CHAPTER 15

Particle Image Velocimetry (PIV) in the Coastal Engineering Laboratory

C.A.Greated[1], D.J.Skyner[2] & T.Bruce[3]

ABSTRACT

The Particle Image Velocimetry (PIV) technique is reviewed, with special reference to its application to coastal engineering problems. The basic principles are described, and strengths and weaknesses of different implementations considered. Examples of some applications of PIV at Edinburgh are described to illustrate the practicalities of its implementation in a hydrodynamics laboratory. Finally, the limitations to the accuracy of PIV in these applications are reviewed and an outline of some new developments in this area is given.

INTRODUCTION

Non-intrusive flow measurement techniques are becoming an increasingly familiar part of laboratory experiments in applied hydrodynamics and coastal engineering. Since its inception in 1966, Laser Doppler Anemometry (LDA) has progressed, and aided by developments in other technologies, notably in lasers, fibre-optics and computing, it is now a most useful tool under a wide range of conditions. LDA is, however, fundamentally a point measurement technique; the time evolution of flow velocity can be measured with great accuracy at a point, but if a map of an area of the flow is to be obtained, it must be built up point by point. Thus the flow must be steady or accurately repeatable. PIV gives a quantitative map of instantaneous flow velocities over a large field. An attractive feature of PIV analysis is that the velocity measurements come out on a regular grid and thus post-processing is relatively straightforward.

PIV is sometimes compared to streak photography where the path lengths of individual marker particles in the fluid are measured. Streak photography suffers two major disadvantages. Firstly, only a very few marker particles can be

---

[1] Director of Fluid Dynamics Unit, Dept Physics, the University of Edinburgh, Edinburgh EH9 3JZ, Scotland

[2] Research Associate, Dept Physics (D.Skyner@ed.ac.uk)

[3] Research Associate, Dept Mech Eng (T.Bruce@ed.ac.uk)

used, otherwise their paths overlap too often and the individual particle motions cannot be separated. Secondly, when used with sheet illumination, the particles may move out of the measurement plane, artificially shortening the streak lengths and causing errors in the velocity measurement.

PARTICLE IMAGE VELOCIMETRY

The basis of PIV is to stroboscopically illuminate a two-dimensional plane of a flow containing small *seeding* particles. A double (or multiple) exposure photograph of this plane is taken. The spacing between the images of each particle on the film gives the local velocity. This photograph is then analysed to obtain the local flow velocities over a grid of points covering the whole field. PIV is normally considered as a two stage process: the acquisition and the analysis of the image.

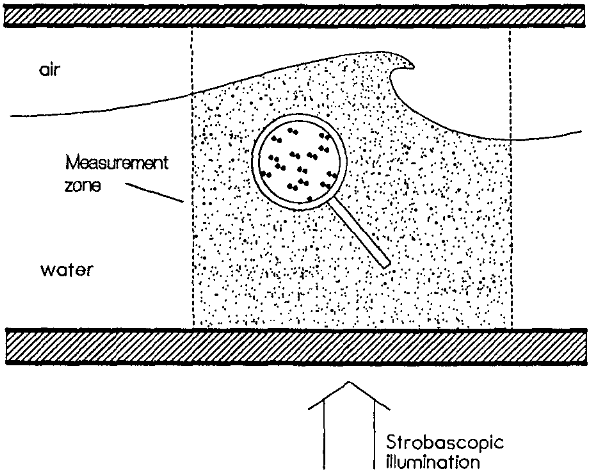

Figure 1: Schematic diagram of PIV acquisition

Figure 1 shows schematically the acquisition stage in a typical application of PIV. The enlarged area shows a portion of the film containing multiple particle images: the optimum form of this area is dictated by the requirements of the subsequent analysis. The acquisition and analysis phases are therefore inextricably linked. The following sections consider the stages of the PIV process: acquisition, ie. illumination and photography, and the automated analysis of the image.

PIV ILLUMINATION

In general, the two-dimensional plane within the flow is defined by a thin sheet of pulsed laser light. Two different techniques are available for producing this light sheet. In the first of these, a continuous wave (CW) laser is deflected off a rotating multi-facet mirror so that the beam is scanned through the measurement area, as illustrated in figure 2. In the illustration, the parabolic mirror collimates the beam to give a more even intensity distribution. The shutter opening time on the camera is normally set to be four to six times the period of the beam scan, in order that this number of exposures are recorded on each negative. In the second technique, a pulsed laser, typically a Nd Yag, is used to produce double pulses, the beam being spread out into a sheet using a cylindrical lens as shown in figure 3.

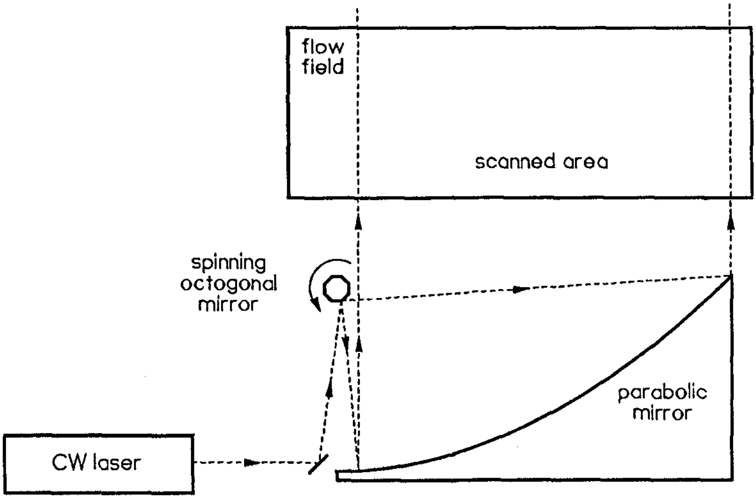

Figure 2: Scanning beam illumination system

A comparison of the illumination techniques shows that the scanning beam method [Gray et al, 1991] is usually the best for use within the velocity ranges normally encountered in the coastal engineering laboratory, i.e. up to a few metres per second. It gives a sufficiently even illumination over a large area and has the advantage that more than two images per particle can easily be obtained, which improves the signal strength in subsequent analysis. For measurement regions of about $1m^2$ and flow velocities up to $2ms^{-1}$ a 10W Argon Ion laser is suitable. The upper limit on velocities which can be measured with the scanning beam system can be proportionally increased by, for example, reducing the width of the scanned area. Above about $10ms^{-1}$ it is necessary to use a pulsed laser; in this case the energy per pulse is almost constant, so that measuring higher velocities does not required a proportional increase in laser power. However, the use of pulsed lasers has certain drawbacks: such systems are inherently more dangerous; it is harder to obtain an even illumination over the measurement zone; the optical alignment is more difficult. Additionally, with a single laser only two pulses per negative

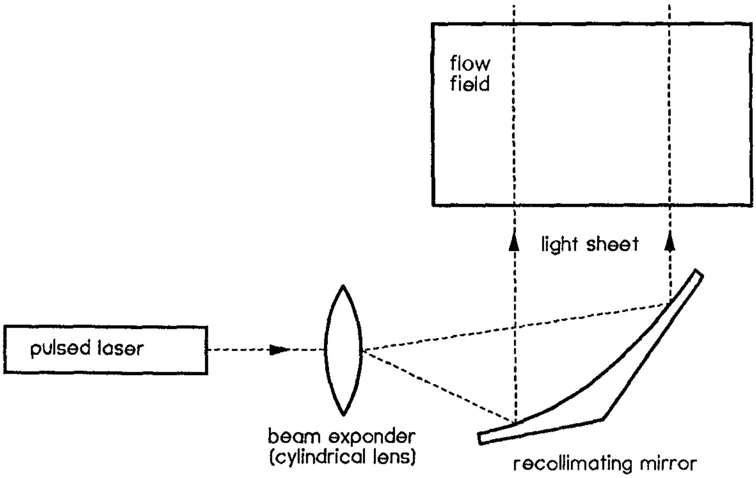

Figure 3: Pulsed sheet illumination

can be obtained easily. It is possible to obtain a four pulse sequence by coupling two lasers, but this is a costly option.

If the surface flow of a fluid is to be measured, then illuminating the measurement zone can be accomplished without the use of lasers. The surface defines the two dimensional plane for measurement and the stroboscopic illumination may be achieved by means of standard flash units. Thus PIV may often be simply and inexpensively applied to such flows.

Practical Implementation of Illumination

In the fluid dynamics laboratories at Edinburgh the wave flumes and water channels have been designed with the application of PIV in mind. All the tanks are made of glass, giving optical access from the sides and below. The beam from a single laser serves three experimental rigs, considerably saving on the cost. Most of the time involved in carrying out PIV experiments is in setting up the apparatus. When the laser is in use for one experiment, another can be in the preparation stage.

For water channel experiments a 15W argon ion laser is used with scanning beam illumination. The beam from such a high powered laser is potentially very hazardous. For safety reasons, the laser is mounted rigidly in an enclosed area and its beam is directed through pipes to the appropriate experiment by a series of mirrors. With this arrangement the beam is totally enclosed up until the point where it enters the water tank and no longer represents a safety hazard. The scanning beam illumination system is mounted on rails beneath each of the water channels so that the position of the measurement area can be easily moved.

PIV PHOTOGRAPHY

The process of acquiring good PIV photographs involves two major stages. The first is selecting appropriate hardware for the application (eg, illumination, seeding, camera and film), and the second is using this hardware to best effect. This section addresses both of the above stages.

Hardware

The hardware considerations may be divided into four broad and interdependent areas: illumination (already discussed), flow seeding, camera and film. Here the practices adopted in the study of water waves in the Fluid Dynamics Unit (FDU) at Edinburgh are detailed. However, most of the issues covered are also relevant to the application of PIV to other flows.

Seeding

The selection of a suitable flow seeding is crucial. The seeding used is conifer pollen which meets the most important criteria: it is almost exactly neutrally buoyant, quite reflective at the wavelength of the Argon Ion laser and small enough to follow the flow patterns (typical particle diameters are $\sim 60\mu m$). It is also quite inexpensive — the cost of seeding a flume containing three tonnes of water is about the same as the cost of the film in the camera.

Camera & Lens

The choice of the camera and lens is also important. A flat-field lens should be used to minimise distortions of the image plane. Choosing a lens of longer focal length will reduce the apparent effect of any out-of-plane motions of particles in the field, but will increase the difficulty of achieving a sharp focus and make the process more susceptible to vibrations. A Hasselblad 500 *EL/M* camera with 80mm lens has proved suitable for many PIV applications at the FDU.

Film

The choice of film is generally straightforward, since there are a number of good, high resolution black and white films available at reasonable cost. Kodak TMax is one such film and is well suited to the green light of the Argon Ion laser.

Taking the Photographs

The important issues remaining are the choice of camera position, achieving the correct seeding density, optimising the focus and optimising the exposure parameters to give good, high contrast photographs. Resolving all these issues involves making compromises, and the aim of the following paragraphs is to outline the most important considerations. Figure 4 shows a well optimised image.

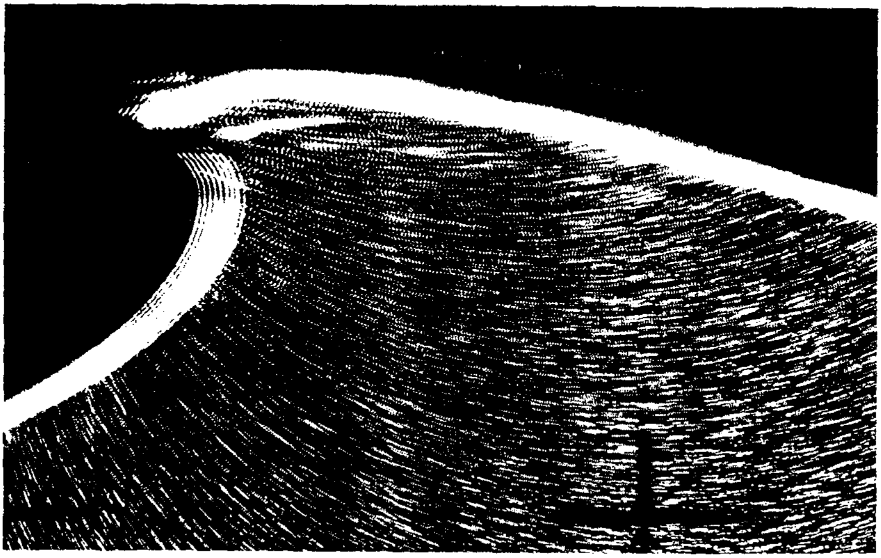

Figure 4: A photographically obtained PIV image

Seeding Density

The optimum seeding density is determined by considering the subsequent analysis of the photograph. In the Edinburgh system, the local flow velocity at a point on the photograph is averaged over a 1mm diameter interrogation region. The seeding density should be high enough that there will always be several particle image pairs in each interrogation area, as shown in figure 1. Experience also plays a large part in getting an optimum level of seeding, and once good results have been obtained, the successful level of seeding can be repeated.

Illumination Interval

As with optimising seeding density, the choice of illumination interval is dictated by consideration of the subsequent analysis. The interval should be set such that the largest velocity in the flow a gives particle image separation on the photograph which is the largest resolvable by the analysis system, 200μm for the FDU system. Normally an estimate of the largest velocity can be made, but its often necessary to take a test film trying different intervals in order to optimise the choice.

Shutter Speed

Once the illumination interval has been chosen, the shutter speed is selected to give the required number of exposures. Typically, four to six exposures are used, but it is important to consider whether the flow changes over the exposure time.

Focus

Achieving a sharp focus is vital if good PIV photographs are to be obtained, and can be quite difficult, especially when a large lens aperture or long focal length lens are being used. Again, a test film in which the focus is varied is the usual route to its optimisation.

Photographic Magnification

The magnification from the measurement zone to the film depends upon the focal length of the lens and the distance from the camera to the measurement zone. Its selection is another compromise. It is often desirable to measure as large a region as possible, but the analysis phase must be considered. The implicit assumption is that particle image displacement over the interrogation region is uniform: if there is a strong displacement gradient present, errors are introduced and the resulting data point may be at best inaccurate and at worst, spurious. Therefore the size of the area imaged onto the photograph is typically limited to that which will result in displacement gradients of less than 3% over any 1mm diameter interrogation area.

Exposure

Once the illumination interval and magnification have been established, the only factors left governing the exposure of the film are the lens aperture, the laser power and the film speed (ie, ASA rating). In general, a test film covering various settings of aperture and laser power is required for a new application. The quality of the resulting photographs may then be assessed on the analysis system and the best settings finalised.

If there is scope for choosing an aperture setting, it should be remembered that the largest apertures (smallest *f* numbers) give the poorest depth of field, so focussing is more difficult. However, if the particle image size is diffraction limited, the smaller apertures (larger *f* numbers) will result in larger particle images. Using *f/4* or *f/5.6* is usually a good compromise.

In general, the lower the ASA rating of the film, the finer the grain, so it is preferable to work with 100 ASA film. However, if the maximum available laser power and the largest aperture still give insufficient film exposure, then faster film may be used — 400 ASA TMax is widely used in the FDU, and has, on occasion, been *push processed* to as much as 3200 ASA, though this does result in a reduction in resolution.

PIV ANALYSIS

There are a number of different approaches to extracting the velocity information from the flow photograph. The objective in this phase of the PIV process is to extract the local velocity at an array of positions on the negative from the most correlated separation of the particle images near that point. For high seeding densities, this is most commonly achieved by obtaining the autocorrelation of the intensity distribution within each interrogation area and locating the strongest correlation peak, which yields the average displacement and hence the velocity.

The autocorrelation function can be conveniently calculated using Fourier transforms (Wiener-Khintchine theorem)

$$A(x,y) = FT^{-1}\{[FT\{a(x,y)\}]^2\} \tag{1}$$

where $a(x,y)$ is the image intensity distribution of the chosen portion of film, FT is a forward Fourier transform and $A(x,y)$ is the autocorrelation function of the area.

Fast Fourier transform (FFT) routines are computationally very efficient. Nevertheless, with a fine grid the amount of computation is still considerable and an alternative approach is to produce optical transforms. Three types of analysis system are therefore possible. Both transforms can be performed digitally, both optically, or one of the transforms can be optical and the other digital. The last approach is the one which has been adopted in the FDU [Gray,1989] , and is commonly known as the Young's fringe method, since the result of the optical transform is diffraction fringes. This gives adequate processing speed, whilst keeping the optical system relatively simple.

If the transformations are to be carried out purely digitally, white light illumination is used and the image of the particles is captured directly on a CCD array. This approach requires two digital transforms it is more computationally intensive than the fringe method. There are certain advantages, one of the most important being that boundaries can be indentified directly during the analysis phase without the need for additional optics.

A complete optical analysis of the PIV records is the most rapid of all, although the techniques for implementing this have not yet been fully developed. The best approach is probably to use an optically addressed spatial light modulator (SLM) as described by Jakobsen et al [1992].

Practical Implementation of Analysis

Figure 5 shows the layout of the Edinburgh analysis rig. The negative is mounted onto a two-axis translation stage in the path of the beam from a low power laser. The illuminated area defines the point where the velocity is to be measured. Behind the negative is the "Fourier transforming" lens which images the fringe pattern onto a CCD array in its back focal plane . For a perfect transform the object should be in the front focal plane of the lens, but in practice it is placed much closer to the lens than this in order to avoid vignetting caused be the finite aperture of the lens.

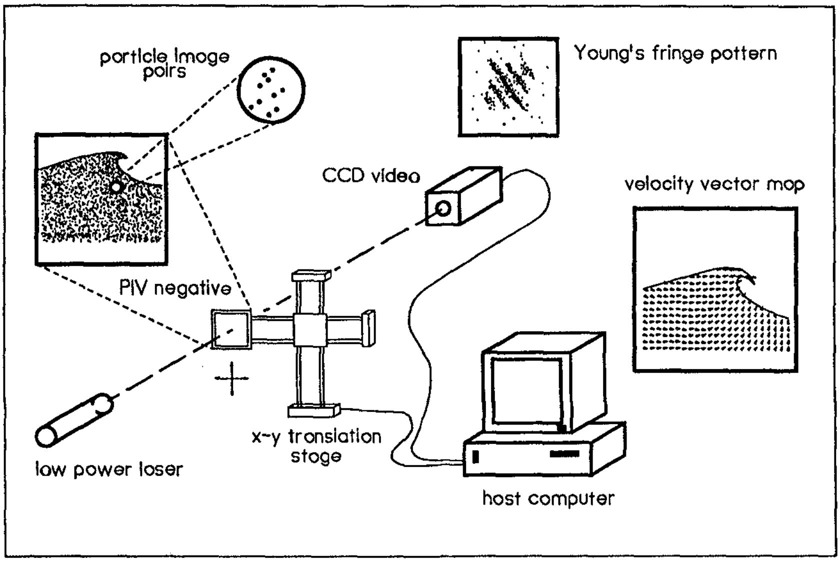

Figure 5: The Edinburgh PIV analysis system

For a typical PIV negative the fringe pattern is shown figure 6. The spacing between the fringes is inversely proportional to the particle spacing, and hence velocity, and their orientation gives the direction of the flow at that point. This fringe pattern is digitised and transfered to the PC. The inverse Fourier transform is then performed using an FFT algorithm.

Figure 6 also shows the computed autocorrelation function, containing two symmetrically placed peaks. The analysis program uses a peak detection routine to locate the position of the correlation peaks; the coordinates of the peaks define the magnitude and direction of the velocity vector at that point. An advantage of the fringe method is that the user can see immediately whether clear fringes are formed at any particular point in the flow. If they are, then the analysis procedure will almost certainly give a reliable velocity reading.

To obtain velocity values on a grid of points in the flow, the translation stage moves the negative relative to the interrogating laser beam. At each point, the analysis takes about 4 seconds, so that a 35mm format negative is analysed on a 1mm grid in under an hour.

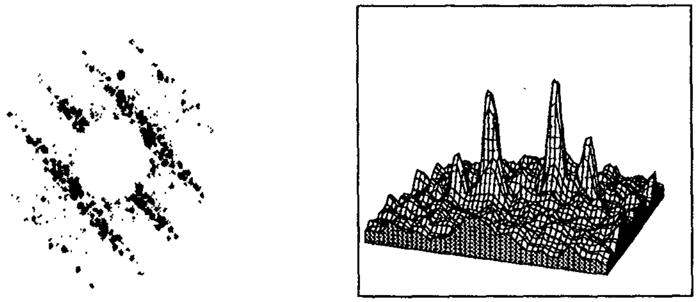

Figure 6: Fringe pattern and autocorrelation function

Data Validation

In the Edinburgh system, each velocity measurement is accompanied by a value which gives a measure of its reliability. This is determined by strength of the correlation peak compared to the background noise level. Velocity vectors whose "quality" factor lies below a specified threshold can either be discarded or checked individually.

PIV Errors

As with LDA, PIV gives an absolute measurement of velocity, provided that the system is calibrated correctly for known particle displacements. Distortions can also occur in the photographic process, for example by refraction at side walls. The extent of these can easily be checked by inserting a grid into the measurement area and photographing this in place of the illuminated sheet.

The accuracy of PIV data is determined by the uncertainties introduced in the two distinct phases of the process; the photographic recording of the flow, and the analysis of this photograph. The specific problems in the case of PIV applied to wave motion have been considered in detail. For the most reliable results, the flow must be well seeded, and the photographic acquisition parameters selected to ensure that the velocity range present corresponds to the full range measurable in the analysis and that the velocity gradients on the film are not too great. If these conditions are satisfied, then the final error bound on measurements can be less than ~2% of the maximum velocity measured.

Examples of PIV Measurements

In this section the results from applications of PIV at the FDU are presented. Other recent studies using the technique are described in [Powell et Al,1992, Morrison & Greated,1992, Bruce & Easson,1992, Skyner & Easson,1992(b)].

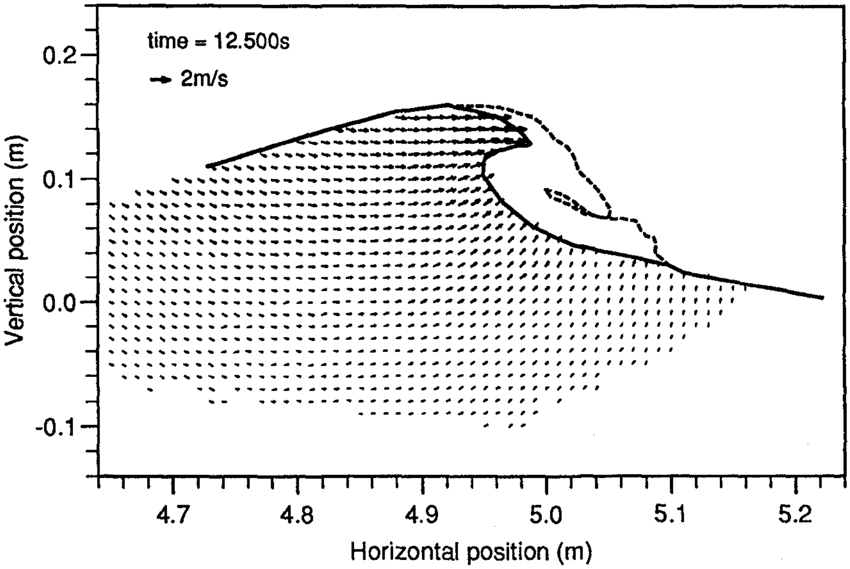

Figure 7: Vector field within a breaking wave

The results of the analysis of a PIV photograph can be displayed as a set of velocity vectors or, alternatively, iso-velocity contours may be plotted. Figures 7 and 8 show matching plots of velocity vectors and iso-velocities for a breaking wave; this type of plot is particularly useful for illustrating wave motions. Further details of this work are given in [Skyner,1992(a)].

For studies of vortex shedding or turbulence, vorticity contours may be the best method of display. Vorticities can be computed at each point by taking central differences in the two coordinate directions. Figure 9 shows a typical vorticity map recorded in the wake of an oscillating cylinder using PIV.

Other Applications

PIV is potentially applicable to the measurement of two-phase flows and has been used successfully to study particle transport in air streams [McCluskey et al,1989]. It is anticipated that it will give good results in the study of sediment movement in water channels when the sediment density is low.

FUTURE DEVELOPMENTS

Two new areas of development in PIV are particularly relevant to the coastal engineering laboratory. One is the use of stereoscopic photography to measure

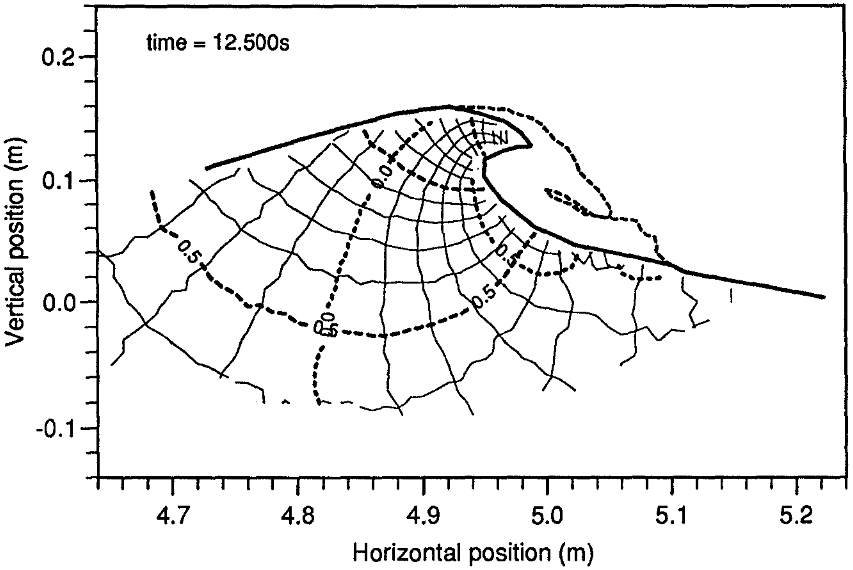

Figure 8: Iso-velocity contours of a breaking wave. Contour interval 0.1ms^{-1}

all three velocity components and the other is the development of on-line PIV using acquisition systems based on a CCD camera.

Stereoscopic Measurements

In stereoscopic photography the normal methods of illumination are used but the illuminated sheet in the flow is photographed from two different positions at the same time. There are two fundamental optical arrangements. In one of these, two cameras are angled to image the same field. The problem here is that different parts of the sheet are at different distances from the lens so it is difficult to achieve a sharp focus. Secondly, there is a distortion of the image. The other approach is to have both cameras aligned normal to the measurement plane. The problem in this case is that only part of the two images overlap and are usable for measurement. Most researchers have chosen the second method.

Two separate photographs may be taken and analysed and the resulting velocity vector plots combined by triangulation to produce the three components of velocity in a plane. For practical purposes the use of two cameras presents alignment difficulties because the resulting negatives need to be matched very precisely. To overcome this difficulty a number of systems have been developed which use just a single lens and produce the two images on a single piece of film. One such system has been devised by Arroyo and Greated [1991]. Instead of having two separate lenses in front of a single piece of film the two images are formed through the same lens. This is achieved by removing the lens from a

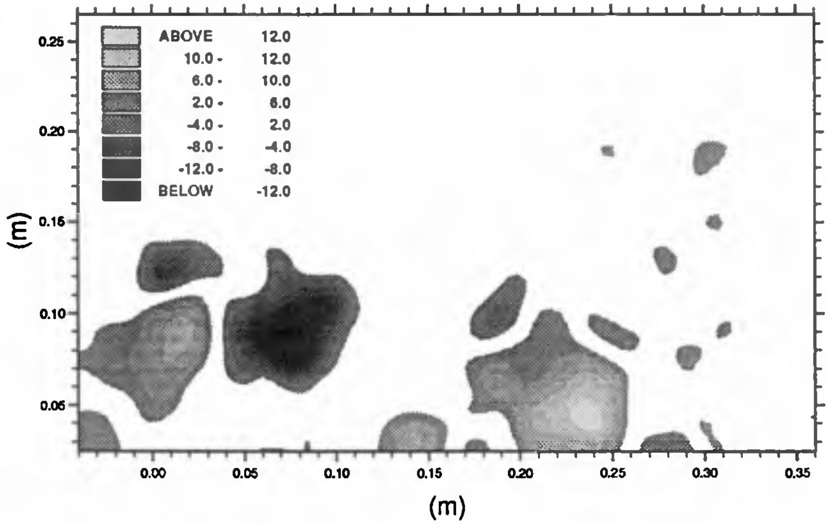

Figure 9: Vorticity in the wake of an oscillating cyclinder

conventional camera and placing two parallel mirrors between the lens and the film. Two further parallel mirrors are also placed on the other side of the lens. After alignment, this system can be operated as a conventional camera.

On-line PIV

It would be attractive if the photographic process could be eliminated completely. This can in fact be done by replacing the film camera with a CCD camera which simply captures an image electronically. Again, with multiple exposures on the same frame, the matrix of pixels needs to be subdivided and correlations carried out digitally. If different frames are used for consecutive exposures or two separate cameras are used, then the autocorrelation can be replaced by crosscorrelation, in which case directional ambiguity is resolved. Systems of this type have been implemented but in the present state of the art they are limited to extremely low velocities and the resolution is extremely poor compared to film.

Fibre Optic Delivery

Considerable developments have been made in fibre optics in recent years and it is now possible to channel powers of up to 4W through monomode fibres. The use of monomode fibres allows the light to be recollimated back into a narrow beam of light at the other end; with the multimode fibres this is difficult to do. Powers of greater than about 4W can in theory be transmitted by splitting the beam. Fibre optic linkage makes the PIV system very much more flexible and may be essential where access to the working area is severely constricted. Fibre transmission of pulsed laser beams is much more difficult due to the very high peak power densities.

ACKNOWLEDGEMENT

The authors wish to acknowledge financial support from the SERC and the EC MAST programme (contract MAST 0035C).

References

[1] Arroyo, M.P. and Greated, C.A. . *Stereoscopic Particle Image Velocimetry.* 1991. Measurement Science Technology, 2, pp 1181-1186.

[2] Bruce T. and Easson, W.J. *The Kinematics of Wave Induced Flows Around Near-Bed Pipelines.* 1992. Proc. 23rd Int. Conf. on Coastal Eng., ASCE.

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

[4] Gray, C., Greated, C.A., McCluskey, D.R. and Easson, W.J. *An Analysis of the Scanning Beam PIV Illumination System.* 1991. J. Phys. Measurement Science and Technology 2, 717-24.

[5] Jakobsen, M.L., Hossack, W.J., Greated, C.A. and Easson, W.J. *PIV Analysis Using an Optically Addressed Spatial Light Modulator.* 1992. Applied Optics and Opto-Electronics, Leeds.

[6] McCluskey, D.R, Easson, W.J., Greated, C.A. and Glass, D.H. *The Use of Particle Image Velocimetry to Study Roping in Pneumatic Conveying.* 1989. 4th European Symposium on Particle Characterization, Nuremburg, Germany.

[7] Morrison, I.G and Greated, C.A. *Oscillating Water Column Modelling.* 1992. Proc. 23rd Int. Conf. on Coastal Eng., ASCE.

[8] Powell, K.A., Quinn, P.A. and Greated, C.A. *The Formation of Shingle Beach Profiles due to Wave Impact, and the Measurement of the Wave Velocities using Particle Image Velocimetry* . 1992. Proc. 23rd Int. Conf. on Coastal Eng., ASCE.

[9] Skyner, D.J. and Easson, W.J. *The Effect of Sheared Currents on Wave Kinematics and Surface Parameters.* 1992. Proc. 23rd Int. Conf. on Coastal Eng., ASCE.

[10] Skyner, D.J. and Greated, C.A. *The Evolution of a Long-crested Deep-water Breaking Wave.* 1992. Proc. 2nd Conf. Int. Offshore & Polar Eng., 3, ISBN 1-880653-03-6, 132-8.