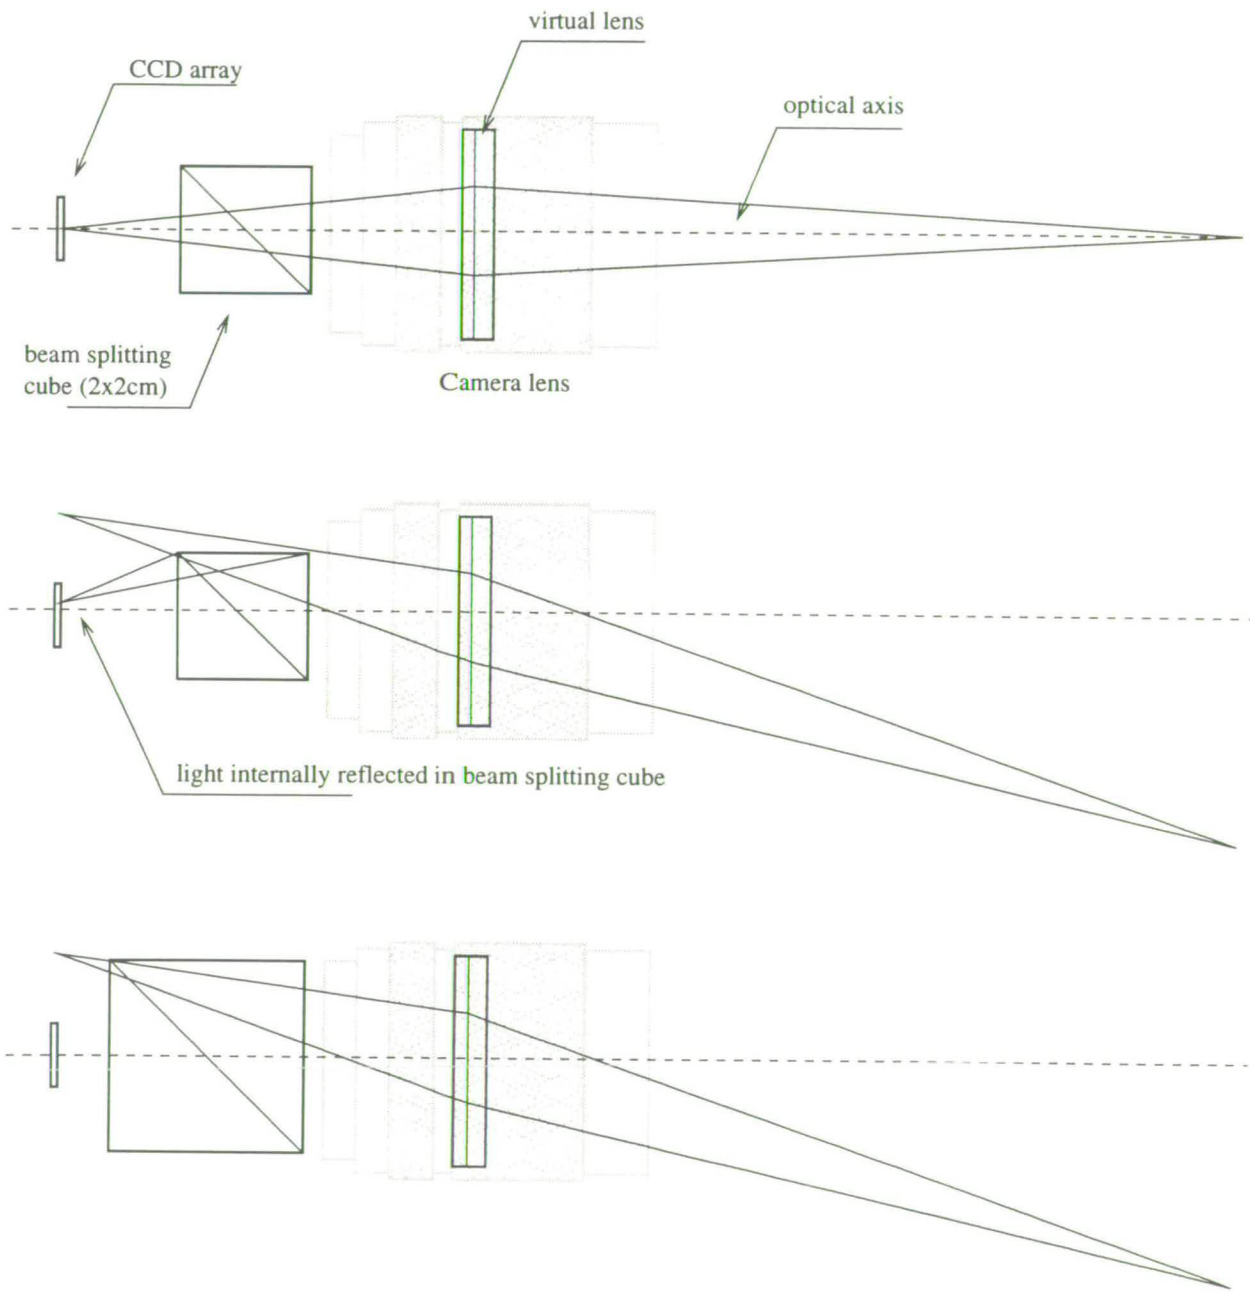

Figure 3.12: Internal reflections caused by off-axis illumination

3.7 CCD Noise

Problems have arisen during the construction of the multiple CCD array systems. Most notable is the problem of clock noise.

The VVL Peach cameras have a pixel clock which operates at 12MHz whilst the Pulnix TM-9701's operate at 14.32MHz. Because of the simple construction of the VVL cameras, access was readily available to the master resonator and synchronising (SYN) inputs. This allowed a clock signal and derivative synchronising SYN signal to be generated easily with standard TTL logic components, enabling the synchronisation of both CCD arrays.

However, the architecture of the Pulnix cameras is much more complicated, and the reluctance of Pulnix to release any information about the synchronisation of the clock signals meant that it has proved impossible to couple the CCD arrays together. This has a serious effect on the quality of the resulting images unless addressed.

The relatively high frequency of the pixel clocks in the cameras being transferred to and from the CCD arrays gives rise to the emission of noise signals of the clock frequency. This can be seen as the effect of passing an oscillating current along any wire. Conceivably, the removal of this effect should be simple in that encasing the clock and all associated electronics in a well grounded metal case should remove any noise emissions. However, if the grounding of the camera electronics has a small capacitance in the path, this can have the opposite effect of making the shielding casing act as an emitter [23]. Added to this is the fact that the most sensitive part of the camera, the CCD array, is exposed to all noise signals.

The result of the noise emission and subsequent interference produced on the CCD arrays is in the form of a set of periodic lines, both horizontal and vertical

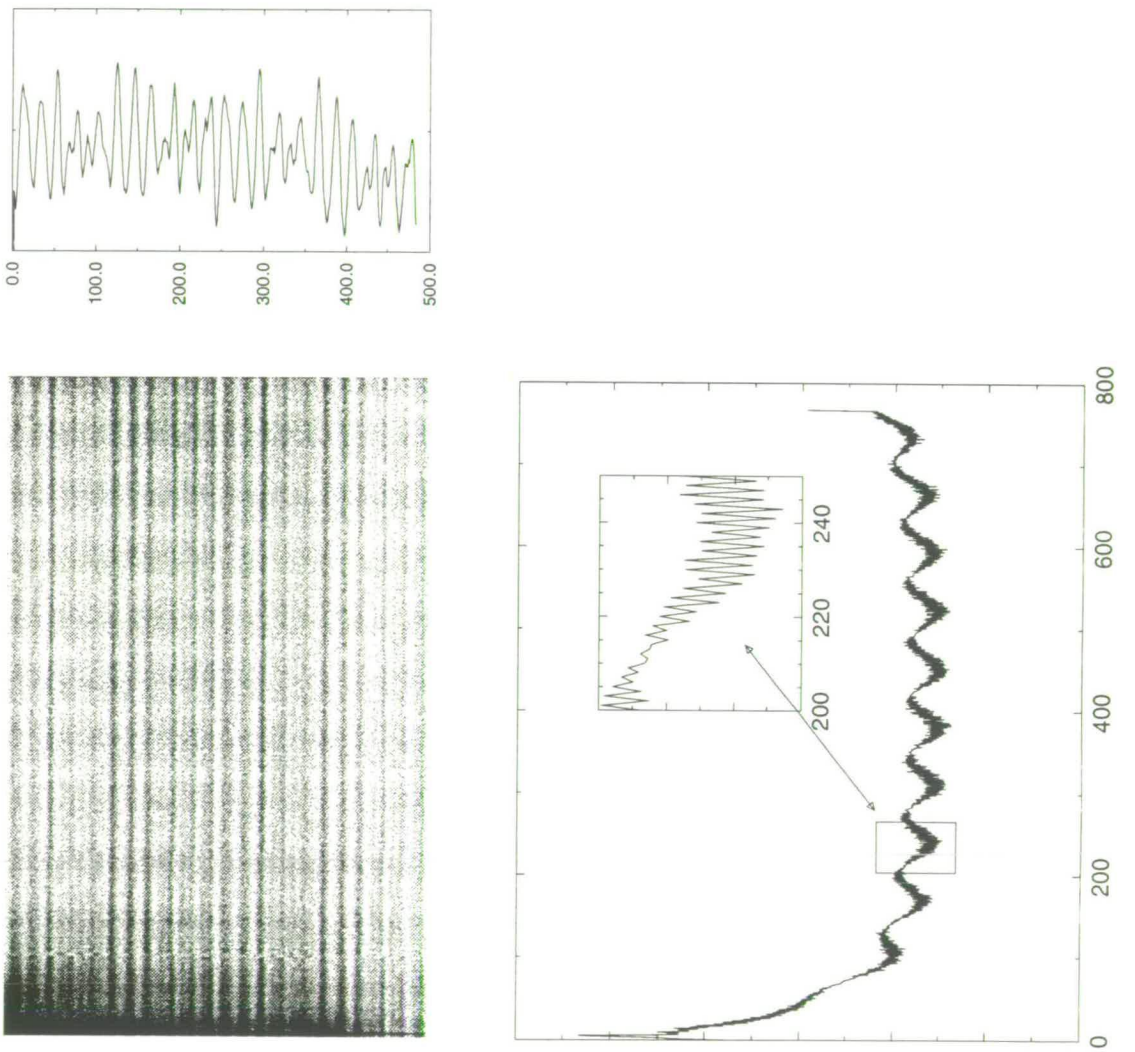

Figure 3.13: Clock noise

(figure 3.13). This image was taken with the lens cap on, the gain of the frame grabber board set to three times its normal gain, and the image has then been normalised to allow the noise to be clearly seen. The original image had a dynamic range of of about 6.5 bits (90 grey levels), and so in the original image the noise would have corresponded to the lower 30 grey values. However, for an 8 bit image, this accounts for 12% of the dynamic range possible. If the output from the two CCD array system is watched, the noise lines can be seen to increase in intensity, and then fade and then transfer to the other camera, indicating a beating effect, most probably due to a fractional difference in the clock frequencies of the two cameras. As the time taken for a complete cycle (i.e. the noise moving from one array to the next and back again) is of the order of 10 seconds, this corresponds to a difference in frequency between the clocks of 0.1Hz, or approximately $7 \times 10^{-9}\%$ of the pixel clock frequency. Obviously, such small effects can only be removed completely by the use of a single clock.

To ensure the removal, or at best, reduction of these signals, very careful screening of all of the electronics has to be taken into account. This can be achieved by screening all of the exposed electronics with a layer of copper sheet, insulated by a paper backing and making sure all of the copper is correctly grounded (figure 3.14, only partly screened to show the linking ribbon cables). The ribbon cable which connects the camera head to the body (containing the rest of the electronics) then also needs to be shielded, and this can most easily be achieved by wrapping all gaps and exposed areas in a layer of crumpled aluminium foil. The total effect is a reduction in the noise signal to a level where they do not influence the correlation process. In the case of the image above, the screening reduces the noise to a level of 3-4 bits, corresponding to between 3 and 6% of the possible dynamic range. The complete removal (or at least to a level when it

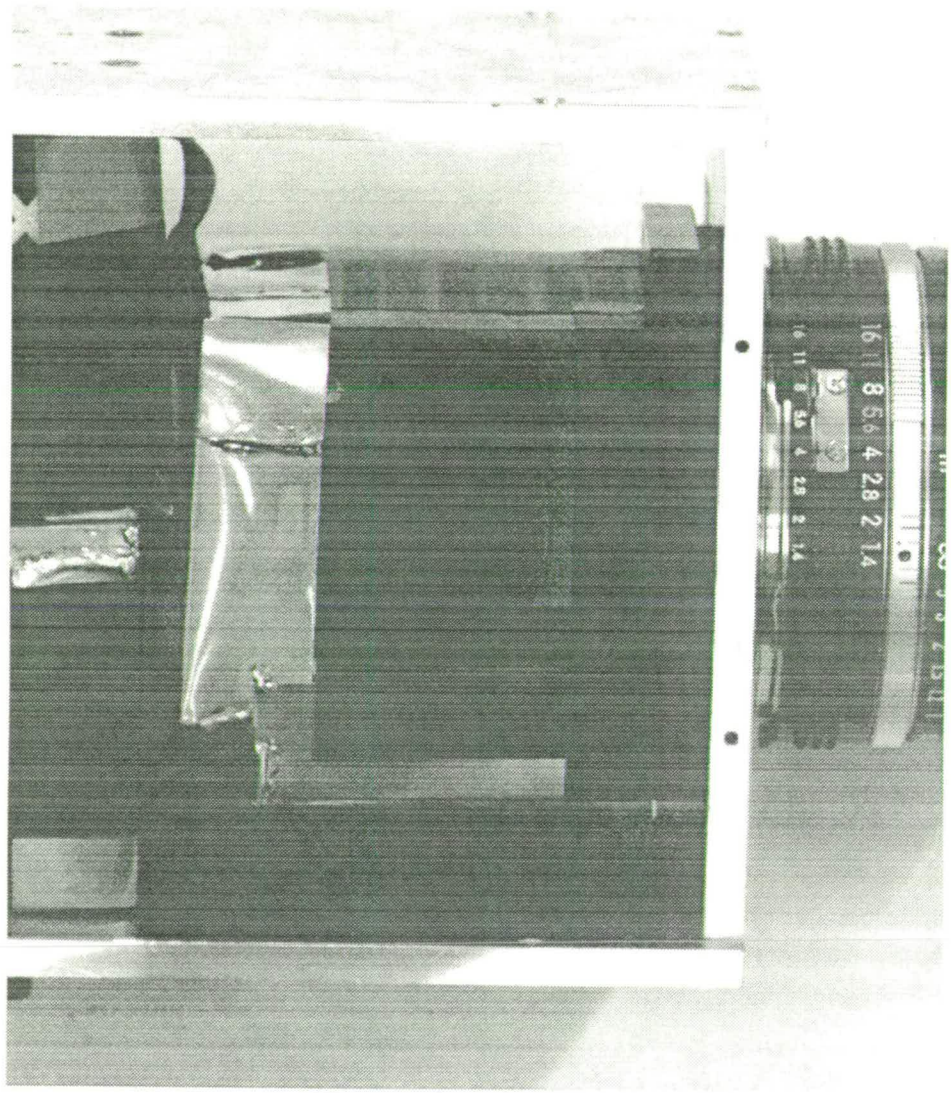

Figure 3.14: View of camera head with partial screening

would not affect the image at all) could be achieved by having a master clock.

3.8 Conclusion

The construction of multiple CCD-arrays is possible, and does create a more flexible imaging system. The resulting system is capable of capturing images separated by μs, and each array is independent of the other.

The system does have problems however. The division of the incident light equally between the arrays means that either the arrays should be very sensitive, or the illuminating laser power should be very high. The availability of arrays which can produce 12 or 16 bits of information is increasing, and coupled with solid state peltier cooling of the arrays the laser power can be reduced to a sensible level [43]. The arrays that are used in this thesis work are , however, not redundant yet. When using our arrays for PIV, we only use a small portion of the available dynamic range. Therefore, if the arrays could be very well screened, or a common clock used, the most dominant noise present in the images should be thermal noise and also the effects of charge leakage. The addition of a cooling device would lower the thermal noise, and so extend the dynamic range to the full 8-bits. As the maximum image intensity is low, the effects of charge leakage should also be low. However, to ensure we can also capture all of the the available information, it would be advantageous to alter the gain of the analogue-to-digital (A-D) converter to reflect the fact that we are only using the lower end of the possible imaging range. This is a feature which is available on the Pulnix TM-9701's without too much difficulty. By combining these two slight modifications, our existing system should have an increased sensitivity at the expense of the total dynamic range of the system.

Chapter 4

Software

4.1 Generation of Artificial PIV Images

The generation of artificial PIV images is of great importance for a number of reasons. It is necessary to have a ready supply of images for testing analysis software, looking at bias effects which depend on seeding concentrations, contrast levels, velocity gradients, to name but a few cases. However, for the results of these simulations to be of any use at all, it is important that the generated images mimic the real PIV images as closely as possible.

There are a number of factors to be included in a program which will generate PIV images, namely:

- Seeding density
- Particle diameter and variance
- Particle position in light sheet
- Background noise
- Real flow properties i.e. local variation in velocity, etc.
- Loss of particles due to out of plane velocity components

- The effect the recording medium has on the image
- Ability to generate "interesting" flows and velocity gradients

The code to generate the images was written in C using a locally available image handling library [1], and consists of under 1,000 lines of code. The software has the ability to read in a flow definition file which contains the coefficients of two polynomials as defined in equation 4.1, with each polynomial being a function of either x or y and giving a displacement in either the y or x direction respectively.

$$d_x = S_x(y) + S_y(y) + sin(y) + cos(y) + a + by + cy^2 + dy^3 + dy^4 \qquad (4.1)$$

$S_x(y)$ and $S_y(y)$ represent the *skew* or spatial gradient in the x or y direction as a function of y respectively.

The position variables are normalised, running from [-1:1] to allow for the easy generation of definition files which are independent of image dimensions.

All of the items listed above are supported by the software, and so realistic simulations can be made which take account of the real factors inherent in PIV. The software has been used in all simulations by the author.

The generation of the particle image is perhaps the most important part of the simulation if the results are to be comparable with a real experiment. It has been generally accepted that the light scattered from the majority of small seeding particles can be closely approximated by a Gaussian intensity profile [13, 1, 44, 62, 27], and so this is the approach taken by the author. This assumption that the particle image is going to be Gaussian in form is a combination of the effects of light scattered from the seeding particle, and the optics present in the camera's lens. However, in most cases this is a very good approximation. Now, however,

[1] This library was written by Dr. W. Hossack of the Department of Physics and Astronomy, University of Edinburgh

we have to take account of the medium onto which the image is being recorded. If photographic film were being used, the spatial resolution of the film would be able to capture the particle image very well, assuming sufficient illumination of the particle and sensitivity of the film. However, with the increasing use of CCD arrays, the *spatially* integrating effects of the pixels have to be taken into account. In section 2.1 it was stated that the optimum size for particle images was of the order of two pixels in each direction. This is due to the fact that we are not trying to reconstruct the particle image, but rather extract the local information about a group or *ensemble* of particle images. Therefore, as the pixel has a finite size, the particle image will in effect be integrated over the dimensions of the pixel. This integration will also take account of the fill factor and the geometry of the pixel.

This integration can be carried out in two ways, either by directly numerically integrating the Gaussian intensity function by using either the *trapezoidal rule* (equation 4.2) or *Simpson's rule* (equation 4.3), or by integrating the Gaussian definitely, making use of the error function $erf(x)$ (equations 4.4, 4.5).

$$\int_b^a f(x)dx \simeq \frac{h}{2}\left[f(a) + f(b) + 2\sum_{k=1}^{n-1} f(a+kh)\right] \tag{4.2}$$

$$\int_b^a f(x)dx \simeq \frac{h}{3}\left[f(a) + f(b) + 2\sum_{r=1}^{n-1} f(a+2rh) + 4\sum_{r=1}^{n} f(a+[2r-1]h)\right] \tag{4.3}$$

$$erf(x) = \frac{2}{\pi}\int_0^x e^{-t^2}dt \tag{4.4}$$

$$\int_{x_1}^{x_2} e^{-ax^2}dx = \frac{1}{2}\sqrt{\frac{\pi}{a}}\left[erf(\sqrt{a}x_2) - erf(\sqrt{a}x_1)\right] \tag{4.5}$$

Both will yield the same result assuming the step sizes in the numerical integration are small enough. For this program the erf approach was used due to the relatively long computation times required for a direct numerical integration.

The generation of the images is done on an additive basis, with each subsequent particle image adding to what was already there. This is the same process that is taking place in a real CCD, and allows the particle images to be superimposed onto a background of noise. The noise is generated by setting each pixel to a normal random value. The maximum value is set by an option flag when running the program. Once the pixels all have a random value, the image is then convolved with a Gaussian a number of times to give a continuous, slowly varying (relative to the size of the particle images) random background. The number of times the image is convolved can be altered to either increase or decrease the average spatial frequency of the noise.

Some further improvements which could be made to the simulation software would be the ability to include a random phase factor with each pixel. This would then allow the generation of specular patterns once the seeding density was increased to a high enough level.

4.2 Removal of Array Misalignment

In the construction of the multiple CCD-array systems, it proves to be impossible to align the arrays perfectly by hand. Commercial multiple CCD-array cameras are aligned by machines, and so can be expected to lie accurately in line with each other, or at least to within 1 pixel. If the arrays are not aligned properly, then an artificial displacement of the imaged flow field will be superimposed on top of the real flow field. This will obviously lead to erroneous velocity measurements, and must be compensated for.

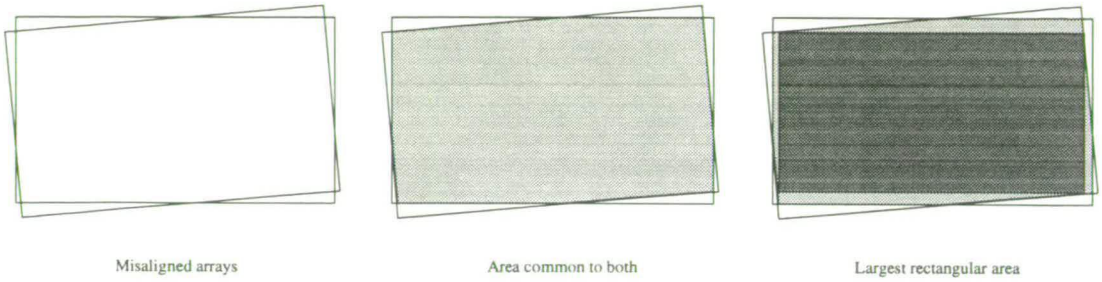

Figure 4.1: Misaligned CCD arrays

In order to find the relative misalignment between the arrays, we shall start by looking at the case of a two-CCD array system. If we consider the first array to be fixed in the correct position, then the second array has to be moved physically until it lies in a position similar to the first array, and typically to within a few pixels. This is a reasonable goal to try to achieve. Once the arrays are in place they are checked to make sure they lie in the same focal plane at all f-numbers of the lens.

We now have a camera system with two slightly misaligned arrays (figure 4.1). The first method for removing the misalignment was to image a random dot or speckle pattern. The image is then captured by both arrays, and the images correlated. This then gives a direct measurement of the misalignment between the arrays, and the vector map produced from this calibration procedure could then be subtracted from the vector map obtained from a real flow.

However, this method has limitations, namely

- lowering of dynamic range
- limitation of analysis grid size

If the misalignment is non-zero, then the maximum measurable velocities in each interrogation area will become non-symmetric, and in the direction of the mis-

alignment the maximum resolvable velocity will be

$$\underline{v}(\underline{x})_{mrv} = \underline{v}(\underline{x}) - \underline{m}(\underline{x}) \tag{4.6}$$

where $\underline{m}(\underline{x})$ is the misalignment and $\underline{v}(\underline{x})$ the maximum velocity measurable without any misalignments at $\underline{x}$. If the misalignment is between 5-10 pixels, a not uncommon value, then this will amount to a significant proportion of the dynamic range of the system, typically 10-12 pixels using an interrogation area of 32x32 which is standard with a medium to high resolution CCD based PIV system.

The second point refers to the fact that in order to use the information in the misalignment vector map, the same interrogation areas have to be used at the same grid locations, and so we are limited in the spatial resolutions that can be acquired.

The solution to these problems is to least-squares fit a polynomial surface to the misalignment data, allowing a continuous representation of the misalignment to be stored as only the coefficients of the polynomial surface, and allowing the misalignment to be calculated at any point in the image plane to sub-pixel accuracy. Because the misalignment can now be calculated at any point in the image plane, the misalignment can be calculated and removed by *dynamically* moving the second interrogation area, and so restore the dynamic range of the PIV system. The order of the polynomial surface can be set by considering how the misalignment can occur.

To the first order, the arrays can be displaced from each other linearly, or they can be rotated relative to each other. Considering the second order effects, such as perspective distortion, these could only be introduced if the arrays did not lie in the same focal plane or if the beam splitting optics were introducing large distortions into the optical system. As both arrays are aligned accurately for focus, and the beam splitting optics consist of high quality optics, flat to typically

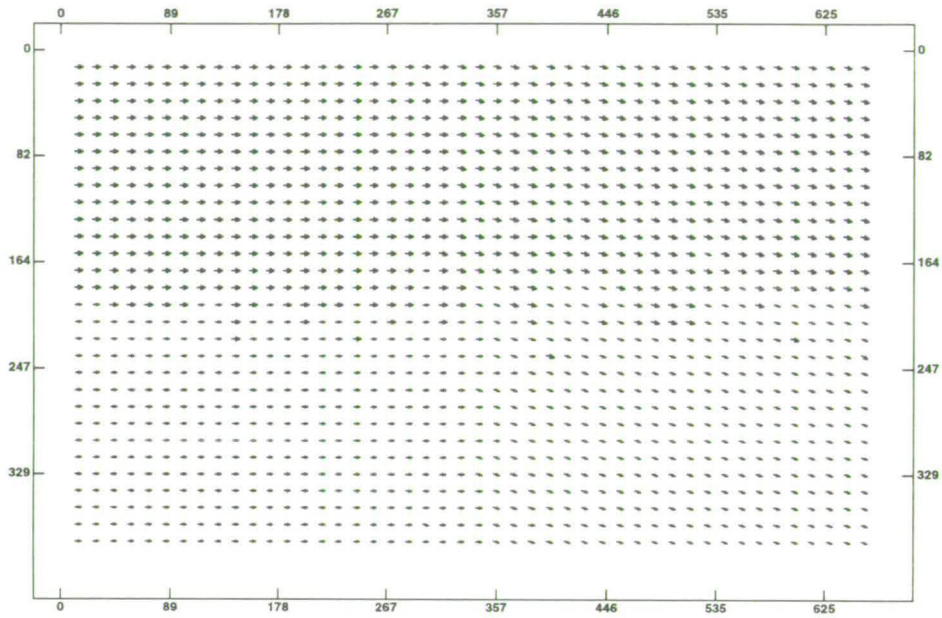

Figure 4.2: Map of array misalignment

$\frac{1}{4}\lambda$, these effects are not likely to occur, and so a simple pair of equations can be used to describe the misalignment.

$$m_x = a_x + b_x x + c_x y \tag{4.7}$$

$$m_y = a_y + b_y x + c_y y \tag{4.8}$$

As the data set from the misalignment map is typically quite small ($\approx$ 1000 points), the coefficients for the surface polynomial can be gained by performing a simple Gaussian matrix inversion upon the normal equations formed from equations 4.7,4.8. The results are then stored to a file and used during the analysis of subsequent images.

As an example, we will look at the misalignment from one of the two-CCD systems at Edinburgh [2]. Both arrays are exposed at the same instance, and the

[2]There are also cameras are Dalian University of Technology, China, and Dept. of Chemical Engineering, Strathclyde University

resulting images analysed. The output vector map can be seen in figure 4.2. Once the image has been analysed, the vector file is fed through a program called *lsf* which follows the procedure explained above to fit a least squares surface to the data. The output from the program is:

```
[tpd@remote_southsider]> lsf out1m.vec off12.cor
a_x: 9.7074561962       b_x: -0.0000627643      c_x:-0.0067696110

a_y: 0.2242905223       b_y: 0.0046936085       c_y: 0.0000105931

[tpd@remote_southsider]>
```

indicating that the misalignment is a combination of a translation and a rotation. Is we now use the coefficients from this least squares surface to remove the misalignment, and then re-analyse the images which gave us the misalignment in the first place, we should get an idea of the residual components of the method. Figure 4.3 shows the residual components. The velocities have been scaled by a factor of 50 to allow them to be seen, and the mean and standard deviations of the residuals are $V_x = -0.007 \pm 0.002$ and $V_y = 0.0002 \pm 0.002$ (both values are in pixels). It can be seen that the standard deviations are of the same size as the mean, and so it would appear that the distribution of these residual values is evenly spread and not clustered around any particular value. In effect, these results should be smaller that the smallest measurement DPIV is capable of making, and so can be neglected. It can be seen from the results that this is the case, and the removal of the misalignment between arrays by this method is very effective.

This work was also carried out simultaneously by *Willert et. al.* [64] during the development of a two-CCD array system for use in high speed air flows. The

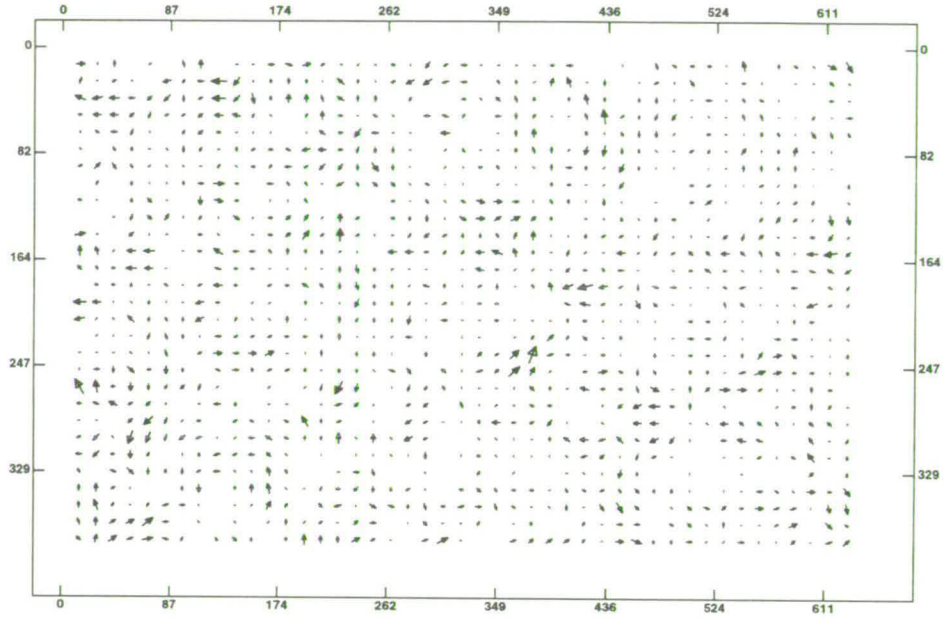

Figure 4.3: Residual components from the removal of array misalignment (vectors scaled up by x50)

motivation for this work was also prompted by the limitations in CCD technology.

4.3 Distortion Removal

In a number of flows which are of interest in the study of fluid mechanics, the optical characteristics of the apparatus within which the flow is contained can introduce a number of distortions into the raw PIV images. This is often the case with flows in pipes, or water flows which involve the camera being placed away from the normal to the illumination plane.

Camera lenses can also be responsible for introducing a number of distortions into the raw PIV images, for example if a short focal length lens is used giving a wide angle view.

The net result however is that all of these distortions must be removed if an accurate image of the flow is to be gained. Traditionally this has not been

possible, and so the reduction or elimination of these distortions has taken first place in the construction of PIV systems and experiments.

We now have a PIV system which is entirely digital, and so the removal of image distortions can be achieved by either mapping the distortion first and then compensating for this effect later, or by warping the raw images to remove the distortions.

The first step is of course to map the distortion, and this can be done in the same way as the removal of array misalignment; we start by imaging a regular grid placed in the plane of the illumination, and an image is captured by the PIV system as it is going to be placed for the rest of the experiment. This is very similar to the procedure carried out to calibrate the PIV measurement, providing a mapping between the real world and pixel coordinates/units.

However, the image of the grid is now analysed by performing a grid wise correlation. As a correlation is simply a pattern matching routine, we need to choose a point in the original image which can be matched with the rest of the image, with the most obvious choice being a crossing point in the grid. As each grid image will be different with each calibration in general, due to changes in the illumination of the grid, the magnification, type of grid, etc., a model of the grid crossing point cannot be stored, and so with each new calibration a new grid crossing has to be chosen.

The correlation is carried out, but the interrogation areas are not overlapped as they are in PIV analysis. The analysis program grid spacing (i.e. location of the interrogation areas) is set to reflect the spacing of the grid, allowing one grid crossing to be present per interrogation. Once the entire image has been analysed we are left with a series of grid points which can then be used to fit a polynomial surface, with the order of the polynomial reflecting the type of distortion. These

can include pincushion and barrel distortions, perspective distortions (similar to the effects introduced by rotating mirror image shifting) and other distortions introduced by the apparatus as previously mentioned.

As distortions commonly encountered with lenses can typically be mapped by polynomials of $\mathcal{O}(x^3)$, the polynomial used to map the distortions should also be of third order, and so by the same process used in the removal of array misalignment, we can form a matrix from equation 4.9 and its normal equations and then perform a Gaussian inversion to yield the coefficients of the equation.

$$distortion(\underline{x})_x = a + bx + cy + dx^2 + ey^2 + fxy + gx^3 + hx^2y + ixy^2 + jy^3 \tag{4.9}$$

As suggested, the distortion can be removed either before or after processing. However, it should be more advantageous to remove the distortion before the raw images are analysed, as typically a distortion will effectively introduce a spatial gradient into the image. Any spatial gradients present within the interrogation area will reduce the height and broaden the correlation peak, increasing the probability of the correlation peak lying below the mean correlation noise level. This technique has yet to be fully evaluated, and it is not known if the spatial gradients introduced into the image due to distortions will be much larger than the spatial gradients commonly found in normal PIV recordings.

4.4 Bias Removal

Analysis of PIV images digitally appears to be very simple, requiring only the correlation of small sub sections of the raw input images. The correlation can be carried out in two ways, either by calculating the correlation directly i.e. by performing a summation, or by performing the correlation in Fourier space, and in

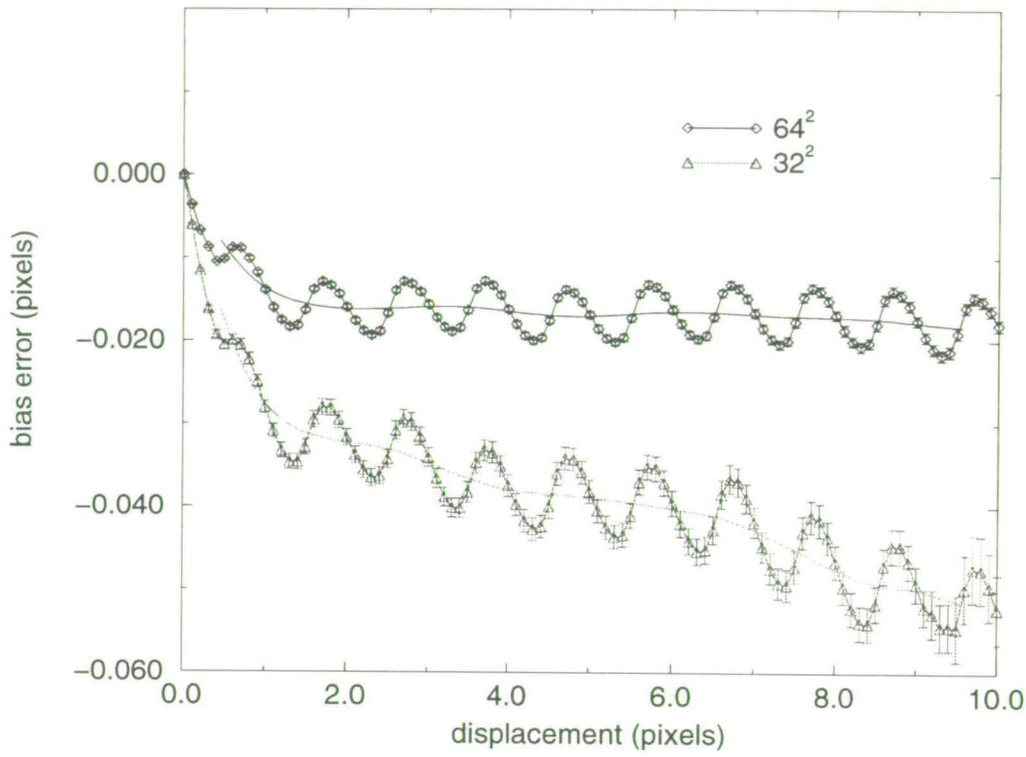

Figure 4.4: Result of simulation to find bias effects, $\diamond = 64^2$ interrogation area, $\triangle = 32^2$ interrogation area.

general the correlation is performed in Fourier space due to the faster performance of the Fast Fourier Transform (FFT) based algorithm ($\mathcal{O}(N^2 ln(N))$ for the FFT based correlation against $\mathcal{O}(N^4)$ for the direct correlation). There are, however, a number of considerations to be taken into account when doing an FFT based analysis as otherwise a bias will be introduced into the velocity measurement.

The kind of biases that can be expected can be seen in figure 4.4, which shows the result from an analysis of simulated PIV images. Each point on the graph is the result of 1296 statistically independent interrogations, allowing a reliable mean and standard deviation to be calculated. All of the PIV images had no

spatial gradients present, and were seeded to allow all measurements to be 95% reliable.

It can be seen that the bias is split into two components, one an oscillating motion with a period of 1 pixel, and the other a plateau which drops steeply and then levels out.

The first component of the bias is the result of having particle images which are too small, and so the three point estimators which are used to calculate the sub-pixel displacement become useless as less information is available to work with. This part is completely separate from the pedestal effect, as can be seen in figure 4.5. The application of different types of centroiding algorithm show it is the assumptions made about the shape of the peak which matter, and so the pedestal effect remains constant. If the particle images are assumed to be approximately Gaussian in shape, then the correlation peak will also be a Gaussian; therefore, if a Gaussian peak fitting routine is used to estimate the centre of the correlation peak, it should perform to a higher accuracy than a parabolic fit, as more information is used in the calculation of the centre position.

If, however, the particle diameters drop below a certain point, the amount of information that can be gained from the correlation plane is reduced until ultimately only integer displacements can be resolved. If a very large interrogation area is used, as is the case when the correlation is done by half-optical, half-digital methods, then this is acceptable, but if we are imaging the flow with a CCD array of limited spatial resolution, small interrogation areas are used, and a resolution of only integer displacements presents a serious limitation in the measurement process.

The cause of the pedestal bias was fully explained by Westerweel in his thesis [62], and can be seen to arise from the assumptions made in the calculation of

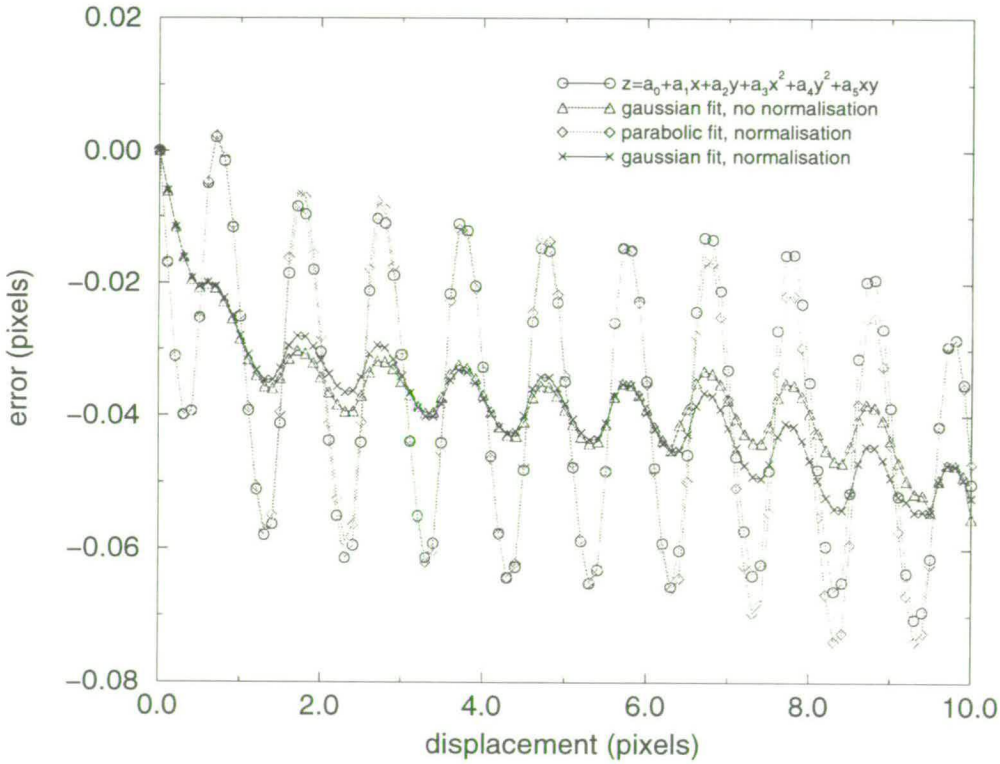

Figure 4.5: Comparison of different sub-pixel displacement estimators

the FFT. The array of data presented to the FFT algorithm is assumed to be periodic, with a period N, the length of the interrogation area. The sample is also considered to extend infinitely in all directions. It is these assumptions which give rise to a biasing of the correlation peak towards a lower velocity, and is termed *displacement bias*. Because the data is periodic, if the particle group which is present in the first interrogation area moves very slightly out of the interrogation area by the time the second exposure is taken (as it must if there is a velocity present), then the number of particle images which are contributing to the correlation peak must decrease. It is this decrease in data which has the effect of biasing the correlation peak towards zero.

These errors can be removed in a number of ways, namely providing an unbiased estimate, re-interrogation, different sized interrogation areas and sub-interrogation area particle tracking.

4.4.1 Correlation Normalisation

Westerweel performed an analysis of the digital PIV system, and concluded that if carrying out the calculation of a co-variance or a correlation, the correlation peak would be biased towards the zero velocity as a result of the diminishing proportion of particle images which were contributing to the correlation. This was termed to give a *biased* estimated of the mean velocity in the interrogation area. The solution to the problem though is relatively simple. If the correlation plane is divided though by a normalising factor, the result should be an unbiased estimate of the velocity. The normalising factor has to counteract the effect of the diminishing number of particle images, and so can be seen to be a square regular pyramid (figure 4.6). The results from Westerweel's thesis show that very good results can be obtained by this method.

Another way of considering this bias is to consider an analogous optical sys-

tem. When analysing the behaviour of lenses in Fourier optics, a lens is considered to be a two dimensional phase function multiplied by a Pupil function. The Pupil function describes the circular symmetry of the lens, and can typically be written as

$$P(\underline{r}) = \begin{cases} 0 & \text{if } \underline{r} > R, \\ 1 & \text{if } \underline{r} <= R \end{cases} \tag{4.10}$$

where $P(\underline{r})$ is the Pupil function, and R is the radius of the lens. We are now interested in finding out how the lens will pass light, and which spatial frequencies will be passed. In the case of coherent light, the Coherent Transfer Function (CTF) is just a scaled version of the Pupil function. However, in the case of incoherent light, we have to calculate the Optical Transfer Function (OTF). This is calculated by performing the autocorrelation of the Pupil function, and typically will have a shape of a triangle in cross section. We now have a function which tells us how much of a signal made up of a certain spatial frequency will be passed by our lens. Of course, if we now consider our interrogation area to be a Pupil function, when we carry out a correlation, we are automatically calculating the OTF of our interrogation area. Therefore, in order to remove this effect we have to divide through the correlation plane by the OTF, and for a square interrogation area we will find this function to be as shown in figure 4.6. This is the same result as found by Westerweel.

If the correlation plane is noisy, then the application of this function to the correlation plane can increase the number of spurious vectors. This is due to the fact that at the corners of the normalising function the correlation plane is divided by 0.25. If the displacement is small, then the signal peak is close to the centre of the correlation plane. This therefore means that the signal peak must be over four times the height of the noise peaks at the edge of the correlation plane. What is seen in practice is that a larger number of erroneous vectors are

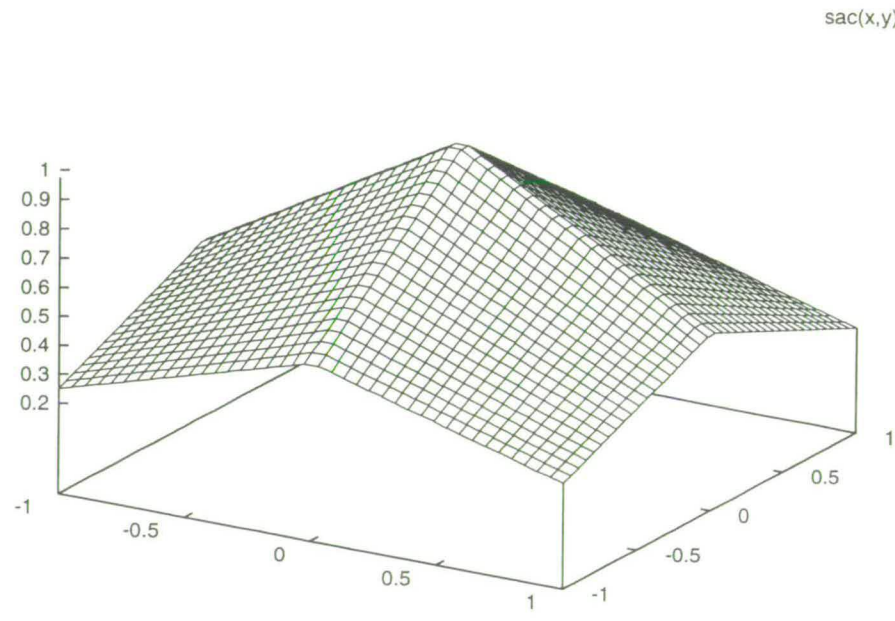

Figure 4.6: Shape of function to be applied to correlation plane data to provide an unbiased velocity estimate

present in the final velocity map, and also these erroneous vectors are gathered around the values which correspond to the four corners of the correlation plane.

If we assume that any signal peaks that are actually placed at the very edges of the correlation plane will be invalid, then we can modify the shape of the function to reflect this. We can now redefine the function to be:

$$\Lambda'(\underline{d}) = \begin{cases} \Lambda(\underline{d}) & \text{if} \Lambda(\underline{d}) \geq 0.5, \\ 0.5 & \text{otherwise} \end{cases} \tag{4.11}$$

and the corresponding normalising function can be seen in figure 4.7. The application of this normalisation function results in the analysis being more tolerant of images with low seeding concentrations or large spatial gradients, all of which flatten and spread the signal peak in the correlation plane. This function will correct exactly all displacements which are less than a quarter of the interrogation area, which corresponds to the limits set by Adrian.

4.4.2 Re-Interrogation

If the interrogation which is biased is used as the input to a second correlation, we now have an *a priori* knowledge of the flow, allowing the interrogation area to be moved dynamically. This insures that the largest displacement within the interrogation area should now be ± 0.5 pixels, and so the bias effects will be very small. This technique also reduces the effect of *gradient bias* which again appears as a biasing of the real velocity towards zero. This is again due to the loss of particle images, but this time due to the effects of spatial gradients being present within the interrogation areas. If there is a spatial gradient present within the interrogation area during the analysis, it is more likely that the particle images which belong to the faster moving particles will leave the interrogation area between exposures. The effect of this larger displacement will then not be included in the average velocity measurement for that grid location, biasing the

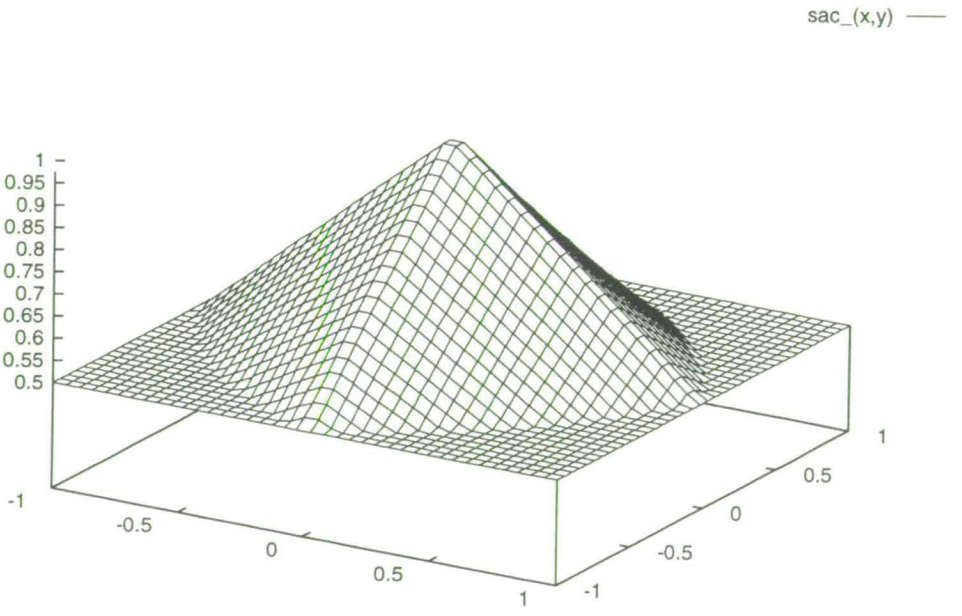

Figure 4.7: Shape of function to be applied to correlation plane data to provide an unbiased velocity estimate

velocity measurement in effect towards a lower value. There are two ways in which to perform this bias removal.

The first method involves analysing the PIV images as normal, i.e. dividing the image up into a grid and performing the correlations. However, at each grid point we perform two correlations, with the input from the first correlation acting as a guide to the second correlation. This method does work very well, but has the drawback that if the first velocity measurement does not return a valid vector, then the second correlation is guaranteed to fail as the interrogation area will be moved in a direction other than the correct direction, again giving rise to another spurious vector.

If the entire image is analysed first to give a complete velocity vector map, then the vector map can be post processed for outlying vectors, and so an accurate representation of the flow can be compiled with a minimum chance of incorrect vectors remaining. If now a high order polynomial surface is fitted to the complete vector map, we have moved from having a discrete representation of the flow to having a continuous approximation. Thus, at every point in the flow we have a knowledge of the flow field, and so the interrogation process can be started over again allowing any biases to be removed by dynamic interrogation areas. This method should be considered as superior, but in itself cannot guarantee the generation of a flawless vector map as the outlying vectors produced in the first pass of the analysis routine were most likely caused by locally low seeding densities. Therefore, even with knowledge of the flow field gained from post processing interpolation of the vector map, the interrogation will fail due to low seeding densities. This is the approach adopted by Cowen [8].

In a recent paper by Westerweel *et. al.* [63] it has been shown that by performing analysis of PIV images using re-interrogation, not only are bias effects

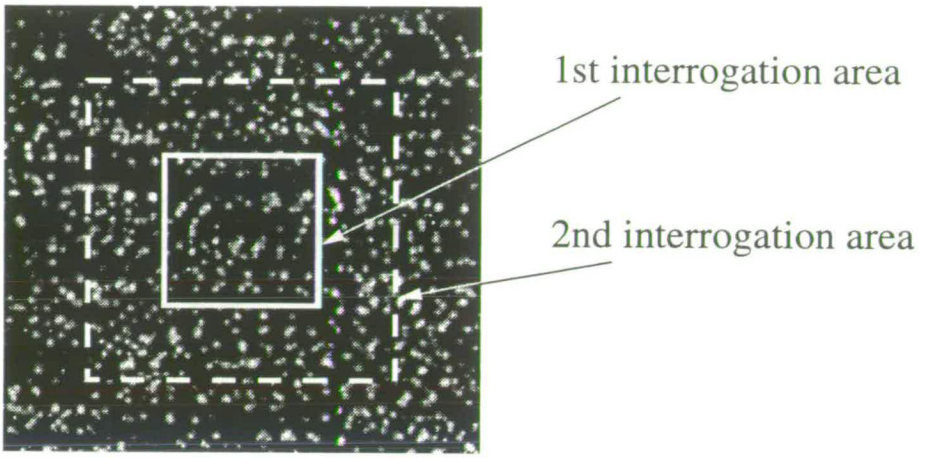

Figure 4.8: Relative locations and sizes of first and second interrogation areas when using zero padding to remove bias.

removed, but a signal-to-noise improvement is noted along with an improvement in the accuracy of very small displacements, giving an increased spatial resolution in their application to homogeneous turbulent flow in a pipe.

4.4.3 Zero Padding

If we have an interrogation area of NxN pixels, we can assume that the maximum resolvable velocity will be less than half of the interrogation area dimensions i.e. if N=32, the maximum particle displacement must lie below 16 pixels in order to be resolved. If we know this fact, then the second interrogation area can be twice as large as the first, and should be centered on the grid location which defines the centre of the first interrogation area to allow velocities in all directions to be resolved (see figure 4.8). If this is the case, then the second interrogation area should therefore contain all of the particle images from the first interrogation area. As the displacement bias is caused by the decreasing number of particle images, this displacement bias should be removed or reduced to a minimum for this case.

However, there are also limitations with this method of bias removal. In general, real PIV images do no have a non-zero background due to the scattered light from out of focus particles. Also in the case of CCD technology, dark noise generated by the CCD chip itself contributes to this background noise. As the correlation is still being carried out by Fourier methods, the interrogation areas have to be the same size, and therefore the first interrogation area has to be packed out with zeroes, or zero-padded to allow the correlation to be performed. This is where the problem arises, as the sudden change from the mean background level of the real PIV image to the zero level in the padded region will give rise to edge effects. These effects appear because the sharp edge alters the spectrum of the image in Fourier space, and will produce unacceptable effects in the correlation plane. This will either ruin the chance of successfully detecting the correct correlation peak, or bias the velocity correlation peak in an unpredictable fashion.

One possible solution to this problem is of course to pack the first interrogation area with the mean image level of the first interrogation area, so that hopefully no edge effects will be produced. However, for the amount of effort consumed in calculating the mean, padding out the interrogation area, and then performing a correlation on a larger area, the better solution is to use a combination of correlation normalising and re-interrogation.

If a kernel is also applied to provide an unbiased estimate, it must also be modified to reflect the geometry of the correlation. The appropriate kernel can be seen in figure 4.9, with the kernel in 4.10 being modified to reduce the number of outliers.

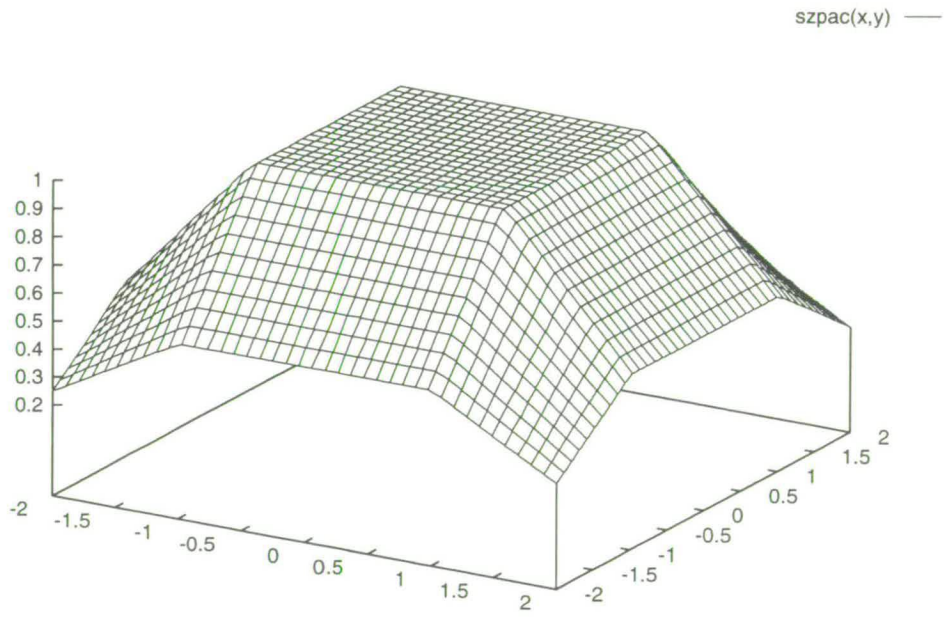

Figure 4.9: Shape of kernel to be applied to correlation plane data to provide an unbiased velocity estimate

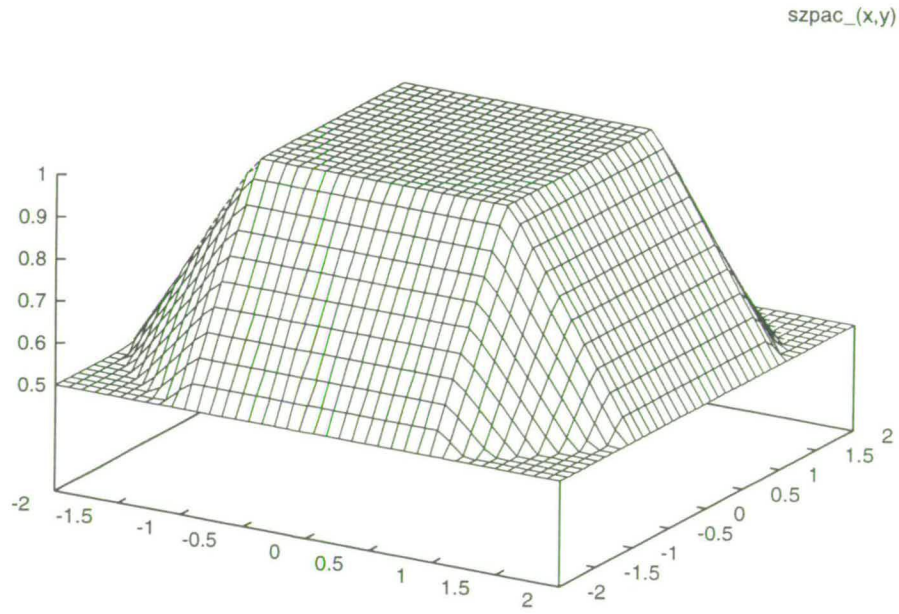

Figure 4.10: Shape of kernel to be applied to correlation plane data to provide an unbiased velocity estimate

4.4.4 Particle Tracking

Particle Image Velocimetry (PIV) and Particle Tracking Velocimetry (PTV) have developed more or less in parallel, although PIV has proved to be the more popular analysis method due to the level of automation that can be brought into the analysis of the flow images. PTV has a greater number of obstacles to overcome, and the algorithmic complexity involved in extracting displacements from low seeding density images is great. However, because PTV is not dependent on the measurement of the local mean or Fourier methods, it does not suffer from the biases described above. If we then follow a procedure similar to that involved in re-interrogation, we should be able to remove the bias from our measurements altogether. The first pass of the correlation routine will yield a knowledge of the flow field structure, then this can then be used to select interrogation areas and perform particle tracking within these local regions.

However, this method will fall foul of the short comings of PTV, namely the particle images must be well defined in order for them to be identified, and hence a result be obtained.

The uniformity of the particle images can be improved by performing a local correlation of the raw PIV image with a computer generated particle image. In order for an accurate correlation to be made, the average diameter of the particles present in the interrogation area must be calculated. This can be carried out in two ways; either by performing a Hough transform [3] on the image to transform the raw image into a suitable parameter space, or by fitting a curve to the correlation peak and finding the roots of the curve. As the particles are going to be of the

[3] A Hough transform is a term which covers nearly all transformations from one parameter space to another. A good example of this is the case of detecting straight lines in computer vision applications. A Hough transform can be constructed which will transform an image from (x, y) cartesian coordinate space into (l, θ), the line length l and angle θ space. A good introduction to Hough Transforms can be found in [32].

order of 2 pixels in diameter, the robustness of a Hough transform would be questionable, and would also involve a large computational overhead.

As the width of the correlation peak is directly proportional to the particle image diameter, if we fit a parabolic or Gaussian curve to the correlation peak, and find the roots of the equation describing the curve, then we have an accurate measure of the peak diameter, and hence the particle image diameter.

A particle image can now be generated by assuming that the particle images are circular and Gaussian in profile. As most particle images are also going to be over 1 pixel in diameter, this correlation also performs a smoothing out of single pixel shot noise present in the real PIV images which would lead to noise in the correlation plane.

In general, this does make the detection of particle images more reliable, but the technique has yet to be fully evaluated by the author.

4.4.5 Smoothing in the Spatial Frequency Plane

If the DPIV image has a high seeding density, and is beginning to border on a speckle image, then the raw image will become more like a random pattern of pixels. Because of this, the upper spatial frequencies contained within the image will become more pronounced. If we allow all components of the spatial frequency plane to remain intact, then the result is a correlation plane in which we have a large amount of noise.

An alternating spatial frequency of period 2 pixels (the highest spatial frequency that can be produced) will produce a frequency peak at 0 and N-1 (for an image that runs from 0 to N-1, not -N/2 to N/2), as can be seen from the transform of a sine function (equation 4.12), and so if a kernel is applied in the spatial frequency domain, all high spatial frequencies can be removed.

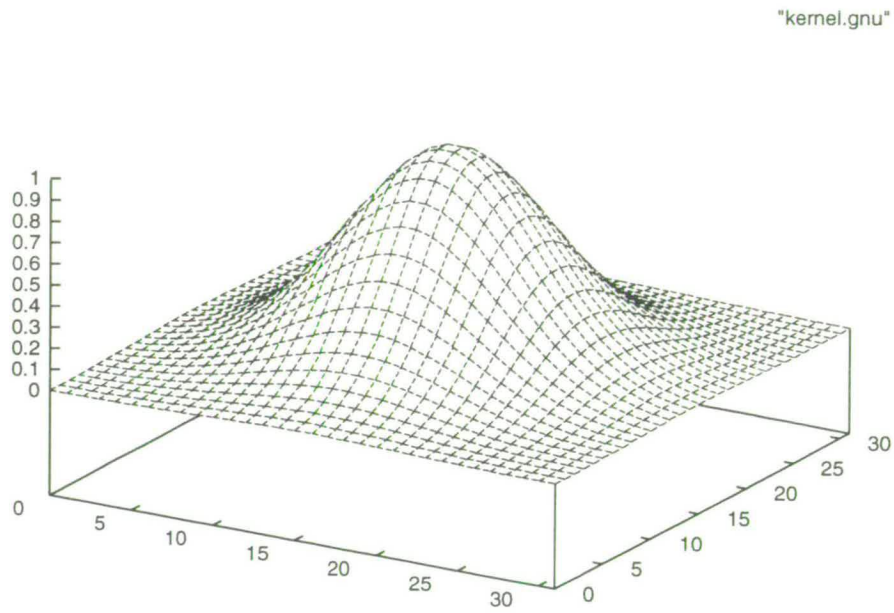

Figure 4.11: Gaussian kernel applied to spatial frequency plane during correlation, width 2σ=N/2

$$\begin{aligned} \mathcal{F}(sin(ax)) &= \frac{1}{2\sqrt{2\pi}} \int_{\infty}^{\infty} e^{-ix(s-a)} dx + \frac{1}{2\sqrt{2\pi}} \int_{\infty}^{\infty} e^{-ix(s+a)} dx \\ &= \frac{1}{2}[\delta(s-a) + \delta(s+a)] \end{aligned} \tag{4.12}$$

The analysis software applies a circularly symmetric Gaussian kernel (figure 4.11) to the spatial frequency plane, with the radius of the kernel, taken as the 2σ point, fixed by a user-defined quantity. The default value for this is taken as N/2, as shown in figure 4.11.

As a multiplication in the spatial frequency domain is equivalent to a correlation in the real space domain, this action can be considered equivalent to applying a smoothing kernel in real space, but the width of the smoothing kernel can be controlled more closely resulting in a finer level of smoothing.

Chapter 5

Application: Two-CCD System

In order to prove that the multiple CCD array systems work, a number of experiments have been carried out.

5.1 Timing Electronics

In order to use a multiple CCD array camera with a scanning beam system, we have to synchronise all events to the scanning beam. This is to ensure that the exposures captured by the arrays are both complete and separate scans of the illuminated area. As most delay or timing boxes cannot cope with counting, either a computer was required, or some simple digital electronics had to be constructed.

The most important events in the measurement of the flow are (in order)

- The event in the flow
- Synchronisation to any other pieces of equipment
- Synchronisation to the illumination system

This will allow the acquisition system to be triggered from the flow, and not vice versa.

The timing system can now be seen to be a straightforward piece of logic (see Figure 5.1). First, we need to acquire the scan rate of the scanning beam

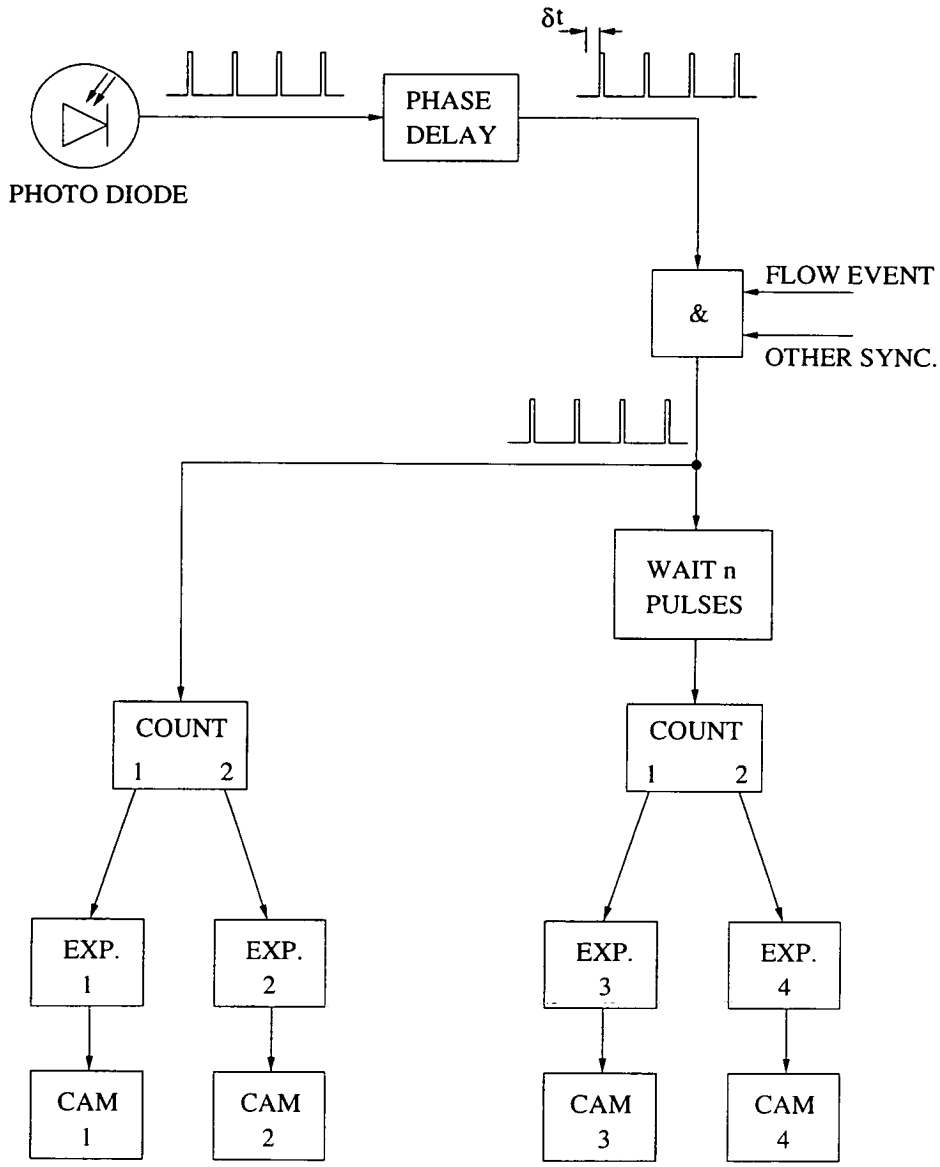

Figure 5.1: Logic behind camera timing

system, and this is most easily achieved by placing a photo diode at one end of the scanning beam system. This photodiode is present in some scanning beam systems already as they used to be essential for accurately setting the scan rate of the system [41].

The next stage is to apply a phase delay to the pulse stream that comes from the photo diode. This is essential as we want to capture a complete scan of the flow, preferably starting at the beginning of the scan. If we manage to achieve this, then there should be no discontinuities due to the exposure of the camera starting in the middle of the light sheet.

After the phase delay has been applied, we now need to wait for the event in the flow which is of interest. This is achieved by simply passing the pulse stream through an AND gate to pass the pulse stream only if the flow event trigger AND the synchronising trigger AND the pulse stream are present.

So far, the timing electronics have been very straight forward, and could have been achieved by any standard pulse or delay generator. However, we now need to count the pulses to allow each array to be triggered on each sequential pulse that arrives at the system after the triggering conditions have been met. The timing electronics have been made flexible enough to allow a four-CCD array system to be used, and in this case we will sometimes want a separation between the first and second pairs of exposures. The reason for this is covered in the next chapter. To allow this to happen, the pulse stream is divided into two, one half triggering immediately the first pair of arrays, and the second half going through a divide-by-n counter and then onto trigger the second pair of arrays. The timing electronics can be seen in figures 5.2 & 5.3.

The timing electronics are based almost entirely on the 74 series of integrated circuits, due to the large number of different functional packages available and

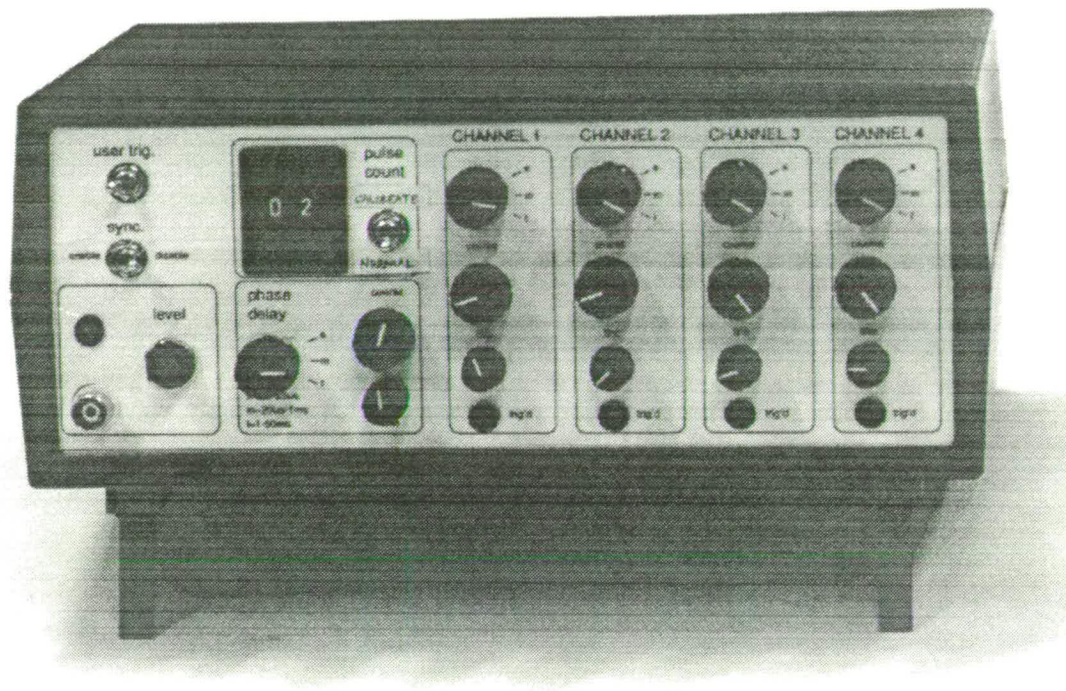

Figure 5.2: Front panel of timing electronics

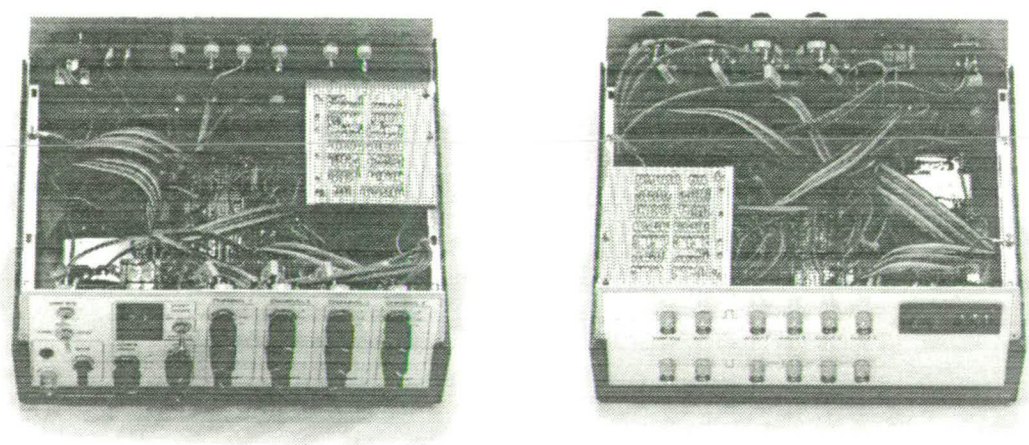

Figure 5.3: Internal structure of timing electronics

the combination of ease of use, relative intolerance to bad circuit design and manufacture, and good speed. The TM-9701 cameras are reset and their exposure length set by the edge and width of a negative TTL (5V) pulse respectively. It was decided that the exposures of the arrays could be set by using monostables [1], as the types of flow which are used with a scanning beam system are relatively slow, being typically no greater than $10ms^{-1}$.

The entire system now consists of a photo-diode, some timing electronics and the multi-CCD array system. This system can now be used as a direct replacement for film based or auto-correlation oriented cameras.

5.2 Wave motion over submerged pipelines

As part of this thesis, some work was carried out in collaboration with Dalian University of Technology, China. The part of the University of Edinburgh in this exchange program was to supply the equipment necessary to make DPIV measurements, and so a two-CCD array system was constructed and taken to China. The analysis system also required that the analysis software be provided, and so the software written by the author for the purpose of thesis work was installed in China.

The experiments that were being run were investigating the effect of wave action over submerged pipelines in shallow conditions. To achieve this the wave flumes have a piston style wave maker installed, as compared to the hinged paddle (figure 5.4) design used in the remaining experiments. The piston wave makers have a rigid front plane, and move in a motion parallel to the long axis of the tank. The resulting wave motion must therefore have a significant motion at the

[1] Monostables are logic devices which will generate a pulse on receiving either a positive or negative going edge. The width of the output pulse is determined by an RC circuit (combination of resistor and capacitor), and so if we use high accuracy analogue components, we can hope to achieve a repeatability of 5-10% in the exposure.

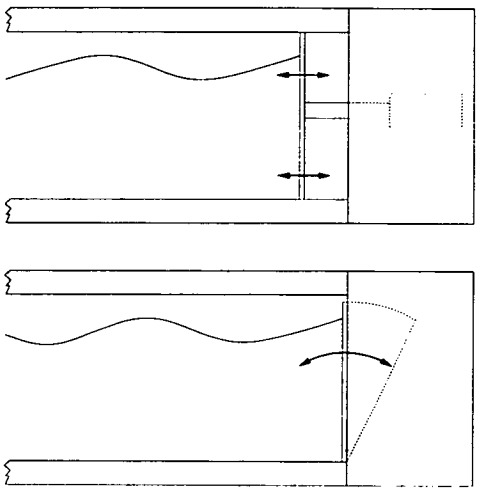

Figure 5.4: Piston and paddle wave makers

bed of the tank, very similar to shallow wave conditions.

The DPIV image shown below shows the effect of a regular wave passing over such a submerged pipeline, showing clearly vortex shedding taking place. The two-CCD array system is exactly the same as described in the previous chapter, and uses two frame grabbing boards in a computer to acquire a time sequence of images. Because the frame grabber boards used are of fairly low bandwidth, between 10 and 12 double images can be captured per second.

5.3 Study of Breaking Waves

Another area of study which has seen a large amount of interest at the University of Edinburgh is the effect of breaking waves. There have been many studies involving water waves, including Gray [18] and Quinn [41], focusing on both practical aspects and comparison with numerical models. In both Gray's and Quinn's thesis' the need for image shifting was not required as there was a large mean flow present. The post breaking turbulence in the body of the fluid is also of interest, but this would require the use of image shifting to allow the directional

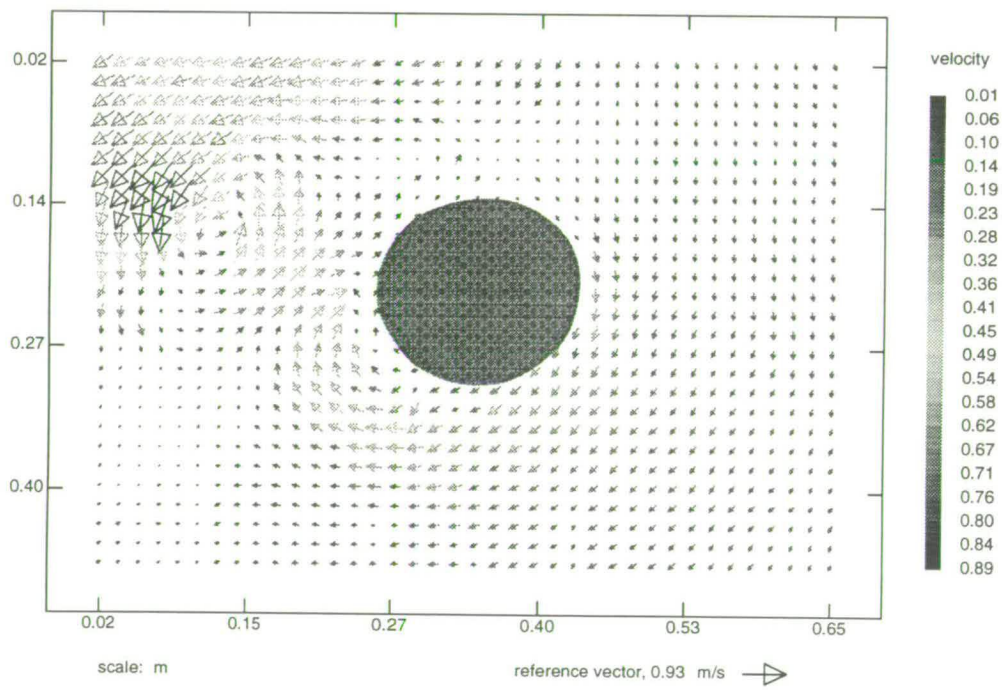

Figure 5.5: Wave over submerged pipeline

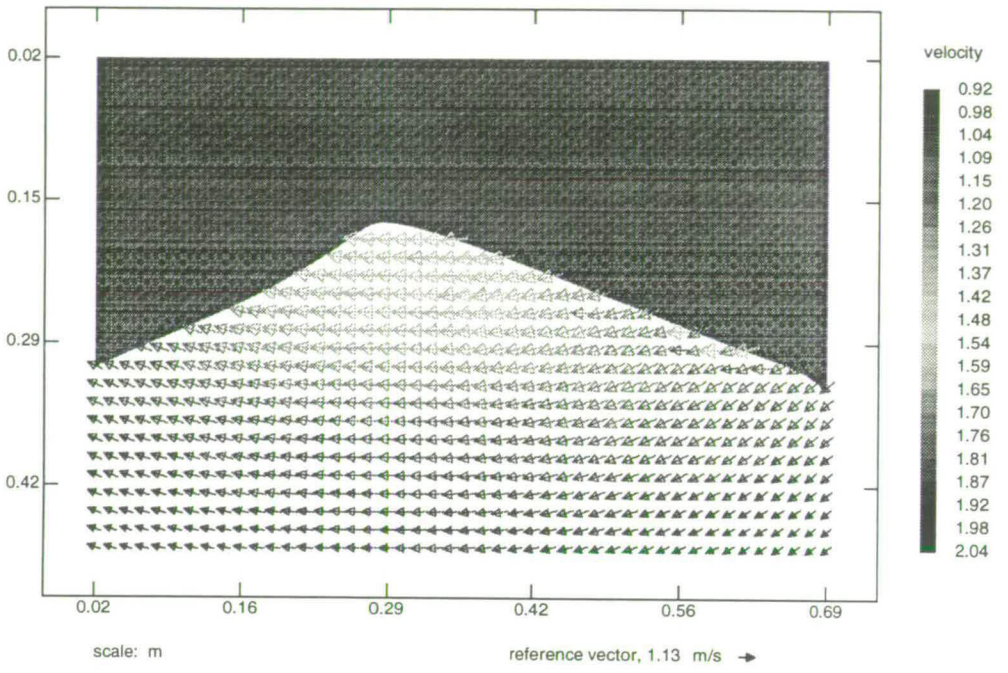

Figure 5.6: Spilling breaker

ambiguities to be resolved. It is in cases like these that the use of a system which is capable of analysis by cross-correlation of separate images is useful.

The removal of the direction ambiguity, the perspective distortions introduced by the traditional image shifting systems, and the larger dynamic range afforded by the technique offset in this case any problems associated with low spatial resolution.

Figure 5.6 shows a spilling breaker formed by the creation of a wave train of different frequencies and amplitudes. The waves have been calculated so that they combine at a pre determined focus point, resulting in the breaking action observed. The equipment used in this study was designed and implemented by Skyner [57] for his thesis work.

5.4 Study of Capilliary Waves

Work being carried out by Pullen [40] at the University of Edinburgh is investigating the effects of surface active films on capilliary waves, and the interaction between wind and wave under these conditions. Capilliary waves are generally considered to be waves in which the predominant resoring force acting is the surface tension, and not gravity. Thus, in the case of water capilliary waves in general have a wavelength of under 17mm. For this study it was decided that PIV should be used to investigate the velocity distribution directly below the surface of the water. This investigation cannot use single frame, multiple image analysis due to the surface of the wave passing over the previously exposed areas, destroying any information that was stored there. Therefore, the two-CCD array system proved to be an ideal solution.

The experimental facility used in this study is essentially a miniature replica of the narrow wave flumes of which there are a number at the University of Edinburgh. The flume is capable of generating a number of harmonic waves, and the velocity maps shown in figures 5.7 & 5.8 are of 3Hz waves, with a wavelength of 10cm and an amplitude of 2cm. The illumination for this experiment was supplied by a 600mW Nd:YAG continuous wave laser and a scanning beam system. The laser is one of a new type of small, powerful lasers which use a Nd:YAG rod pumped by a laser diode. This allows a good beam profile and long coherence lengths to be generated in a small and efficient package.

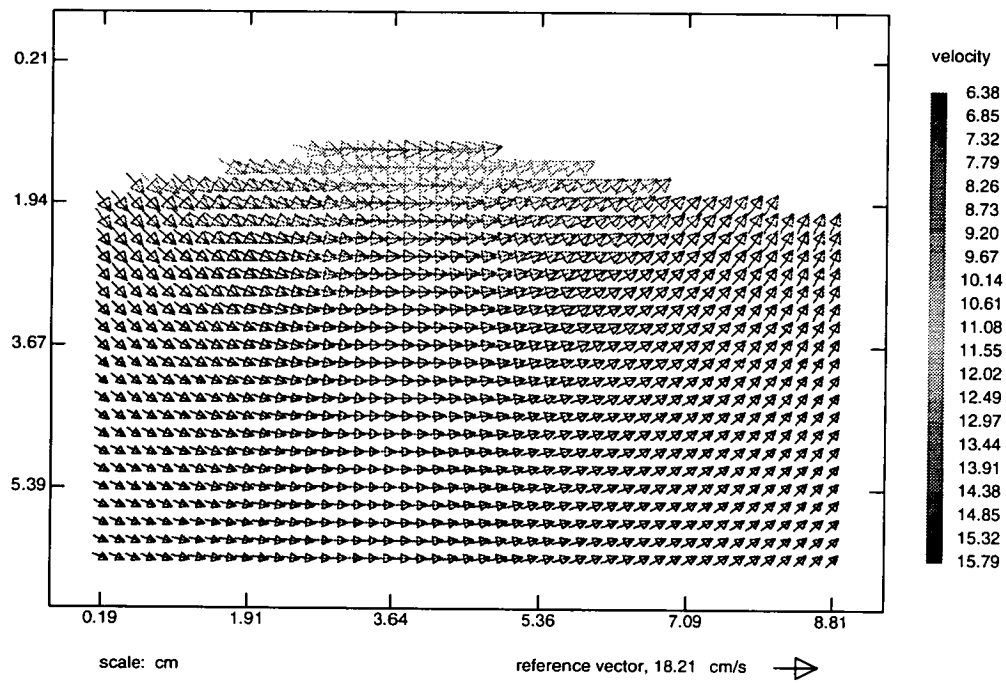

Figure 5.7: Small amplitude wave, no surface film

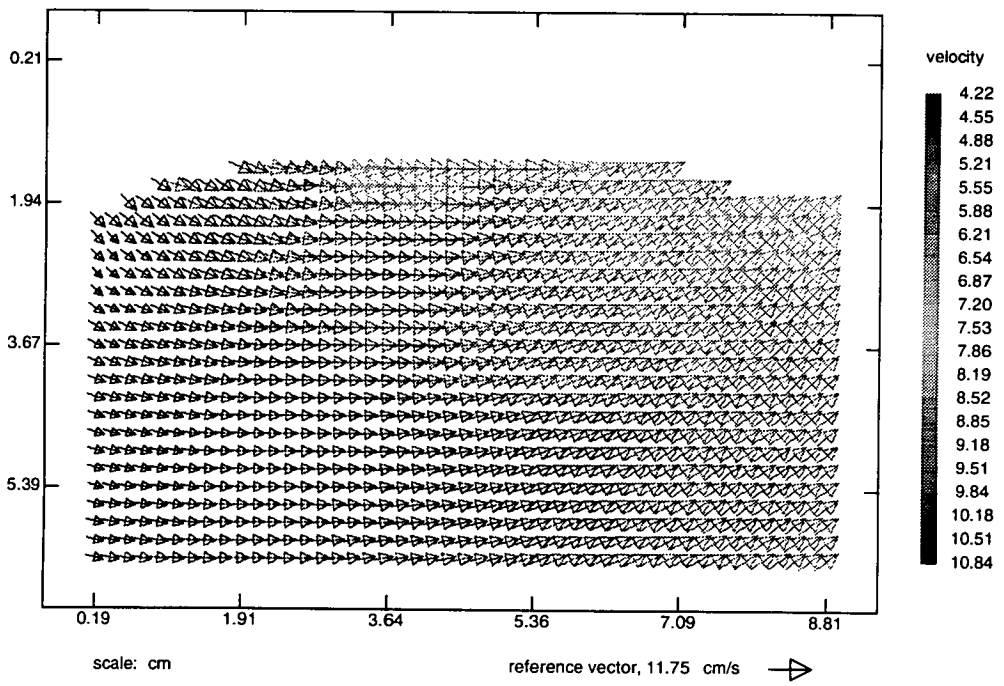

Figure 5.8: Small amplitude wave, surface film of detergent

Chapter 6

Acceleration Measurements

6.1 Motivation

> We notice that the [Navier-Stokes] equation is a non-linear partial differential equation in $\underline{u}$. The non-linearity arises from the dual role of the velocity in determining the acceleration of a fluid particle. This non-linearity is responsible for much of the mathematical difficulty of fluid dynamics, and is the principal reason why our knowledge of the behaviour of fluids in motion is obtained in many cases from observation (both of laboratory experiments and of natural phenomenon) rather than from theoretical prediction.
>
> *Physical Fluid Dynamics, D.J. Tritton, pp59*

The concept of non-dimensional scaling is a familiar one in fluid mechanics, with a number of non-dimensional parameters existing to help class fluid flows based on geometric similarities (Reynold's number), frequency of vortex shedding (Strouhal number), convection studies (Grashof and Prandtl numbers), oscillatory flows (Keulegan-Carpenter number) and free surfaces on which waves can develop (Froude number).

The reason these non-dimensional quantities are of such importance is that

they allow the prediction of some of the fluid mechanical effects which will be present when, for example, a large scale offshore structure is transferred from model to production. One of the key areas of interest however is in the prediction of forces acting on the offshore structure induced by the action of fluid flow, and this is an area which has seen a tremendous amount of research and growth since the 1940's.

It is in this area that one of the most important uses of Particle Image Velocimetry can be found. Designers often use the velocity fields present around model structures to predict the real pressure distributions and overall forces, rather than pressures or forces measured directly from the model structure itself as the highly non-linear nature of the force and pressure measurements creates difficulties in scaling from model to prototype. An example of this would be in the estimation of the wave forces acting on ocean structures [10].

Morison's equation (Eq. 6.1) is perhaps the most widely used formula for estimating the forces acting on objects submerged in a fluid flow, with the result being a force per unit length adjusted by the two empirical *force coefficients* C_m and C_d.

$$f = C_m \rho \Pi R^2 \dot{u} + C_d \rho R u|u| \tag{6.1}$$

It is known that simple Froude scaling laws can be used to predict velocities under waves in the ocean from small scale model experiments [53] whereas prediction of the force directly requires very large scale experiments and in doing so the choice of many of the scaling parameters adds to the sources of uncertainties [7], in particular the choice of C_m and C_d are still the subject of numerous studies and papers [36, 38, 2], with, for example, an investigation by Allen *et. al.* using a basin 21'x16'x2968' to tow a cylinder 79"x3.5" at speeds of up to 85'/second,

thereby realising Reynolds numbers in water of $2x10^6$. These experiments were carried out in order to test the effects of tidal currents acting on vertical oil platform drilling risers.

If a large scale experiment is conducted, then a series of force transducers can supply the necessary force information that the designers require. However, the production of a full scale prototype is in most cases not feasible. Laser Doppler Anemometry (LDA) has been used in the study of forces acting on submerged structures, but LDA can only be reliably used in cases where we have either a flow with a convective element that can be neglected, or a precisely repeating flow, in which case the LDA measurement point can be traversed across the entire experimental section. Applying LDA in this fashion is of course extremely time consuming and difficult, as entirely repeatable flow situations are difficult to create. The use of LDA would also exclude the investigation of transitory events, which are often of more interest to the designer.

It is the ability to measure the *total* acceleration and velocity simultaneously, then, which is desirable, as it is the total acceleration which is of most interest in the field of offshore structures; it has been shown that if only the temporal derivative of the velocity is used, the acceleration measurement can be 10-15% lower than the total derivative of the velocity [54], which if following the Eulerian flow definition is defined as

$$\frac{D\underline{u}(\underline{x},t)}{Dt} = \frac{\partial \underline{u}(\underline{x},t)}{\partial t} + \frac{\partial \underline{u}(\underline{x},t)}{\partial \underline{x}} \cdot \underline{u}(\underline{x},t) \tag{6.2}$$

As it is the total derivative of the velocity which is used in the computation of Morison's equation, the inclusion of the convective terms is essential if the correct results are to be obtained.

The construction of a camera which can capture four images of a flow in rapid succession makes it possible to make two velocity measurements, and therefore

estimate the acceleration present in the flow. This is the aim of the rest of the chapter.

6.2 Prediction of Force

The force or pressure which acts on a structure can be gained from the Navier-Stokes equation describing (in the case of a water wave) an incompressible fluid (equation 6.3).

$$\rho \underline{a}(\underline{x},t) \equiv \rho \frac{D\underline{u}(\underline{x},t)}{Dt} = -\nabla p(\underline{x},t) + \mu \nabla^2 \underline{u}(\underline{x},t) + \underline{F} \tag{6.3}$$

with ρ the density of the fluid, μ the viscosity, $p(\underline{x},t)$ the pressure, $\underline{F}$ the volume force (in the following case gravity) and $\frac{D\underline{u}(\underline{x},t)}{Dt}$ or $\underline{a}(\underline{x},t)$ the *total* acceleration.

As $\mu \nabla^2 \underline{u}(\underline{x},t)$ is the viscous force per unit volume, this will become significant only in the boundary layer next to the vertical wall. While PIV has been used in investigations of boundary layers, in the experiments which follow the spatial resolution was not sufficient to resolve the velocities present in the boundary layer, and so $\mu \nabla^2 \underline{u}(\underline{x},t)$ can be neglected in the calculation of the pressure. This now means that the pressure in the fluid will be related to the acceleration by a first order linear differential equation (equation 6.4).

$$p(z_1,t) - \bar{p}(z_2,t) = \rho \int_{z_1}^{z_2} (\underline{a}(\underline{x},t) - \underline{g}) dz \tag{6.4}$$

In order to provide some experimental validation to the acceleration measurements, the pressure derived from the PIV images and the pressure measured from the pressure transducers placed in the front face of the wall will be compared.

6.3 Calculation of Accelerations

6.3.1 Background

In fluid mechanics there are two distinct approaches to the measurement of flow quantities; the *Eulerian* and *Lagrangian* specifications.

The first, the Eulerian specification, requires that the observer remains at a fixed point in space relative to the flow under investigation. This is the more commonly used coordinate system, as it allows the formulation of the equations of flow more easily, for example yielding the instantaneous velocity immediately, and it is this notation which is used throughout the rest of this chapter.

The second, the Lagrangian specification, requires that the observations of the flow are taken while traveling with a well defined particle of the fluid itself. This coordinate system is inherently linked with a number of flow visualisation schemes, such as dye or smoke streaks [58], but it is difficult to relate the Lagrangian approach to the equations formed by the Eulerian method.

The total acceleration is calculated as in equation 6.2 for the Eulerian specification and by

$$\frac{D\underline{u}(\underline{x},t)}{Dt} \simeq \frac{\Delta\underline{u}(\underline{x},t)}{\Delta t} \tag{6.5}$$

in the Lagrangian specification. Both should yield the same result as Δt approaches zero.

The motivation to use the Lagrangian method can be clearly seen, as the calculation of the total derivative of the velocity can be achieved by making two velocity measurements on the same fluid particle, and hence to the first order the acceleration is the difference between the velocity measurements. It is the simplicity and inherent accuracy of the Lagrangian method which influenced the decision to develop the algorithms required to track a fluid volume. The Eulerian

method can be seen on the other hand to be simple to implement, but it has two areas of approximation present, the temporal *and* spatial gradient calculations, making the Eulerian implementation perhaps less accurate and more susceptible to rounding errors than the Lagrangian. However, the requirement to track a fluid particle has been removed, and so the Eulerian approach should be more robust.

6.3.2 Bias effects

In PIV, one of the defining criterion for the accurate measurement of velocities is that the time separations should be small enough to make the accelerations (i.e. spatial gradients) present within interrogation areas negligible. A great deal of work has been done into the limiting cases, and a study of the implications of increasing spatial gradients has been carried out most notably by Keane and Adrian [27, 28, 29]. As the analysis of PIV images is a statistical process, the introduction of large spatial gradients in the interrogation area changes the shape of the correlation peak to a broadened, flattened shape. This not only has the effect of increasing the probability of measuring a rogue noise peak instead of the correct correlation peak, *detection bias*, but it also gives rise to an effect termed *gradient bias* following the terminology of Keane and Adrian. This has the effect of biasing the velocity measured towards a lower value due to the effects of loss of image pairs.

In order to allow the measurement of acceleration, we need two accurate velocity measurements, which will have to be sufficiently different to allow the calculation of an acceleration without large errors being introduced. This need to measure the velocity accurately implies that the time separating the velocity measurements should be large in comparison to the time required to gain a velocity measurement.

6.3.3 Calculating Accelerations

As described in section 1.2, it is typical for the raw PIV images to be divided up into a grid of interrogation areas, with each grid location being analysed by performing a local correlation. The selection of the size of the interrogation area depends on the highest spatial frequencies present within the flow, as the presence of any large velocity gradients will reduce the reliability of the measurement. Therefore the size of the interrogation areas has to be a trade off between sufficient particle image density and minimising spatial gradients, and to avoid the introduction of large biases, the displacement is also kept below a quarter of the interrogation area [27].

If the above guidelines are followed, the interrogation areas which are selected during analysis can essentially be considered as being fluid volumes in the Lagrangian sense, as the velocities present within the area should be essentially constant. Thus, the act of making a PIV measurement in intrinsically Lagrangian, as we are essentially using pattern matching to track a fluid volume. However, the grid wise analysis produces a regular set of velocity vectors, and in the majority of cases this is taken to be the same as the Eulerian velocity field. This cannot be correct though, as only a first order approximation to the velocity can be made from two images. This requires therefore that an adjustment be made to the location of the calculated velocity vector.

6.3.4 Timing

As previously stated, the time separation between the velocity measurements has to be large to allow the accelerations to be measured reliably, while allowing accurate velocity measurements to be made. To allow two velocity measurements to be made, either three or four images of the flow have to be taken in quick succes-

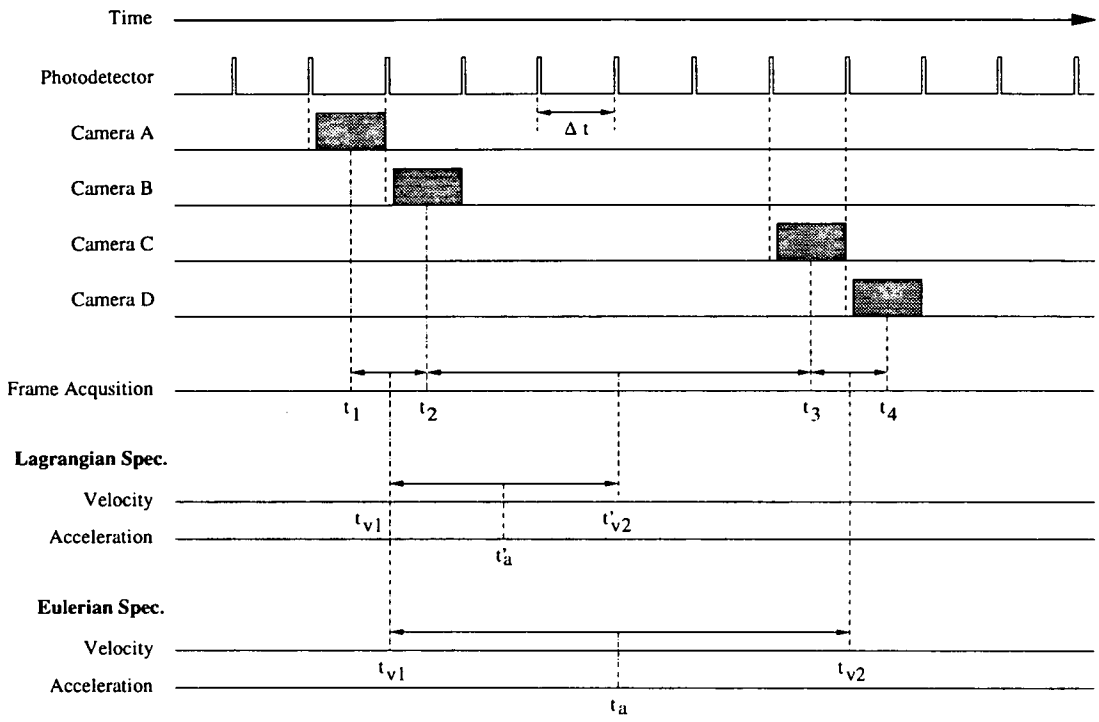

Figure 6.1: Timing of camera exposures for acceleration measurements

sion, depending on which flow specification we are following. As the preliminary experiments were carried out in water flows, the flow was illuminated with a scanning beam system. This allowed the synchronisation of the entire experiment to the scan frequency of the rotating polygon mirror by use of a photo diode placed on top of the scanning beam box. It was therefore convenient to define the time separation between the velocity measurements in terms of an integer number of scans of the scanning beam system.

The acceleration field was calculated by acquiring four single exposure images of the flow for both the Eulerian and Lagrangian specifications. The four frames are from now referred to as A, B, C and D and were acquired at times t_1, t_2, t_3 and t_4 respectively. The first velocity measurement, which is the same for both the Eulerian and Lagrangian specifications, is made by cross correlating between images A and B, and so the time separation between these images will be one

scan period of the illumination system, Δt. The relationship between the four images and their times of acquisition are show in figure 6.1 and are characterised by equation 6.6

$$\Delta t \equiv t_2 - t_1 = t_4 - t_3 \ll t_3 - t_2 = n\Delta t \tag{6.6}$$

The four camera system (Chapter 3) allows the image acquisition, and the exposure time of the arrays is set such that each array only receives one complete scan of the flow.

Some more definitions now need to be made to allow reference to the times at which the velocities measurements were made. As we are making a measure of the local velocity with only a single displacement, we have made a first order approximation to the instantaneous velocity at that point in the flow. Thus, the only information we can gain about the time which this velocity belongs to is that it must lie exactly between the time of the first and second images. Therefore, the time of the first velocity measurement is t_{v1} which is:

$$t_{v1} = \frac{t_1 + t_2}{2} \tag{6.7}$$

As this velocity is the same for both of the specifications, this term will be used in both of the resulting equations describing the total acceleration.

The next velocity measurements now are strongly dependent on the flow specification, and the interrogation areas have to be defined accordingly.

Eulerian

In the case of the Eulerian specification, each interrogation areas is kept at a constant grid location throughout the analysis. This is illustrated by figure 6.2. The dotted squares represent the location of the particle groups that were present in the interrogation area at the time of capture. As we are now going to make a

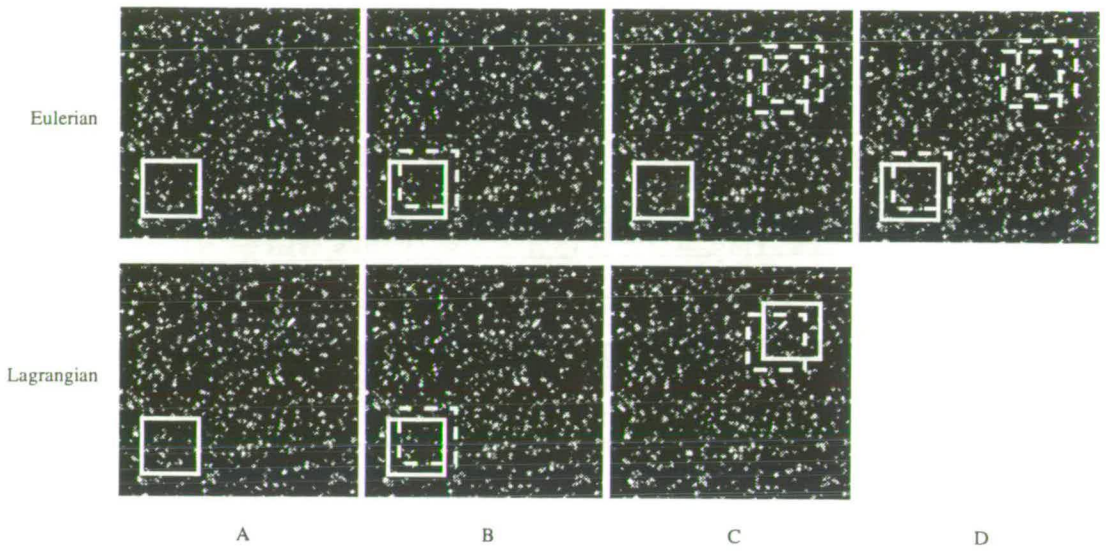

Figure 6.2: Location of interrogation areas for acceleration measurements

second velocity measurement at the same grid location, we use images C and D, again with a time separation of Δt, giving t_{v2}:

$$t_{v2} = \frac{t_4 + t_3}{2} \tag{6.8}$$

$$t_{v2} - t_{v1} = (n+1)\Delta t \tag{6.9}$$

The inter velocity separation is as defined in equation 6.9, and so now the Eulerian acceleration field can be calculated.

Lagrangian

The Lagrangian specification requires that we attempt to follow the same group of particles throughout the entire measurement, and so we have to move the interrogation area accordingly. In order to attempt this, we have to use the velocity measurement gained from the first two images A and B to move the interrogation area dynamically, allowing us to then correlate images B and C. As long as the acceleration has not caused the particle group to dissipate completely, the correlation should then reveal the true location of the original particle group.

The second velocity measurement, t'_{v2} occurs at:

$$t'_{v2} = \frac{t_2 + t_3}{2} \tag{6.10}$$

$$t'_{v2} - t_{v1} = \frac{(n+1)}{2}\Delta t \tag{6.11}$$

with the dash indicating that this is the time relating to the Lagrangian specification. Again, figure 6.1 shows the relative position of the acquisition of the images.

Time of Acceleration Measurement

In both specifications, the time of the acceleration measurement is assumed to be between the first and second velocity measurements, as again we are making a first order approximation to the instantaneous acceleration within the flow at either that grid point (Eulerian) or the group of particles that is being followed (Lagrangian). Again, using a dash to denote a Lagrangian quantity, the time of the acceleration measurements is found to be:

$$t_a = \frac{t_{v1} + t_{v2}}{2} \tag{6.12}$$

$$t'_a = \frac{t_{v1} + t'_{v2}}{2} \tag{6.13}$$

6.3.5 Location of Particle Group

The location of the particle group in the Lagrangian specification is also of importance, as at the time which best describes the acceleration measurement, the particle group will have moved a considerable distance from its original location. The centre of the particle group will be defined as $\underline{r}_{i,j}(t)$ at time t, with (i,j) denoting the location of the interrogation area in which the particle group was in image A. This corresponds well to the centre of mass of the fluid volume, and assuming the dissipation of the particle group is kept to a minimum over the

measurement, this should continue to be the case throughout the measurement. In the same fashion as the definition of the location of the particle group, the velocity is defined as being $\underline{v}_{i,j}(t)$

6.3.6 Definition of Accelerations

Eulerian

As we now have two velocity measurements at the same grid location in the fluid, we can now move onto calculating the acceleration at that spatial location. Because we have only two velocity measurements, the total acceleration (equation 6.2) can only be calculated to the first order, and so the accelerations must be kept reasonably linear over the measurement to ensure the correct result.

If we split the total acceleration into its two component parts, the spatial and convective terms, we can find their approximations. Assuming the temporal acceleration is approximately constant during the measurement, we can set a bound on the magnitude of the acceleration:

$$\left| \frac{\partial \underline{u}(\underline{x}, t_a)}{\partial t} - \frac{1}{2} \left(\frac{\partial \underline{u}(\underline{x}, t_{v2})}{\partial t} + \frac{\partial \underline{u}(\underline{x}, t_{v1})}{\partial t} \right) \right| \ll \left| \frac{\partial \underline{u}(\underline{x}, t_a)}{\partial t} \right| \tag{6.14}$$

and if this is held, the temporal acceleration is simply:

$$\frac{\partial \underline{u}(\underline{x}, t_a)}{\partial t} = \frac{\underline{u}(\underline{x}, t_{v2}) - \underline{u}(\underline{x}, t_{v1})}{t_{v2} - t_{v1}} \tag{6.15}$$

where $\underline{u}(\underline{x}, t_{v1})$ and $\underline{u}(\underline{x}, t_{v2})$ represent the measured fluid velocities at location $\underline{x}$ at times t_{v1} and t_{v2} respectively.

Now considering the spatial or convective term of the acceleration measurement, if we assume that the grid spacing of the velocity measurements are small enough to allow the spatial gradients to vary only slightly between neighbouring locations, we can therefore assume the variation will be linear, and so we can

define another bound for the spatial terms:

$$\left|\frac{\partial \underline{u}(\underline{x},t_a)}{\partial x_j} - \frac{1}{2}\left(\frac{\partial u_i(\underline{x}+\Delta\underline{x}_g,t_a)}{\partial x_j} + \frac{\partial u_i(\underline{x}-\Delta\underline{x}_g,t_a)}{\partial x_j}\right)\right| \ll \left|\frac{\partial \underline{u}(\underline{x},t_a)}{\partial x_j}\right| \quad (6.16)$$

with $\Delta\underline{x}_g$ being the spacing between neighbouring grid locations. The indices run over i,j,k=1,2,3. The spatial gradient term can now be approximated to the first order as:

$$\underline{u}(\underline{x},t_a).\nabla\underline{u}(\underline{x},t_a) = \frac{1}{2}\sum_{i=1}^{N}\underline{e}_i\sum_{j=1}^{N}\left(\frac{\partial u_i(\underline{x},t_{v1})}{\partial x_j}.u_i(\underline{x},t_{v1}) + \frac{\partial u_i(\underline{x},t_{v2})}{\partial x_j}.u_i(\underline{x},t_{v2})\right) \quad (6.17)$$

The partial velocity gradients in equation 6.17 are calculated from each of the velocity maps, and are then averaged together to give a measure of the spatial gradients present at time t_a. The velocity gradients are calculated by fitting a least squares surface to the vector map locally, and then differentiating in each direction, to give a result which should be more tolerant of small differences in vectors than for example a central differences scheme.

As we are only considering an image plane, the number of vector directions we have to sum over is 2, and so N=2 in the above equations.

Lagrangian

The Lagrangian acceleration (equation 6.5) can now also be calculated from the two velocity measurements, and we are assuming we have succeeded in following a group of particles. Again, if we assume that the acceleration of the fluid volumes is constant within the time of the measurement, we can place a bound on the accelerations present within the flow:

$$\left|\underline{a}_{i,j}(t'_a) - \frac{1}{2}\left(\underline{a}_{i,j}(t'_{v1}) + \underline{a}_{i,j}(t'_{v2})\right)\right| \ll \left|\underline{a}_{i,j}(t'_a)\right| \quad (6.18)$$

If we then assume equation 6.18 is held, the total acceleration of the fluid

volume will be:

$$\underline{a}_{i,j}(t'_a) = \frac{\underline{v}_{i,j}(t'_{v2}) - \underline{v}_{i,j}(t_{v1})}{t'_{v2} - t_{v1}} \tag{6.19}$$

As we have assumed that the acceleration is going to be nearly constant over the measurement time, we can now find the location of the centre of mass of the fluid volume at time t_a:

$$\underline{r}_{i,j}(t'_a) = \frac{1}{8}\left(\underline{v}_{i,j}(t'_{v2}) + 3\underline{v}_{i,j}(t_{v1})\right)(t'_{v2} - t_{v1}) + \underline{r}_{i,j}(t_{v1}) \tag{6.20}$$

with the position of the fluid volume at $\underline{r}_{i,j}(t_{v1})$ being:

$$\underline{r}_{i,j}(t_{v1}) = \underline{v}_{i,j}(t_{v1}).\frac{(t_2 - t_1)}{2} + \underline{r}_{i,j}(t_1) \tag{6.21}$$

6.4 Uncertainty in Acceleration Measurements

The calculation of quantities based on the subtraction of two nearly equal measured values implies a potential for having a high sensitivity to both random and systematic errors which can be present in the PIV measurement technique. Generally several noise sources contribute to the uncertainty present in the predicted acceleration measurements, and these have been covered in great detail by a number of authors [27, 62]. These can include the effects of noise in the correlation plane due to random correlations , and noise introduced in the imaging process, for example optical aberrations, noise in the CCD sensors, non-uniform light response in the CCD sensors, charge leakage and dark current. The majority of these effects, however, only become significant under very low illumination levels, when these non-uniform effects of the CCD sensors can begin to dominate over the captured image of the flow field.

If we assume that the net result of all of these effects is an uncertainty Δs in the measurement of the displacement of a particle group, and that Δs is isotropic

and also independent of the displacement [1] then we can construct two expressions for the uncertainty of the acceleration measurements. The uncertainty of the *i'th* component of the Lagrangian acceleration is therefore:

$$\Delta a_i = \frac{2\sqrt{1+n^2}}{(n+1)n}\frac{\Delta s}{\Delta t^2} \tag{6.22}$$

while the uncertainty of the *i'th* component of the Eulerian acceleration is:

$$\Delta a_i = \sqrt{\frac{2}{(n+1)^2} + \frac{1}{2}\left(\frac{\langle|\underline{u}|\rangle\,\Delta t}{\Delta \underline{x}_g}\right)^2}\frac{\Delta s}{\Delta t^2} \tag{6.23}$$

with $\langle|\underline{u}|\rangle\,\Delta t$ being the mean displacement measured from the two velocity maps and $\Delta\underline{x}_g$ the distance between neighbouring vectors. Furthermore, it has been assumed that the values of $\langle|\underline{u}|\rangle\,\Delta t$ will be much larger than the spatial velocity gradient. The expressions give both of the uncertainty measurements for the calculated accelerations a hyperbolic relation to the time interval between the two velocity measurements (i.e. n), while the Eulerian expression appears to have a minimum noise floor.

6.5 Validation of Methods

The criteria which have been developed in the previous sections to ensure accurate and reliable acceleration measurements rely in their basic form on a high level of knowledge of the flow situation which is present in the experiment. As one of the main attractions of PIV is its ability to capture instantaneous, non-repeatable events, *a priori* knowledge of the flow is more often than not not available. In this case where only experimentally obtained data is available additional methods must be developed for validation of the vector maps.

In order to validate the acceleration prediction techniques described, an analytical model of standing water waves in a cavity has been developed. The model

[1] This assumption does not hold under idealised simulations, however, for the measurement of real flows, it is considered to be appropriate

is based on linear wave theory [54], which is restricted to waves of small amplitude, compared to their wave length and the still water depth, and does not include viscosity. Considering a co-sinusoidal wave of amplitude A and wavenumber k in water of depth h, we can find the free surface elevation above the still water level, η by:

$$\eta = A\cos(kx - \omega t) \tag{6.24}$$

with ω being the angular frequency of the wave. The horizontal and vertical velocity components can be found from this expression by applying potential theory, leading to:

$$u = \frac{Agk}{\omega}\frac{\cosh k(z+h)}{\cosh kh}\cos(kx - \omega t) \tag{6.25}$$

$$v = \frac{Agk}{\omega}\frac{\sinh k(z+h)}{\cosh kh}\sin(kx - \omega t) \tag{6.26}$$

with z being measured from the still water level. If we now place a plane, vertical, perfectly reflecting wall at $x = b$ we can make the assumption that the amplitude of the reflected wave will be the same as the incident wave. The wavenumber and frequency of the reflected wave are also taken to be identical to the incident wave. This then gives a surface elevation of:

$$\eta_s = 2A\sin(kb - \omega t)\cos(kx - kb) \tag{6.27}$$

with the horizontal and vertical velocity components of this standing wave equal to:

$$u_s = \frac{2Agk}{\omega}\frac{\cosh k(z+h)}{\cosh kh}\cos(kb - \omega t)\sin(kx - kb) \tag{6.28}$$

$$v_s = \frac{-2Agk}{\omega}\frac{\sinh k(z+h)}{\cosh kh}\cos(kb - \omega t)\cos(kx - kb) \tag{6.29}$$

PIV images are simulated based on the velocity fields produced by the model. The parameters for the model are matched to the experimental set-up for the

standing wave measurements, being described in section 6.6. The standing wave model considers only single harmonic solutions for the simulation of the PIV images, while a solution containing the second harmonic as well (table 6.1) is used for comparison with the experimental results.

Two phase points are of special interest to be considered in both the simulations and experiments. The first of these is immediately after the surface elevation has passed the still water level. The conditions present during this case give a minimum acceleration and maximum temporal changes in the acceleration field, and so this should give a good test of the resolution of the method and indicate how far the criteria, of having a constant acceleration during the velocity measurements, can be pushed. The second phase point of interest is immediately before the surface elevation reaches its minimum value. In this case the conditions provide a maximum acceleration and minimum temporal changes of the same. This will give an indication as to when the method based on the Lagrangian specification breaks down due to distortion of the fluid volume. The phase points will be specified further in section 6.7.

6.6 Experimental Setup

As described above, the initial experimental setup involved the measurement of acceleration fields under waves impinging against a flat, vertical, perfectly reflecting wall.

The wall itself is made of perspex, and housed three pressure transducers which were placed at three equally spaced locations down the wall, at 30mm above the mean water level, 20mm and 70mm below. All of the pressure transducer amplifiers have low pass filters with a cut-off frequency of 100Hz, and so no components of the experiment should be lost. There was also a flat film wave

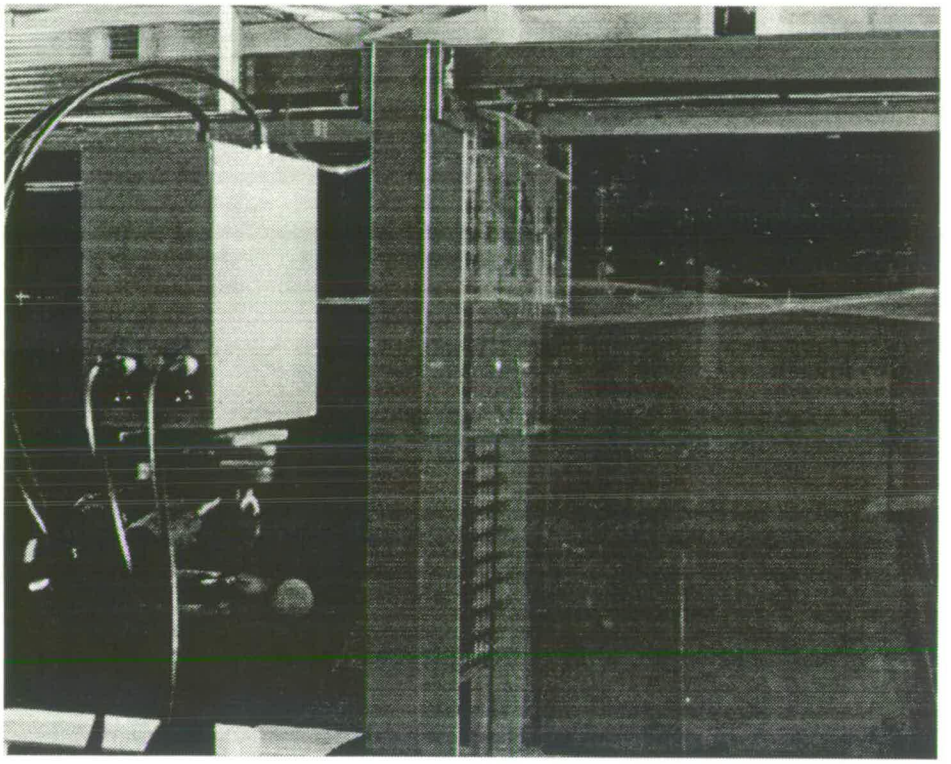

Figure 6.3: Experimental setup showing vertical wall

gauge placed on the front of the wall to allow an accurate record of the wave height to be recorded. All of the outputs from the transducers were recorded by the camera controlling system, with the results being sampled at 800 samples per second and logged to a PC. The instance at which the cameras were triggered was also recorded, allowing a temporal reference point between PIV data and transducer data. The camera, wall and wave flume can be seen in figure 6.3.

The timing equipment for the four-CCD array system is identical in logic to the system described in Chapter 5. In this case however we need also to acquire data from the transducers previously mentioned. For this a computer is required, and so as a result of this it was decided to build a system which contained all of the components for the four-CCD array system, and also modules for data acquisition.

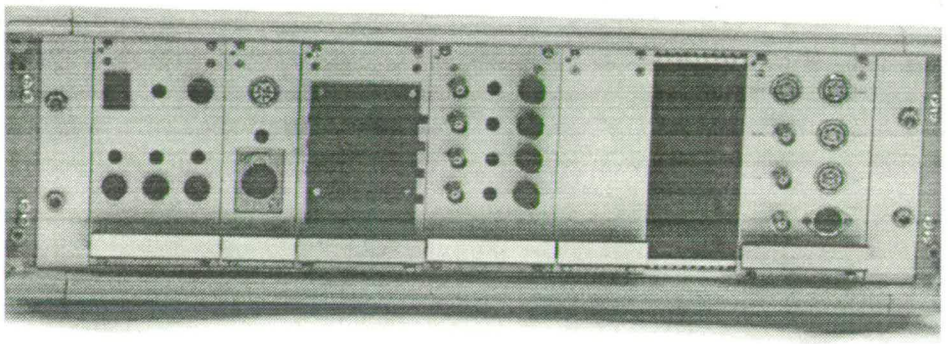

Figure 6.4: Timing electronics

The resulting system has twelve 16-bit counter/timers, giving three for each array. The timers are programmed by the general purpose digital and analogue I/O card used for logging the transducer data. The counter/timer chips used allows the system to be used with either a pulsed laser system or a scanning beam system. The resulting system is housed in a 19" rack, and can be seen in figure 6.4. The modular design allows any type of data acquisition to be configured, with the results then being broadcast onto the data bus and subsequently acquired by the I/O card required to run the system. An example of the modules can be seen in figure 6.5. As this timing system logs the transducer data, it was a simple task to trigger the entire system from a particular fluid dynamic event.

The bulk of this timing and data acquisition system was designed and built by Michael Jakobsen, with the author contributing only to the design and construction of a few plug in modules.

The experiments were carried out in a wave flume of dimensions 10m x 48cm with a designed water depth of 75cm. At one end of the wave flume there is a hinged paddle type wave maker which is computer controlled to allow the

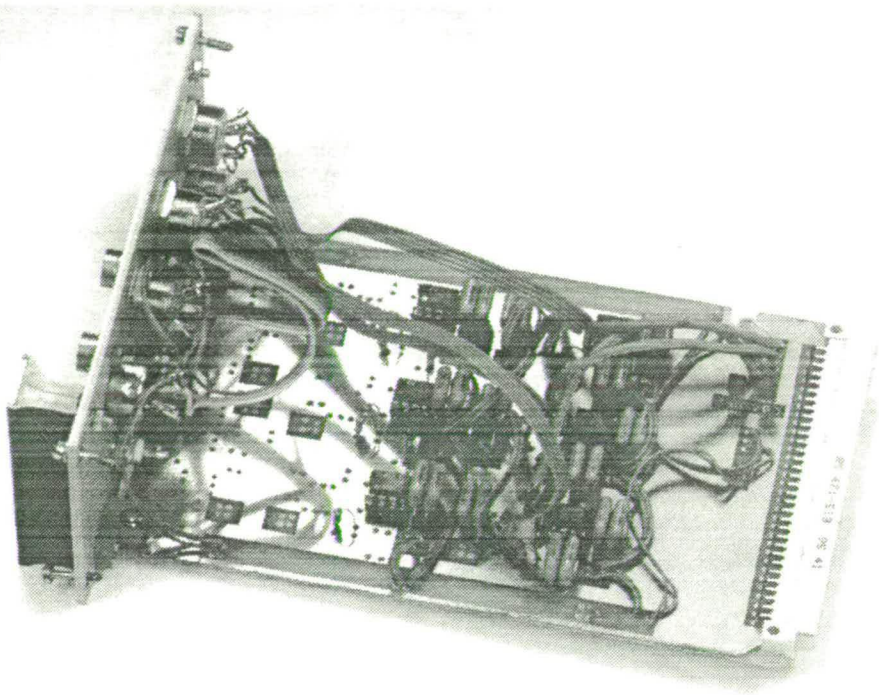

Figure 6.5: Plugin modules

production of a wide range of frequencies and amplitudes, with the ability to produce both shallow and deep wave conditions. The paddle system can build up various types of wave structures by superimposing a number of different frequency components allowing both regular and irregular waves to be produced. The wave maker also acts as an active wave damper, allowing all of the reflections produced by the vertical wall to be absorbed. At the other end of the flume a number of options can be placed, including absorbent foam, beaches of various slopes, or indeed a vertical wall. For these experiments, the wall was located 6.38m from the paddle.

Two types of wave were investigated during this experiment, a harmonic wave and a freak wave. The freak wave is a wave train composed of different frequencies, the phases of which have been timed to superimpose constructively just in front of the vertical wall.

The camera used in these experiments was the four-CCD array system described previously, fitted with a Nikon Micro-Nikkor 55mm $f2.8$ lens. As the light scattered from the seeding in the flow is divided equally between the arrays, the resulting light intensity at each array is low. However, the lens could be used at its maximum aperture without any noticeable distortion or defocusing effects.

The flow was illuminated with a scanning beam system, fed by a Spectra Physics 15W Argon-Ion laser. The scanning beam box used supplied a light sheet 1m wide, however only a tenth of this possible field of view was used. The scan rate of the system was set to 5ms per scan.

The camera was arranged such that the long edge of the arrays were vertical. For the experiments with a harmonic wave, the cameras field of view lay immediately below the still water surface and next to the vertical wall. For the freak wave experiments, the field of view was moved upwards, with the still water level now lying 85% of the way down the array. The physical area contained within each PIV image is 0.11m(H)x0.15m(V). Figure 6.6 shows the relative positions of the wall, pressure transducers and wave gauge, and the location of the cameras field of view.

The timing controller was setup up to monitor the signal from the wave gauge placed on the front of the wall, and set to trigger when the water surface passed the still water level moving from trough to crest. This should be the point during the wave cycle at which the surface is moving the fastest, and so should provide an accurately repeatable trigger point. In order to allow images to be acquired at any point during the wave phase, an adjustable phase delay was added to the trigger signal before it was passed onto trigger the CCD arrays. Once the system was triggered, the first image, A was taken during the next occurring scan, with image B being taken on the second scan, and images C and D being taken during

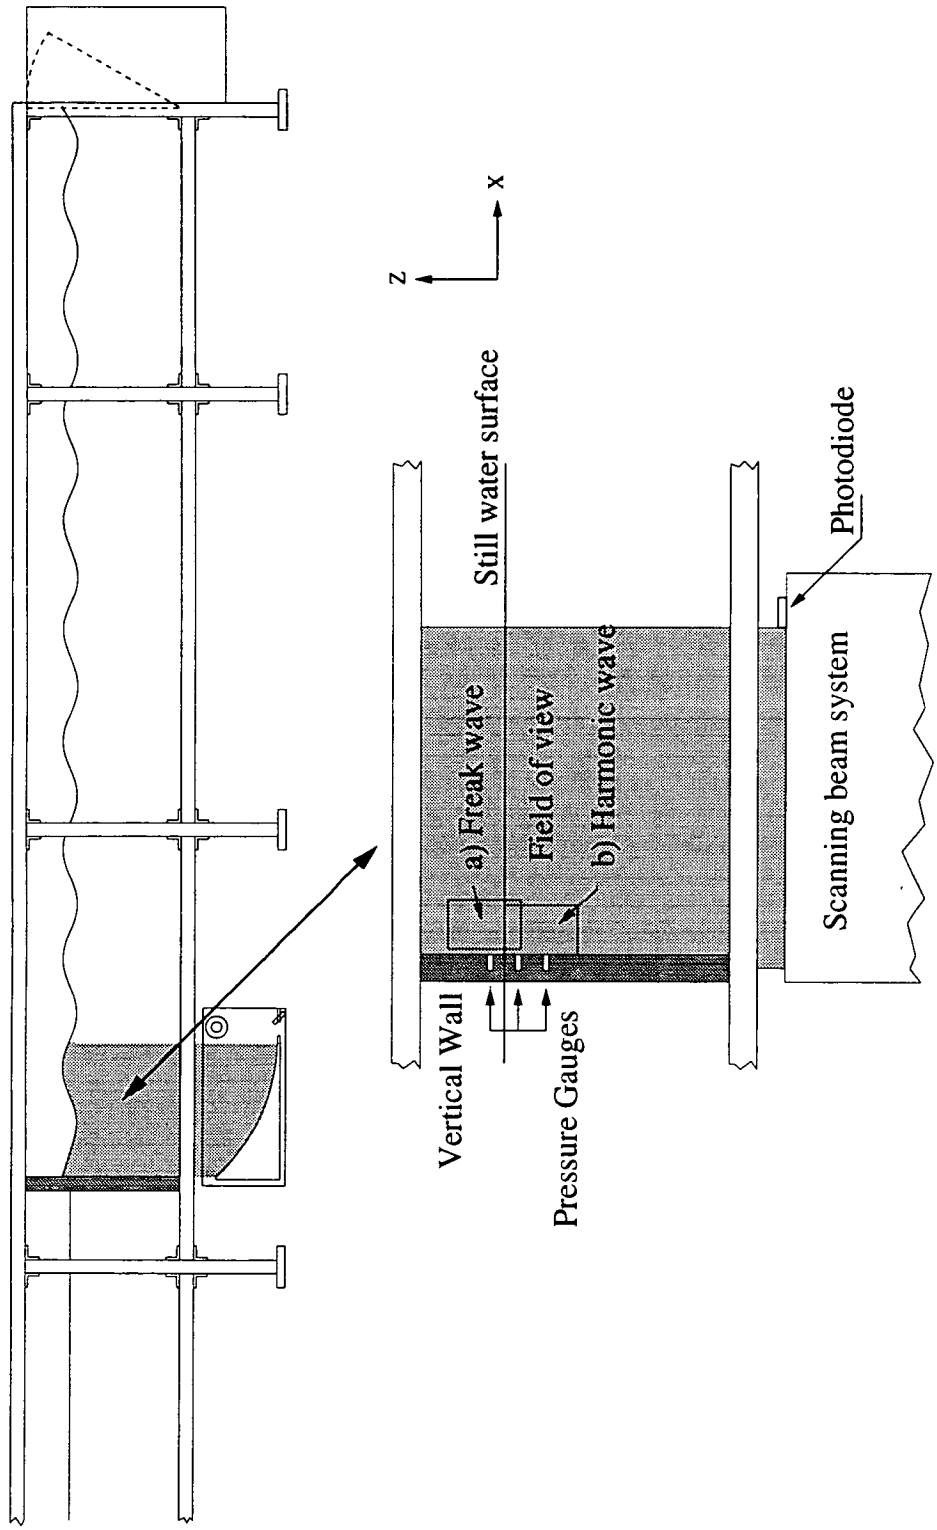

Figure 6.6: Schematic of vertical wall showing relative locations of fields of view and pressure transducers

Harmonic	Amplitude Normalised	Phase Delay
First Order	1.0	0.0
Second Order	0.081	2.76
Third Order	0.014	-0.70

Table 6.1: Components of standing wave

scans $n+1$ and $n+2$ respectively. The parameter n was set between 1 and 24, depending on the wave structure.

6.7 Harmonic Results

In the harmonic case, the presence of the vertical reflecting wall set up standing waves in the wave flume. The period of the waves was set to $1.00s \pm 0.2\%$ giving a wavelength of 1.56m and an amplitude of 60.7mm$\pm$0.5%. The PIV images are acquired at two phase points of the standing wave. The first phase point, a, is located 5ms after the triggering event, with the surface elevation still close to the still water level. The second phase point, b, is taken 630ms after the triggering event, and is immediately before the surface elevation reaches its first trough.

The images are acquired with an uncertainty of $\pm 5ms$, due to the synchronisation with the scanning beam. This corresponds to $\pm$0.5% of the wave period.

When the surface height records were analysed, it was found that the standing waves were not composed of a single harmonic, but had non-negligible second and third harmonic components also. The effect this had on the comparison between real PIV data and simulation will be described later. These components are listed in table 6.1

6.7.1 Results

While the comparison of the model and real PIV data is carried out by averaging the PIV data, the vector maps are displayed to illustrate both the progression of

the wave at each phase point, and a typical acceleration field. Figures 6.7,6.8,6.9 show the progression of the wave, and figures 6.10,6.11,6.12,6.13,6.14 show the resulting acceleration fields, both Eulerian and Lagrangian for phase point a. Figures 6.15,6.16,6.17,6.18,6.19,6.20,6.21,6.22 show the same data, but for phase point b.

The acceleration maps, in comparison with the velocity maps, all show a high level of noise. This is mainly due to the fact that the acceleration data is more often than not subject to not only the noise in the PIV recordings, but is calculated from velocity vector maps that have only slight differences between them.

A way to reduce this noise is to spatially average the acceleration data to give a smoothed representation of the data. This method of course has its disadvantages, namely a loss of spatial resolution, but the increase in the accuracy of the acceleration measurement should counter this. The noise could also be reduced in the vector maps by performing a phase averaging, whereby a highly repeatable flow situation is repeated a number of times, with the data from the same flow phase being averaged together. This has the advantage that the spatial resolution of the original measurement is preserved, however, the effect of averaging will remove the instantaneous structures within the field. As with conventional PIV, it is these instantaneous structures that are often of the most interest.

6.8 Freak Results

The freak wave is generated by ramping different wave frequencies up from 0.6Hz to 1.6Hz, with the fundamental wavelength as the wave meets the wall is approximately 3.7m. For the freak wave three different phase points are used to trigger the system, and only the Eulerian acceleration was calculated. These

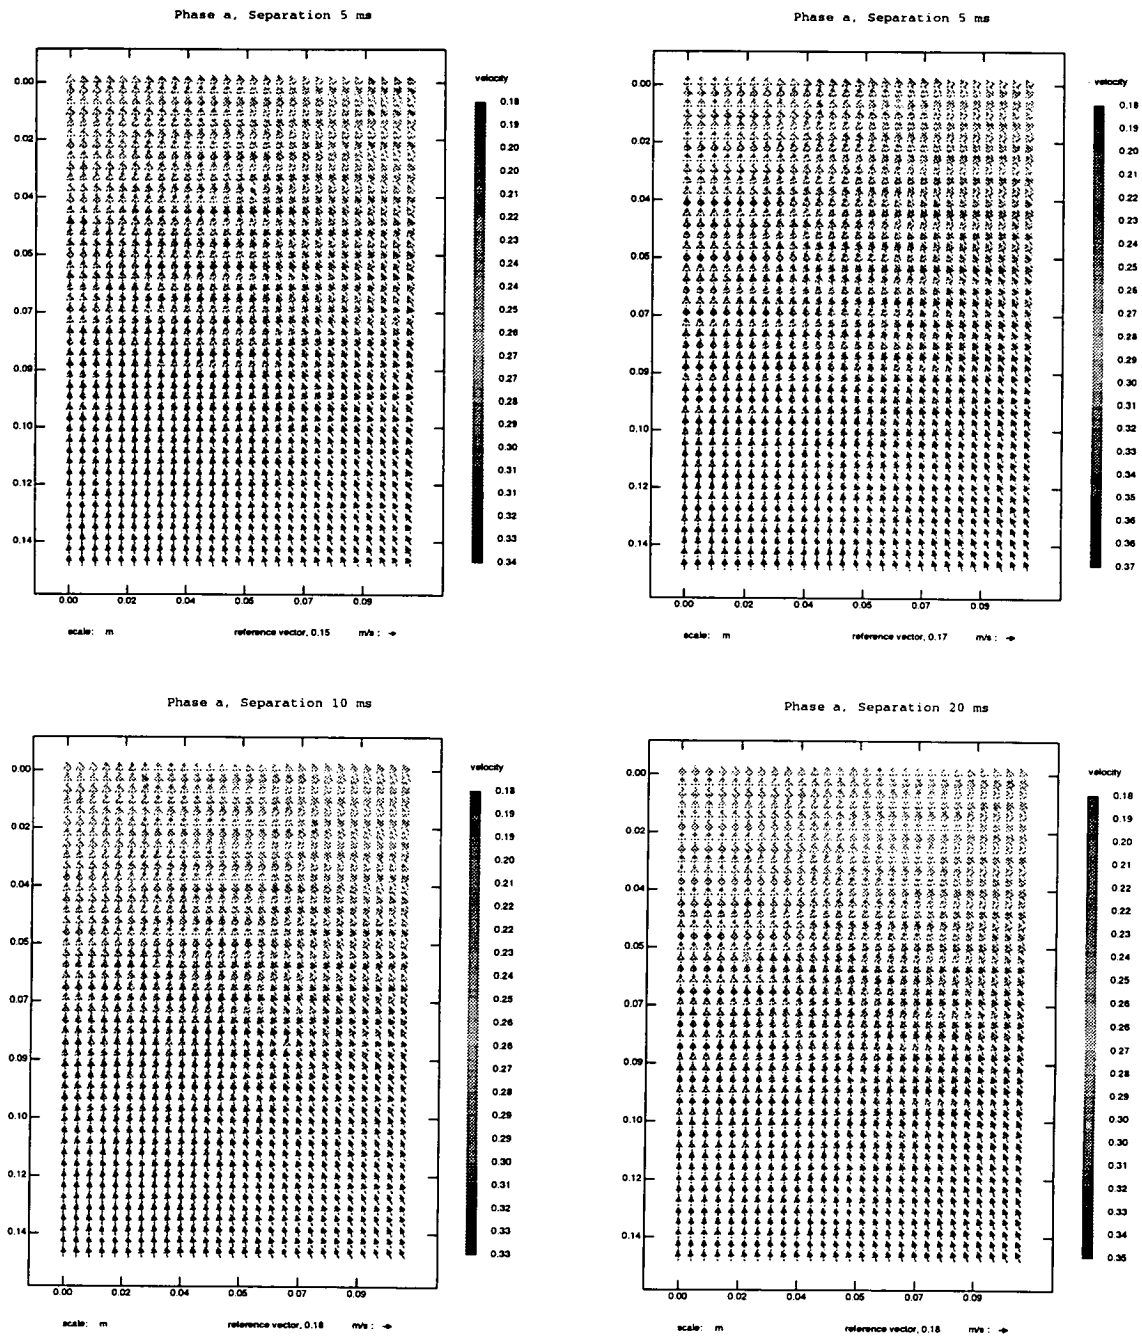

Figure 6.7: Velocities, phase point a, n=0,1,2,4

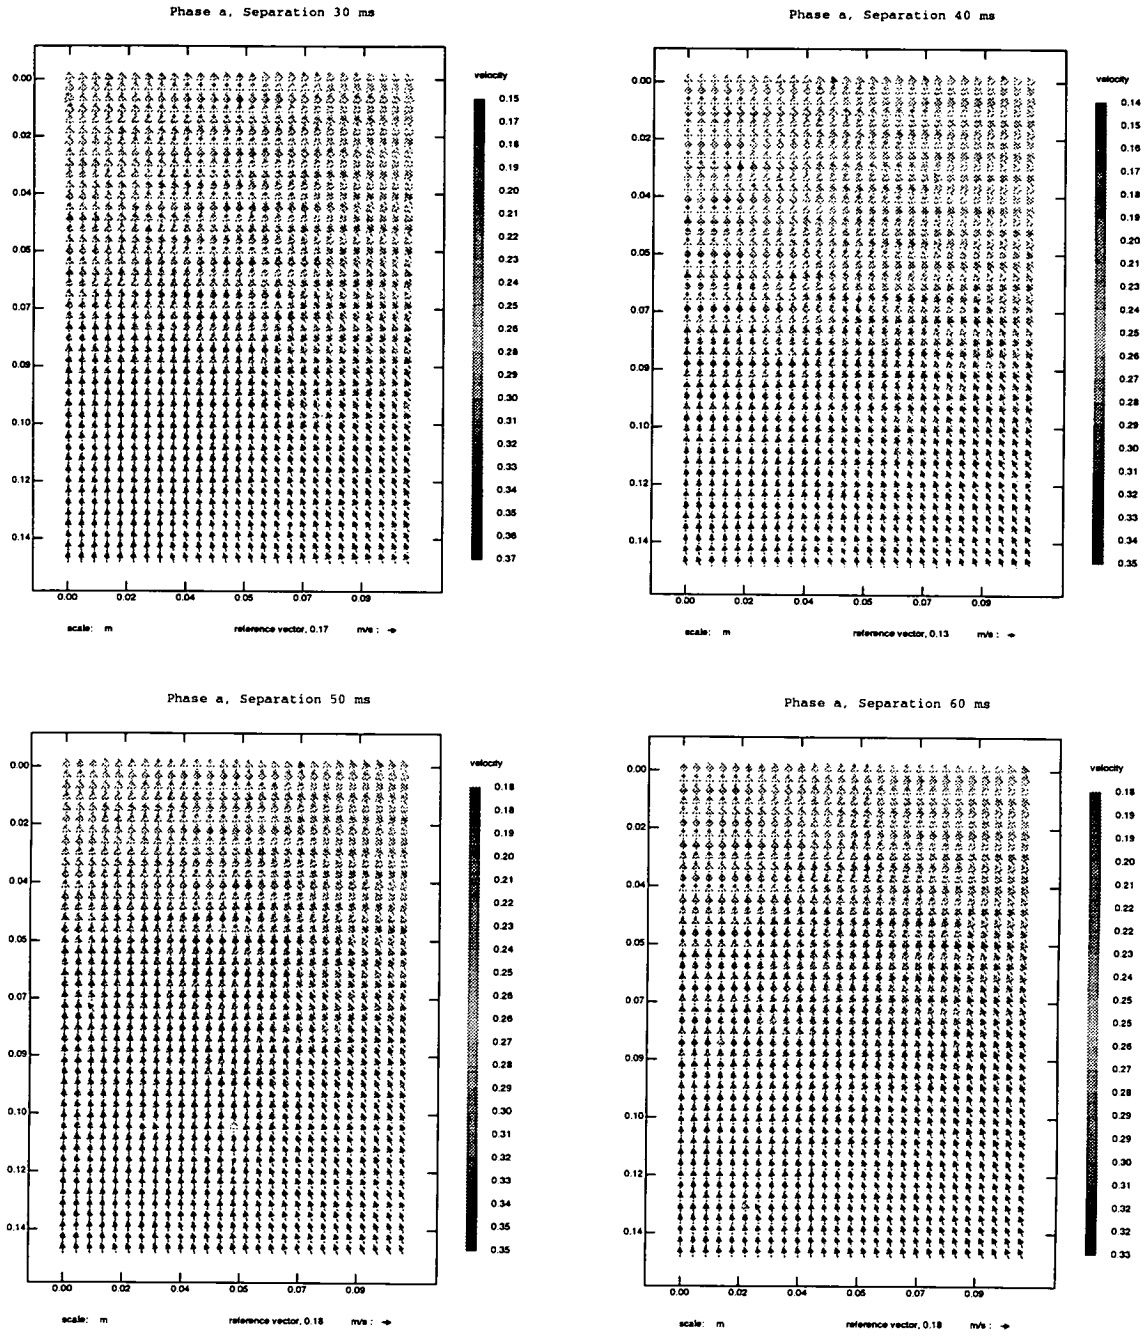

Figure 6.8: Velocities, phase point a, n=6,8,10,12

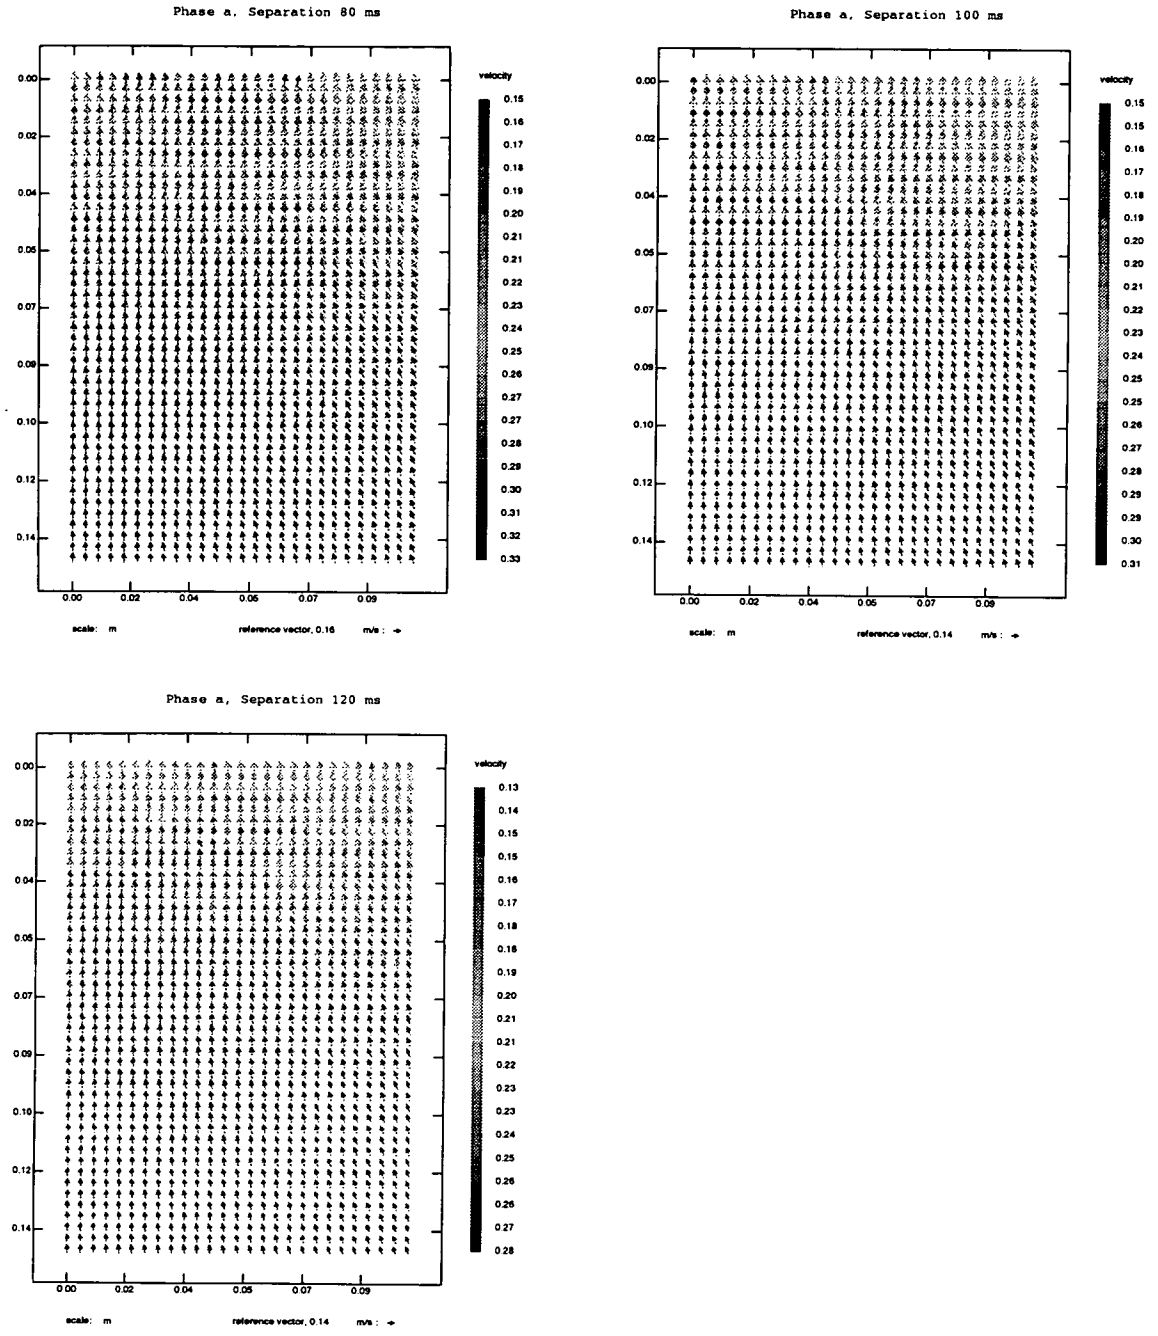

Figure 6.9: Velocities, phase point a, n=16,20,24

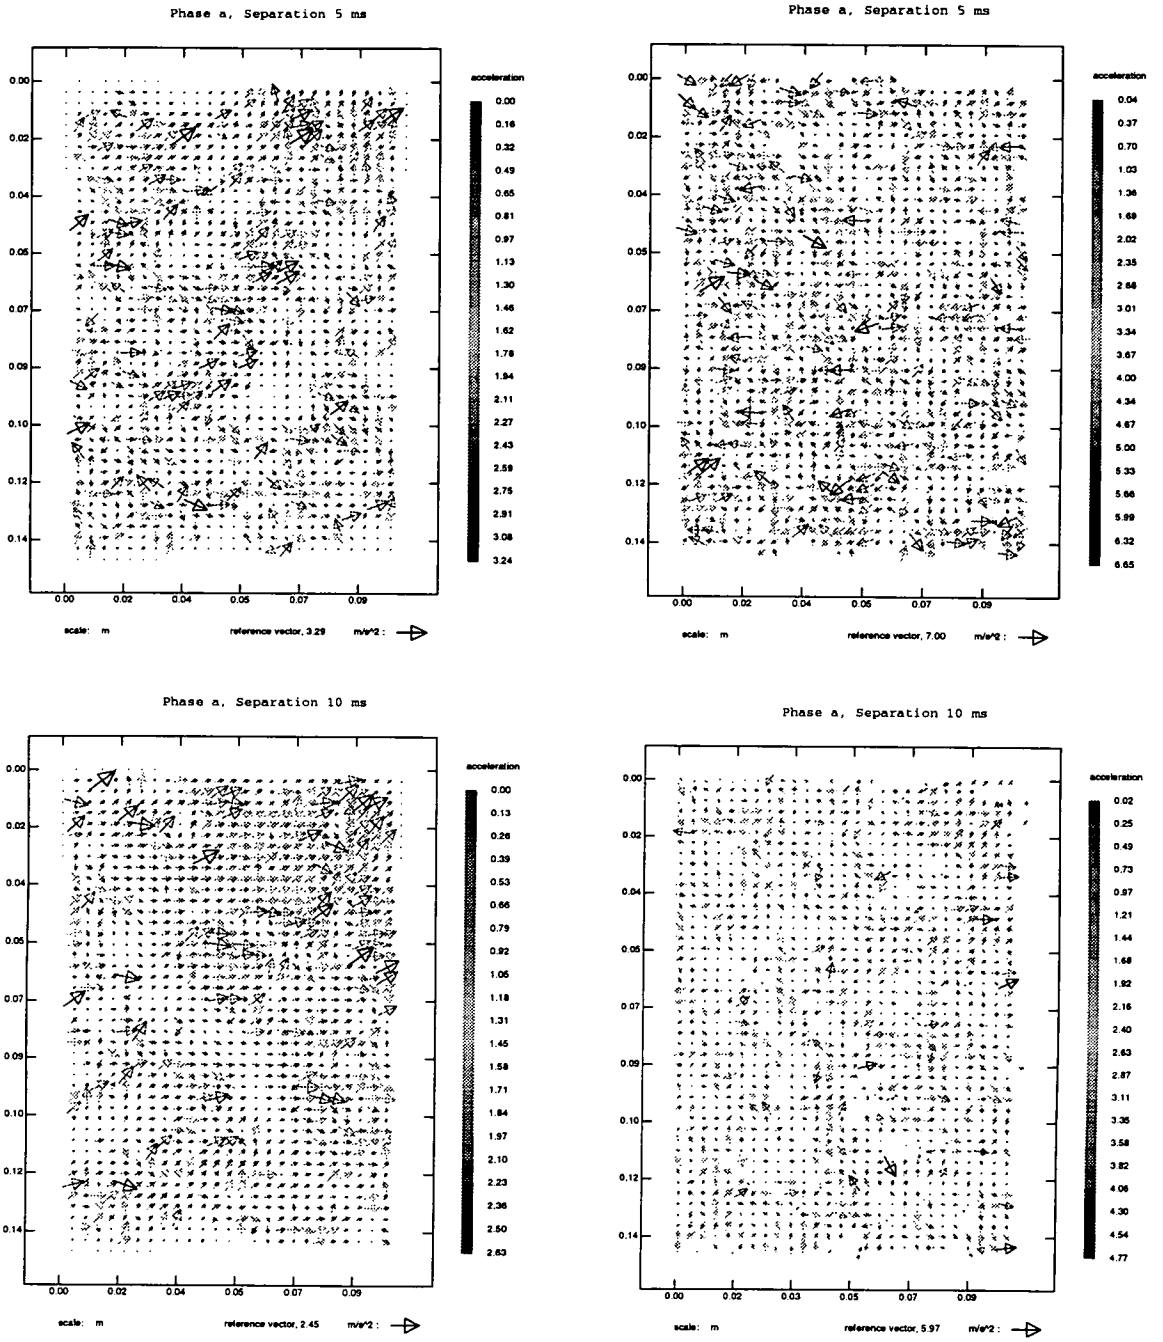

Figure 6.10: Accelerations (Eulerian-left, Lagrangian-right), phase point a, n=1,2

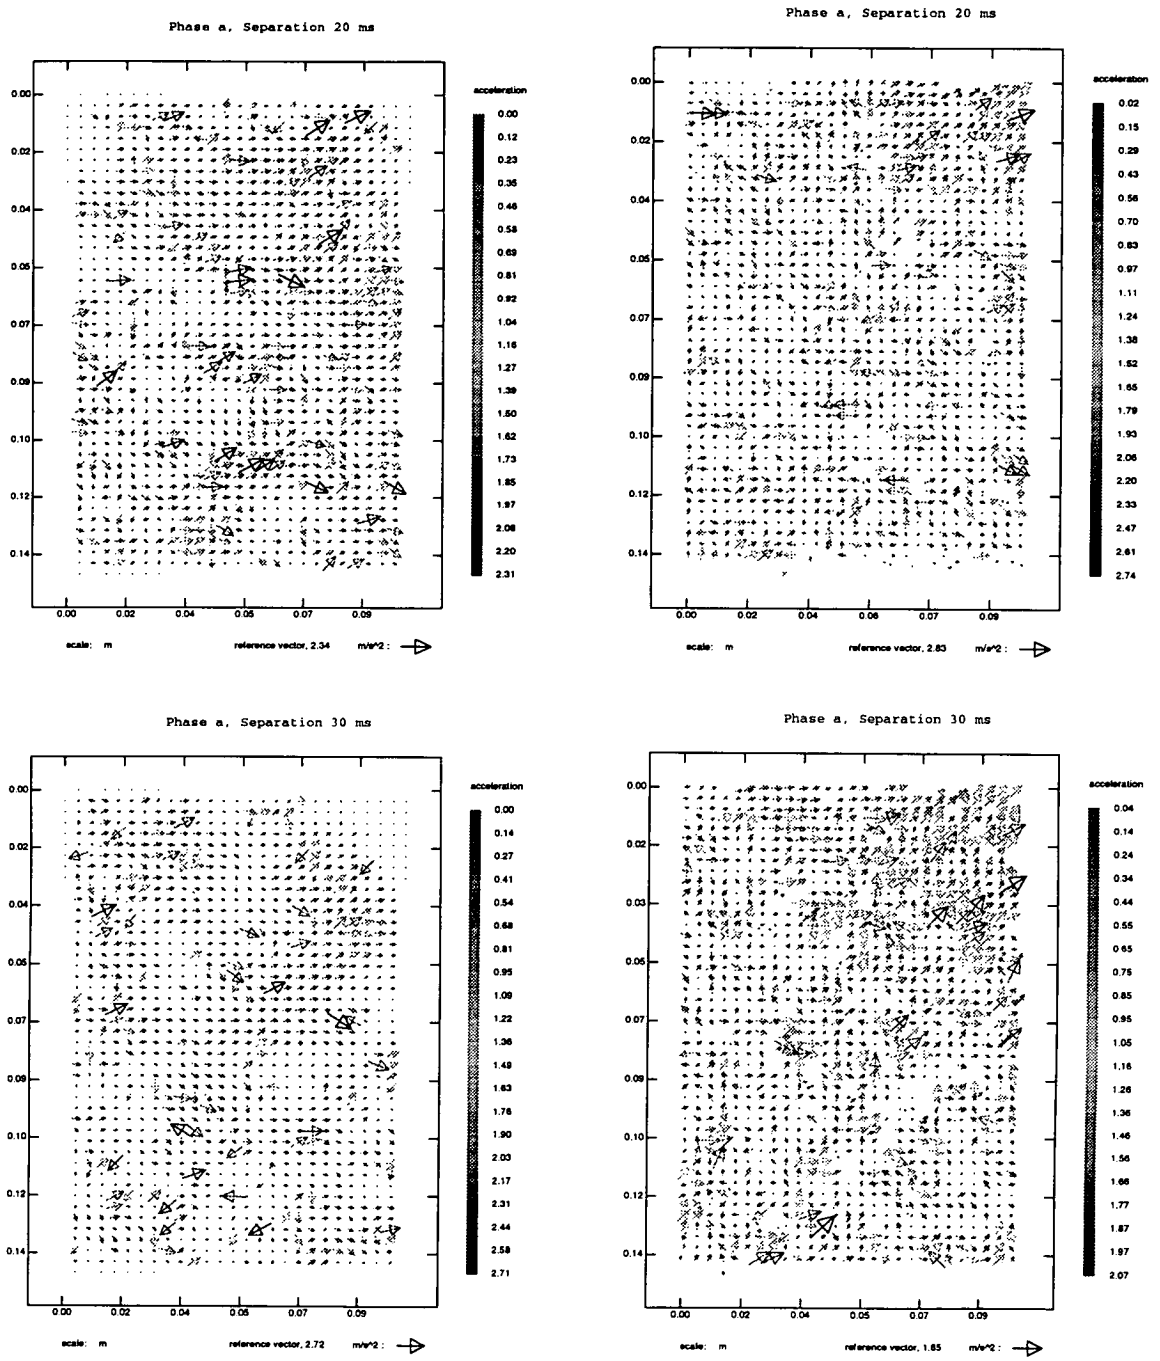

Figure 6.11: Accelerations (Eulerian-left, Lagrangian-right), phase point a, n=4,6

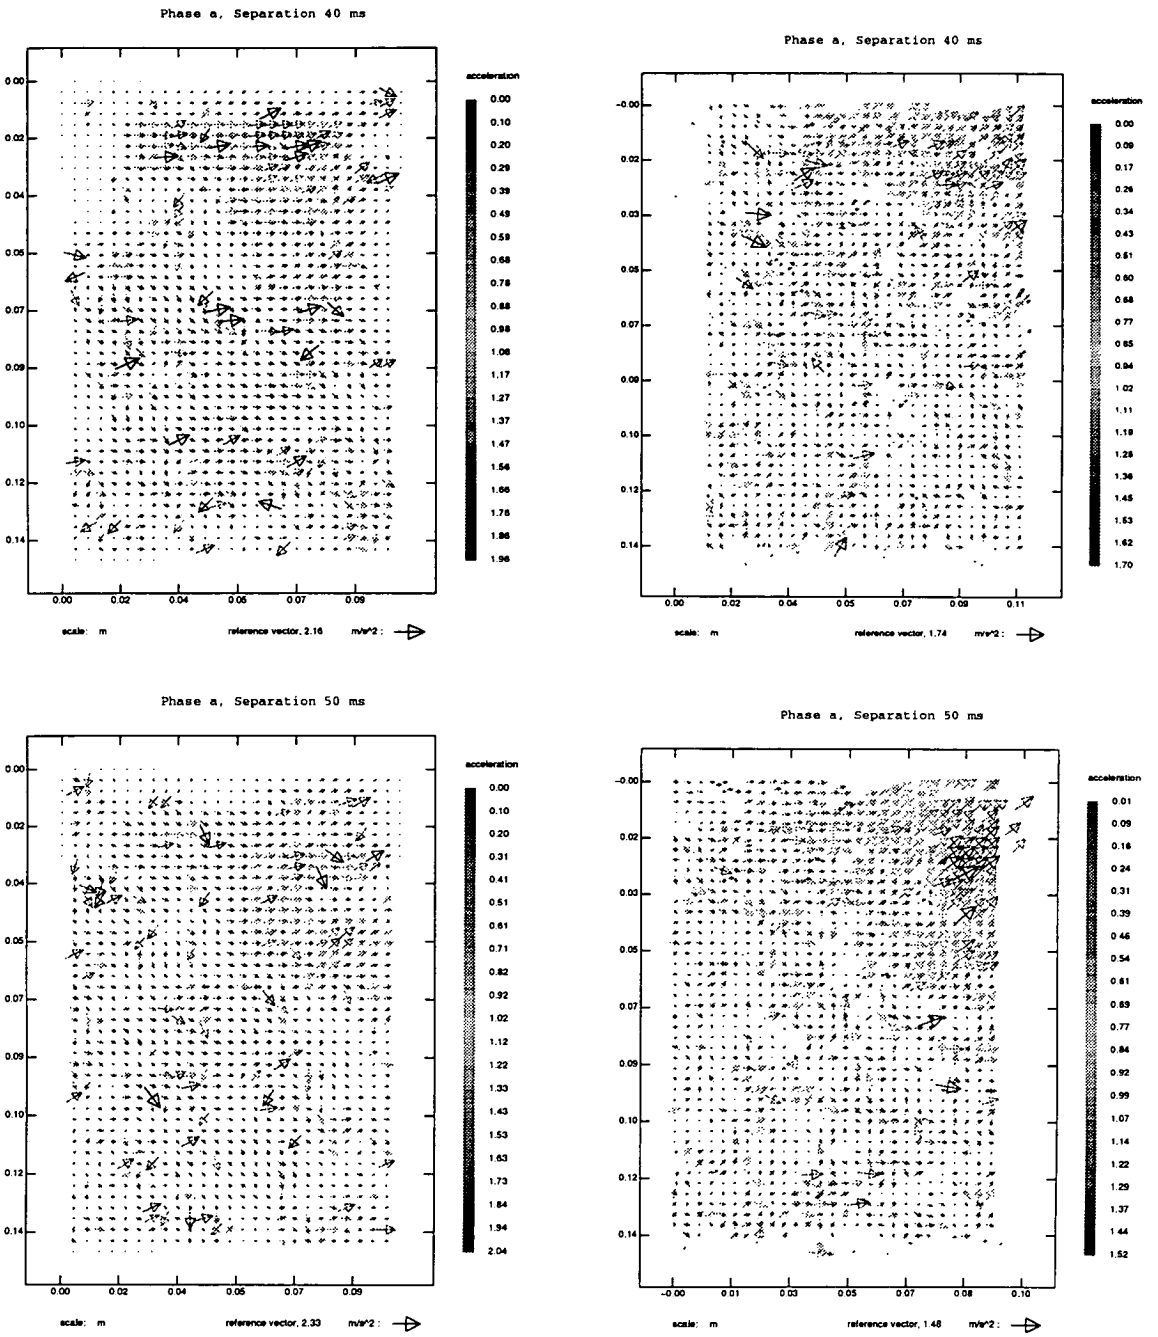

Figure 6.12: Accelerations (Eulerian-left, Lagrangian-right), phase point a, n=8,10

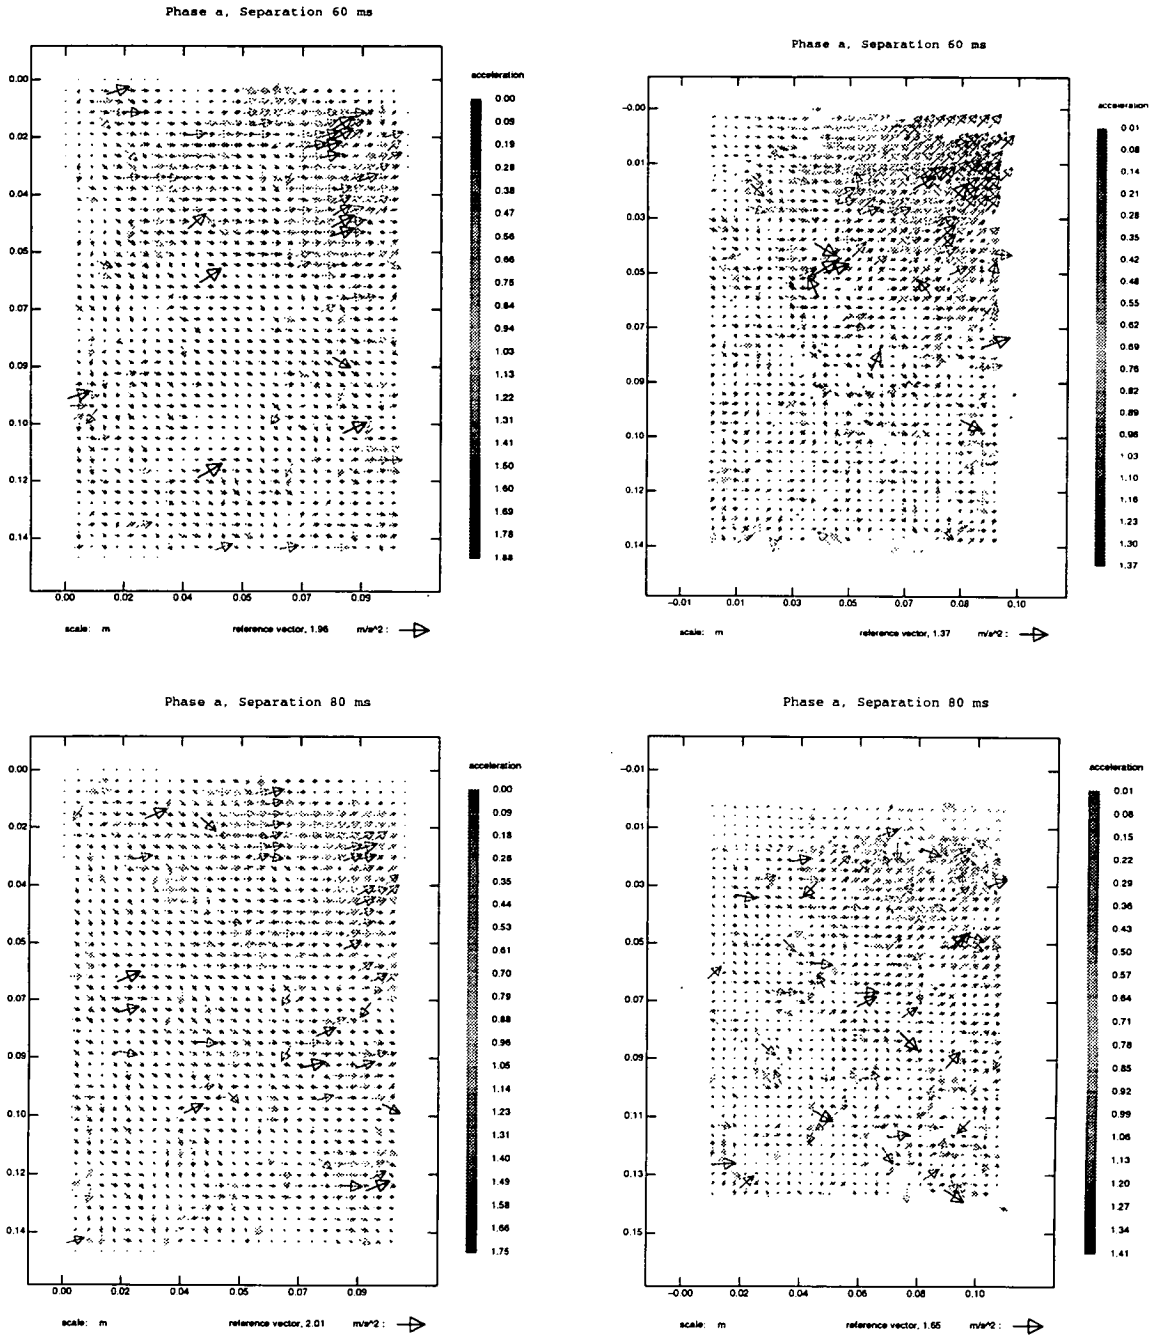

Figure 6.13: Accelerations (Eulerian-left, Lagrangian-right), phase point a, n=12,16

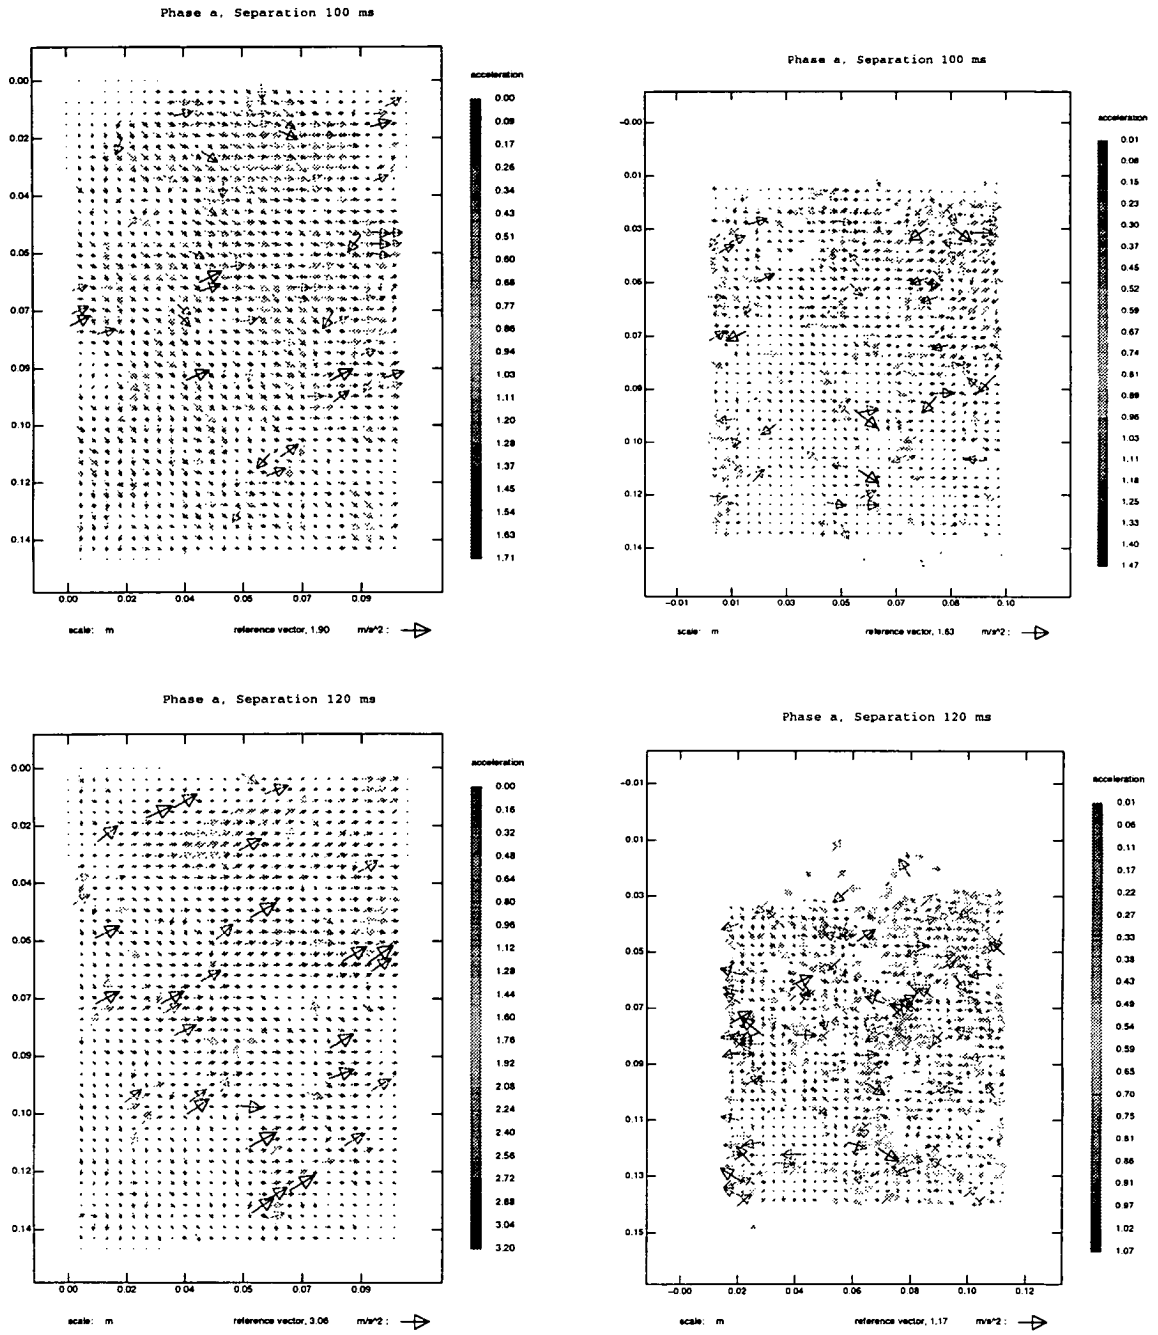

Figure 6.14: Accelerations (Eulerian-left, Lagrangian-right), phase point a, n=20,24

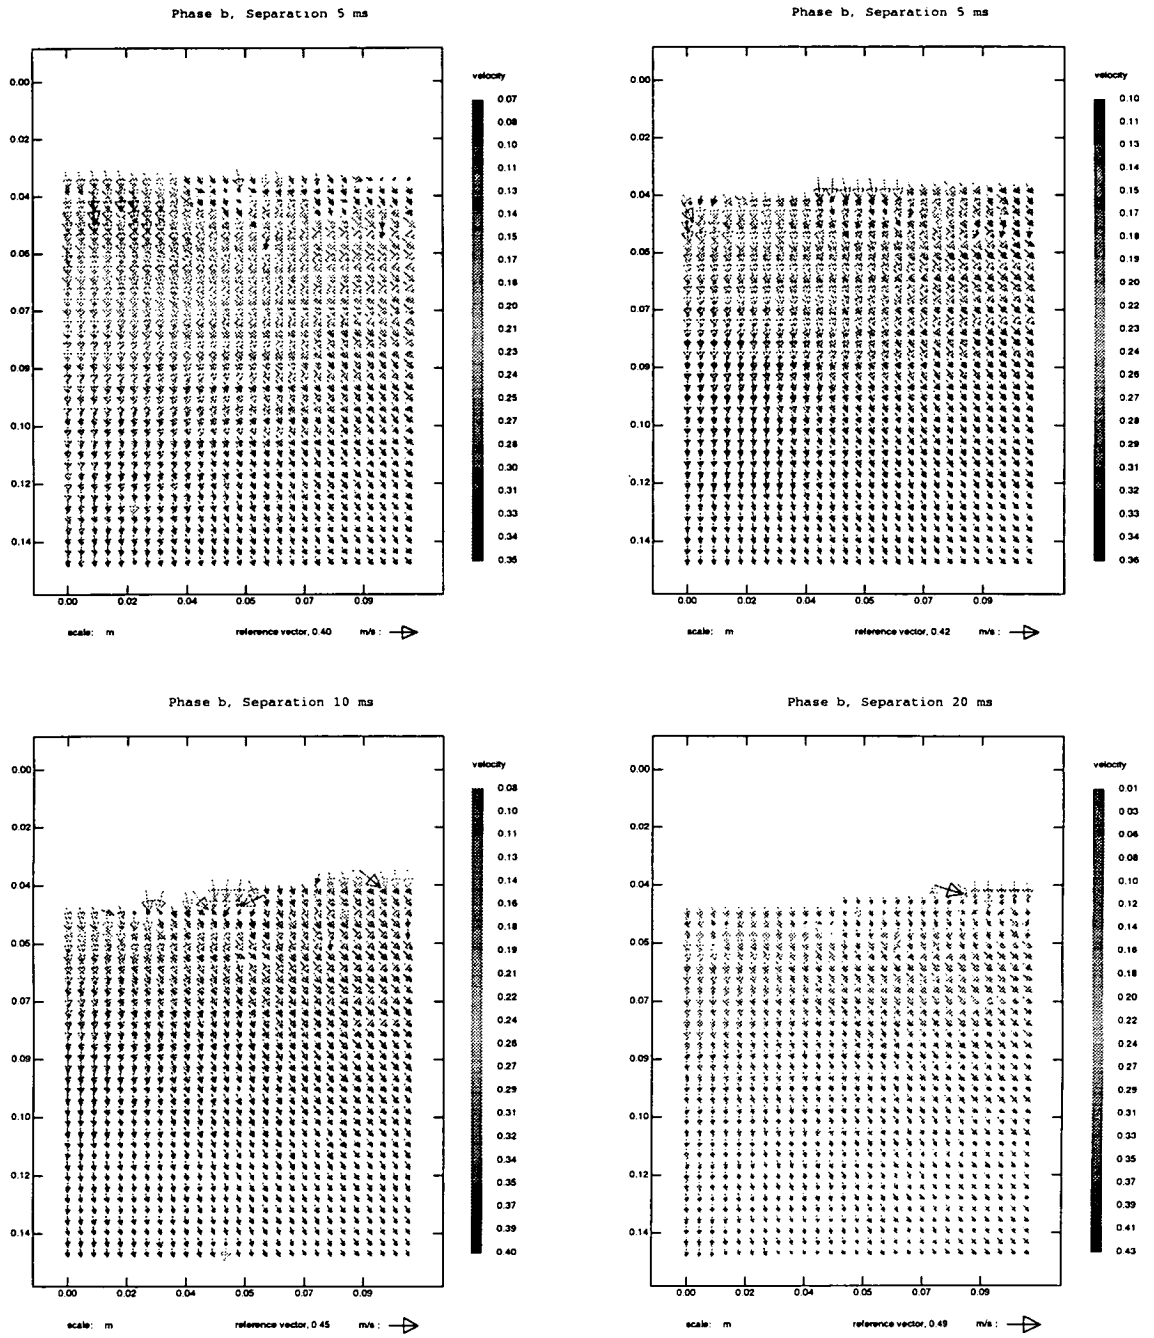

Figure 6.15: Velocities, phase point b, n=0,1,2,4

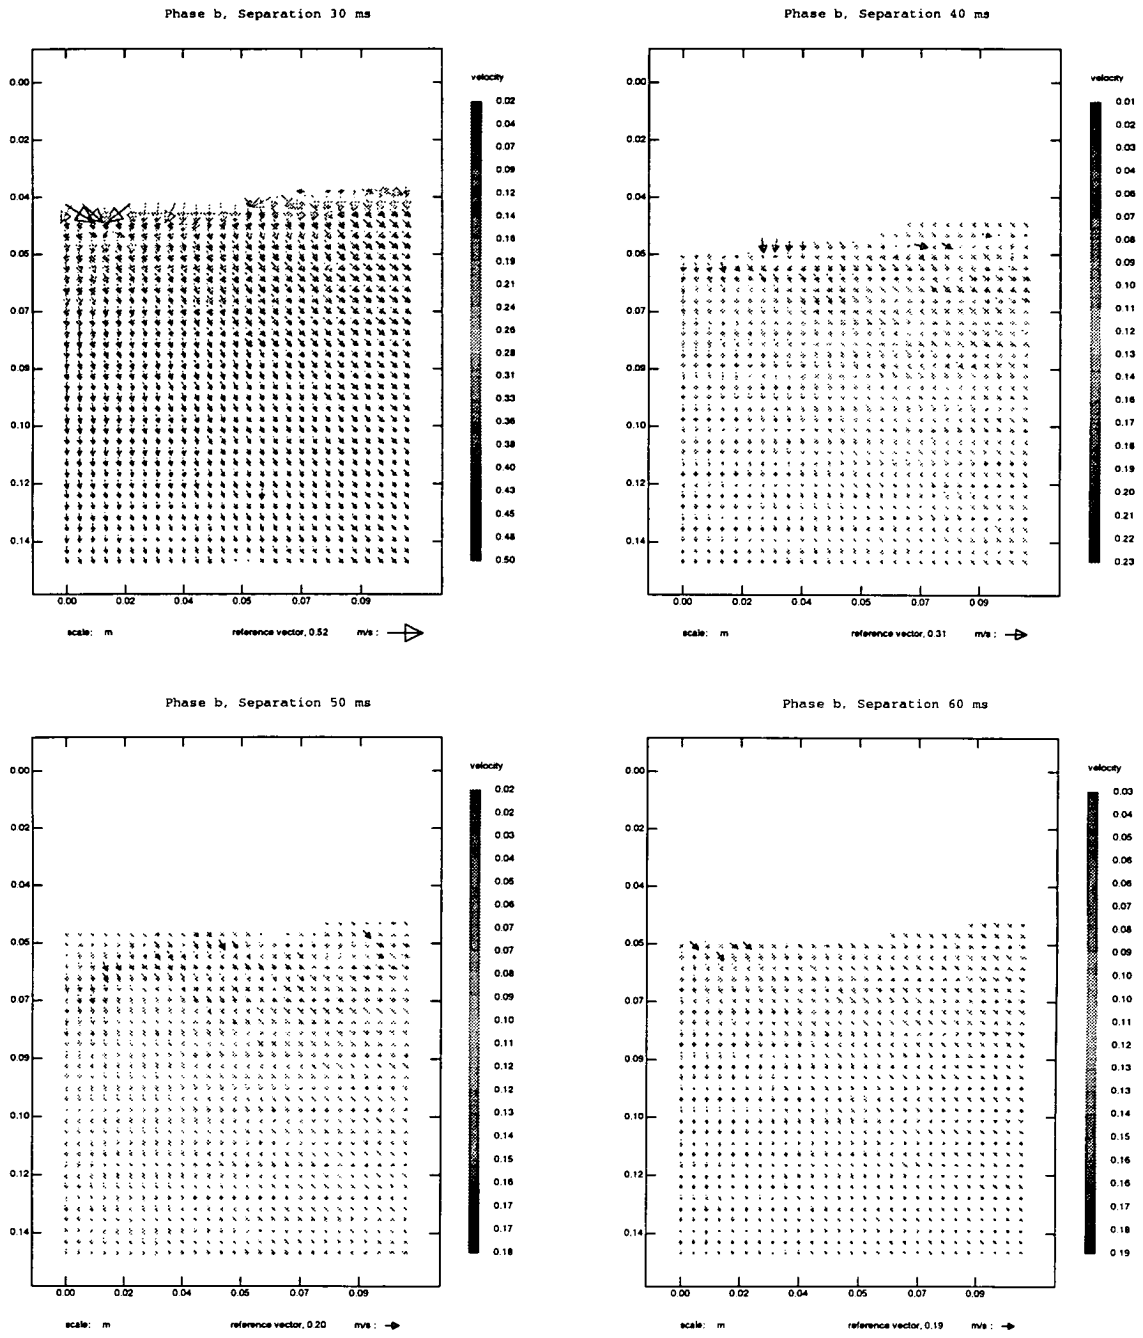

Figure 6.16: Velocities, phase point b, n=6,8,10,12

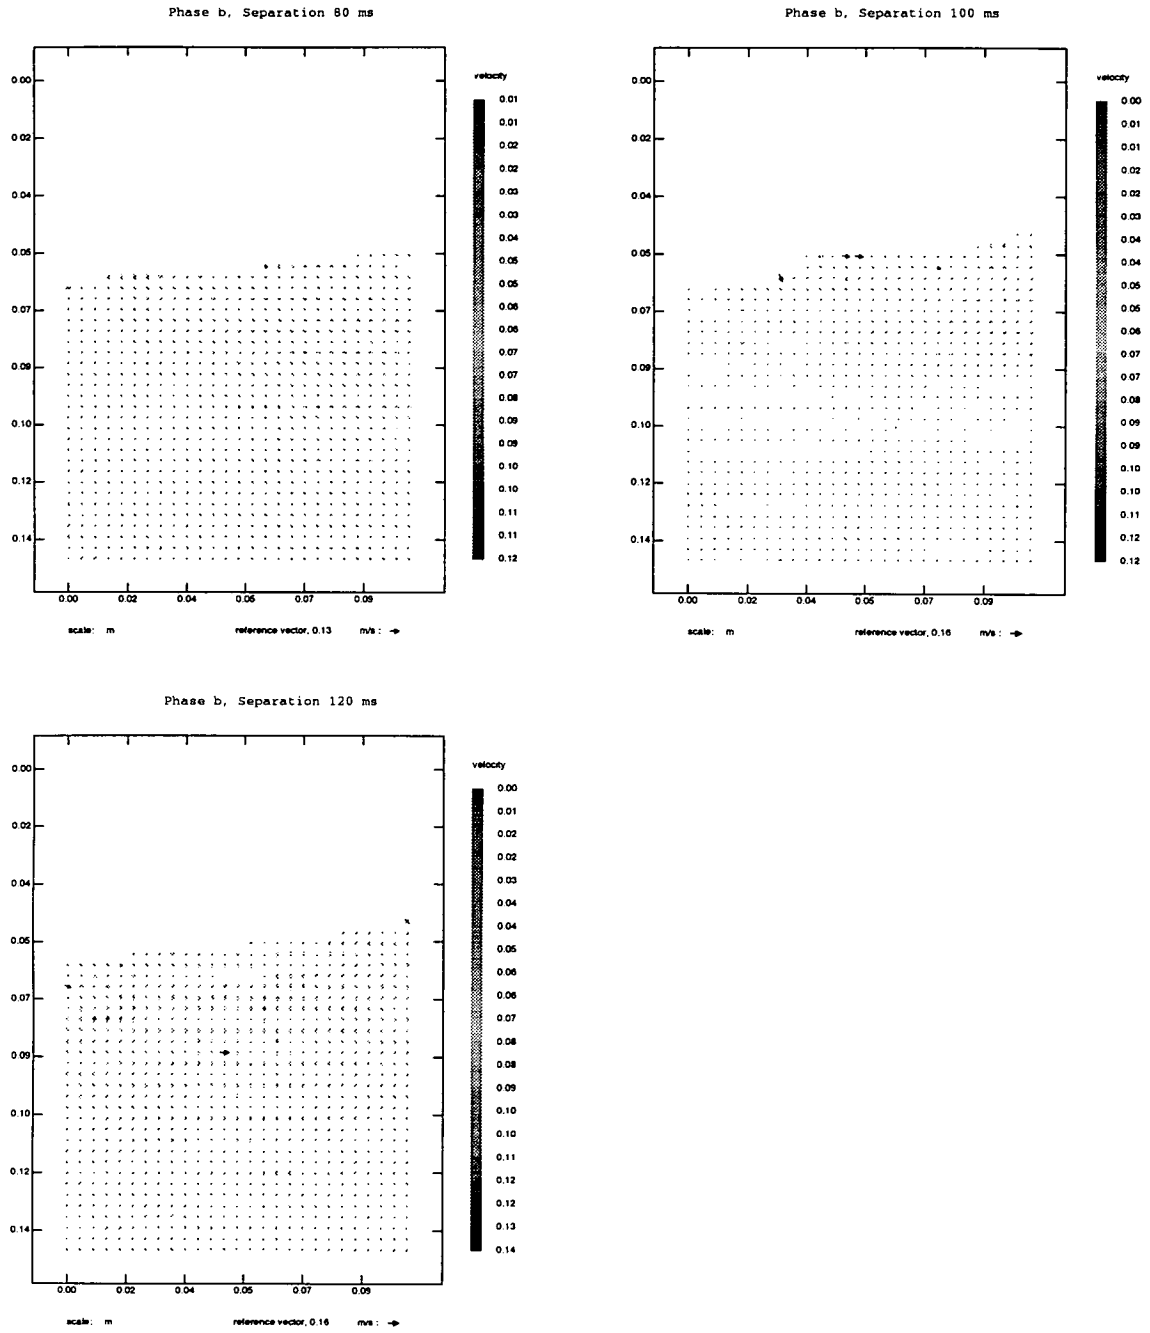

Figure 6.17: Velocities, phase point b, n=16,20,24

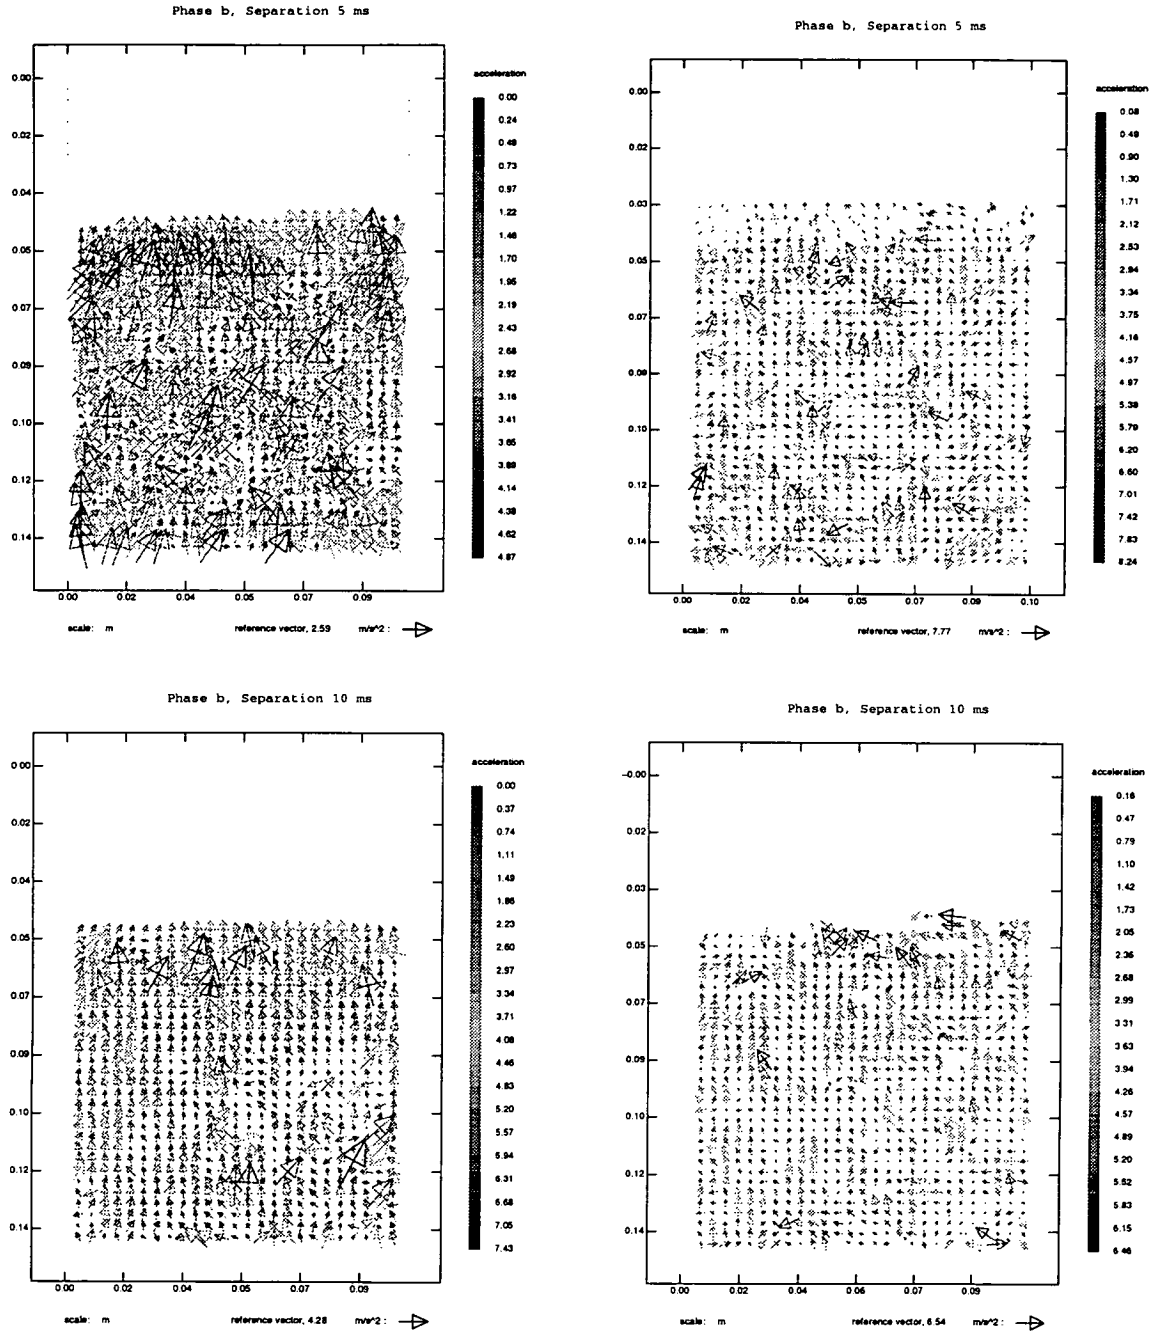

Figure 6.18: Accelerations (Eulerian-left, Lagrangian-right), phase point b, n=1,2

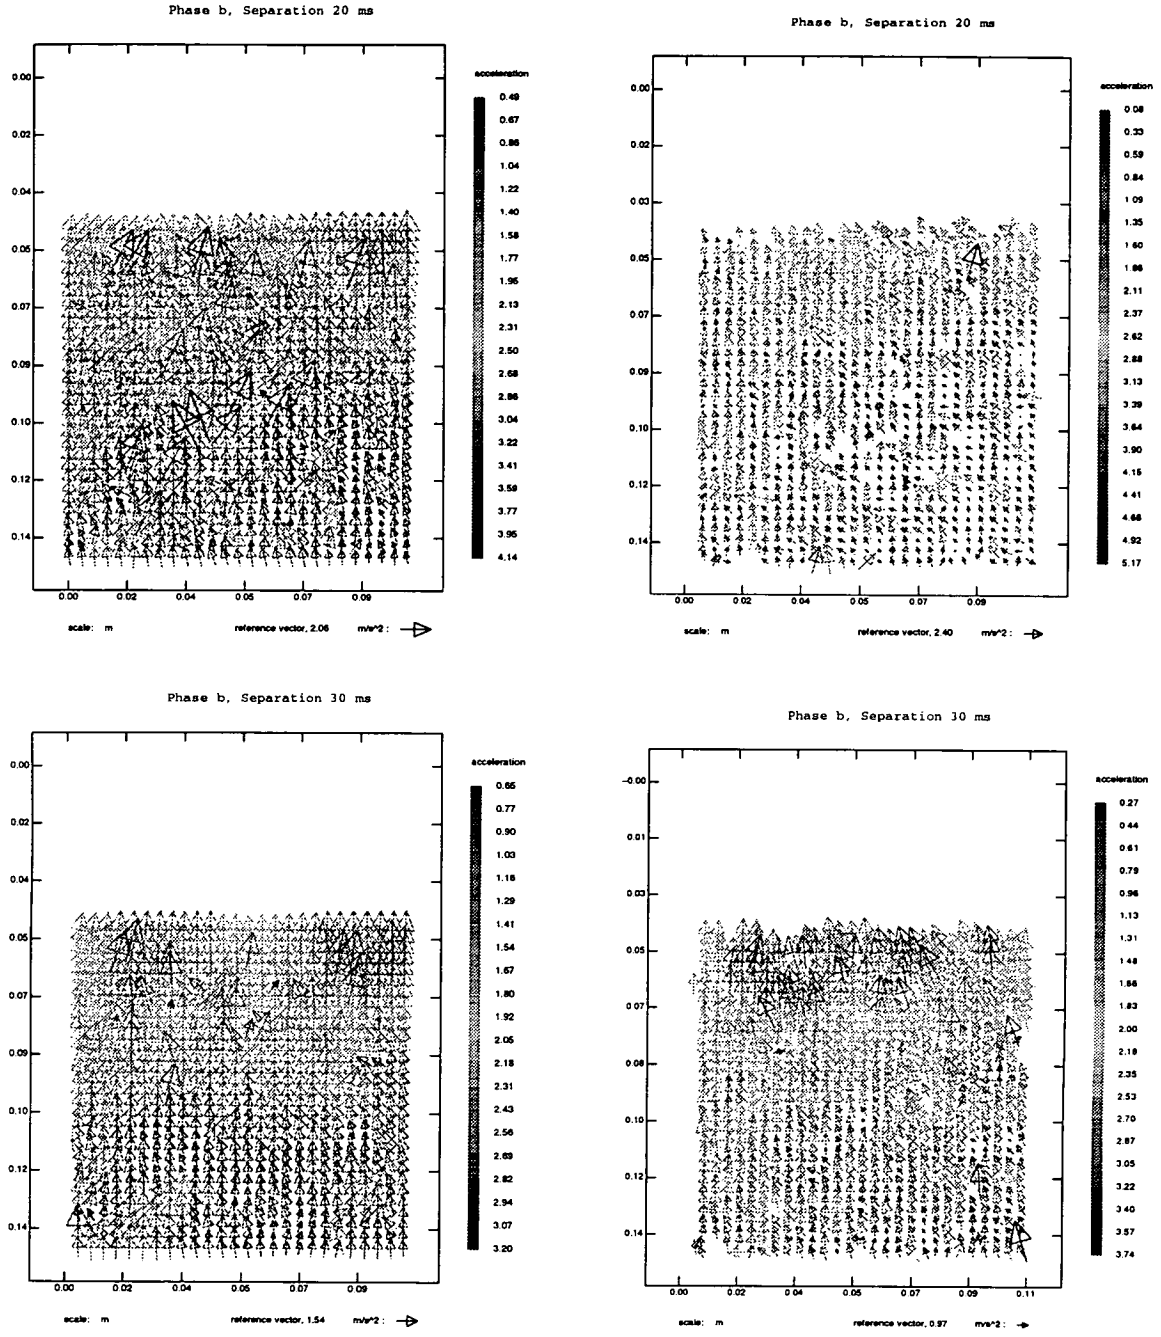

Figure 6.19: Accelerations (Eulerian-left, Lagrangian-right), phase point b, n=4,6

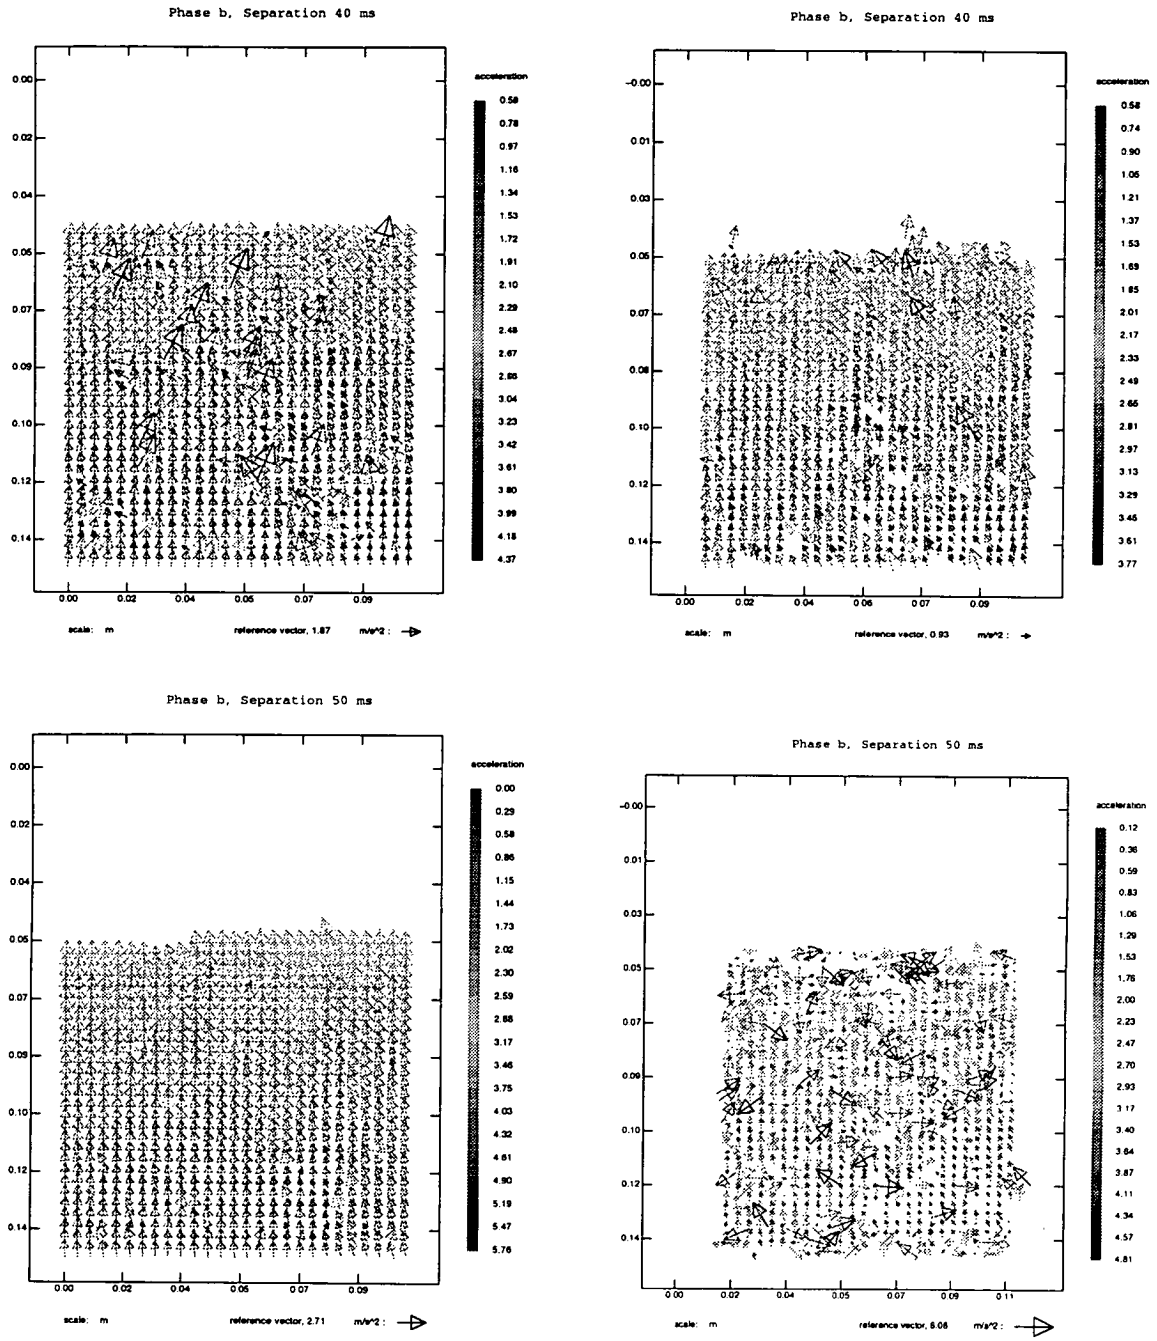

Figure 6.20: Accelerations (Eulerian-left, Lagrangian-right), phase point b, n=8,10

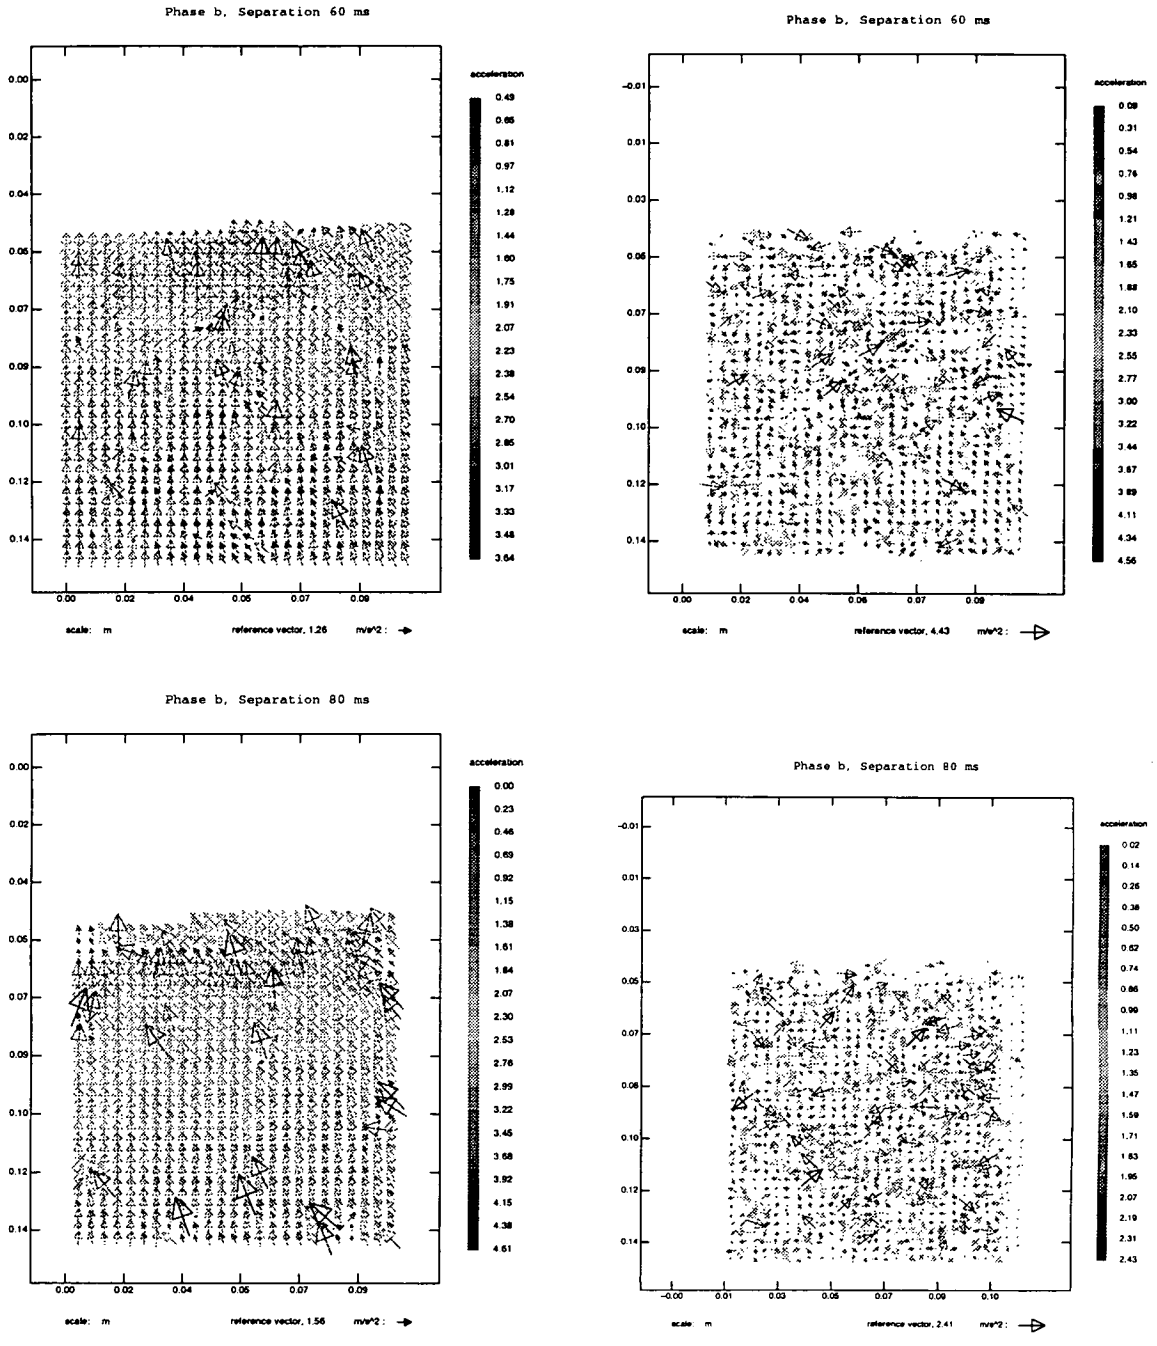

Figure 6.21: Accelerations (Eulerian-left, Lagrangian-right), phase point b, n=12,16

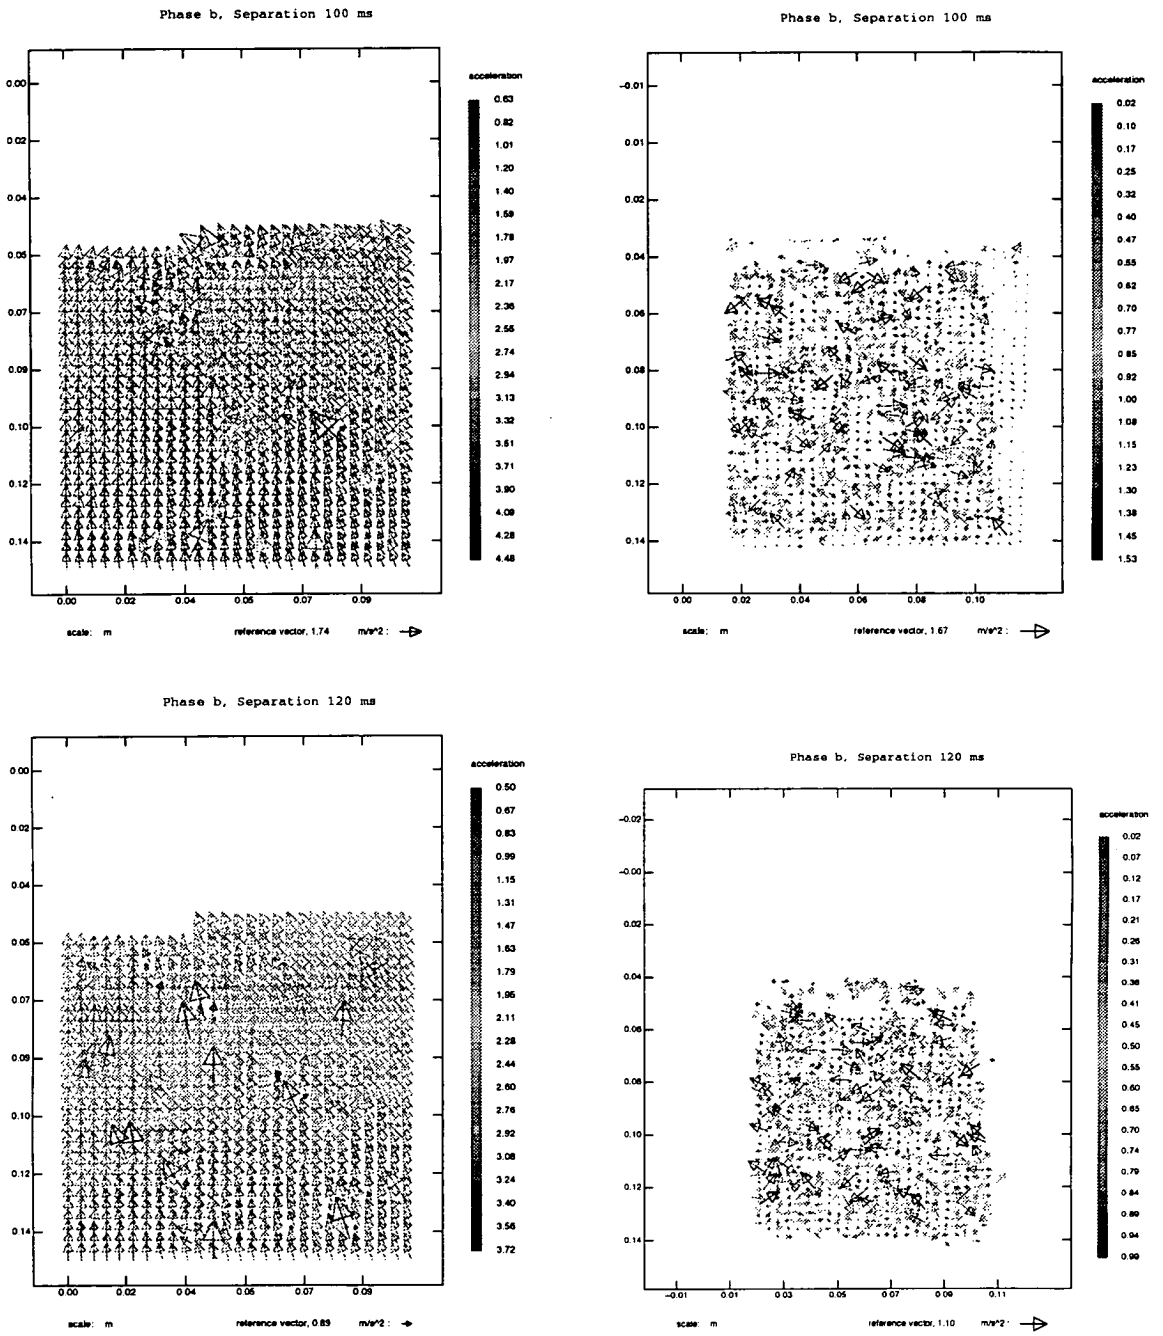

Figure 6.22: Accelerations (Eulerian-left, Lagrangian-right), phase point b, n=20,24

phase points are 700ms after the triggering event (point a), 800ms after (point b) and 1230ms after (point c). Phase point a corresponds to the formation of the main upward moving jet, phase point b to immediately before the crest reaches the wall, and phase point c to the wave withdrawing from the wall.

Because of the high accelerations present in this wave, the range of separating times was limited from $n = 1..12$. For the images taken in phase c, the analysis was carried out with an interrogation area that was twice as large as normal (64x64 as opposed to 32x32) due to the high spatial gradients present during this phase of the wave cycle.

6.9 Comparison of Simulated PIV Images and Analytical Solution

As previously mentioned, a model of the harmonic wave conditions was constructed from linear wave theory. From this PIV images were made with the same parameters are were found in the wave flume experiment. The simplest way to make sense of the sometimes very noisy vector maps was to apply a spatial averaging, and in fact for the comparison between the model and the simulated PIV images, the entire velocity map was averaged. As each map contains typically between 400 and 1000 vectors, this gives a reduction in the uncertainty of a factor of 20-30. The maximum spatial frequency that can now be measured from the averaged acceleration maps is $9 \times 6.66m^{-1}$, corresponding to a wave with frequency 2.6Hz. The results are shown in figures 6.35 and 6.36.

The Eulerian accelerations show an excellent match with the analytical data points for both phase points. The Lagrangian accelerations show a relatively small but consistent bias, introducing a systematic error of up to 10% of the maximum acceleration within a wave cycle. This error seems to increase as the

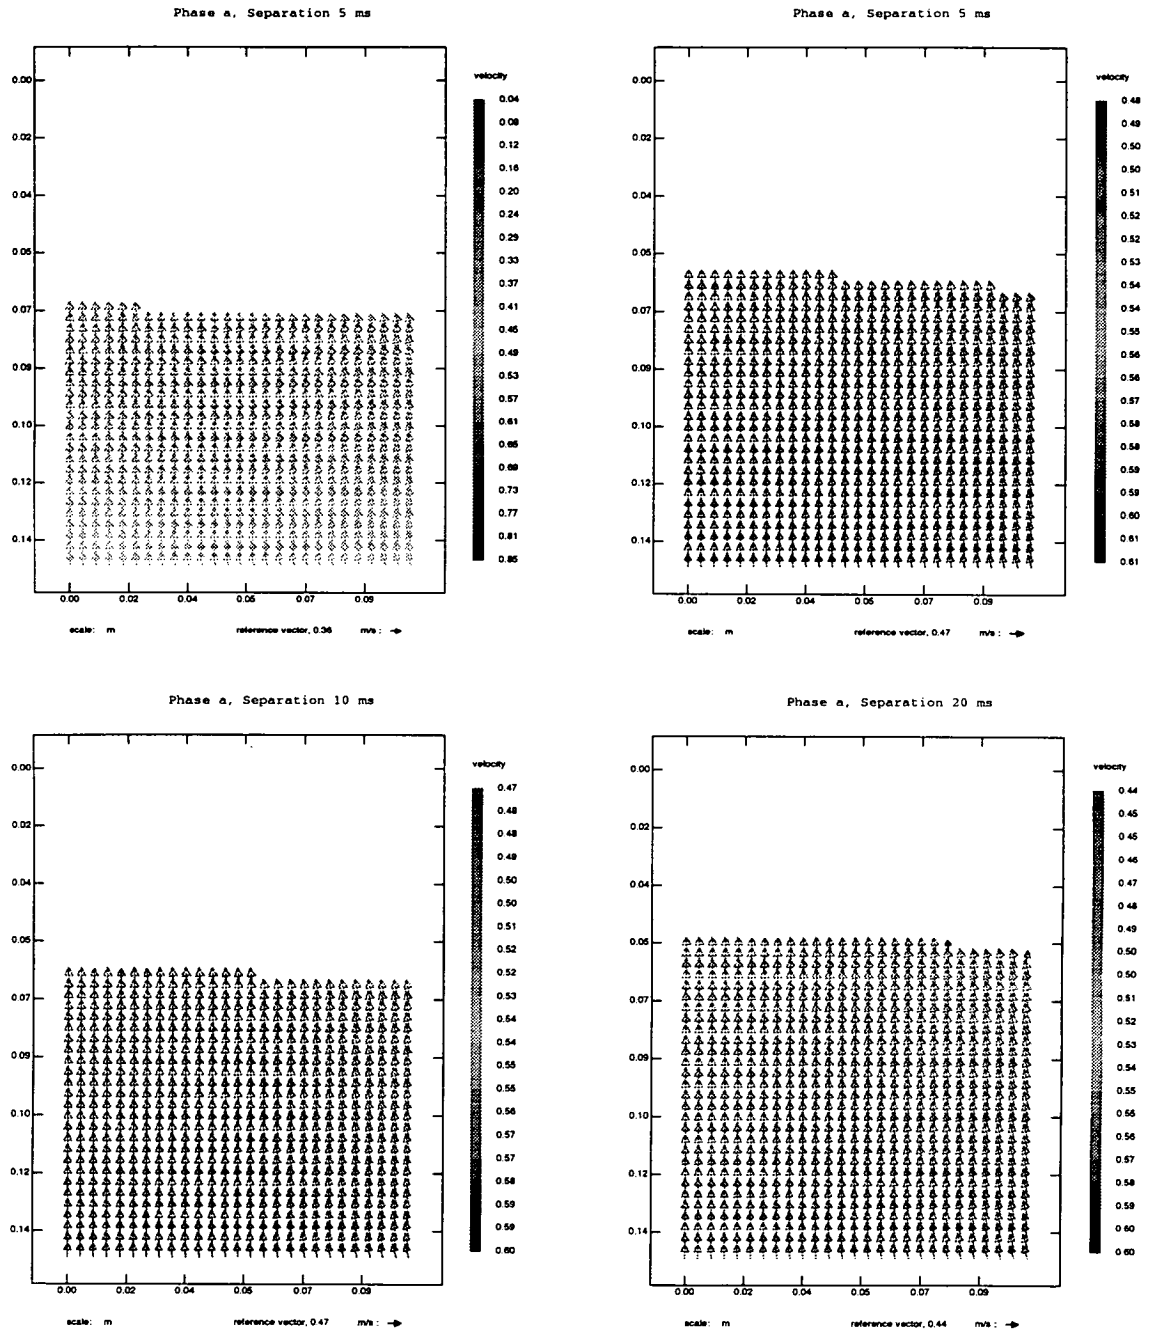

Figure 6.23: Freak velocities, phase point a, n=0,1,2,4

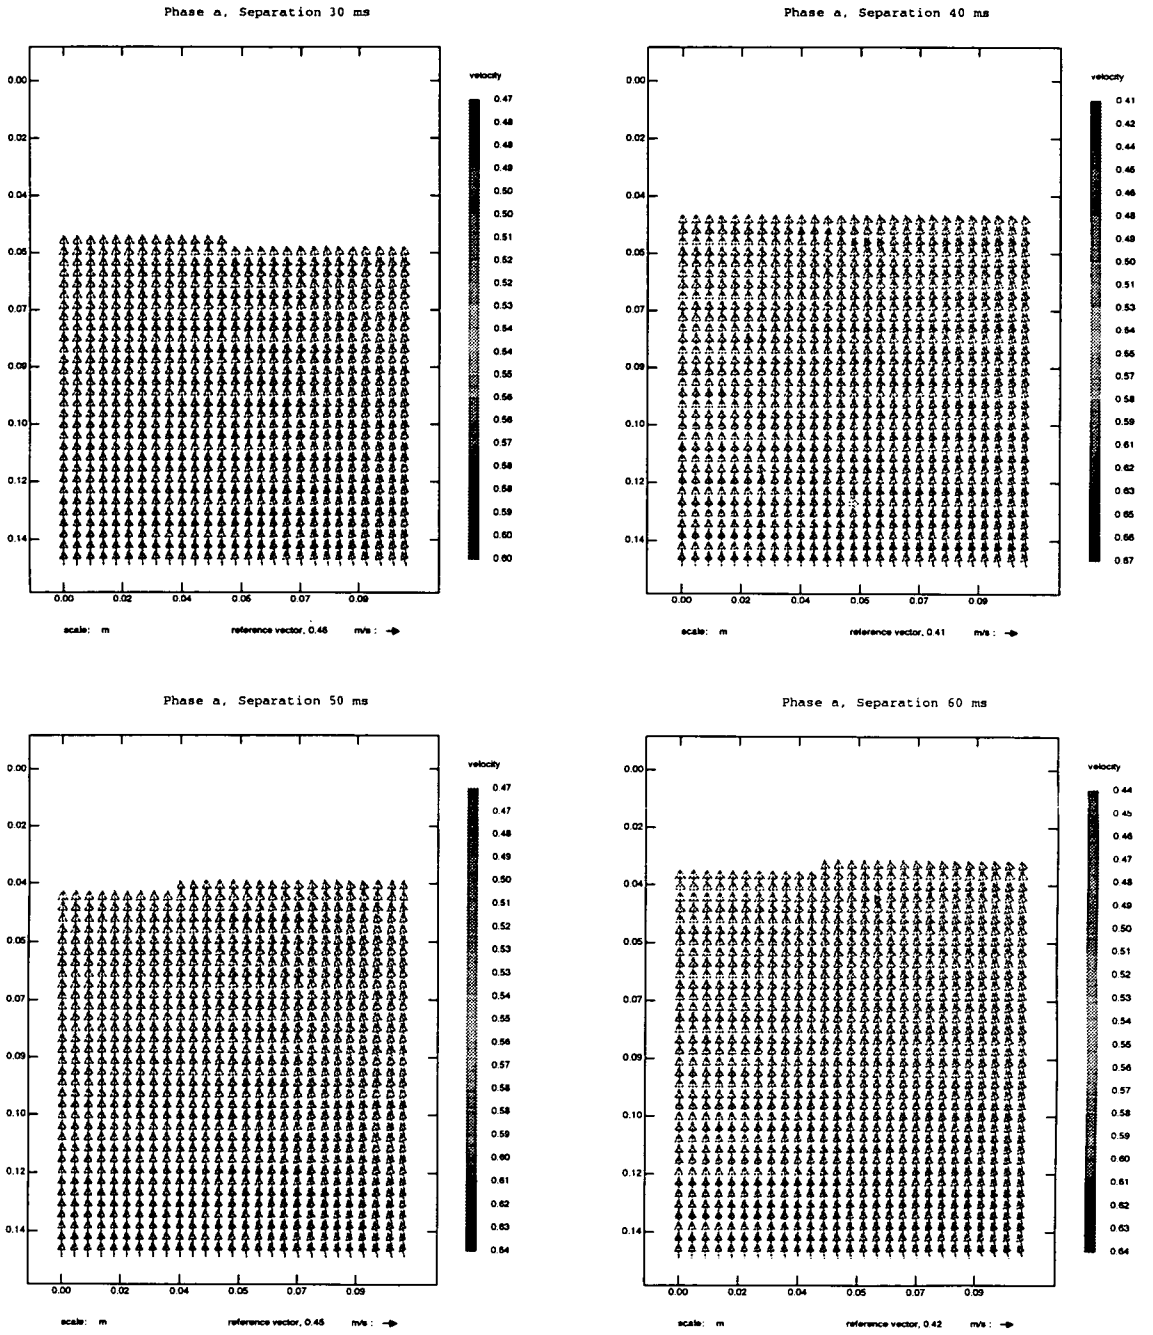

Figure 6.24: Freak velocities, phase point a, n=6,8,10,12

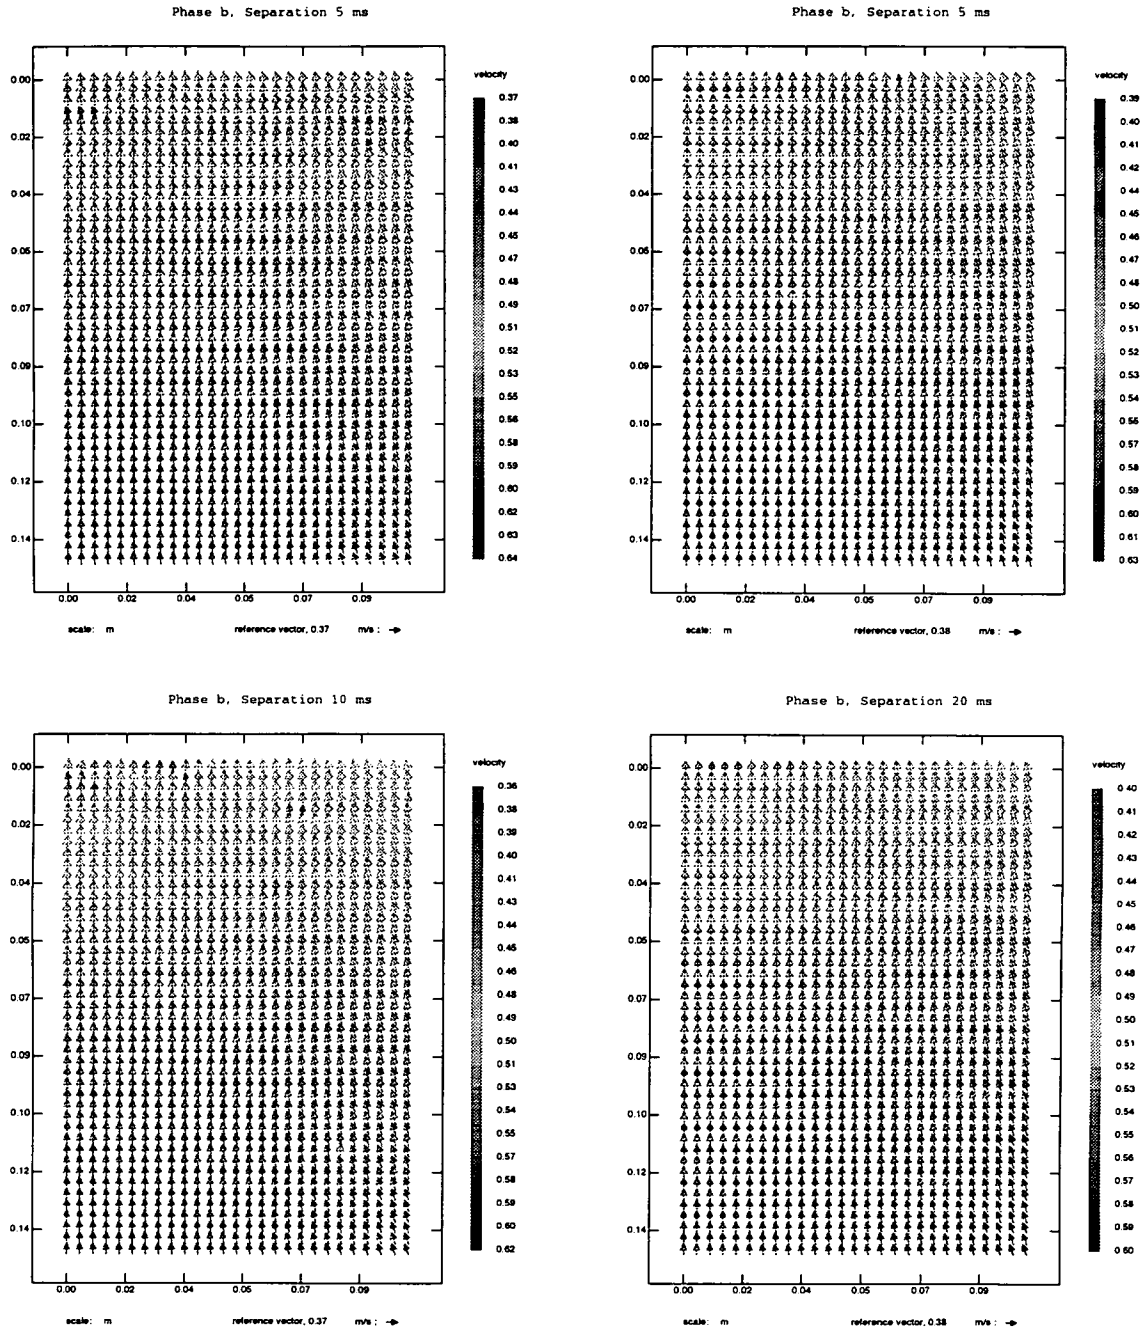

Figure 6.25: Freak velocities, phase point b, n=0,1,2,4

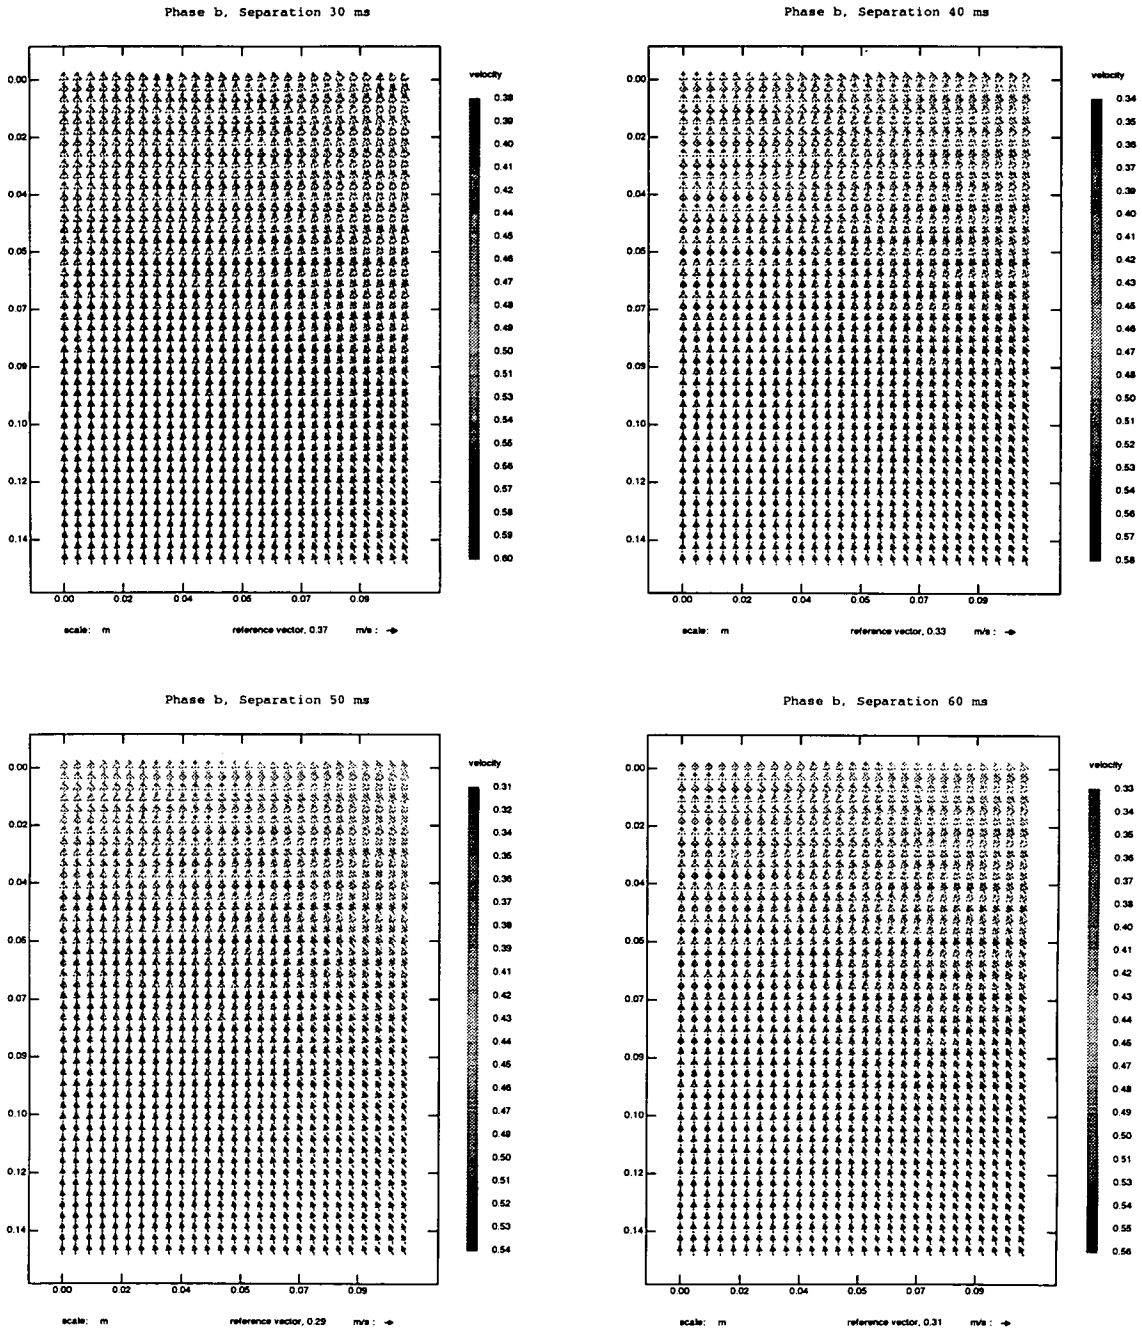

Figure 6.26: Freak velocities, phase point b, n=6,8,10,12

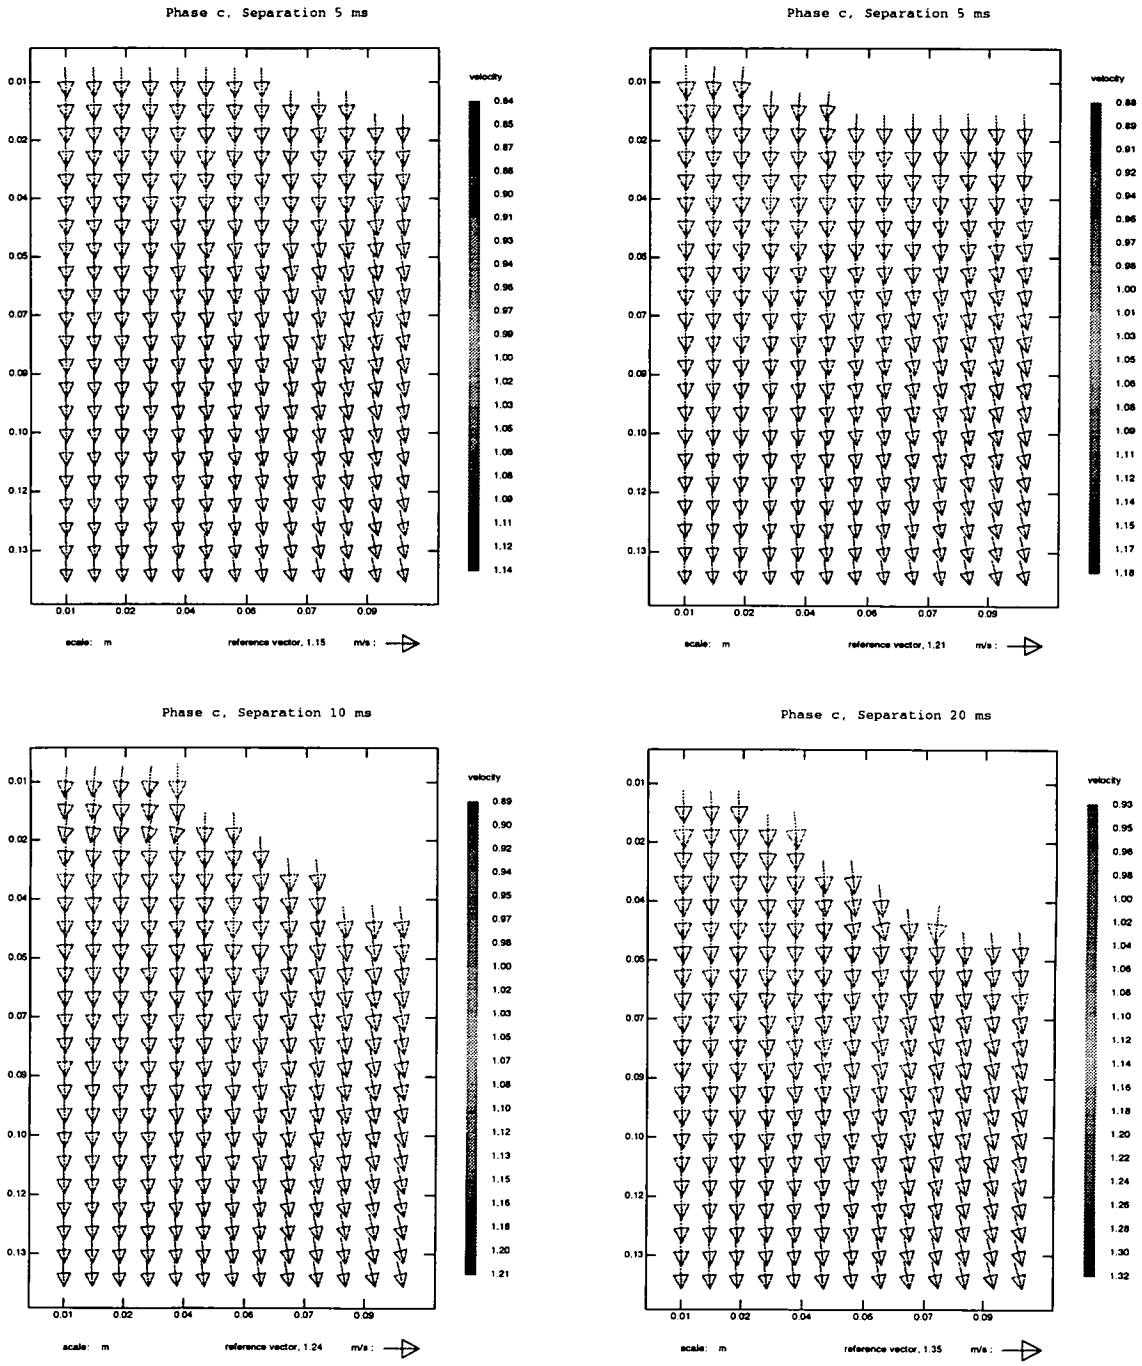

Figure 6.27: Freak velocities, phase point c, n=0,1,2,4

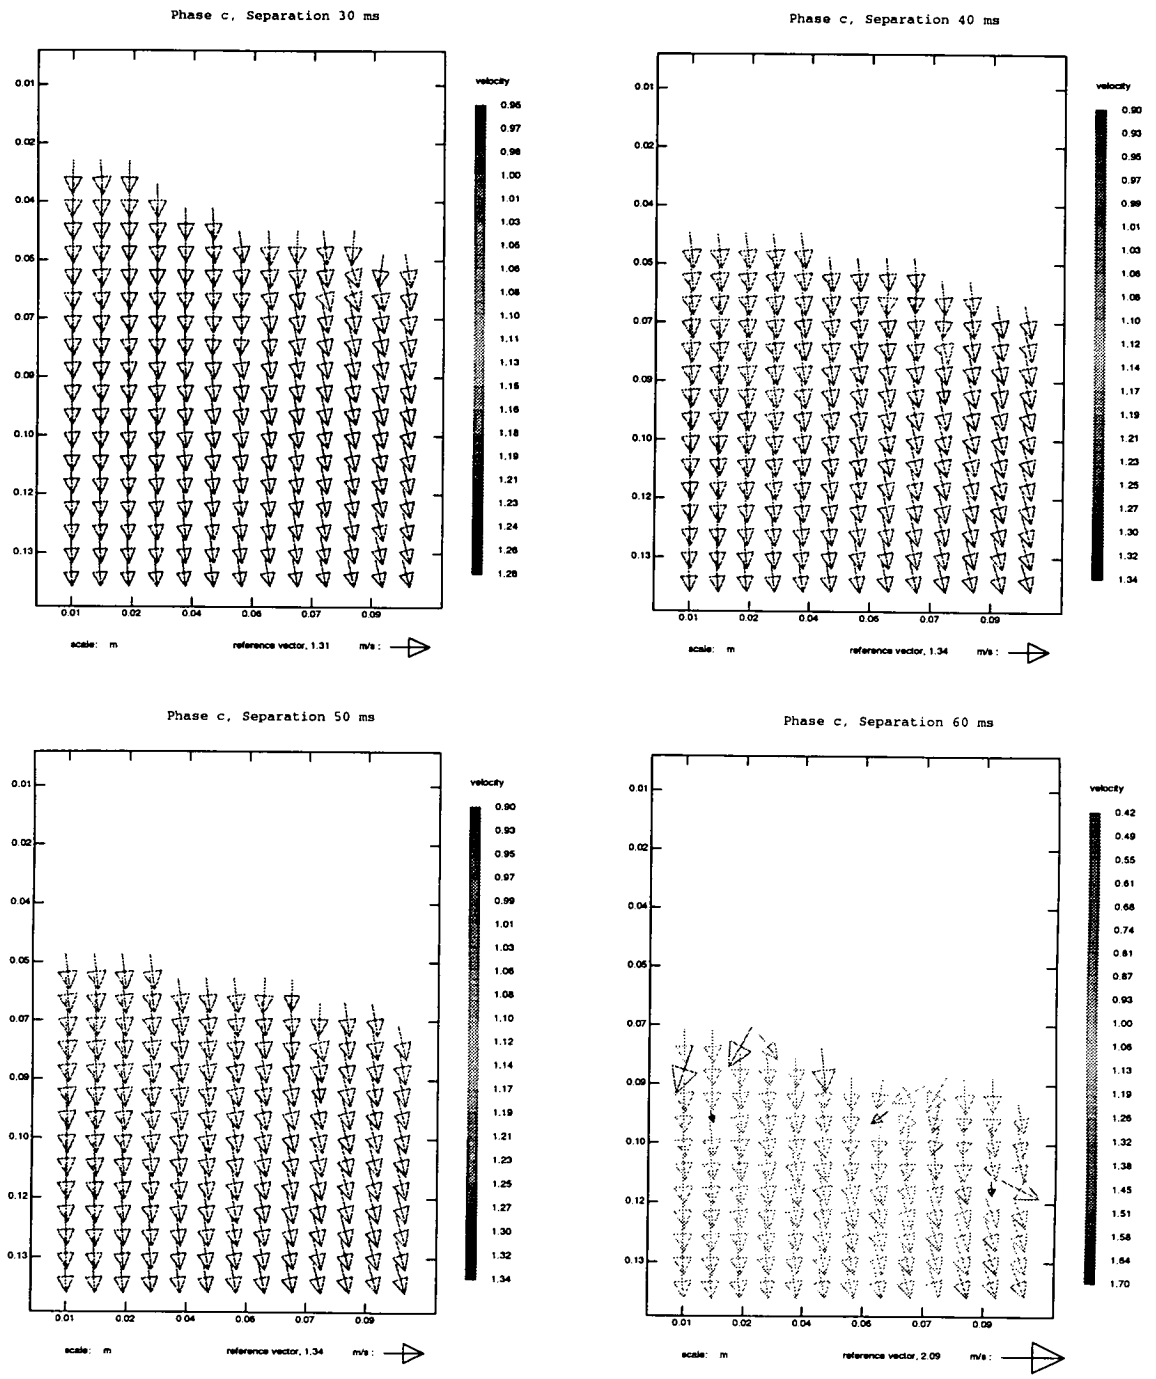

Figure 6.28: Freak velocities, phase point c, n=6,8,10,12

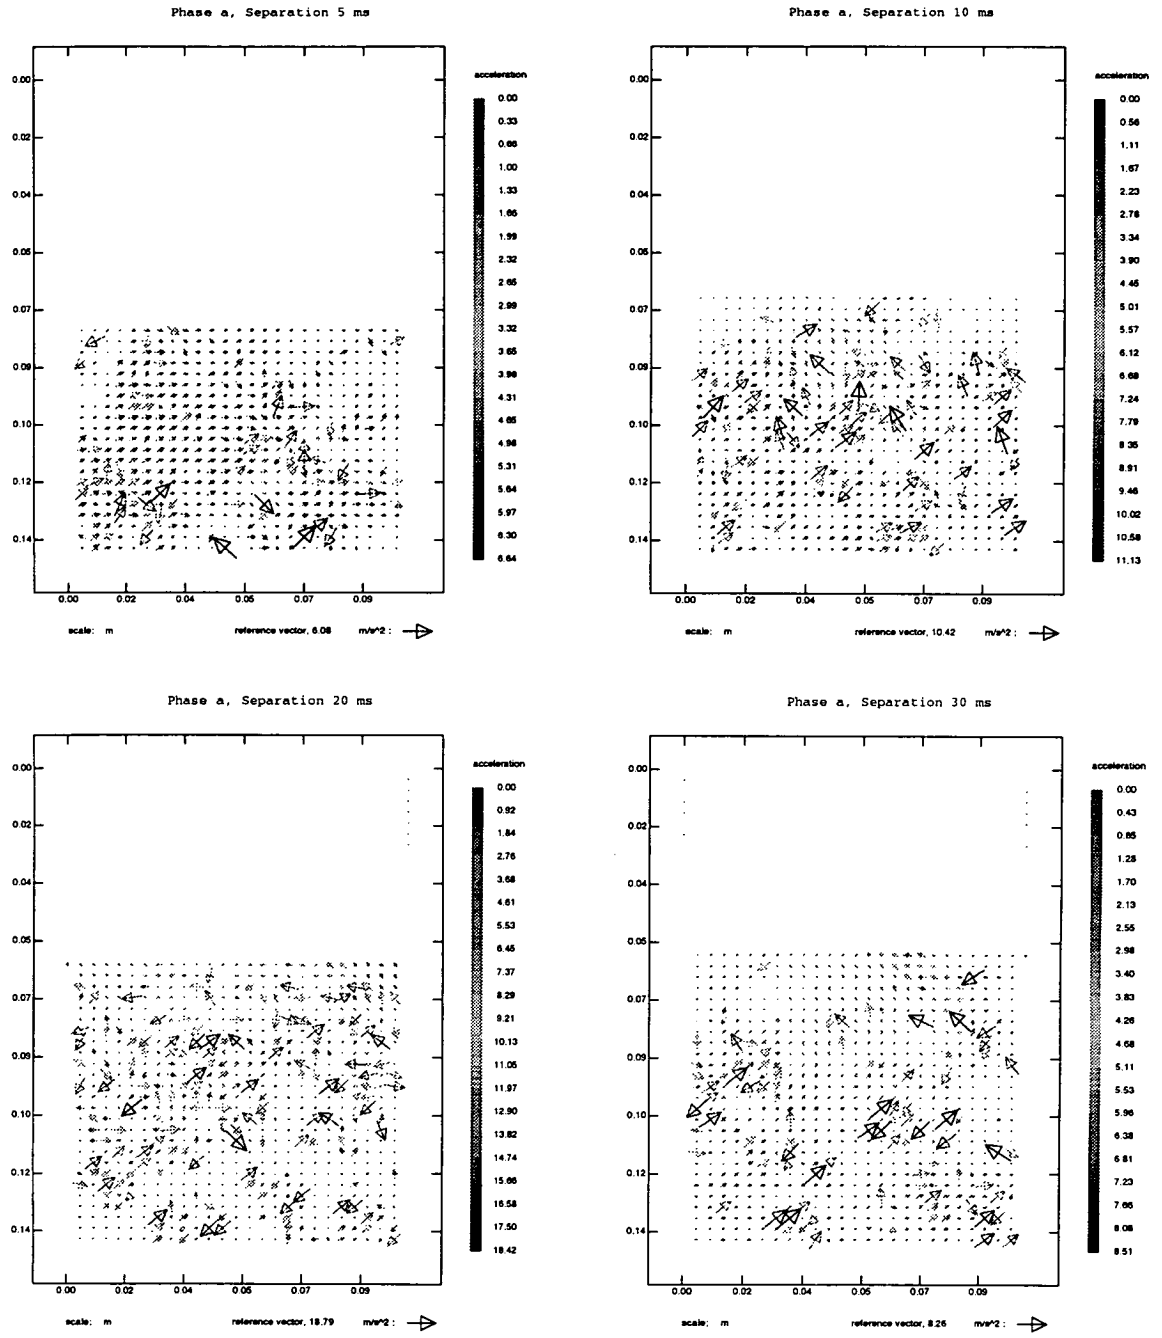

Figure 6.29: Freak accelerations, phase point a, n=1,2,4,6

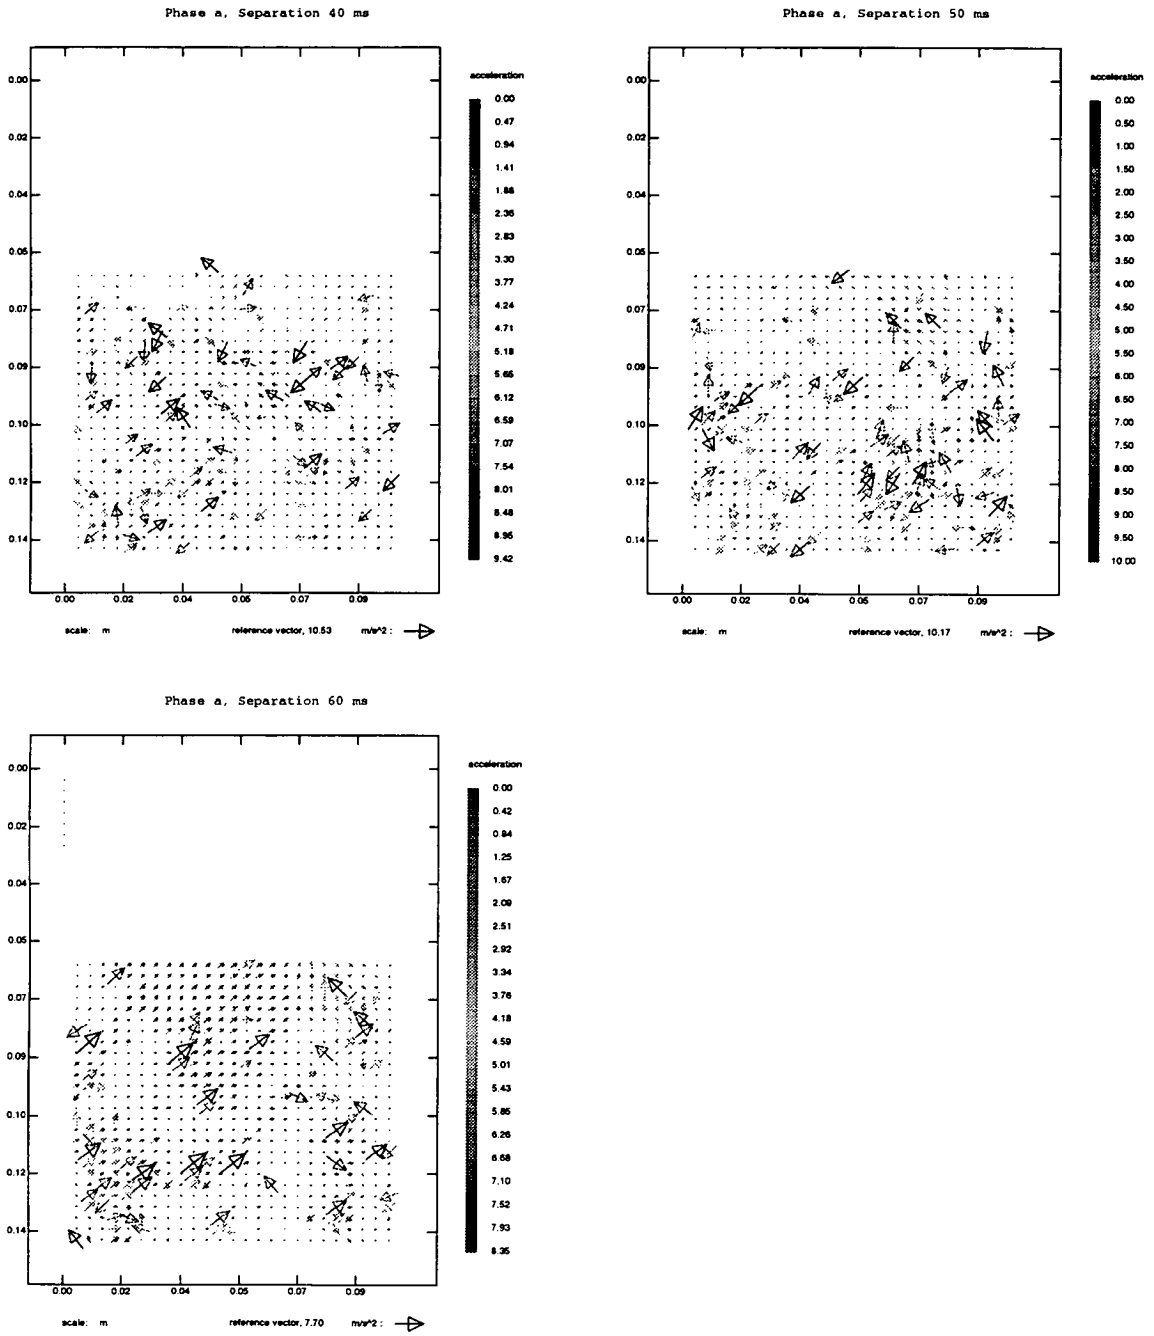

Figure 6.30: Freak accelerations, phase point a, n=8,10,12

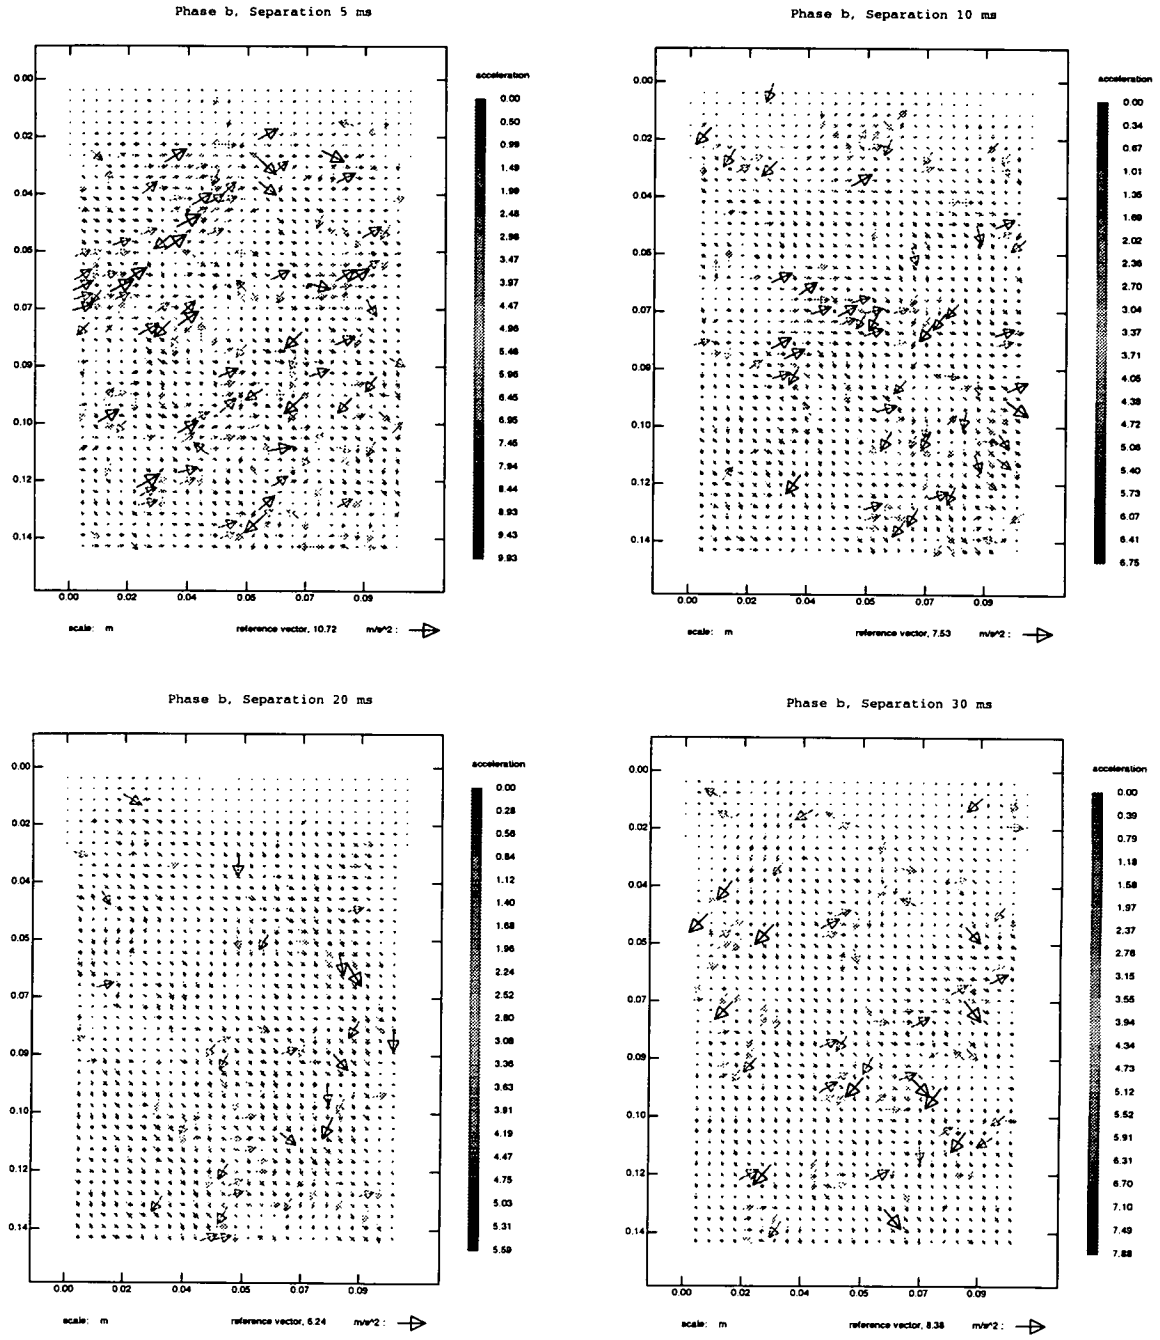

Figure 6.31: Freak accelerations, phase point b, n=1,2,4,6

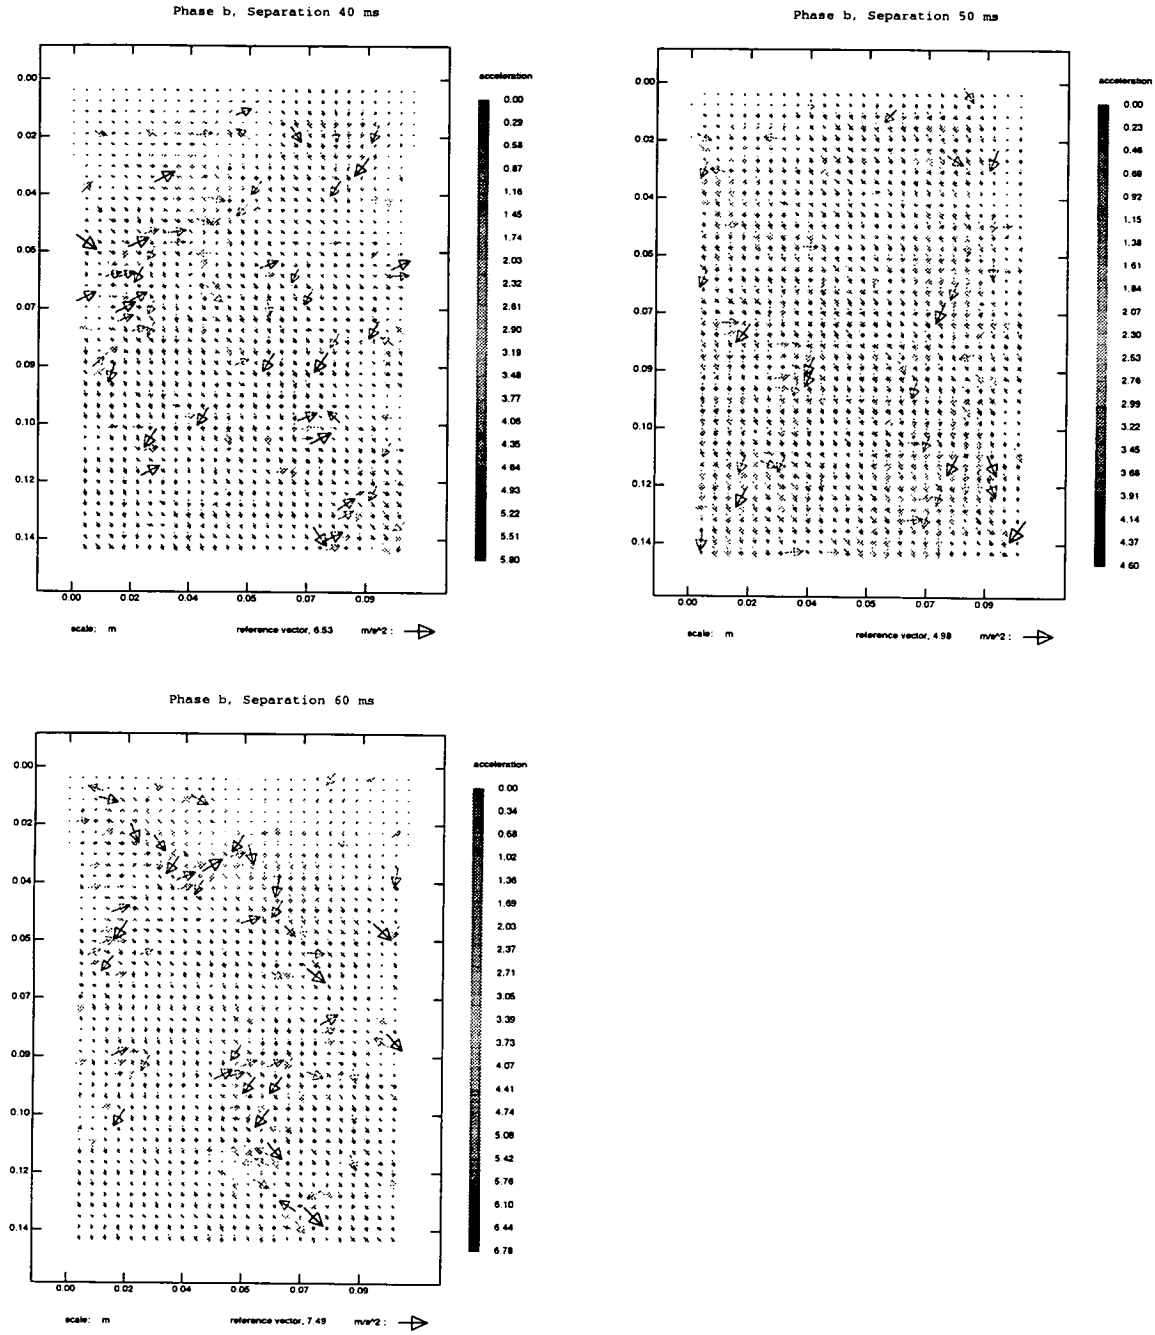

Figure 6.32: Freak accelerations, phase point b, n=8,10,12

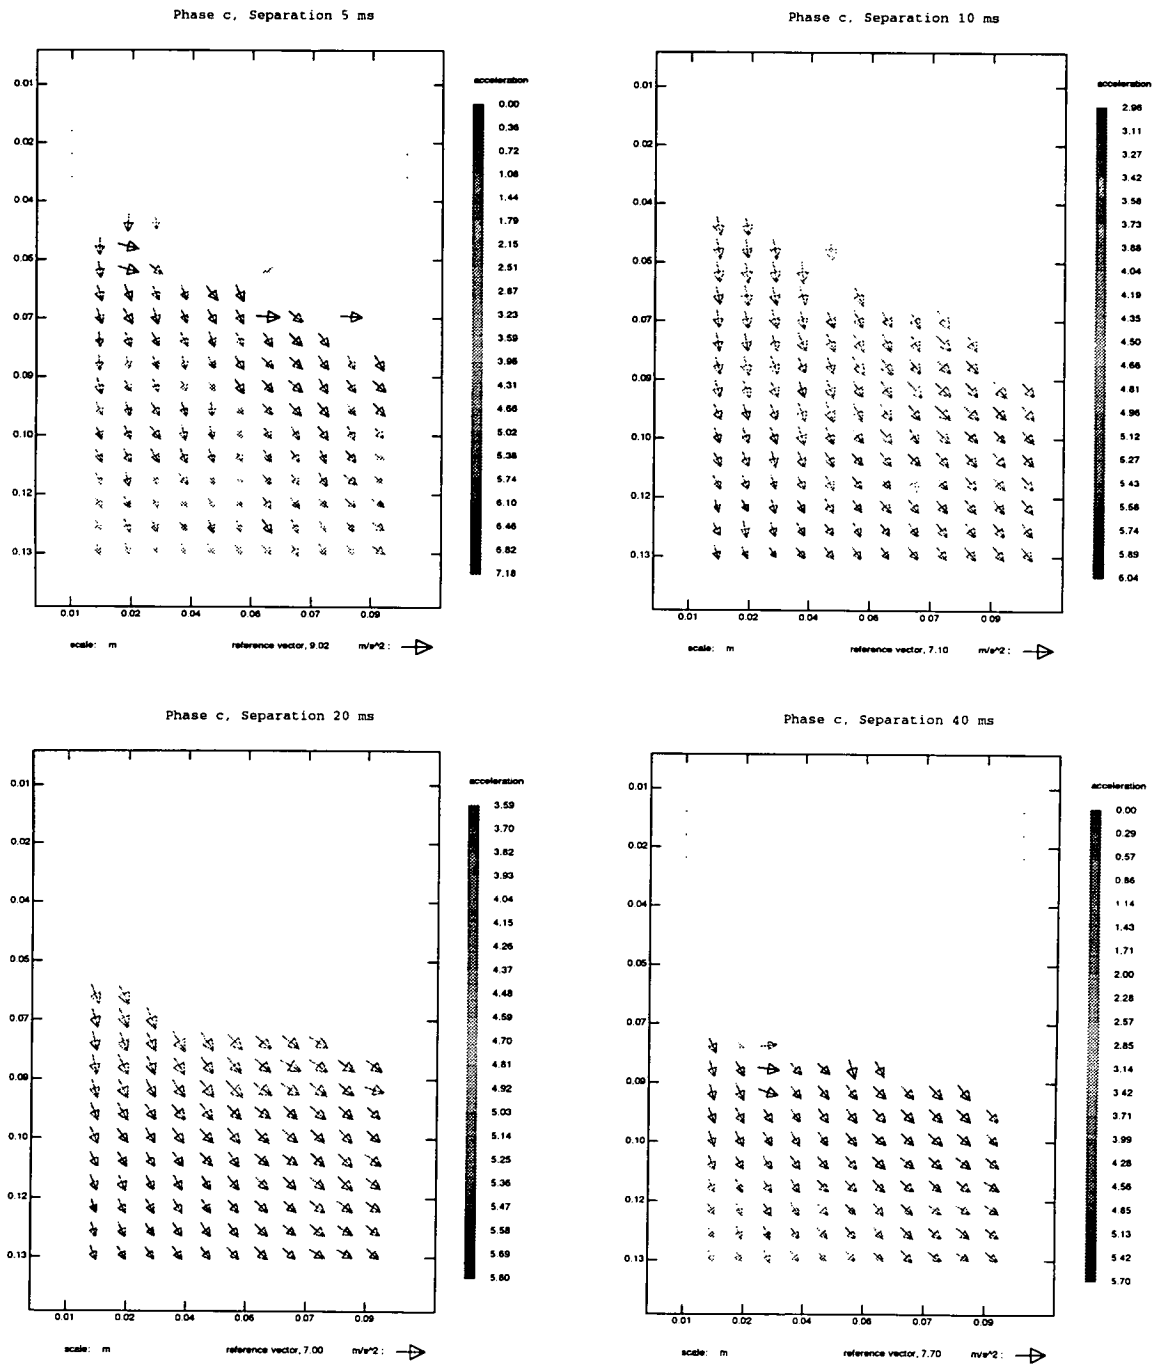

Figure 6.33: Freak accelerations, phase point c, n=1,2,4,6

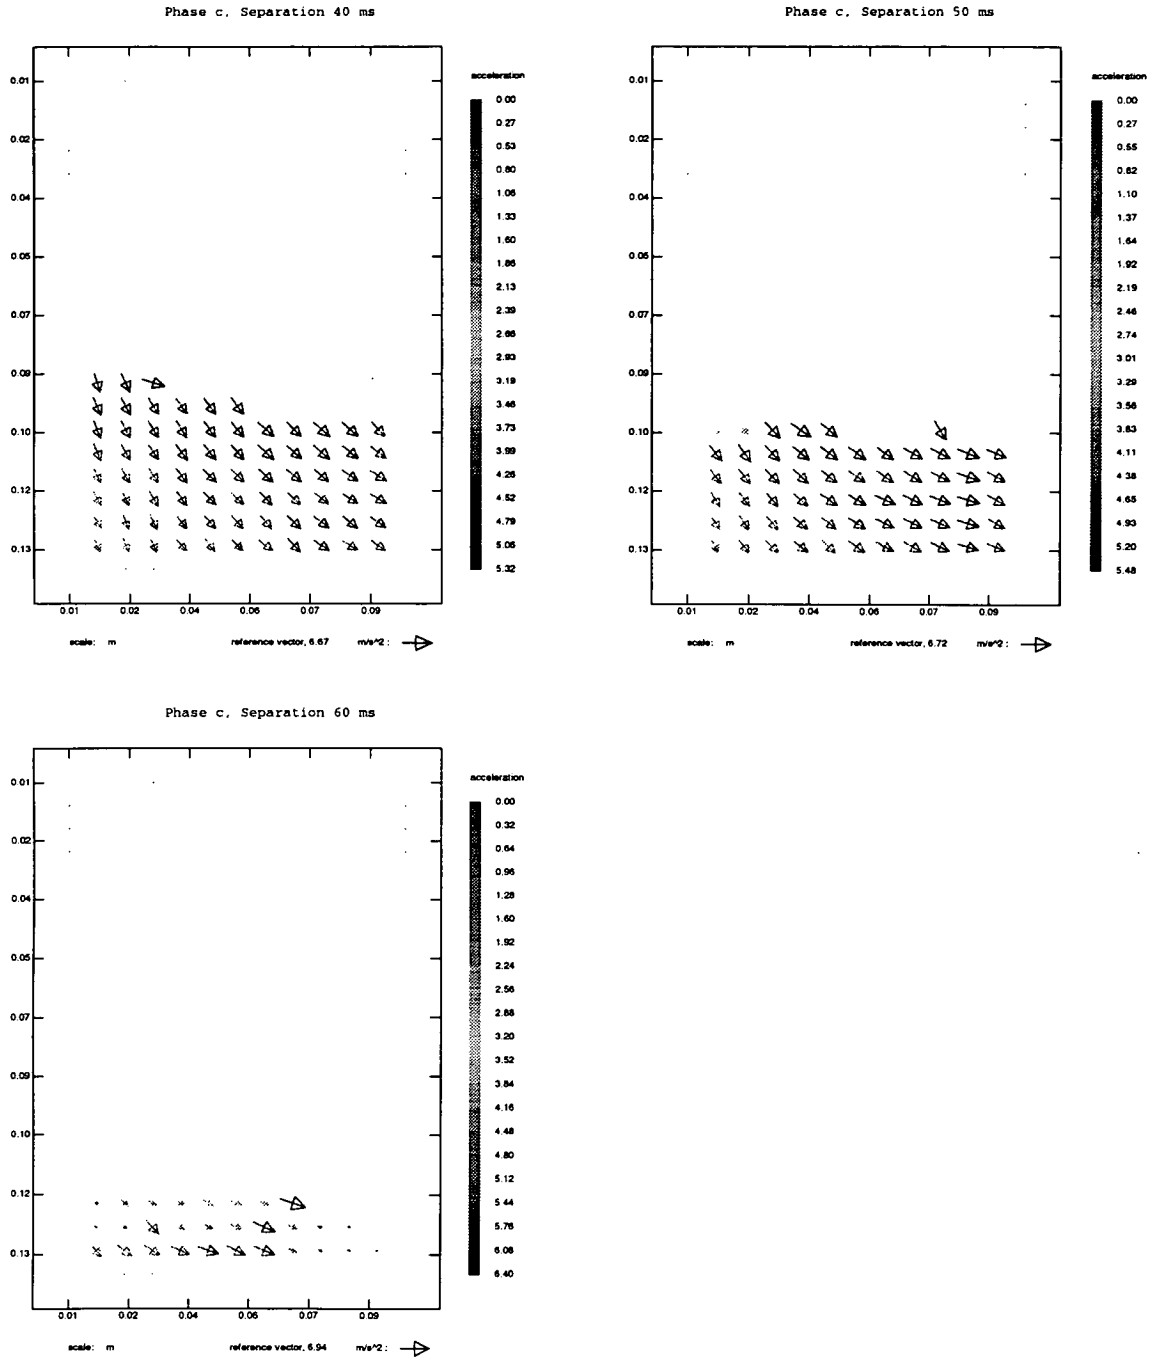

Figure 6.34: Freak accelerations, phase point c, n=8,10,12

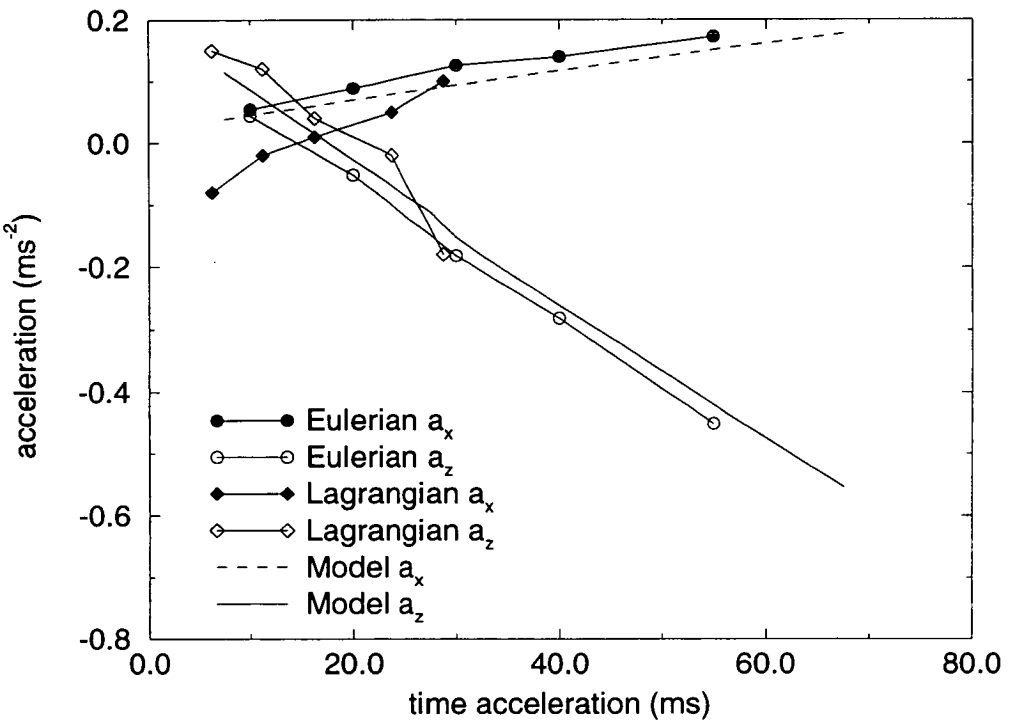

Figure 6.35: Comparison of model and simulated PIV images, phase point a

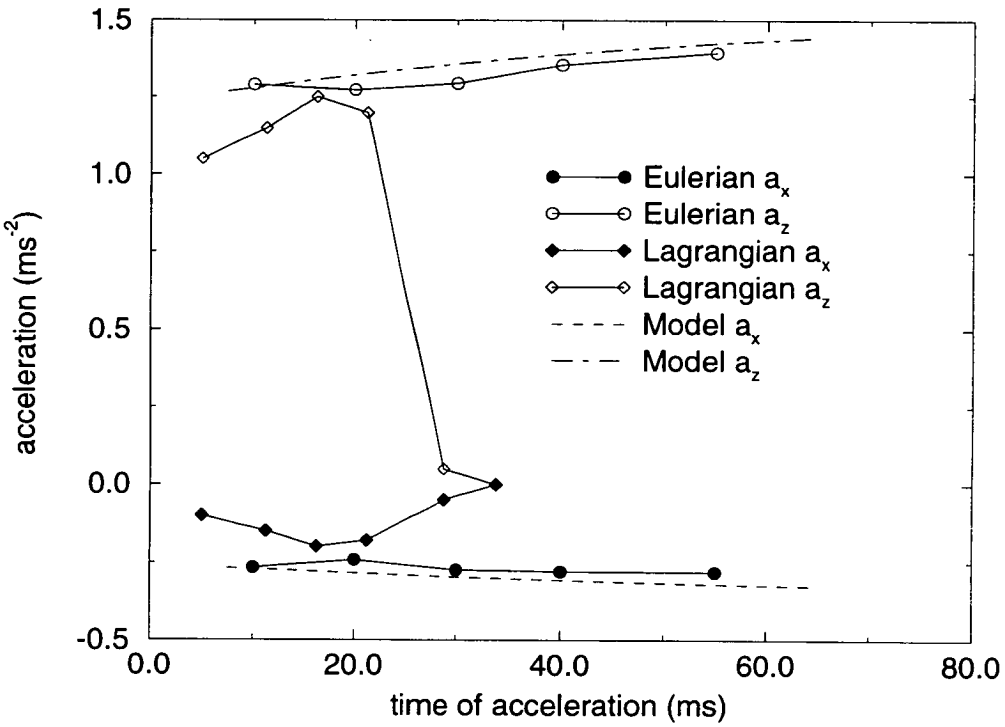

Figure 6.36: Comparison of model and simulated PIV images, phase point b

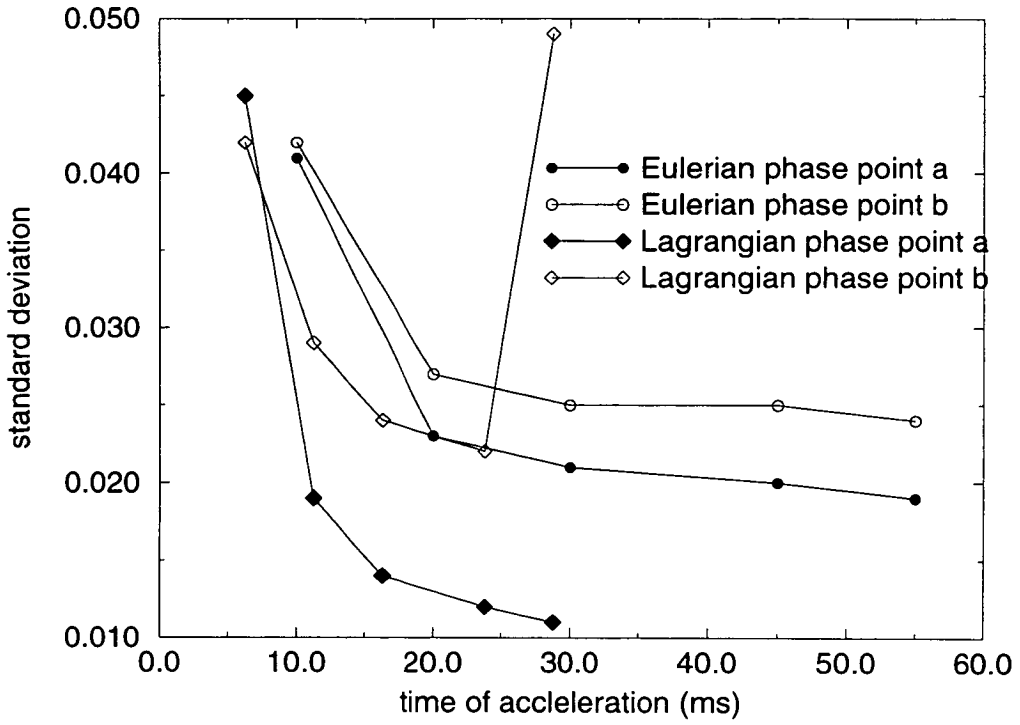

Figure 6.37: Measure of the standard deviation of the acceleration measurements

time separation between the two velocity measurements decrease, and in the case of phase point *a* becomes smaller with increasing time separation.

However, the effect of constantly increasing the inter velocity time can be seen to have a large effect on the resulting Lagrangian acceleration measurement. Looking at phase point *b*, it can be seen that the acceleration measurement fails as the particle group becomes deformed under the action of the acceleration, until finally the correlation is lost, resulting in a random measurement. While these spurious vectors could be removed, with n=14, 50% of the acceleration vectors are already lost, mainly from the top of the vector map.

Figure 6.37 illustrates the standard deviations gained from the spatially averaged acceleration maps. The form of the standard deviation can be seen to

approximate the hyperbolic behaviour that was expected from section 6.4. Further, it can also be seen that the standard deviation of the Lagrangian analysis increases as the fluid group becomes dissipated.

As the time interval between the velocity measurements increases, the assumption of having a linear acceleration starts to fail. By comparing the accelerations calculated from the model for comparison with the PIV data and the instantaneous acceleration at that time, we find for the two phases a 1% error is introduced when the time interval between the velocity measurements exceeds n=6 for phase point a and n=17 for phase point b.

The hyperbolic behaviour of the standard deviation leads us to conclude that there is an optimum time separation between velocity measurements. This can be seen to be a compromise between the minimum uncertainty due to the effects of correlation noise ($n \leq 6$ for both phase point a and b) and the assumption of the acceleration remaining linear being broken ($n \leq 6$ for phase point a, $n \leq 17$ for phase point b) or losing track of the fluid particle using the Lagrangian specification ($n \leq 14$).

6.10 Comparison of Real PIV Images and Analytical Solution

The results from the experiments carried out in the wave flume have been processed and analysed in the same manner as the simulated PIV images.

It can be seen in from figures 6.38 and 6.39 that at both phase points the modeled vertical acceleration under estimates the experimental prediction. While the simulation software was altered to include the higher order terms present in the real wave, the results are still under estimated. It could be the case therefore that the choice of a linear wave theory model was insufficient to accurately describe

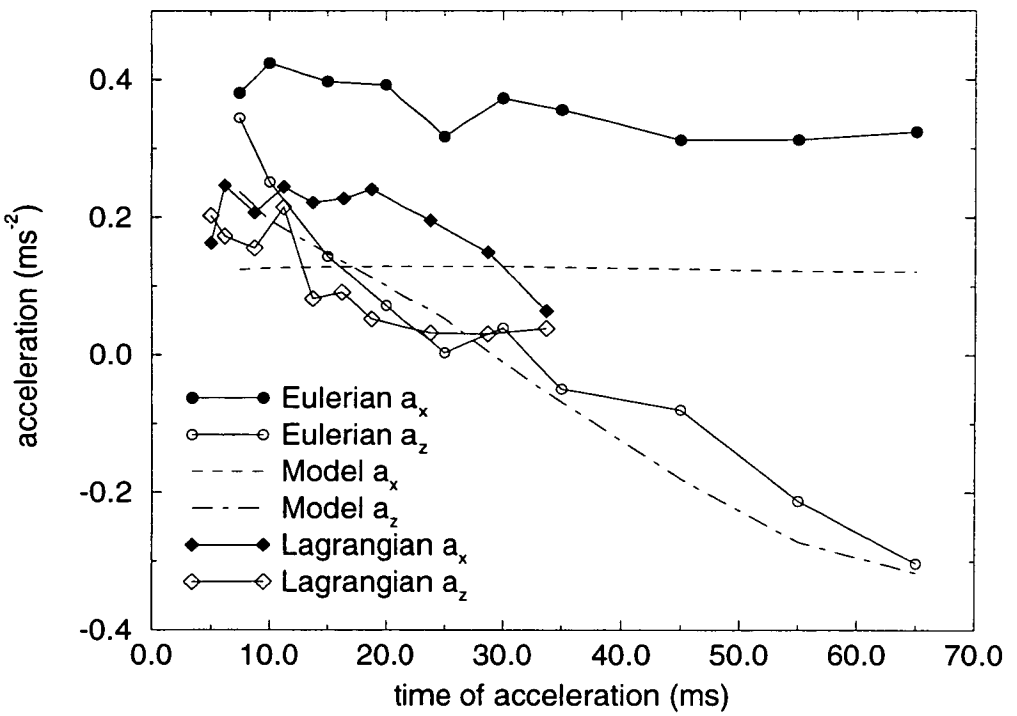

Figure 6.38: Comparison of model and real PIV images, phase point a

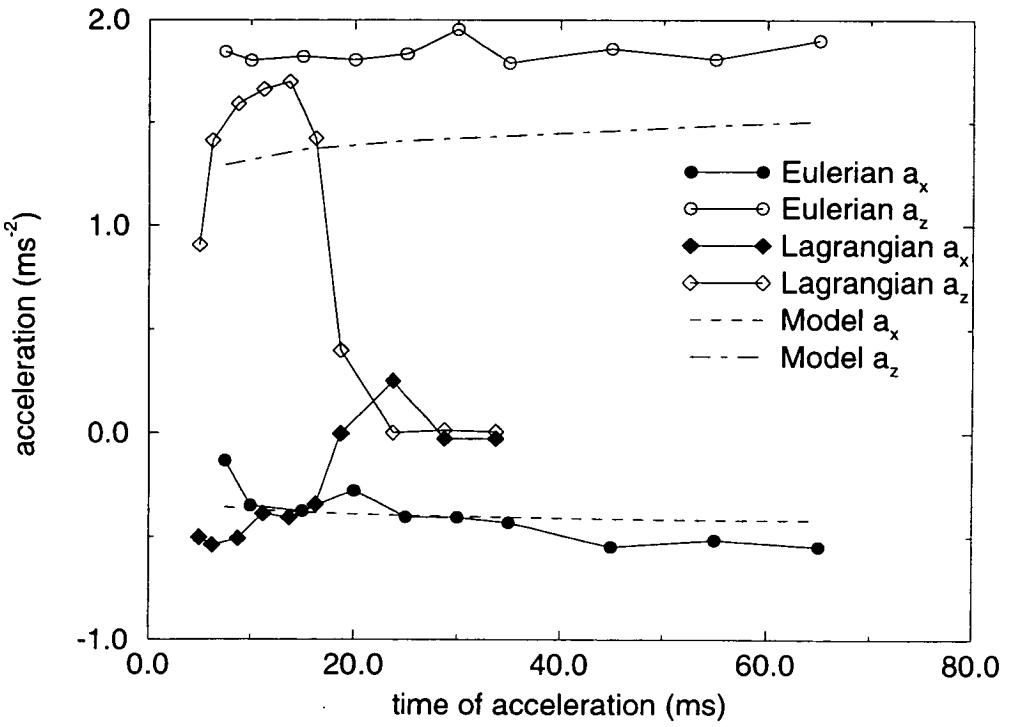

Figure 6.39: Comparison of model and real PIV images, phase point b

the real world situation, with perhaps a model based on Stokes finite amplitude wave theory being more appropriate.

Splitting the Eulerian prediction into its temporal and convective terms, an indication of the relevance of acquiring both temporal and spatial information from the flow under investigation can be gained. The two phase points for the standing wave both give a maximum convective contribution of +10% relative to the maximum acceleration. For the freak wave, the first two phase points give a maximum convective contribution of +15% of the maximum acceleration, while at the third phase point up to 50% of the maximum acceleration is due to the convective component. However, with such large accelerations present in the flow, the validity of the assumption of a linear acceleration have to be brought into question.

6.11 Pressure Measurements

A further way to validate the acceleration fields calculated from the PIV measurements is to predict the pressure drop in the fluid flow along a solid boundary where simultaneous pressure readings can be made for comparison. Making use of the equations in section 6.2, and especially equation 6.4, we can integrate in a narrow section next to the wall, where it is assumed that the velocity is almost constant and uniform. In this case, the integration was performed by averaging the acceleration maps over an area bounded by the pressure transducers, the wall and a vertical line 0.075m into the flow.

The result of this integration can be seen in figure 6.40. The dots represent the difference between the middle and top pressure transducers during the freak wave event. The pressure difference excludes a static pressure difference between the pressure transducers, being 490Pa.

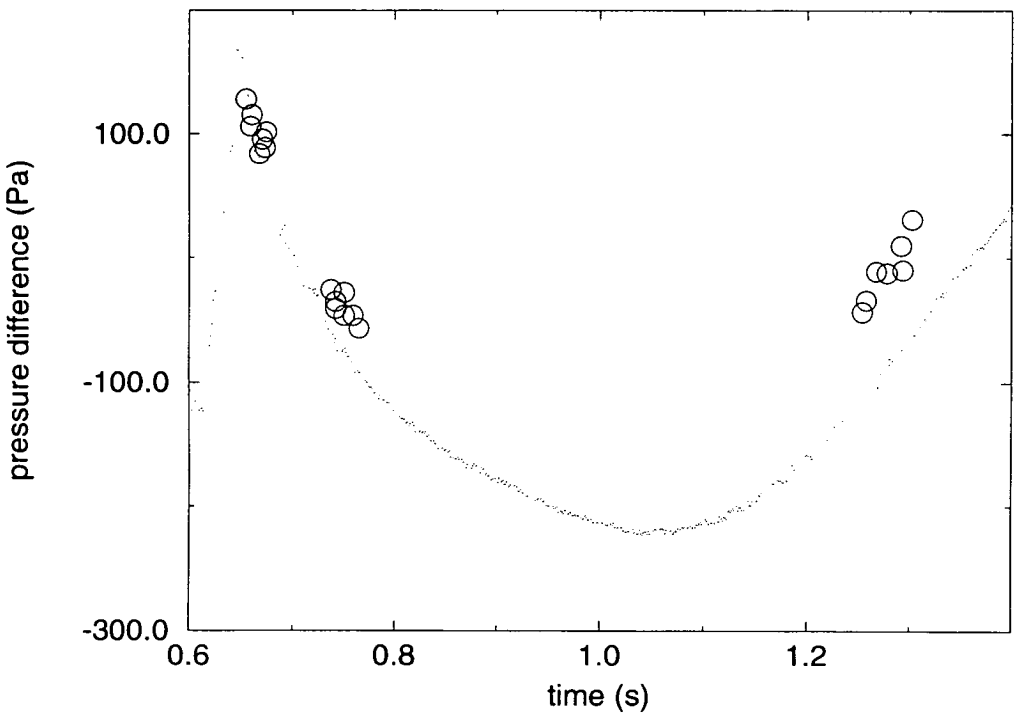

Figure 6.40: Comparison of pressure difference gained from transducers and PIV

The pressure data gained from the PIV measurements are plotted as the large circular points, and for the first two phase points there is a quite good match. The third point, however, differs quite significantly from the transducer values. It is believed that this can be explained due to the large non linear accelerations that are acting at this phase point. Another possible reason is due to the larger interrogation areas that were used during the analysis of this phase point. These reduce the maximum wave frequency measurable to 4.6Hz, while the Fourier analysis of the surface profile reveals significant contributions from frequency of up to 9Hz.

6.12 Conclusion

In this chapter, the four CCD-array camera has been used to capture a sequence of images in close succession of two different types of wave. These PIV recordings have then been analysed by two different methods to reveal the total acceleration present within the fluid flow.

A model of the harmonic wave motion was used to generate simulated PIV images, as well as providing a method of validating the results from both the simulated PIV recordings, and the real PIV recordings.

Of the two flow specifications, the Eulerian has shown itself to be the most robust, providing acceleration measurements which lie within 2-3% of the real acceleration value for an optimal velocity separation of n=6. The limitations of the Lagrangian specification have been demonstrated.

The improvement of this technique could finally yield a tool for measuring the total accelerations present within fluid flows. However, the ever increasing capabilities of technology will undoubtedly finally allow a flow to be completely spatially and temporally resolved. Once this is achieved, the calculation of tem-

poral quantities like acceleration will become much simpler.

The Lagrangian implementation could also be improved, mainly by use of an iterative process, allowing a greater tolerance to the effects of particle group dispersion. However, the robustness of the Eulerian method allows acceleration measurements to be made without the computational overhead of an iterative process.

Chapter 7

Further Applications of the Multiple-CCD Array Cameras

The four-CCD array system was originally designed to be used to capture a sequence of four images with short time separations, a feature not available with conventional CCD cameras. In the measurement of accelerations this is obviously essential, but the system can be used for other applications. Some of these applications are outlined now.

7.1 Dual Plane 3-D DPIV

The basis of the PIV technique is a correlation, with the volume of the correlation peak being directly proportional to the number of particle images which contributed to that correlation peak. In ideal circumstances (i.e. small velocities which are 2-dimensional and not subject to large amounts of shear) this holds true, but in real flows there are always spatial gradients in the velocity field present, along with out of plane or three dimensional components. In conventional PIV all of these factors would combine to reduce the detectability of the correlation peak, increasing the probability of detecting a noise peak in the correlation plane resulting in an incorrect velocity measurement for that interrogation area.

However, this reduction in the height of the correlation peak can be used to

measure the out of plane component of the flow.

In order to measure the out of plane component of the velocity, the illuminating sheet of light is moved very slightly in the out of plane, or z, direction (see figure 7.1). Typically the movement will be of the order of half the light sheet thickness. The light sheet can now have two different axial positions, and to allow information about the out of plane motion to be retrieved three images are acquired, two in the first axial position and a further image in the second axial position.

Correlations are then carried out to find not only the in plane velocities, but also the ratio of the correlation peak heights at different z-positions of the sheet, and combined with the knowledge about the intensity profile of the illumination sheet this can then be used to give a measure of the z-component of the displacement. This is the basis of a technique developed by Raffel *et. al.* [44, 45, 48, 49], and this technique has been successfully applied to several cases including a vortex ring in a water tank ($\approx$ 5-10cm/s), a rotating flow in a glycerine tank ($\approx +/-$ 5cm/s, and computer simulations.

All of these flow cases involve only low velocities, due to the limitations of the CCD technology used in the initial experiments. However, with the four-CCD array system, the minimum time separation between the images which can be acquired is small, allowing it to be used in flows of up to $\mathcal{O}(100ms^{-1})$.

7.1.1 Apparatus

The device for moving the illumination sheet is itself very simple, comprising a rotating disk with three apertures, one of which contains a piece of thick ($\approx$ 5mm) glass. The rotational axis of the disk is placed at an angle to the incoming laser beam, and so as the section with the glass lies in the path of the incoming laser light, the beam is deflected sideways. This deflection is easily calculable, and the

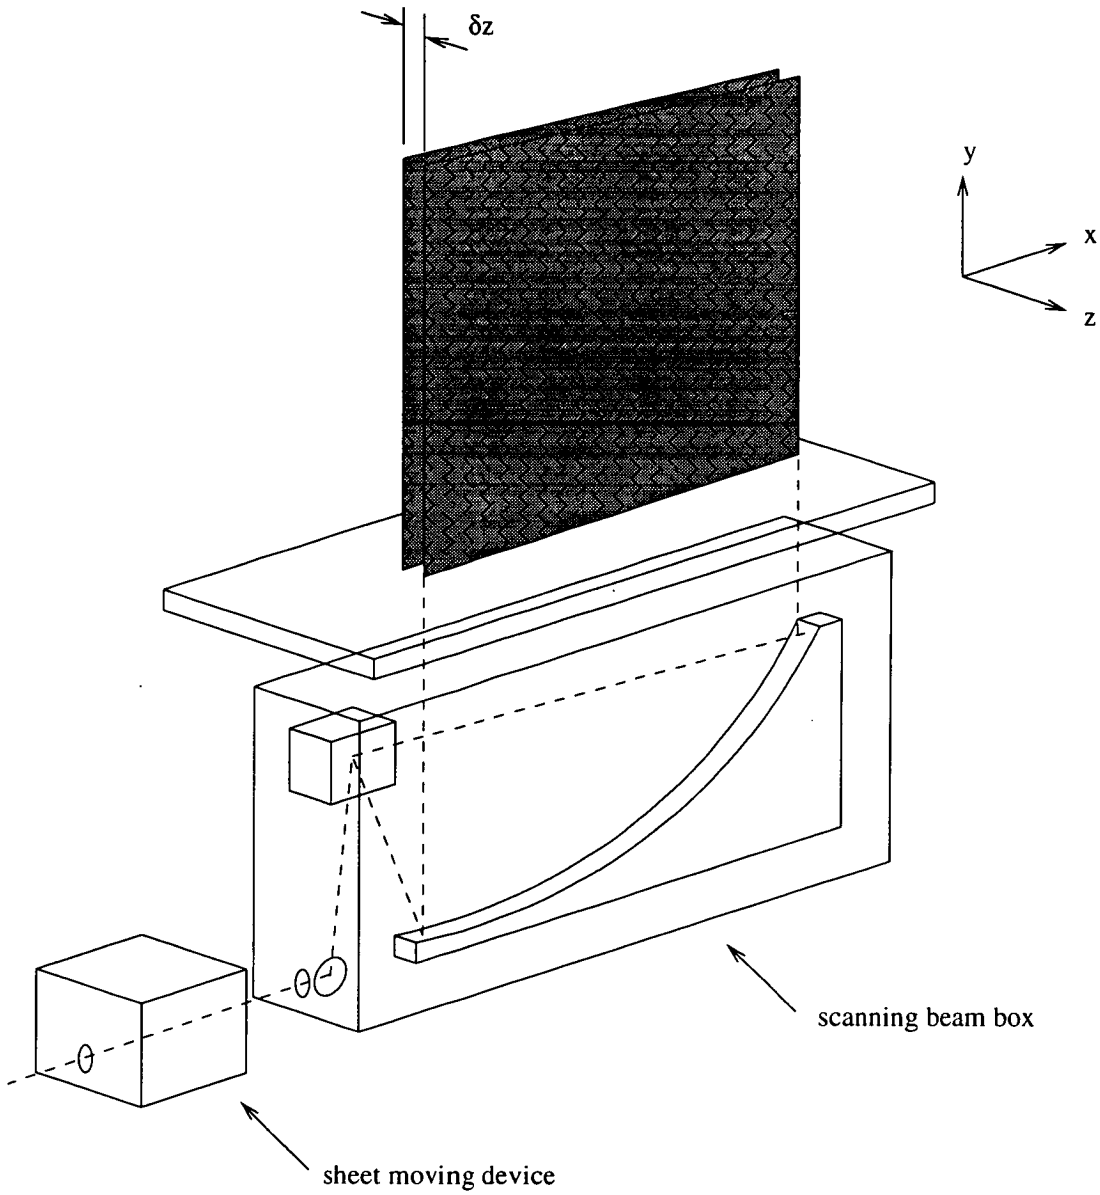

Figure 7.1: Arrangement of the sheet moving device

entire rotating disk and associated drive motor are mounted on a rotating column so that any angle may be set.

The original design of the system was based upon the assumption that the sheet moving device was to be used in conjunction with the scanning beam system which has found widespread use at the University of Edinburgh. As previously mentioned, in order to successfully use a multiple camera system with a scanning beam box, the triggering of the cameras has to be synchronised to the scan rate of the illumination system. This is easily achieved by placing a photo diode at one end of the pseudo-static light sheet. When properly filtered this will give a pulse stream which gives the frequency and phase of the illuminating beam.

The pulse stream which is output is then used in two ways; the first to synchronise the rotating disk to the light sheet to allow individual scans to be shifted, and the second to be used in conjunction with a signal output from the sheet moving device to trigger the cameras.

The advantage of creating a sheet moving device which is synchronised to the rest of the illumination system is that it allows the entire system to be used with a number of illumination schemes with little or no modification required to either the controlling electronics of the sheet moving device or the illumination system. It also provides an easy experimental system to work with as once the system is configured, if the scan rate of the illumination is altered, all other components in the chain are self-synchronising.

The sheet moving device is shown in figure 7.2.

7.1.2 Application

The apparatus described above was built with the aim of carrying out dual-plane PIV with the four-CCD array system. However, due to the time limits of the Ph.D. it proved impossible to actually conduct any experiments. It is hoped that

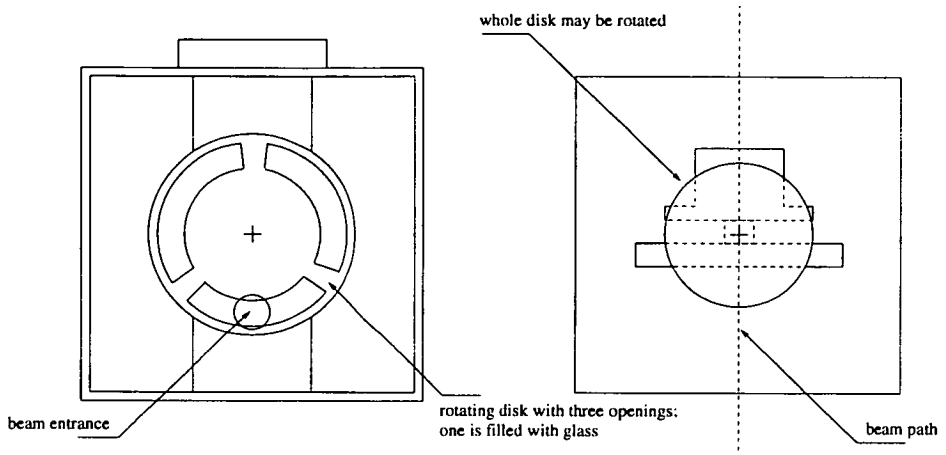

Figure 7.2: Sheet moving device

the equipment that has been built will be used in the future by another student.

7.2 Doppler Global Velocimetry

Doppler Global Velocimetry (DGV) is similar in concept to Laser Doppler Anemometry (LDA), however while LDA is a point measurement, DGV is a full field measurement technique more suited to higher flow velocities.

Doppler Global Velocimetry is again a conceptually relatively simple technique. The flow under investigation is seeded with particles, and a plane of the flow is illuminated with monochromatic light. The light scattered from the particles is Doppler shifted, with this Doppler shift giving the velocity of the fluid flow. However, in order to measure the Doppler shift an additional piece of equipment is required.

The insertion of an iodine cell in between the flow and the recording medium will convert the frequency shift of the scattered illumination light into an amplitude shift. This amplitude shift can then be recorded and processed to yield a velocity map. However, the effects of non uniformities in the illumination process, experiment, etc. have to be accounted for, and it is not often that a completely

uniform flow image will be present. Therefore a reference and a signal image have to be taken, the reference signal being taken without the iodine cell being present.

The multiple CCD array system is ideal for this application, as the capture of the reference and signal images can be achieved simultaneously and using the same set of imaging optics. The placement of the iodine cell in between the output side of the beam splitter and one of the CCD arrays will provide the signal image. A neutral density filter placed on the other output side of the beam splitter will compensate for the losses due to the iodine cell. Thus, both the reference and signal images can be acquired by one system.

7.3 High Speed Photography

Another application for the four-CCD array system could be high speed photography. As the optics of the multiple array systems have been designed to allow the use of standard 35mm format lenses, the system can be used for normal imaging of any scene. If each camera is operated in a free running mode (i.e. allowed to produce 30 frames per second) but staggered equally with respect to each other, the result is a system which can produce 120 full frames per second.

Another option is to capture only a burst of four images. This seems at first to be an unattractive option until it is realised that the cameras can be triggered with an uncertainty of maximum $5\mu s$. The internal storage in the Pulnix TM-9701's can then be used to store the images indefinitely until they are captured and stored by a computer. This mode of operation allows the recording of high speed events to be captured.

7.4 Conclusion

While the multiple CCD array systems were developed for use in PIV, the main reason for their construction was the limitations of commercially available CCD cameras. The ability to capture a sequence of images in rapid succession has allowed the measurement of accelerations within fluid flows. However, this ability to acquire a sequence of images is not limited to applications in PIV, but can be used in any situation which has a need for a temporal resolution which is higher that possible from commercially available CCD cameras.

Appendix A

Papers Published & Conferences Attended

The following papers have been either published in journals or presented at conferences during the course of my Ph.D. studies.

Journals

- Jakobsen, ML, Dewhirst, TP, Greated, CA (1997) "PIV: Acceleration Fields and Force Prediction within Fluid Flows" *Measurement Science & Technology* Vol. 8, No. 12, pp1502-1516

- Arnott, A, Dewhirst, TP, Pullen, J, Greated, CA (1996) "Underwater CCD Cameras for PIV Flow Mapping" *Journal of Optical and Laser Technology* Vol. 29, No. 1, pp45-50

Conferences

- Dewhirst, TP, Jakobsen, ML, and Greated, CA (1996) "Multiple CCD camera DPIV for Force and Acceleration Measurements in High Speed Flows" *Proc. IMechE Optical Methods and Data Processing in Heat and Fluid Flows*, pp251-258

- Dewhirst, TP, Jakobsen, ML, Greated, CA (1997) "The Measurement of Acceleration Fields Under Waves by DPIV" *Proc. 7th Int. Offshore and Polar Engineering Conference*, Vol. III, pp390-394, Honolulu, USA, May 25-29, 1997

- Dewhirst, TP, Jakobsen, ML, Greated, CA (1997) "Multiple CCD camera DPIV" *Underwater Optics: Light in the Sea II*, Aberdeen, UK, February 1996

- Arnott, A, Dewhirst, TP, Pullen, J, Greated, CA (1996) "Underwater CCD Cameras for PIV Flow Mapping" *Underwater Optics: Light in the Sea I*, Ketering, UK, February 1995

Appendix B

Multiple CCD Array Camera Plans

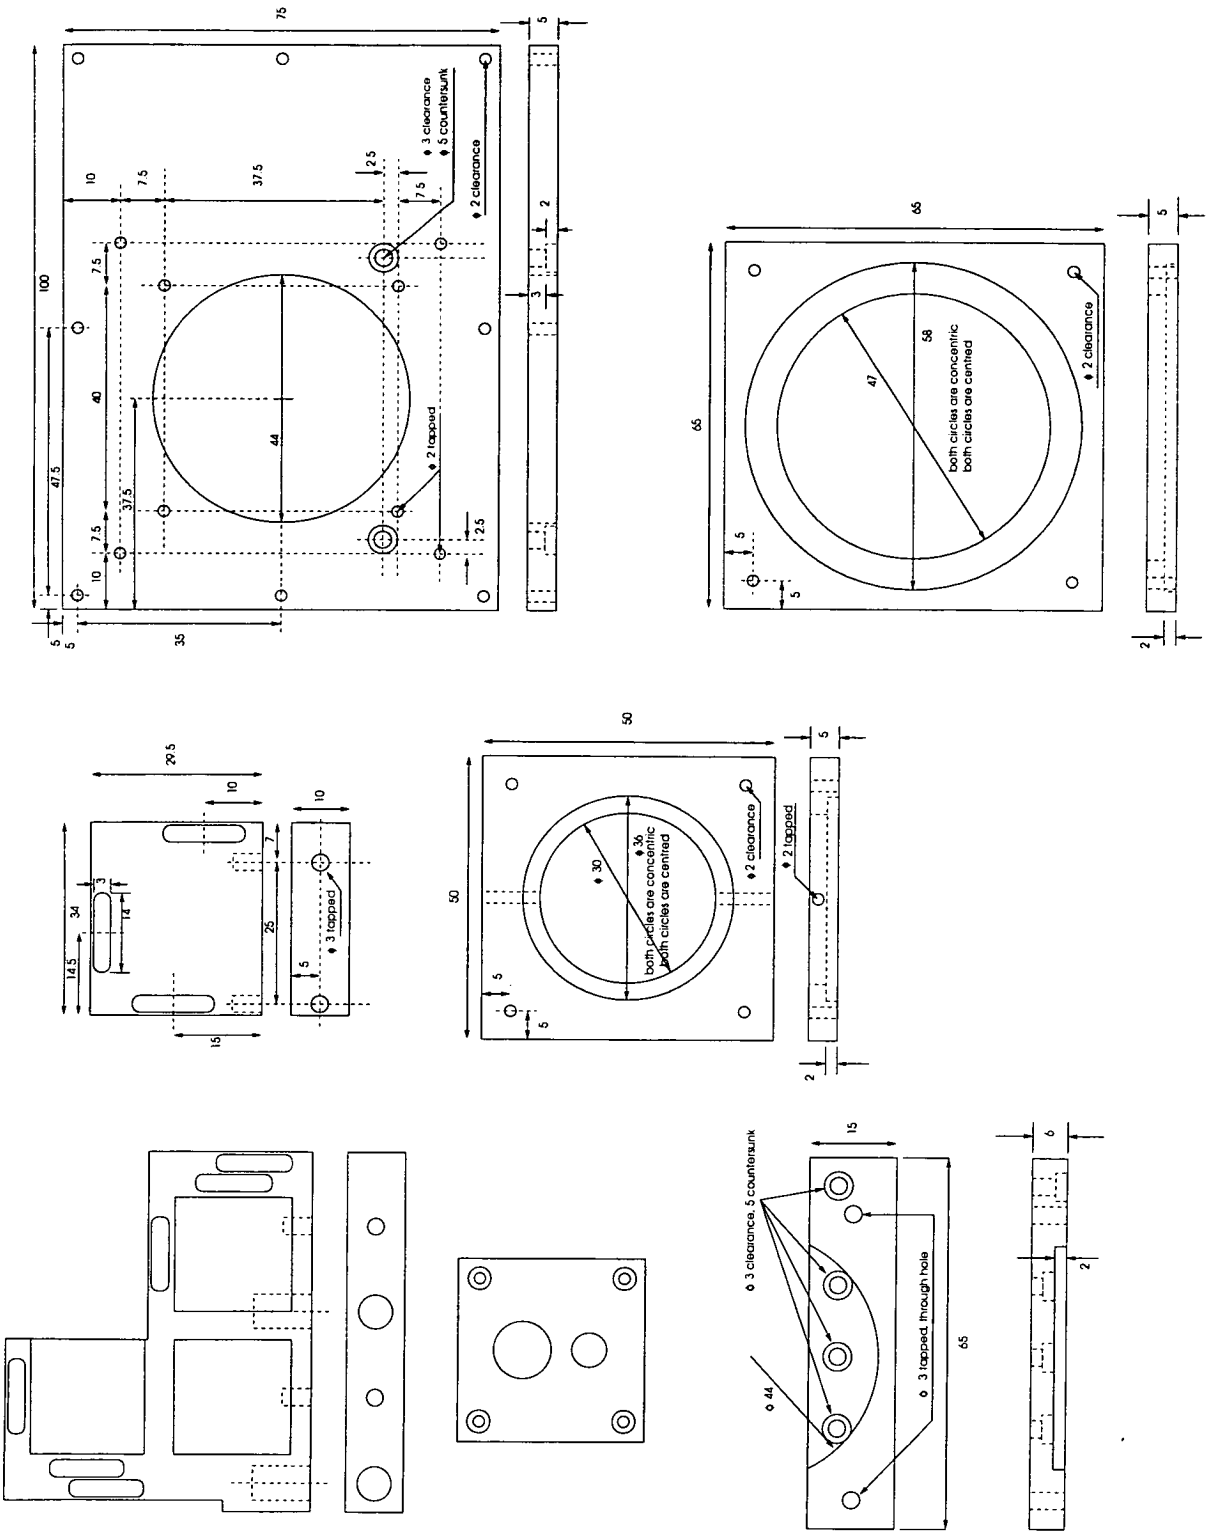

Figure B.1: Component parts of VVL prototype camera

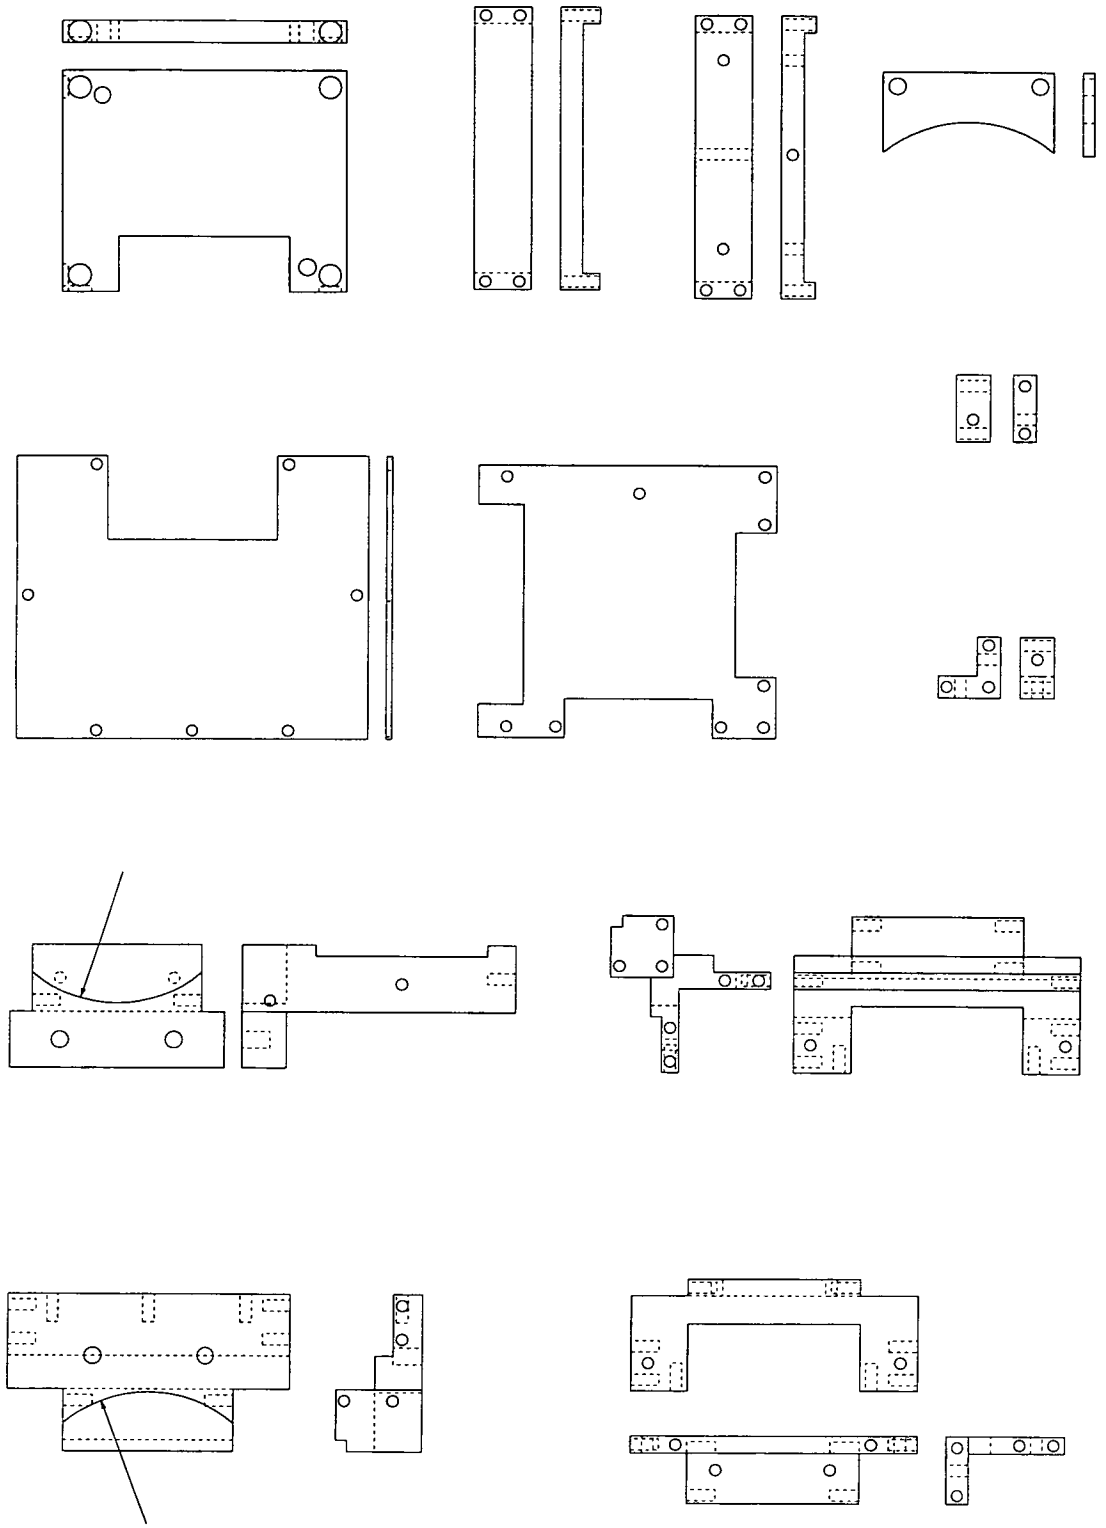

Figure B.2: Components of the mounting for beam splitting cube and CCD arrays

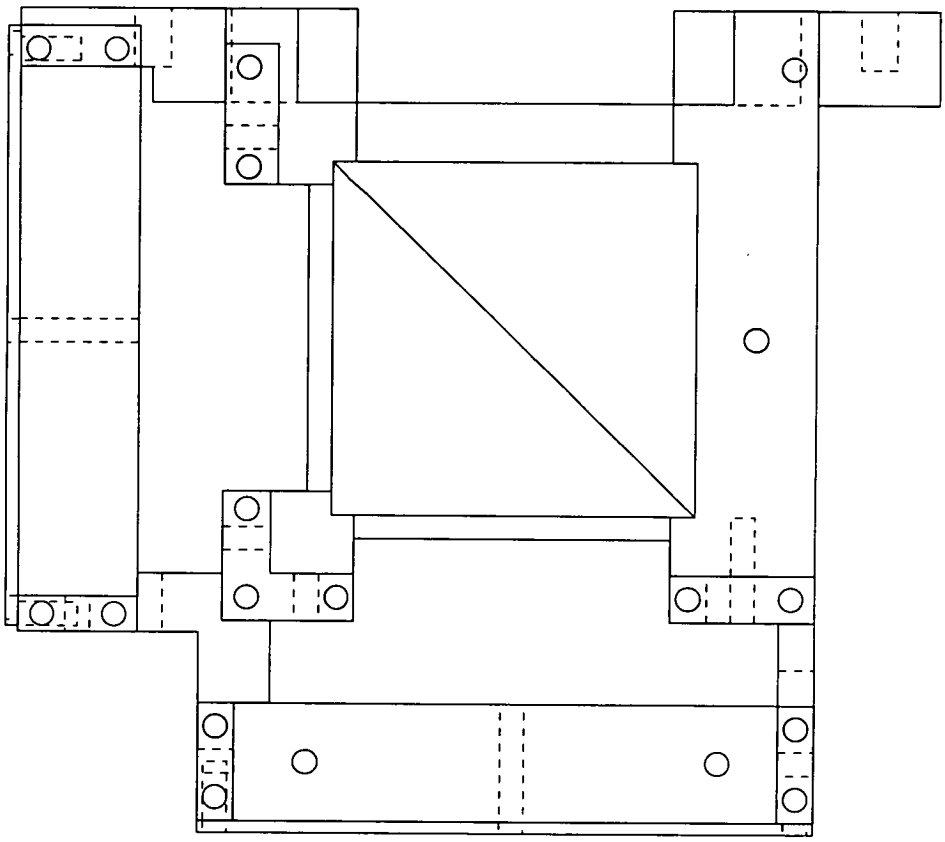

Figure B.3: Assembled components from above

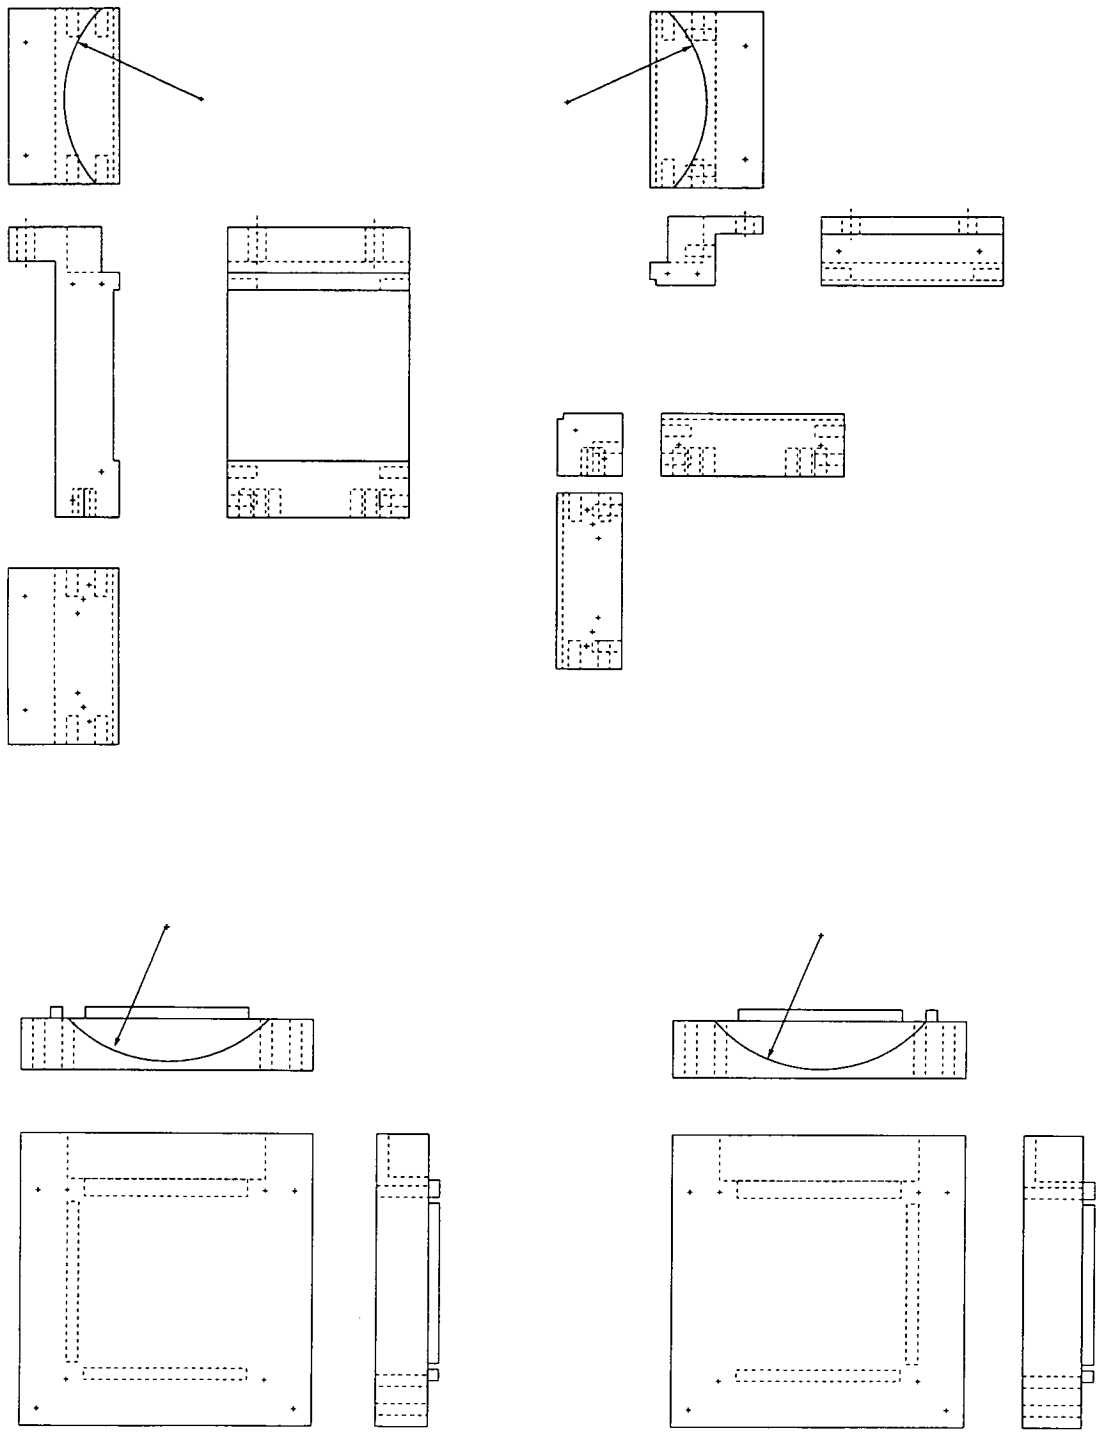

Figure B.4: Components of the mounting for beam splitting cube and CCD arrays

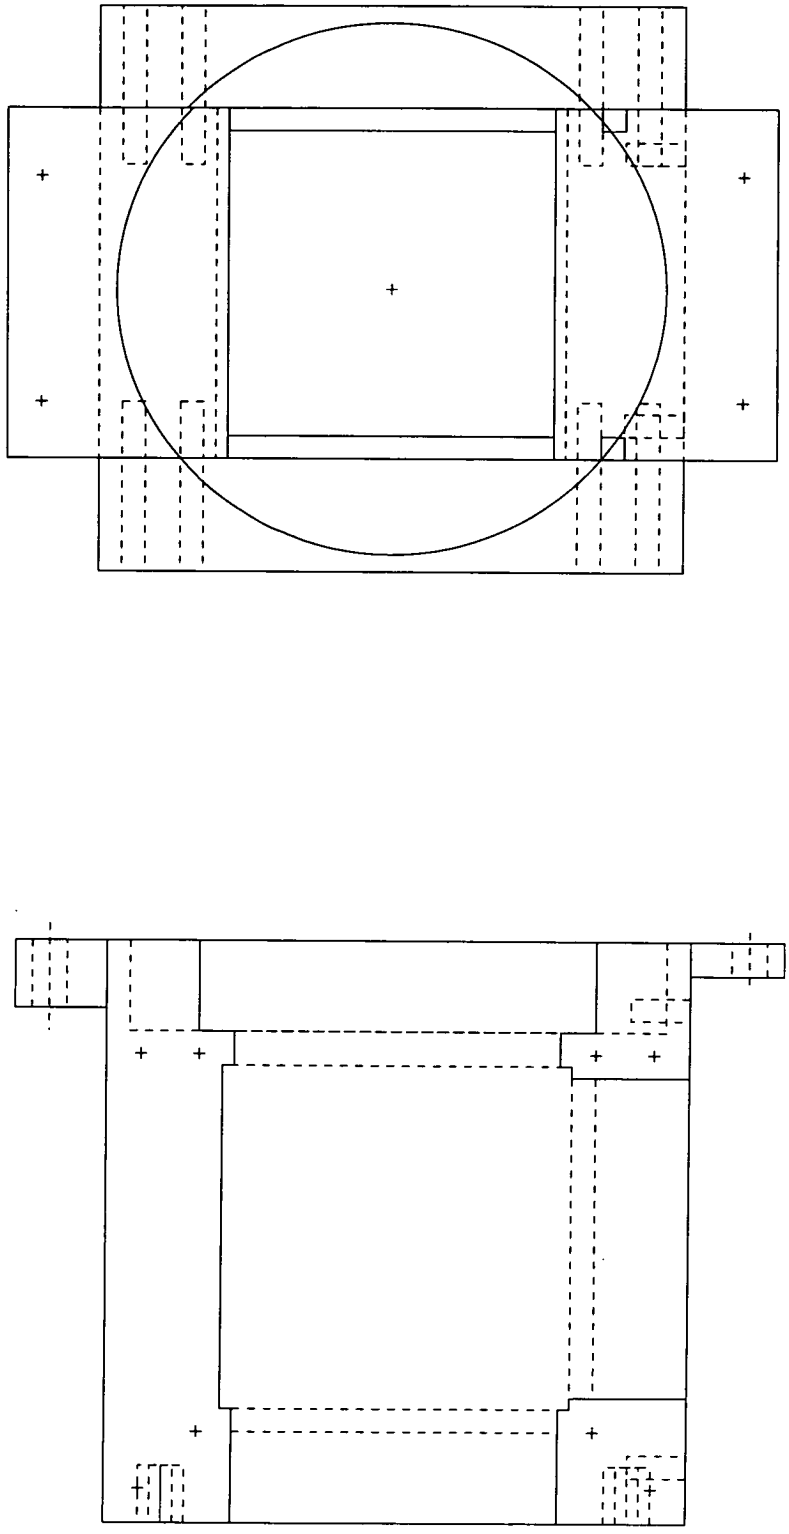

Figure B.5: Assembled components from above and front

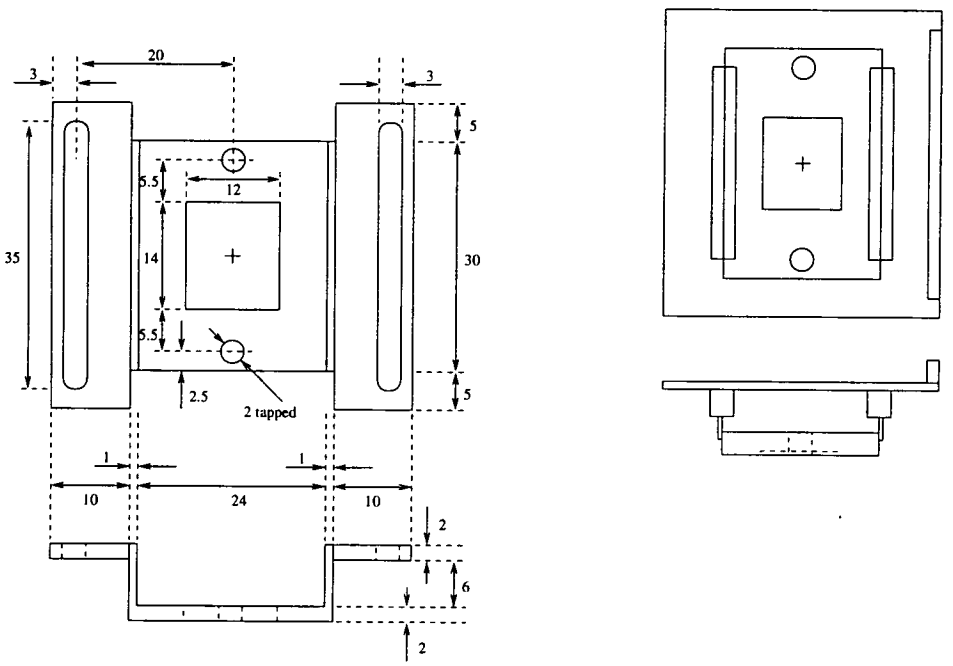

Figure B.6: Mounting plates for CCD arrays

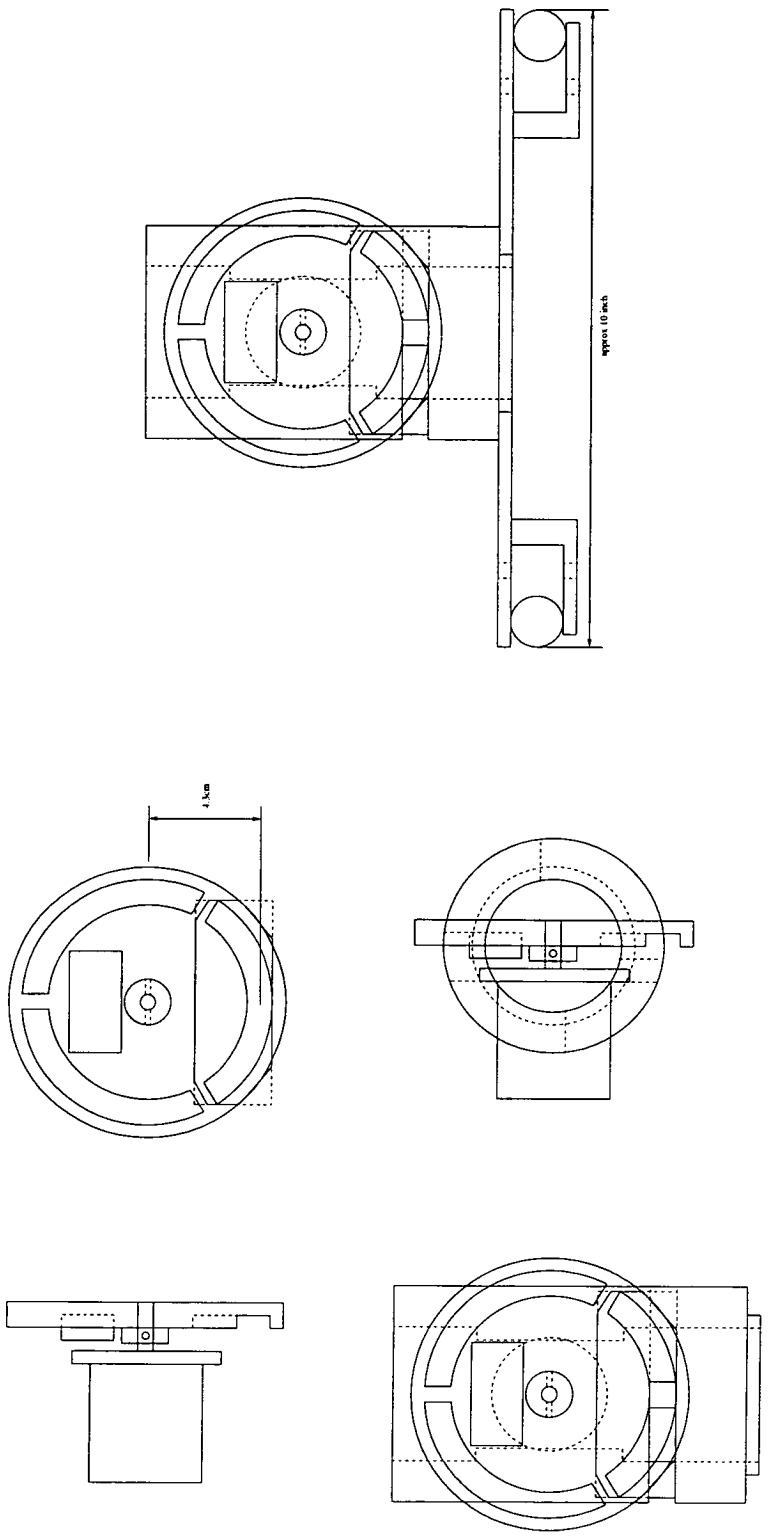

Figure B.7: Components of sheet moving device

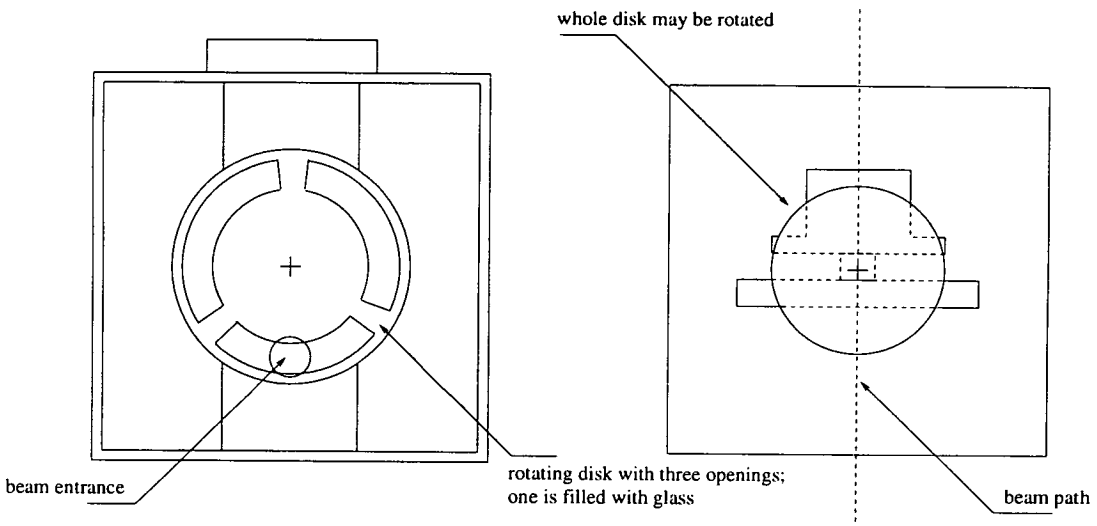

Figure B.8: Assembled sheet moving device

Appendix C

PIV Analysis Software

To allow acceleration measurements to be made from the images captured by the four-CCD array system, software had to be written which was not available commercially. However, as velocity measurements are required before an acceleration measurement can be made, the basis of the analysis software is essentially the same algorithmically as every other commercial and non-commercial DPIV analysis package.

While the algorithms behind the analysis are commonly enough known, the real effort was in creating a relatively large piece of software from scratch. The software also became more complex due to the necessity of a Graphical User Interface (GUI). It was felt that a GUI was required as the analysis of PIV images is time consuming enough without the extra difficulty of trying to optimise the processing parameters blind. Displaying the velocity vectors as they are derived from the PIV images allows a quick check on the quality of the end result, and can be used while in the lab to check on the parameters of the experimental setup.

C.1 EdPIV

The analysis software can be seen in figures C.1, C.2, C.3 and C.4.

The software can be split into two sections, one concerned with the analysis

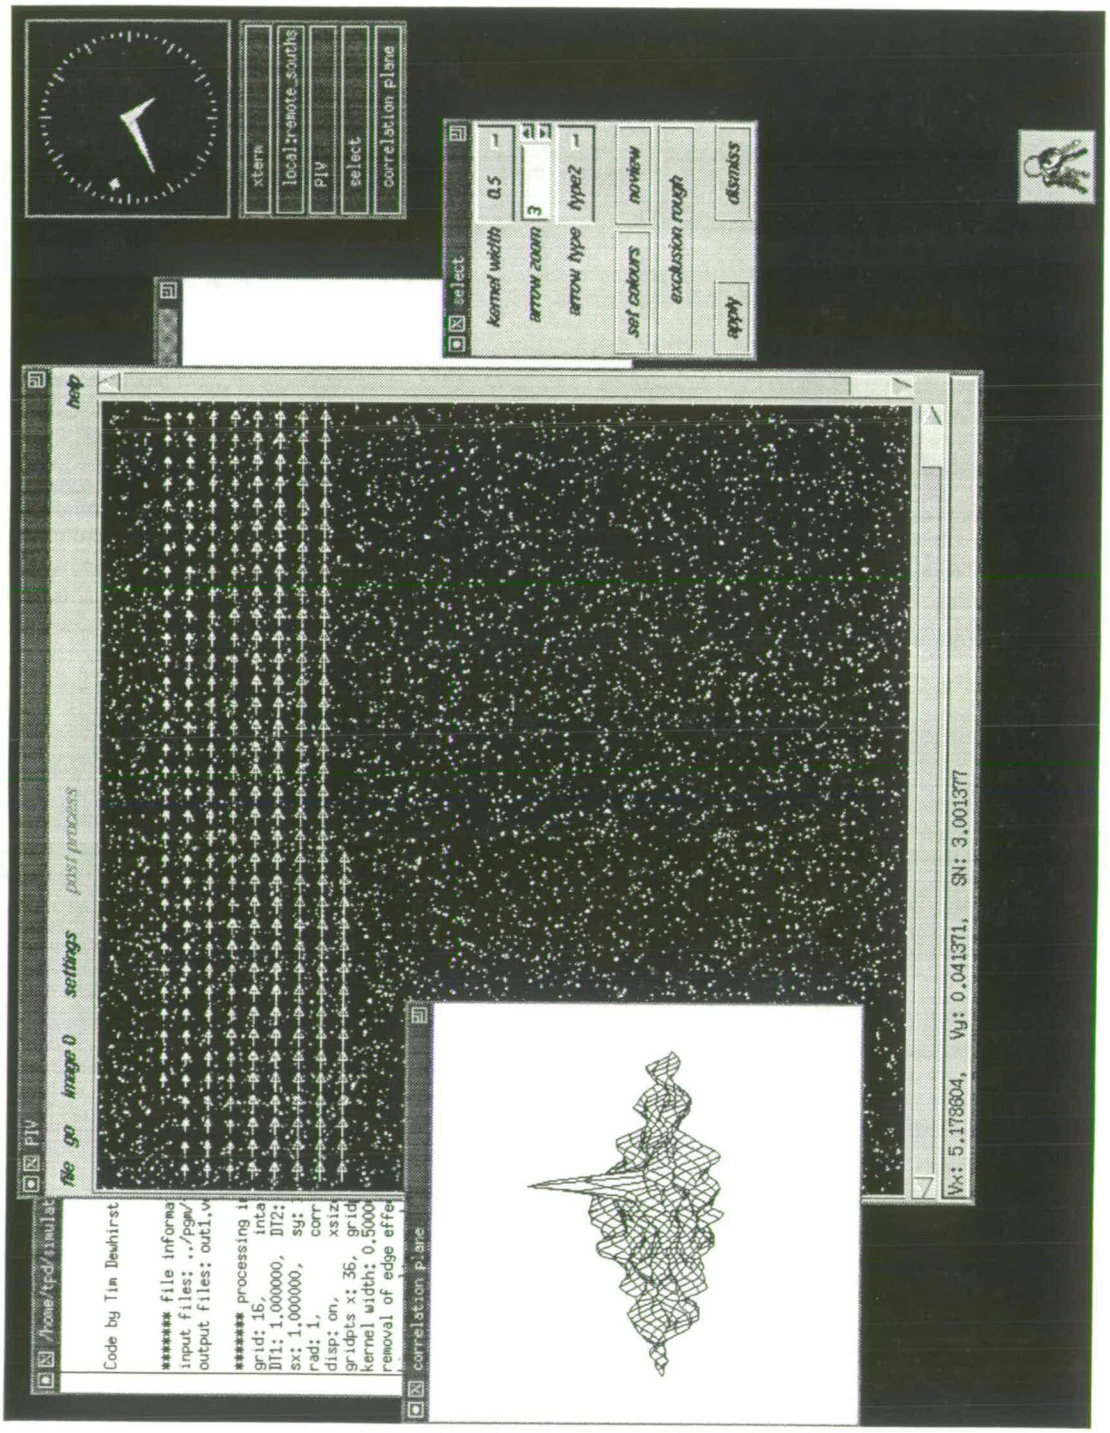

Figure C.1: EdPIV carrying out an analysis

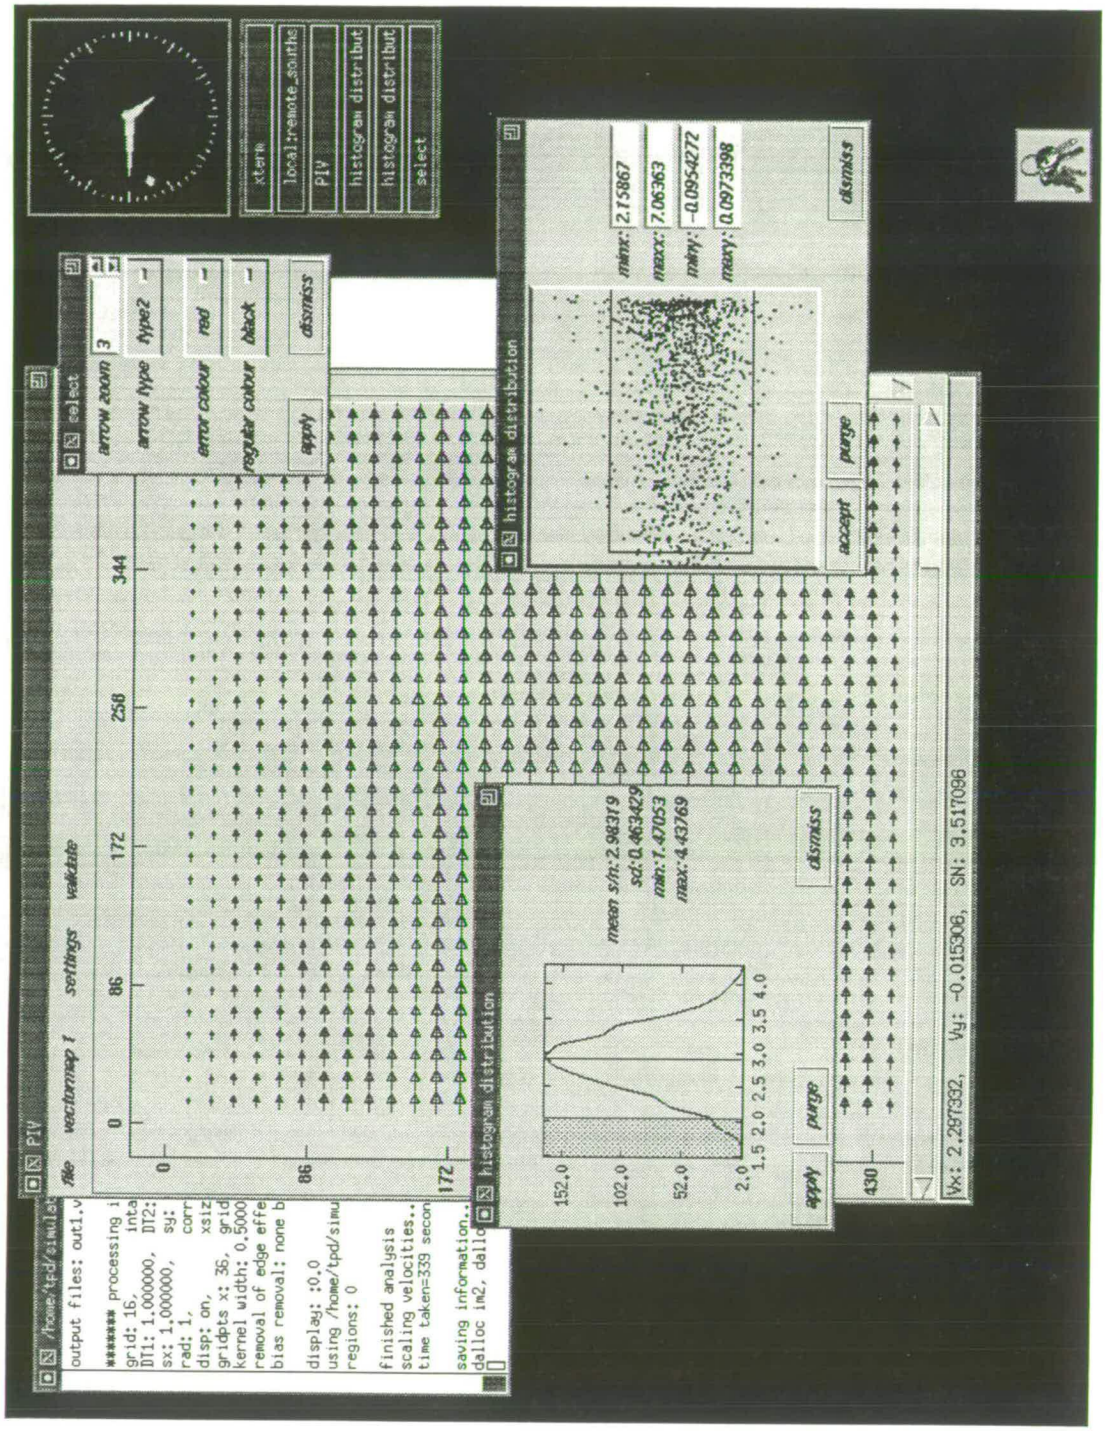

Figure C.2: EdPIV in post processing mode

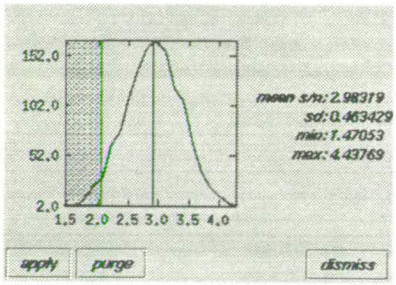

Figure C.3: Signal to noise histogram window

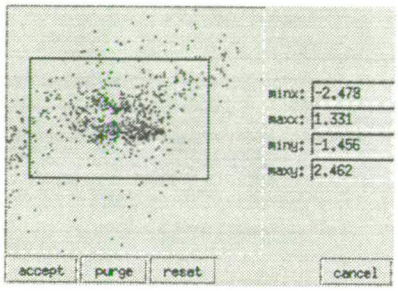

Figure C.4: Global histogram window

and the other concerned with displaying the data and construction the GUI.

C.2 Analysis

The heart of the analysis routines is an image handling library, written in C by Dr. W. Hossack, Department of Physics, University of Edinburgh. The library contains a number of routines to allow the easy manipulation of *PGM* (Portable Grey Map) images, including the expected functions of extraction, copying, etc.

As Dr. Hossack is a member of the applied optics group, this library was developed with the needs of optical simulations in mind. The resulting internal format for the images is capable of supporting either one byte per pixel, one floating point number per pixel or one complex number per pixel. The availability of the complex format simplified the calculation of the correlation immensely.

The non-GUI portion of the code is also responsible for peak detection rou-

tines, validation of the vector file, storage of the vector information, calculation of acceleration and parsing of the input parameters, to name a few functions. All together, this analysis code comprises some 6,000 lines of C code.

C.3 Graphical User Interface

One of the first difficulties that arises when writing a GUI based piece of code is which library to use to generate the *widgets*, or graphical elements which form the interface. This choice depends very much on the computing platform which is chosen. If DOS/Windows is used, so the windows libraries must be used exclusively. This is also the same for MacOS. However, if the operating system is a Unix, there are a number of libraries available for building GUI's.

Most of these libraries are based on C or C++, and as such generally require a great deal of code to be written to produce even a simple program. Libraries which fall into this category include Openlook and Motif. However, there is a high level language which allows the simple and rapid generation of GUI's. The language is Tcl, and its extension, Tk, handles all graphical code. Tcl/Tk is platform independent, powerful, allows rapid code generation and can be integrated with normal C code without too much difficulty.

The downside of using Tcl/Tk is that it is slow relative to e.g. Openlook. However, with the current desktop computer rapidly becoming equivalent to a low end workstation, this overhead is not noticeable unless large amounts of graphics are required to be updated frequently. If this is the case (as in the correlation window in the analysis program), a different library specialising in primitive graphics operations can be used instead [1].

The remaining bulk of the code written is devoted to the generation of the

[1] To draw the three dimensional correlation plane, the SRGP (Simple Raster Graphics Package) library was incorporated into the analysis code

GUI, and this amounts to another 4,000 lines of code.

The speed of execution can be increased further by "compiling" the high level Tcl/Tk language into its C equivalent, and this is achieved by using another package, *ET2C* (Embedded Tk 2 C). This also allows the integration of Tcl/Tk code and C to be made simpler, easing the job of passing variables, etc..

C.4 Post Processing

Once the PIV images have been analysed, they can be viewed by the second half of the analysis program. This allows the vector map to be plotted and validated, and then finally saved or printed.

The validation of the vector map is an integral part of a PIV measurement. While the experimental parameters are optimised to produce as few erroneous vectors, or outliers, as possible, the generation of a completely valid velocity vector map is not always possible. If this is the case then the vector map has to be filtered somehow to remove these vectors, and preferably replace the now missing vector with a vector which accurately represents the flow.

The filtering of velocity vector maps has been a subject of study for quite some time, but the results have not always been very promising. The most commonly used routines for detecting these outliers automatically are now considered to be a *s/n histogram*, a *local median filter* and a *global mean filter.*

C.4.1 S/N Histogram

During the analysis, a value of the quality of the correlation peak is usually assigned to the vector by comparing the height of the signal peak to the next highest noise peak. This is commonly referred to as the *detectability.* Once the complete vector map has been analysed, the mean and standard deviation σ of

the s/n values can be found. If we then assume that the distribution is normal, then removing the values which lie outside the lower 2σ of the distribution should result in $\approx 97\%$ of the vectors being deemed valid. The histogram of s/n values of an analysed vector map can be seen in figure C.3. It can be seen that the histogram is relatively symmetric, with the mean of the distribution marked by the central blue line, and the hatched rectangle on the left hand side of the histogram displaying the part of the distribution to be discarded.

This is possibly the least successful of the methods in removing outliers. The simple reason is that even if the vector map is completely free of outliers, this method will judge a certain proportion of those vectors to be defective, and remove them. However, in certain cases this filter can be useful.

C.4.2 Local Median Filtering

The local median filter is one which has been used in image processing since the seventies, and is especially effective at removing single isolated outliers.

The principal is simple. For each vector, a fixed number of its nearest neighbours are chosen. The velocities of each of these vectors are then entered into a one dimensional array, and the array is sorted in order of ascention (or descention). To find the median of any distribution it is simply a case of choosing the value placed at the centre of this array. The use of the median filter makes an assumption about the data which is being filtered, namely that the data is only slowly varying relative to the size of the sample taken around the vector in question. If this is the case, then finding the median of our vectors will produce a meaningful statistic about the flow locally, and can be used to reject any outliers.

The rejection of the outlier can be done in two ways. Once we have found the median we can then compare our original vector to it, and if the two vary by a significant amount, the vector is replaced by the median value of the local

distribution.

Secondly, each vector can be replaced directly by the median value of the distribution. This can at first sight appear to be a rash and unwise approach. However, if we have adhered to the primary principles of PIV, then our flow should have only small spatial velocity gradients in each interrogation area, and so the vector distribution should indeed be only slowly varying around any one vector. Therefore it should also be safe to directly replace each vector by its median value. Furthermore, it is quite probable that the median value of the distribution is going to be the original vector, thereby not altering the measurement at that point at all. This has been confirmed directly by Westerweel [62]. His research showed that applying a local median filter to correctly sampled PIV data results in the smallest change to the statistics of the flow of any of the validation methods.

This method is also the method adopted as standard in the field of image processing, where the local median filter is used to remove single pixel noise from images.

C.4.3 Global Mean Filtering

Global mean filtering is again simple in concept. This time however the filter acts on the entire data set as opposed to local regions, hence the name.

To perform the global mean filter, the data is binned into a two or three dimensional (for either two or three dimensional measurements) array. Once this is done the mean and standard deviation σ of the distribution is once again calculated and the same selection rules applies, namely discarding any vectors further than 2σ from the mean. This is not a fully automatic process, as in the majority of cases the velocity distribution is not normal in the flow. Therefore manual intervention is the most reliable method for removing the outliers. A typical global histogram window can be seen in figure C.4, showing each vector

marked as a point, and a rectangle placed around the data representing the 2σ bounds of the distribution. This rectangular area can be altered by human interaction, and as previously noted this is quite often necessary. In the case shown the flow had already been filtered to remove a few erroneous vectors, and so the data rejected would again have been valid data. However, for strongly normal distributions (for example the velocity map derived from an image shifted PIV recording) with a small standard deviation, global mean filtering is an ideal method of removing outliers.

C.4.4 Local Mean Filtering

Local mean filtering was until recently one of the more commonly used filtering methods for the removal of outliers in PIV data. However, this method has a number of features which make its use undesirable.

The local mean filter works in exactly the same manner as the local median, except of course the mean is calculated instead of the median. Once this mean has been calculated, then what can be done with it? The options are as for the local median case, namely comparison of the central vector with the mean value, or direct substitution of the mean value.

The second case cannot be used for the following reason. The application of a local mean at each vector location, followed by the direction substitution of this mean value will be the same as applying a low pass filter to the data. This is not usually an acceptable option unless the PIV data is noisy, in which case the low pass filtering may serve to bring out structures in the flow above the noise.

The first case was the most commonly used, but also has its problems. The first is the setting of the threshold value over which the central vector is replaced with the mean value. This can only be determined experimentally, and in general will need to be the result of a large number of PIV analyses. The second fault

is that in the calculation of the mean the value of the outlying vector will be included, and so the mean will be altered accordingly. If this is the case, the the effect of the outlier is to drag the mean away from a valid vector, and so replace it with a biased value. This can be avoided by first of all tagging the dubious values and performing a second pass of the data using only the valid vectors. However, the effect of the mean biasing can still be present resulting in the replacement of valid original data with the biased mean value.

This issue has been dealt with by Raffel [47] by applying a locally varying threshold value which is derived from the experimental results and the locally calculated standard deviation of the data. The application of this filter to flows with discontinuities has lead to good results.

C.5 Vector Interpolation

Once the velocity vector map has been filtered for outliers in the data set, there are going to be a number of missing vectors. It is often desirable to have a complete vector map, for example to enable the calculation of derivative quantities. Therefore it would seem obvious to interpolate the missing vectors.

The choice of scheme to replace the vectors can take on a number of forms, including linear interpolation, weighting functions and surface fitting. A summary of some of the available methods is given by Wang [61], although no mention was made of Delaunay triangulation. However, the conclusion of Wang was that a weighting scheme was the most computationally efficient and robust method.

The interpolation of the missing vector is carried out by making a weighting function or kernel which for two dimensional data is typically either a circular inverse square function or a circular Gaussian function. A fixed number of nearest neighbours are then chosen and summed after being multiplied by the weighting

function:

$$f'(0) = \frac{\sum_{i=-\underline{a}}^{\underline{a}} W(\underline{r}).f(\underline{r})}{\sum_{i=-\underline{a}}^{\underline{a}} W(\underline{r})} \tag{C.1}$$

with $W(\underline{r})$ being either

$$W(\underline{r}) = \frac{1}{\underline{r}^2} \tag{C.2}$$

or

$$W(\underline{r}) = e^{(-\alpha \underline{r}^2)} \tag{C.3}$$

The resulting interpolated vectors were shown to be in good agreement with the data that was removed (for the test simulations), although the procedure to obtain the data is not based upon any physical assumption.

C.6 EdVEC

Finally, a general purpose vector maps display and editing program was required to allow display, printing and review of the vector data. The program had to have a number of features incorporated, namely:

- editing and validation of the vector map
- display of the vector map
- display of derived quantities (i.e. vorticity, strain rates,...)
- the ability to make arbitrary translations of the map
- addition of scaling and annotation information
- production of a hardcopy

- the ability to display non-grided vector maps

The final point is one of the most important, as in general velocity vector maps, although analysed on a regular grid, do not produce results which lie on a regular grid. This is especially true when considering the analysis of an acceleration map following the Lagrangian specification. Of course, the data could be re-grided by means of either Delaunay triangulation or an adaptive Gaussian window, but this is another step away from the original data set, and if possible should be avoided.

The result of these specifications is *EdVEC*, a vector display and editing program written entirely in Tcl/Tk. The code itself is short at just over 4,000 lines considering that all vector file parsing, editing, validation and display is handled by this code.

The output from this program can be seen throughout this thesis, and the functions performed are listed and also described above. A screen shot of EdVEC can be seen in figure C.5

C.7 Computing Platform

When performing a large amount of data transfer and processing, it is essential that the computing platform is stable and computationally efficient. To achieve this, Linux was chosen as the operating system of choice for a number of reasons. Linux is a POSIX compliant Unix clone, is free, and comes with a great number of software packages for all aspects of computing. However, it is also extremely stable, and makes good use of Intel based processors which were conventionally used in conjunction with DOS/Windows.

As a comparison, a DEC Alpha 166 based workstation and an Intel Pentium 166 PC were compared, both with 32MBytes of RAM and running Linux. The results are listed in table C.1.

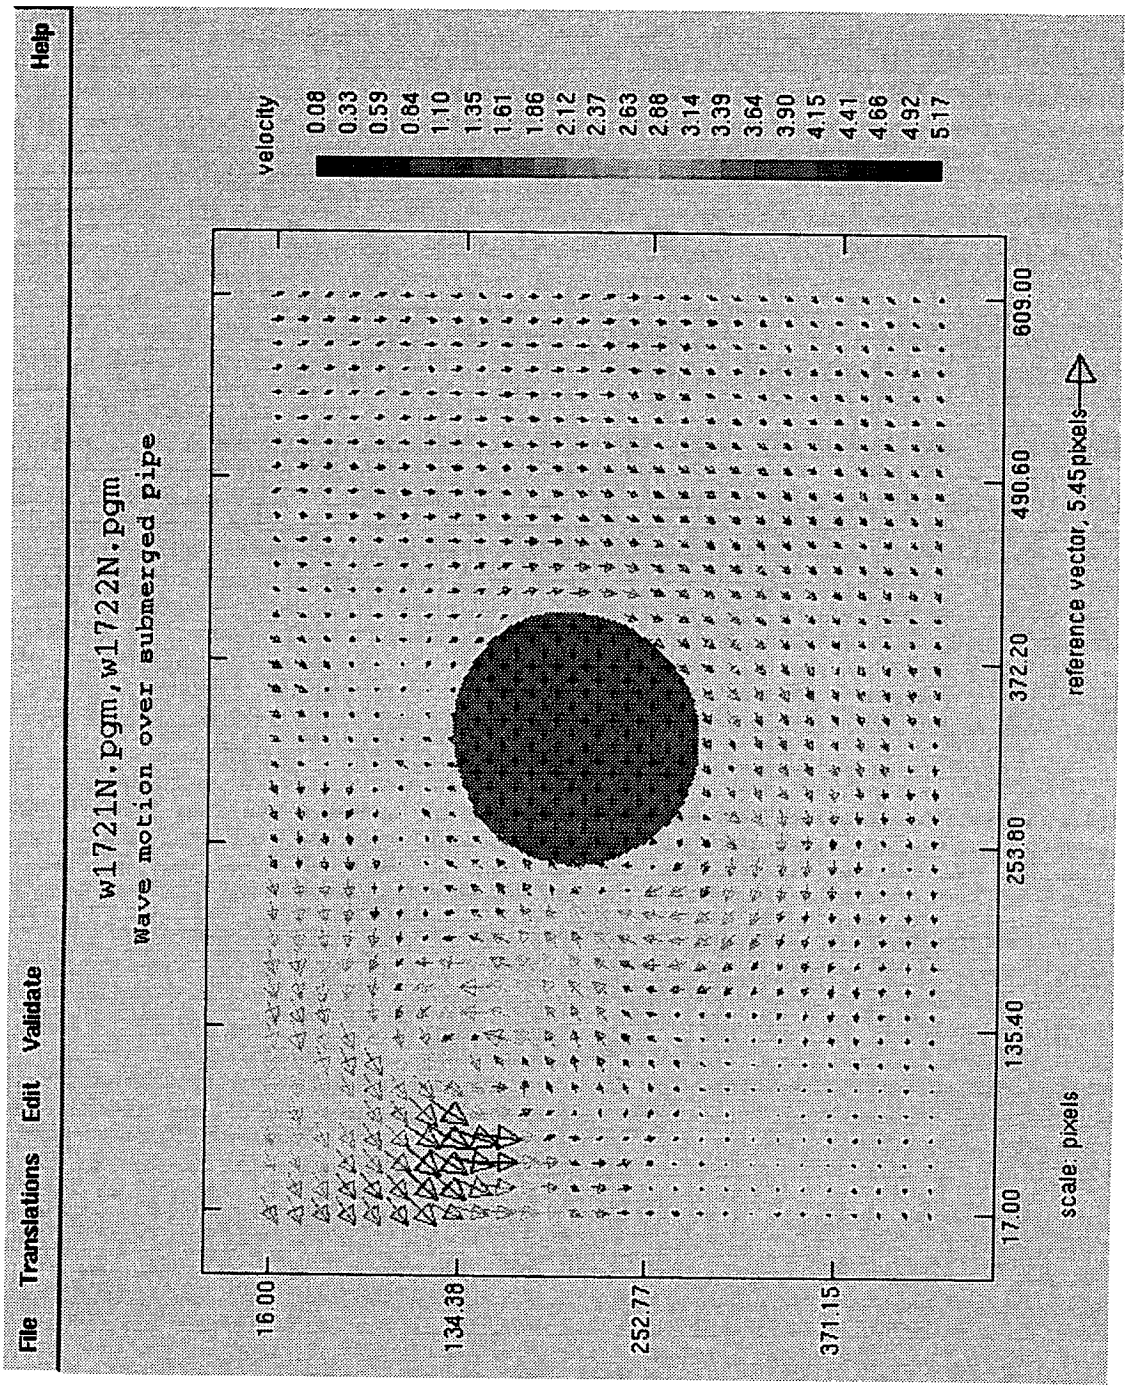

Figure C.5: EdVEC

	DEC Alpha 166	Intel Pentium 166
repetitions	time(s),kflops	time(s),kflops
8	1.50,8199.005	0.69,18311.111
16	3.01,8199.005	1.36,18780.627
32	6.02,8199.005	2.79,18621.469

Table C.1: Comparison of processors

The comparison is calculated by a program called *LINPACK* which performs large array manipulations on floating point numbers. It can be seen, that for an equivalent clock speed, the Intel processor is faster than this particular (and low end) Alpha. However, the Alpha costs substantially more that the Intel based machine, and with the continuing acceptance of Intel based machines in the domestic market, the performance per pound spent is increasing at a phenomenal rate.

With this in mind, and that fact that Unix also supports useful features such as TCP/IP, NFS, and a whole host of network based utilities as standard, Linux was chosen for use throughout the group.

Linux was tested and found to be much more stable that DOS/Windows, allowing continuous processing for days and weeks as opposed to hours. Linux also has advantages from the programming point of view. DOS is an 8-bit operating system which was created from an old 4-bit operating system. This history means that DOS has some odd limitations, especially when it comes to allocating memory. This can be worked around, but at the expense of extra complexity in the software. Linux, on the other hand, is a completely multi-user, multi-tasking, 32-bit operating system, which can easily be optimised to take advantage of the latest processor developments.

Bibliography

[1] R J Adrian. Scattering particle characteristics and their effect on pulsed laser measurements of fluid flow: Spekle velocimetry vs particle image velocimetry. *Applied Optics*, 23(11):1690–1691, 1984.

[2] D W Allen and D L Henning. Vortex-induced vibration tests of a flexible smooth cylinder at supercritical reynolds numbers. In *Proceedings of the 7th International Offshore and Polar Engineering Conference*, volume 3, pages 680–685, 1997.

[3] B Bretthauer, G E A Meier, and B Stasicki. An electronic cranz-schardin camera. *Rev. Sci. Instrum*, 62(2):364–368, 1991.

[4] Ch Brücker. Digital-particle-image-velocimetry (DPIV) in a scanning light-sheet: 3-D starting flow around a short cylinder. *Experiments in Fluids*, 19:255–263, 1995.

[5] Ch Brücker. 3-D PIV via spatial correlation in a colour-coded light sheet. *Experiments in Fluids*, 21:312–314, 1996.

[6] Ch Brücker. High-speed video: An upcoming chance for time-resolved 3-D velocimetry based on light-sheet tomography. In *Laser Anemometry Advances and Applications*, pages 427–435, University of Karlsruhe, Germany, 1997.

[7] J R Chaplin. Nonlinear forces on a horizontal cylinder beneath waves. *Fluids Mech.*, 147:449–464, 1984.

[8] E A Cowen and S G Monismith. A hybrid digital particle tracking velocimetry technique. *Experiments in Fluids*, 22:199–211, 1997.

[9] T P Dewhirst, M L Jakobsen, and C A Greated. Multiple CCD camera DPIV for force and acceleration measurements in high speed flows. In *Proc. IMechE Optical Methods and Data Processing in Heat and Fluid Flows*, pages 251–258, 1996.

[10] W J Easson and C A Greated. Breaking wave forces and velocity fields. *Coastal Engineering*, 8(3):233–241, 1984.

[11] L J M Esser and F L J Sangster. Charge transfer devices. In C Hilsum, editor, *Handbook on Semiconductors*, volume 4 of *Device Physics*. North Holland Publishing Co., 1981.

[12] N Farrugia, S Kanne, and D A Greenhalgh. Three pulse digital particle image velocimetry. Private Communication, 1995.

[13] J W Goodman. *Introduction to Fourier Optics*. McGraw Hill, New York, 1968.

[14] James Gourlay. *Ferroelectric Liquid Crystal Spatial Light Modulators: Devices and Applications*. PhD thesis, The University of Edinburgh, 1994.

[15] I Grant. Particle image velocimetry: a review. In *Proceedings Institute of Mechanical Engineers*, pages 55–76, 1997.

[16] I Grant and A Liu. Directional ambiguity resolution in particle image velocimetry by pulse tagging. *Experiments in Fluids*, 10:71–76, 1990.

[17] I Grant and X Wang. Directionally-unambiguous, digital particle image velocimetry studies using a image intensifier camera. *Experiments in Fluids*, 18:358–362, 1995.

[18] C Gray. *The Development of Particle Image Velocimetry for Water Wave Studies.* PhD thesis, The University of Edinburgh, 1989.

[19] C Gray. Private Communication, 1997.

[20] C Gray, C A Greated, D R McCluskey, and W J Easson. An analysis of the scanning beam PIV illumination system. *Meas. Sci. Technol.*, 2:717–724, 1991.

[21] E Hecht. *Optics.* Addison-Wesley Publishing Company, 2 edition, 1987.

[22] A K Hind and J R E Christy. Digital PIV applied to flows around artificial heart valves: analysis by autocorrelation. In *Proceedings of Euromech Colloquium 279.* Kluwer Academic Publishers, 1992.

[23] P Horowitz and W Hill. *The Art of Electronics.* Cambridge University Press, 2 edition, 1989.

[24] W. J. Hossack. *Modern Optics Lecture Course.* The University of Edinburgh, The University Of Edinburgh, 1994.

[25] W M Humphreys, S M Bartram, and J L Blackshire. A survey of particle image velocimetry applications in langley aerospace facilities. In *Proc. 31st Aerospace Sciences Meeting*, Reno, Nevada, 1993.

[26] M L Jakobsen, W J Hossack, and C A Greated. Particle image velocimetry analysis using an optically addressed spatial light modulator: Effects of nonlinear transfer function. *Applied Optics*, 34(11):1757–1762, 1995.

[27] R D Keane and A J Adrian. Optimisation of particle image velocimeters. part i: Double pulsed systems. *Meas. Sci. Technol.*, 1:1202–1215, 1990.

[28] R D Keane and A J Adrian. Optimisation of particle image velocimeters: Ii. multiple pulsed systems. *Meas. Sci. Technol.*, 2:963–974, 1991.

[29] R D Keane and A J Adrian. Theory of cross-correlation analysis of PIV images. *Flow Visualisation and Image Analysis*, pages 1–25, 1993.

[30] C C Landreth and R J Adrian. Electrooptical image shifting for particle image velocimetry. *Applied Optics*, 27(20):4216–4220, 1988.

[31] N J Lawson, N A Halliwell, and J M Coupland. Particle image velocimetry: Image labeling by use of adaptive optics to modify the point-spread function. *Applied Optics*, 33(19):4241–4247, 1994.

[32] V.F. Leavers. *Shape detection in Computer vision using the Hough Transform.* Springer–Verlag, London, 1992.

[33] B Lecordier, M Mouqallid, S Vottier, E Rouland, D Allano, and M Trinite. Ccd recording method for cross- correlation PIV development in unstationary high speed flow. *Experiments in Fluids*, 17:205–208, 1994.

[34] L M Lourenço. *Particle Image Velocimetry: Historical Background, Processing Techniques.* von Karman Institute for Fluid Dynamics, Lecture Series 1996-03.

[35] W Merzlirch and T Wagner. PIV with two syncronised video cameras. In *Proc. IMechE Optical Methods and Data Processing in Heat and Fluid Flows*, pages 237–242, 1996.

[36] G Moe and O T Gudmestad. Prediction of morison type forces in irregular, high reynolds number waves. In *Proceedings of the 7th International Offshore and Polar Engineering Conference*, volume 3, pages 399–406, 1997.

[37] Iain George Morrison. *The Hydrodynamic Performance of an Oscillating Water Column Wave Energy Converter.* PhD thesis, The University of Edinburgh, 1995.

[38] I A R Neill and J B Hinwood. Hydrodynamic forces acting on a bottom mounted circular cylinder. In *Proceedings of the 7th International Offshore and Polar Engineering Conference*, volume 3, pages 415–422, 1997.

[39] K Prasad, R J Adrian, C C Landreth, and P W Offutt. Effect of resolution on the speed and accuracy of particle image velocimetry interrogation. *Experiments in Fluids*, 13:105–116, 1992.

[40] J. Pullen. Private Communication, 1997.

[41] Paul Anthony Quinn. *Breaking Waves on Beaches.* PhD thesis, The University of Edinburgh, 1995.

[42] R G Racca, O Stephenson, and R M Clements. High-speed video analysis system using multiple shuttered charge-coupled device imagers and digital storage. *Optical Engineering*, 31(6):1369–1374, 1992.

[43] M Raffel. Private Communication, 1997.

[44] M Raffel, A Derville, C Willert, O Ronneberger, and J Kompenhans. Dual-plane correlation for three- dimensional particle image velocimetry on planar domains. In *8th International Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon*, 1996.

[45] M Raffel, M Gharib, O Ronneberger, and J Kompenhans. Feasibility study of three-dimensional PIV by correlating images of particles within parallel light sheet planes. *Experiments in Fluids*, 19:69–77, 1995.

[46] M Raffel, J Kompenhans, B Stasicki, B Bretthauer, and G E A Meier. Velocity measurement of compressible air flows utilising a high-speed video camera. *Experiments in Fluids*, 18:204–215, 1995.

[47] M Raffel and J Kompenhnans. *Particle Image Velocimetry: Theoretical and Experimental Aspects of PIV Recording Utilizing Photographic Film and Mechanical Image Shifting.* von Karman Institute for Fluid Dynamics, Lecture Series 1996-03.

[48] M Raffel, J Westerweel, C Willert, M Gharib, and J Kompenhans. Analytical and experimental investigations of dual-plane PIV. In *SPIE International Symposium on Optical Science, Engineering, and Instrumentation*, pages SPIE Proceedings Article No. 2546–13, 1995.

[49] M Raffel, J Westerweel, C Willert, M Gharib, and J Kompenhans. Analytical and experimental investigations of dual-plane particle image velocimetry. *Opt. Eng.*, 35(7):2067–2074, 1996.

[50] M Raffel, C Willert, and J Kompenhans. *Particle Image Velocimetry: A Practical Guide.* Springer-Verlag, 1998.

[51] O Ronneberger. *Messung aller drei Geschwindigkeitskomponenten mit Hilfe der Particle Image Velocimetry mittels einer Kamera und Zweier paralleler Lichtschnitte.* PhD thesis, DLR, Goettingen, Bunsenstrasse, 10, Goettingen, D-37073, 1998.

[52] E Rouland, S Vottier, B Lecordier, and M Trinite. Cross-correlation PIV development for high speed flow with a standard ccd camera. In *Proc. IMechE Optical Methods and Data Processing in Heat and Fluid Flows*, pages 9–20, 1994.

[53] H Rouse. *Advanced Mechanics of Fluids.* Wiley, 1959.

[54] T Saprkaya and M Isaacson. *Mechanics of Wave Forces on Offshore Structures.* Van Nostrand Reinhold Company Inc., 1981.

[55] G X Shen and G Y Ma. The investigation on the properties and structures of starting vortex flow past a backward-facing step by WBIV technique. *Experiments in Fluids*, 21:57–65, 1996.

[56] Rajpal S Sirohi, editor. *Speckle Metrology.* Marcel Dekker Inc., 1993.

[57] David John Skyner. *The Mechanics of Extreme Water Waves.* PhD thesis, The University of Edinburgh, 1992.

[58] D J Tritton. *Physical Fluid Dynamics.* Oxford Science Publications, 2 edition, 1988.

[59] G Van der Graaf and H E A Van den Akker. PIV measurements on flow phenomena during mould filling. In *Optical Methods and Data Processing in Heat and Fluid Flow*, pages 249–253, City University, London, 1994.

[60] J T Verdeyen. *Laser Electronics.* London: Prentice-Hall International, 3rd edition, 1995.

[61] J. Wang, F. Barnes, and C. Gray. A comparison between interpolated algorithms for post processing of particle image velocimetry data. In *Optical*

Methods and Data Processing in Heat and Fluid Flow, pages 291–300, City University, London, 1996.

[62] J Westerweel. *Digital Particle Image Velocimetry. Theory and Application.* PhD thesis, Delft University Press, Delft, 1993.

[63] J Westerweel, D Dabiri, and M Gharib. The effect of a discrete window offset on the accuracy of cross-correlation analysis of digital PIV recordings. *Experiments in Fluids*, 23:20–28, 1997.

[64] C Willert, B Stasicki, M Raffel, and J Kompenhans. A digital video camera for application of particle image velocimetry in high-speed flows. In *SPIE International Symposium on Optical Science, Engineering, and Instrumentation*, pages SPIE Proceedings Article No. 2549–19, 1995.

[65] C E Willert and M Gharib. Digital particle image velocimetry. *Experiments in Fluids*, 10:181–193, 1991.

[66] M Wirth and T A Baritaud. A cross correlation PIV technique using electro-optical image separation. *Experiments in Fluids*, 21:410–416, 1996.

[67] D C Wormell and J L Sopchak. A particle image velocimetry system using a high resolution CCD camera. In *SPIE 2005, Optical Diagnosics in Fluid and Thermal Flow*, pages 648–654, 1993.

[68] D C Wormell and J L Sopchak. A particle image velocimetry system using a high resolution CCD camera. In *Optical Methods and Data Processing in Heat and Fluid Flow*, pages 21–28, City University, London, 1994.