CHAPTER 1

QUANTITATIVE IMAGING TECHNIQUES AND THEIR APPLICATION TO WAVY FLOWS

J. Kristian Sveen[1] and Edwin A. Cowen[2]
[1] *Mechanics Division, Department of Mathematics, University of Oslo*
Box 1053 Blindern, N-0316 Oslo, Norway
E-mail: jks@math.uio.no
[2] *School of Civil and Environmental Engineering, Cornell University*
Ithaca, NY 14853-3501, USA
E-mail: eac20@cornell.edu

Quantitative imaging (QI) techniques are a general class of optically based laboratory measurement techniques used in the field of experimental fluid mechanics, which have seen rapid growth over the last two decades. They are particularly well suited for the study of wavy fluid flows which are characterized by unsteady free surfaces and internal motions. This paper presents an overview of QI techniques in general, with a particular focus on particle image velocimetry (PIV). We present QI methods in the context of the broader fields of pattern recognition and image processing techniques, which are currently used in a wide range of fields. In this review QI methods and their fundamentals are described in detail and recent developments, targeted at increasing accuracy and resolution, are described and put into perspective. More specifically we identify QI techniques as a digital data analysis (through software) set of issues built upon general principles of pattern matching. Throughout the paper we address the aspects that are particular to wavy flows although these issues can be argued to be important for any unsteady fluid flow of interest, e.g. turbulence. This review article thus serves as a general reference for the neophyte and experienced fluid mechanics experimentalist.

1. Introduction

Perhaps the greatest challenge of making measurements in wavy free surface flows is measuring flow field characteristics near the dynamically moving free surface – i.e., between the trough and crest. Quantitative imaging (QI) techniques are a robust solution to this problem, as demonstrated in Fig. 1, and hence QI techniques are becoming the methods of choice when

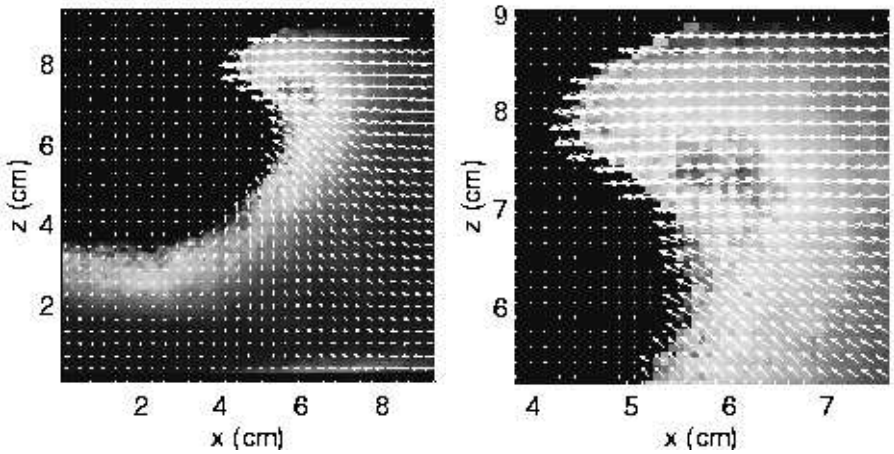

Fig. 1. Incipient shoaling breaking wave. Left image with every 4^{th} determined vector shown in each coordinate direction; right image is a magnified subregion of left image with every 2^{nd} determined vector in each coordinate direction.

attempting to interrogate laboratory wavy free surface flows. An important aspect to the emerging dominance of QI techniques is their ability to capture *whole field* properties – e.g., $\vec{u}(x,z)$ (the velocity field), $\vec{\omega_y}(x,z)$ (the vorticity field), $\overline{u'w'}(x,z)$ (the Reynolds stress field), and $\varepsilon(x,z)$ (the turbulent dissipation field). While many wavy free surface flows are periodic the reality is that due to reflections and to the difference between the phase velocity and the energy propagation velocity (group velocity) flows are often quasi-periodic at best. Researchers interested in spatial gradients have traditionally attempted to employ single-point measurement technologies at two or more spatial locations but on different experimental runs. Variability among experimental runs will lead to variability of induced rms velocities which is often on the order of the turbulence intensity itself making it extremely challenging to employ data from different runs for the determination of gradients. While QI techniques do not eliminate wave-to-wave variability they do capture a spatial field under the identical free surface conditions allowing accurate instantaneous gradients to be determined. Presently the most cited reason to work with single-point measurement technologies is a need for improved temporal resolution relative to QI techniques. However, over the last decade computers and image capture technologies have progressed sufficiently that QI techniques are now capable of reasonably high temporal as well as spatial resolutions, further

propelling the popularity of QI techniques.

A challenge for these techniques is optical access, particularly for existing facilities that were not designed with PIV in mind. Recent advances with boroscopes and index-of-refraction material matching, however, demonstrate that optical access issues can usually be overcome with a bit of ingenuity.

In order to visualize some of the methods presented in this paper, we have chosen to use the free, open-source PIV program MatPIV[1] and two example images taken from the paper by Jensen *et al.*[2].

This paper is organized as follows. Section 2 presents a brief introduction to some of the more commonly used QI techniques in the perspective of pattern matching. Section 3 is written to provide a basic overview of the fundamentals of Particle Image Velocimetry (PIV), while section 4 targets the more fundamental aspects of the technique. Section 5 gives a brief overview of Particle Tracking Velocimetry (PTV) and section 6 focuses on higher order measurements from velocity fields. Finally, Section 7 contains a short conclusion. and also provides an overview of a few other areas where the same principles of pattern matching are applied.

2. Quantitative Imaging Techniques

QI techniques can be broken immediately into several fundamentally different types of techniques – flows seeded with discrete particles, flows seeded with continuous tracers (e.g., fluorescent dyes) and unseeded flows (to look at density differences). The former are generally employed for the determination of velocity while the second and third are generally used to determine a scalar field quantity (e.g., concentration or temperature). The literature is rife with various acronyms for these types of techniques and we will briefly introduce some of the more popular QI nomenclature here.

2.1. *Particle Based QI Techniques*

- *Particle Streak Velocimetry* (PSV) -A general class of techniques where the image exposure time is long relative to the time a particle occupies a point in space. The result is images of particle streaks. The length of the streak can be calculated to determine the velocity based on the known exposure time.
- *Laser Speckle Velocimetry* (LSV) -A general class of techniques where the seeding density is sufficiently high that an image captures predominantly overlapping and interfering particle images, which can be thought of as an intensity texture or speckle field.

Essentially no discrete particle images are seen. The velocity is extracted by correlating the speckle pattern in a small subregion with that in another subregion, either optically (Young's fringe analysis) or digitally (auto or cross correlation analysis).

- *Particle Image Velocimetry* (PIV) -This term is sometimes taken to mean the entire broad class of discrete particle based techniques, however, its preferred definition is a general class of techniques where the seeding density is moderate such that the nearest neighbor distance of particle images is on the order of a few to perhaps ten times the particle diameter ensuring that all small subregions have several distinct discretely imaged particles within them and relatively few particle images overlap. The velocity field can be extracted in a number of ways, the most popular of which are digitally via auto or cross correlation analysis.
- *Particle Tracking Velocimetry* (PTV) A general class of techniques where the seeding density is sufficiently low that an image captures predominantly non-overlapping or interfering particle images and the velocity can be extracted by tracking the motion of individual particles over known times.

Other common names for particle based techniques include: pulsed light velocimetry (PLV), particle image displacement velocimetry (PIDV), particle displacement velocimetry (PDV), digital particle image velocimetry (DPIV), digital particle tracking velocimetry (DPTV), correlation image velocimetry (CIV), spatial correlation velocimetry (SCV) and large-scale particle image velocimetry (LSPIV). We note that three dimensional QI techniques have seen rapid growth recently and are based on one of four fundamental approaches: holography[3], stereoscopic imaging with multiple cameras[4], depth-of-field with a single camera[5], and scanning a light sheet through a volume[6]. We will restrict our review of QI techniques to two-dimensional implementations. Ron Adrian, in a survey through 1995[7], showed that the number of publications per year on particle based QI techniques grew exponentially. All indications are that the trend continues.

The common principle to all particle based techniques is that the instantaneous fluid velocities can be measured by recording the position of images produced by small tracer particles, suspended in the fluid, at successive instants in time. The techniques listed above, as well as others, fall into two broad categories, each with different development paths: PIV and PTV.

LSV and PIV are different operating modes of the same technique. The velocity field is generally determined on an interrogation grid and each velocity vector is the average velocity over many tracers contained in a small volume of fluid. These techniques have their roots in solid mechanics. They were originally used to determine the in-plane displacement and strain of solids with diffusively scattering surfaces.

PTV and PSV can also be thought of as different operating modes of the same technique. In contrast to PIV, velocity vectors are determined from the individual particle images or streaks produced by a single particle at random locations. These techniques have their roots in the field of flow visualization; particle streak photography and stroboscopic photography. Prandtl[8] was an early developer of particle tracking techniques, although not the first.

For the remainder of this manuscript we will use PIV to indicate the class of particle based QI techniques where the velocity is extracted by looking at the movement of an ensemble of particle images. PTV will be used to indicate the class of particle based QI techniques where the velocity is extracted by looking at the movement of a single particle image.

2.2. *Tracerless QI Techniques*

This class of techniques relies on measuring either the concentration or displacement of a chemical tracer substance added to the flow.

- *Laser induced fluorescence* (LIF) -A general class of techniques where a flow is seeded with a dilute fluorescent chemical tracer that will fluoresce proportionally with its local concentration. The local concentration field is often the objective of LIF measurements but careful choice of dye can allow the measurement of temperature, and pH as well as other scalar properties of the flow.
- *Scalar image velocimetry* (SIV) -Really a specific analysis applied to LIF images where the gradient in intensity information recorded in images is treated like speckle in LSV and correlation based analysis is used to extract velocity field information.

Another common name for chemical tracer based QI techniques is planar laser induced fluorescence (PLIF). As with the particle based techniques we note that three dimensional scalar QI techniques have seen growth recently and are generally based on scanning a laser light sheet through a volume[9,10].

2.3. *Other QI techniques*

Other QI techniques do exist that do not rely on the use of passive tracers. In fact some of the oldest flow visualization[11] techniques do not, such as shadowgraphy, schlieren imaging and the Mach-Zendner Interferometer. Today many researchers utilize so-called Synthetic Schlieren or Quantitative Shadowgraph techniques using digital cameras and the very same principles of pattern matching as are used in for example Particle Image Velocimetry. We shall briefly return to this aspect in section 7.

2.4. *QI Techniques – Image Processing and Pattern Recognition by a Different Name*

Many newcomers to PIV in particular, and QI techniques in general, are often confused by the relatively large number of details that need to be addressed in order to apply these techniques to real fluid flow experiments. One of our primary goals in this review is to focus on the foundations upon which these techniques are built. PIV in particular has received a lot of attention within the last 10 to 15 years. As implied in section 2.1, authors often use different nomenclature for what are essentially identical approaches, perhaps with subtle implementation differences. We would like to stress that PIV relies on image processing and pattern recognition analysis and as such it should more properly be viewed as an interdisciplinary field between the experimental fluid mechanics research community and the image processing and pattern recognition (IP&PR) research communities. QI techniques are widely known and used by IP&PR researchers and in fact, experimental fluid mechanicians often re-invent analysis techniques as they ignore the previous efforts documented in the IP&PR literature. The basics for understanding PIV and pattern recognition may actually be found in most introductory books on image processing[12]. The details of applying QI techniques to fluid flows essentially consists of imaging a temporally and spatially varying pattern within the flow, generally by adding discrete tracer particles although we should not feel restricted to this particular case, illuminating the flow in a nearly two-dimensional slice as the majority of imaging devices capture two-dimensional information, and using IP&PR algorithms to extract the displacement in a known time of the imaged tracers. The entire QI measurement process can be divided into two fundamental components:

(1) a hardware problem of experimental techniques, including illumination, seeding and image recording, and

(2) an analysis (software) problem of applying IP&PR techniques to extract displacement or other information of interest from the images.

The latter point should be recognized as the fundamental application of QI methods, while the former is merely an experimental problem (which in many cases is the hardest part as the majority of QI experimentalists will agree the challenge is getting 'good' images – with 'good' images in hand it is simply a matter of finding the right IP&PR technique to extract the information one desires!). Hence we contend that QI analysis in general, and PIV as a particular example, is simply a sub-specialty within the broader fields of IP&PR. For our purposes QI/PIV are defined as extracting information of interest to fluid mechanicians by image processing and pattern recognition means.

There are subtleties that we avoid by this generalization, but as a starting point this view is effective and allows us to break down a large problem into smaller, solvable pieces. We can view the more recent developments in QI and PIV analysis techniques as emerging to solve problems that arise in terms of accuracy and resolution in displacement estimates – to date the fundamental quantity that the vast majority of QI velocimetry techniques seek to extract from images.

The reader should note that applications of pattern matching is currently an area of active research. For example, looking at the motion picture industry, the MPEG standards[13] for sound and video compression, storage and transfer are built on many of the same ideas we use in PIV. As an example we can consider the DVD standard (MPEG2), where local motion estimates are used in order to limit the amount of storage required. Instead of saving every single frame in a movie, only a few (typically every 8-16 frames) are stored as full frames and only the local motion is stored for the remaining images. There are many advances from this industry that have yet to be applied to PIV. For example, we should be able to perform motion estimation in the frequency domain within the PIV framework, much like is done in image processing[14]. In this way we should be able to avoid one fast Fourier transform (FFT), and hence save about one-third of the calculation time. In section 7 we present a brief overview of other research fields that use pattern matching approaches.

3. PIV – a General Overview

Let us begin the description of the basics of PIV with a fluid flow that is seeded with particles that can be considered passive tracers, perhaps because they are very small or if they are not so small they are near neutrally

buoyant with respect to the flowing fluid. A light source, generally a laser light source, is shaped with optical components into a light sheet to illuminate the particles. The light sheet is quasi-two-dimensional in the sense that it is thin in the direction orthogonal to the plane of motion that contains the two components of velocity we are interested in measuring while it is broad in the other two directions. An imaging device, in this case a digital camera, is equipped with appropriate optics (e.g., a lens) to collect images of the particles as they pass through the light sheet. In general, computer controlled timing signals are sent to the digital camera and laser light source (or the optical components that shape the light source into a light sheet) to synchronize the light source to the camera such that discrete images of particles (e.g., short time exposure images) are captured at desired times within each collected image.

There are, of course, myriad ways in which particle images can be collected, which fall into two general categories: the single exposure of multiple images and the multi-exposure of single images. Again there are myriad ways that the mean displacement of particles in any sub-region (henceforth referred to as a *subwindow*) can be extracted from the images but the fundamental technique used in multi-exposed single images is autocorrelation analysis while the fundamental technique used in single-exposed multiple images is cross-correlation analysis. As cross-correlation analysis is more straightforward as well as more accurate, we will restrict our discussion of PIV basics to cross-correlation analysis of image pairs – that is the exposure of two sequential images, each individually, specifically for analysis by cross-correlation.

Let's assume we have collected an image pair where the second image was captured a known time, Δt, after the collection time of the first image. The most straightforward approach to cross-correlation analysis is to define a square subwindow with side length $N = 2^n$ where n is an integer. N is typically taken as a power of 2 to take advantage of determining the cross-correlation in the frequency domain via the fast Fourier transform (FFT). As will be described later N need not be restricted to these discrete values but as this restriction is frequently employed we will assume it for now. Let us assume a value of $N = 32$ and an image size of $NR \times NC$ pixels (an acronym for *"picture element"* that describes the smallest discrete unit of scattered light intensity measured by a digital camera, sometimes also referred to as a "pel"), where NR and NC are the number of rows and columns in the digital image, respectively. The simplest algorithm is to divide each image of the pair into non-overlapping $N \times N$ subwindows and to then perform the two-dimensional cross-correlation of each subwindow

pair in the image pair. The cross-correlation function is defined as

$$R(s,t) = \frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} F'_{I,J}(i,j) F''_{I,J}(i+s,\ j+t) \tag{1}$$

where R is the cyclic cross-correlation between subwindows I, J in the first image of the image pair (F') and the second image of the image pair (F''), i, j is the pixel location within subwindow I, J, and s, t is the 2-D cyclic lag at which the cross-correlation is being computed. As indicated above R is often calculated in the spectral domain and hence equation (1) can be found as:

$$R(s,t) = \mathcal{F}^{-1}\left[\mathcal{F}^*\left\{F'_{I,J}(i,j)\right\} \mathcal{F}\left\{F''_{I,J}(i+s,\ j+t)\right\}\right] \tag{2}$$

where $\mathcal{F}$ and $\mathcal{F}^{-1}$ are the Fourier and inverse Fourier transform operators and the star denotes complex conjugate. This basic feature is known as the correlation theorem and can be found in most introductory level books on image processing[12]. For the purpose of visualization we consider the example images in Fig. 2a–b, which shows the image pair F' and F'' along with the non-overlapping subwindow interrogation grid. Fig. 2c and 2d shows subwindow $I = 10, J = 15$ for each image in the pair along with panel e, which shows the resultant cyclic cross-correlation of this subwindow pair. The correlation plane contains a peak which has a maximum at $(s,t) = (26, 25)$. The displacement is measured from the center of the correlation plane to this peak. The integer displacement in our example is estimated to be $dx = -6$ and $dy = -7$.

Our simple introduction here raises several issues. First we note that the maximum unambiguous displacement that can be resolved is $N/2$ pixels. If the displacement is larger than $N/2$ pixels (but less than N pixels, so at least a few of the particles in the first image subwindow remain in the second image subwindow) the correlation peak will alias to the location $-(N - \xi)$ where ξ is the actual displacement. If the displacement is larger than N pixels then R represents the cross-correlation of two uncorrelated subwindows and the returned displacement estimate will be the result of a random noise peak (i.e., the lag where a maximum number of particles randomly align themselves between the two subwindows). Secondly, there is an implicit assumption that the particles are being translated without rotation or shear. If the particles undergo rotation and/or shear over the time Δt between the capture of the two images then we must be concerned about the effect of this rotation and/or shear on the existence of a usable correlation peak. Thirdly, if the flow is not two-dimensional where the out-of-plane motion is identically zero then there is a finite non-zero probability

of a particle appearing in the first subwindow moving out of the light sheet (and hence the image) by the time the second image is captured. This out-of-plane motion can clearly lead to a measurement bias if the out-of-plane motion is correlated to the velocity itself. And lastly, if $N = 32$ our maximum resolvable displacement is just 16 pixels. If we cannot extend this or resolve the displacements to sub-pixel accuracy (e.g., estimate the location of the correlation peak to fractional pixel values) then our maximum accuracy is greater than 3% and our typical accuracy is considerably higher than this. The above example also raises many other questions. With the above primer in PIV as our starting point we will now turn to the details of PIV.

4. PIV – the Fundamentals

There are essentially three typical implementations of PIV: single exposure multiple image (cross-correlation based), double-exposure single image (auto-correlation based) and multi-exposure single image (auto-correlation based). There are several excellent fundamental references on QI techniques in general and PIV techniques in particular and the reader should explore these references for further details and perspectives on PIV[15,16,17,18,19].

4.1. *Displacement Estimates – Correlation Approaches*

PIV techniques often rely on either auto or cross-correlation of subregions to extract the mean displacement of particles contained within the subregion. This methodology thus relies on estimating the auto or cross-covariance function between the subregions to extract the displacement estimate. For details on auto-correlation see Adrian[20] and Keane and Adrian[21]. The covariance can be determined in a number of ways but it is generally determined by either Fourier transforms[22] or the direct determination of the two-dimensional covariance function[23] (also known as the correlation coefficient[12]). The expected value of the cross-covariance $R(s,t)$ (see equation 1) is shown by Westerweel[17] to be

$$E\left\{R(s,t)\right\} = \left(1 - \frac{|s|}{N}\right)\left(1 - \frac{|t|}{N}\right) R(s,t) \tag{3}$$

and it is seen immediately to be *biased* (note the bias vanishes as $N \to \infty$). The bias occurs because the shift over (s,t) results in only a part of F' correlating with F''.

Adrian[24] reported this as the result of in-plane loss-of-pairs. In the presence of strong gradients the interpolation region contains more particle pairs

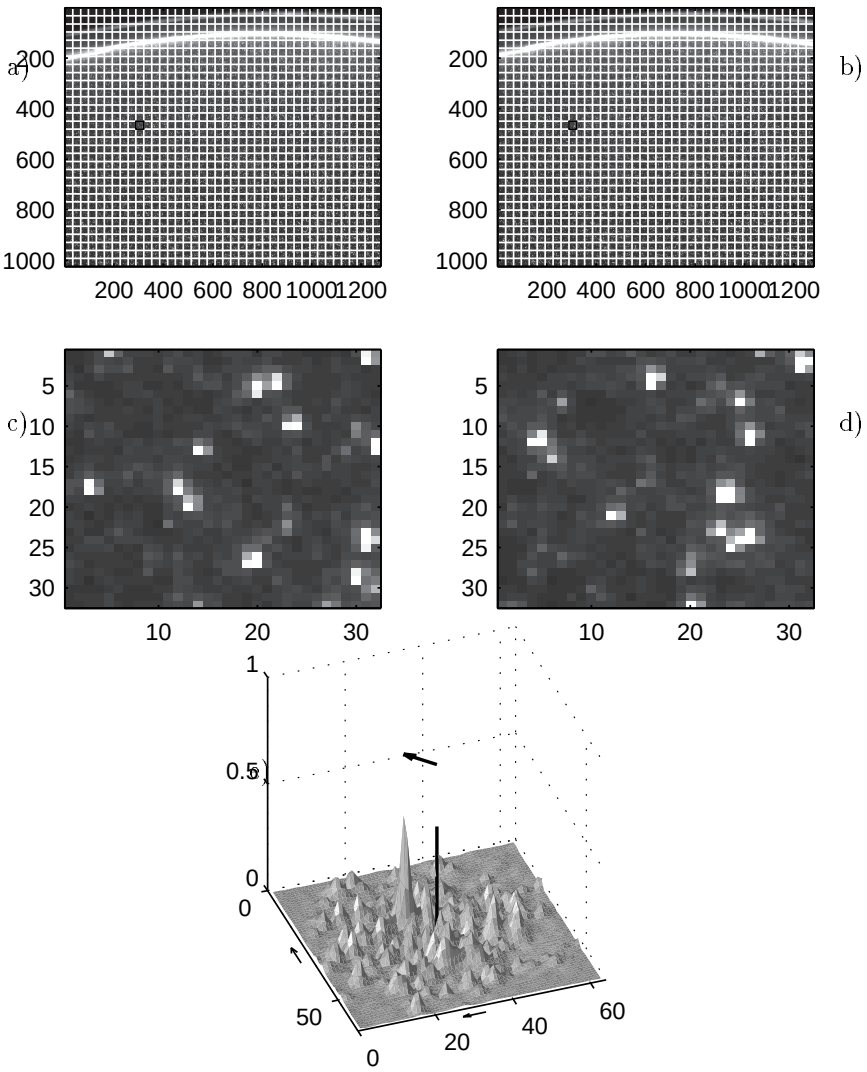

Fig. 2. Images with superimposed non-overlapping grid. Black cells marks the sub-windows shown in c) and d). Correlation plane shown in e). Upper arrow shows the displacement. Lower arrows denote the displacements along each axis. Black vertical line marks the center of the plane.

with *small* displacements, therefore, this is an *under* bias. From the above equation it is apparent that the bias grows linearly with the shift size. Westerweel shows that two conditions arise from looking at the variance of the expected value of R. He finds that the noise in R due to random correlations

is approximately a uniform random field and goes as $1/N^2$. The probability that the noise peak is greater than the correlation peak grows for increasing displacement (s, t).

4.1.1. *Removing Effects of Correlation with the Mean Background*

As shown by Adrian[24], the correlation between two images may be split into 3 contributions:

(1) correlation of the mean background intensities, $R_c(s, t)$,
(2) correlation of the mean and the fluctuating intensities, $R_f(s, t)$, and
(3) correlation of the fluctuating intensities, $R_d(s, t)$.

In this way the correlation may be written as $R = R_c + R_f + R_d$. The latter part of the correlation plane will contain the displacement peak, and to avoid the other parts of the correlation, the mean is normally subtracted from each image prior to computing the correlation. The process is shown in Fig. 3 and also included in equation (4).

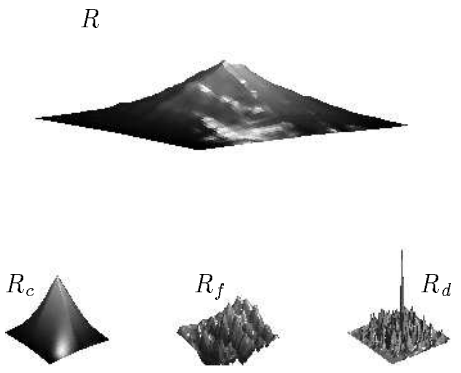

Fig. 3. Schematic figure showing the different contributions to the correlation plane.

4.1.2. *Normalizing the Correlation Plane*

In many cases it may be favorable to be able to evaluate the degree of correlation between two subwindows. The maximal value in the correlation plane is a relatively important parameter to estimate since it is a direct

measure of the degree of match between the subwindow patterns. The ratio between the highest peak and the rest of the correlation plane will furthermore provide information of the quality of the match compared to the noise. In order to evaluate these quantities we need to normalize the correlation in equation (1). One function that calculates a normalized correlation is the full correlation coefficient or covariance function, which is defined as[12]

$$\hat{R}(s,t) = \frac{1}{N^2} \frac{\sum_i \sum_j [F'(i,j) - \bar{F}'][F''(i+s,\ j+t) - \bar{F}'']}{\sum_i \sum_j ([F'(i,j) - \bar{F}']^2 [F''(i+s,\ j+t) - \bar{F}'']^2)^{1/2}}, \qquad (4)$$

where $\bar{F}' = \text{mean}(F'(i,j))$ and $\bar{F}'' = \text{mean}(F''(i+s, j+t))$. We will return to this function later, but for now we note that the numerator easily can be calculated by the use of Fourier transforms. The denominator, however, is usually approximated by $N^2 \text{std}(F'(i,j)) \cdot \text{std}(F''(i,j))$, where std() denotes standard deviation, which is defined as $\text{std}(F') = [\sum (F' - \bar{F}')^2/(N-1)]^{1/2}$, and is calculated prior to $\hat{R}$. In this way we can use the peak height in the correlation plane as a measure of how good the correlation was. A value close to 1 will indicate a very high degree of match between our windows, while a value of 0 will indicate the opposite. The underlying assumption for making this approximation is that our pattern is evenly distributed, meaning that the standard deviations do not change if we vary s and t during the calculation of $\hat{R}$. An uneven distribution of the pattern may easily reduce the accuracy of our measurements, as we shall see later in the paper.

4.2. *Minimizing Covariance Estimate Induced Bias*

The source of the bias in estimating the displacement of an ensemble of particles can be minimized by insuring that the majority of particles captured in the first subregion are also found in the second subregion. There are two traditional approaches to minimizing this bias – dynamic subwindow location and adaptive subwindow sizes.

The goal of dynamic subwindow location[22,25] is simply to insure that the mean displacement of particles contained within two subwindows is, to the nearest integer, zero pixels. Thus the second image subwindow (or both subwindows as will be discussed shortly) is dynamically located so that to the nearest integer the covariance estimate of the displacement is zero pixels. Cowen and Monismith[22] used this concept to remove spurious analysis regions. They argued that if a dynamic estimate of the integer pixel displacement does not converge to zero pixels in three iterations, the analysis of the subwindows is problematic and the subwindow pair is discarded.

Fig. 4 (middle pair) shows the concept of this dynamic subwindow location. The interrogation region in the first frame is fixed in space, while in the second image it is shifted in integer steps found after an iterative process.

Based on the information from two subwindows, the optimal estimate of displacement that can be obtained is second-order accurate. This is achieved by considering the location of the displacement estimate to be at the midpoint between the center of the first and second subwindows (this assumes the center of mass of the particles for each subwindow is at the center of the subwindow – a reasonable model but as subwindows get small there may be significant deviations from this assumption). Allowing the second subwindow to move dynamically with respect to the first subwindow results in a random perturbation to the originally intended interrogation grid due to the potentially random motion of the second subwindow. Wereley and Meinhart[26] eliminate this random perturbation by allowing both subwindows to move dynamically as shown in Fig. 4 (lower pair). This avoids the need to interpolate velocity vectors and ensures a truly second-order accurate spatial velocity field.

The goal of an adaptive subwindow size[21] is to expand the second subwindow in area so as to ensure that all particles contained in the first subwindow will remain within the bounds of the second subwindow. If this is done in the absence of discrete subwindow offsets, the potential for the second subwindow to become large exists, and the signal to noise ratio of the image (ratio of the magnitude of the true correlation peak to the rms correlation peak) will decrease causing dropouts. An improved approach is likely to couple the two techniques, handling the mean displacement with a discrete offset and the effects of local turbulence induced shear by a slightly larger second image subwindow. An optimal approach involves the consideration of the image 1 subwindow as a system in the thermodynamic sense – a collection of mass of fixed identity. Tracking the system, and thus its boundaries, from its location in image 1 to its new location in image 2 results in no in-plane loss of particles but requires that the image 2 subwindow boundaries become deformed. Several researchers have developed such optimal methods, which have become known as particle image deformation, or PID, methods [27,28,29,30,31]. PID methods are described in more detail in section 4.7.3.

4.3. *Subwindow Size*

An important consideration is the size of the subwindow to be used for PIV analysis as this leads directly to issues of resolution. Prasad *et al.*[32]

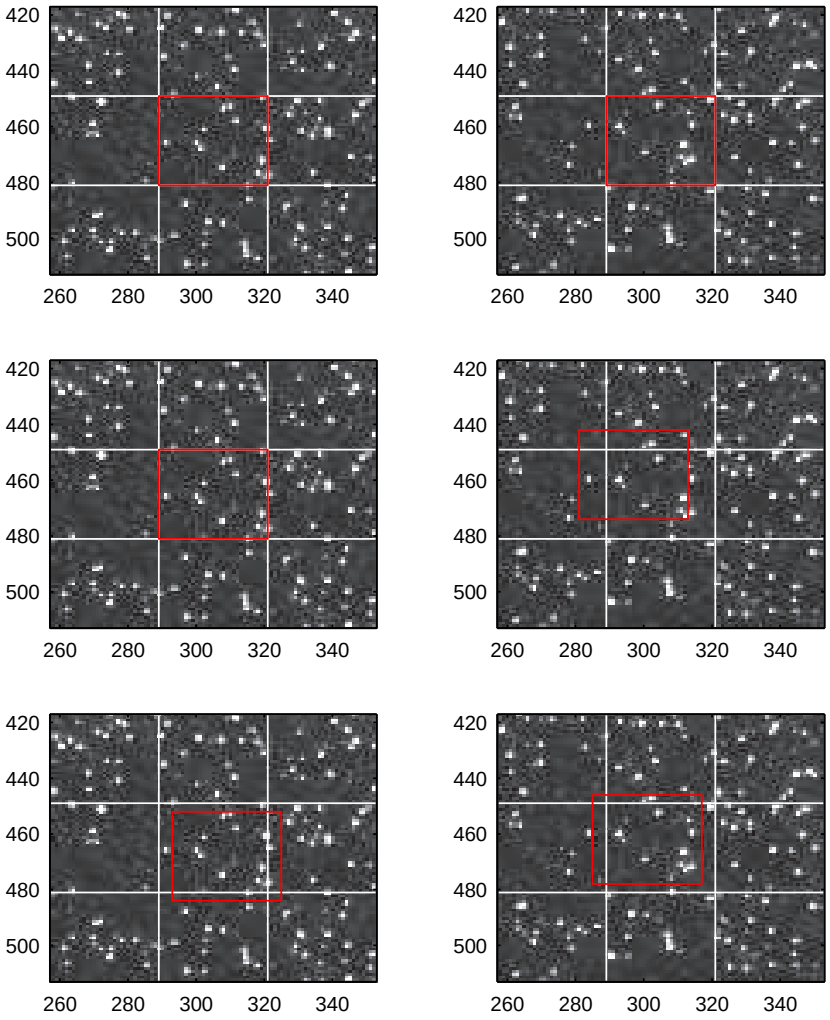

Fig. 4. Schematic view of the three most basic dynamic subwindow location techniques. The upper pair shows fixed interrogation windows. The middle pair shows first order window shifting[22,25] and the lower pair shows second order accurate window shifting[26].

investigated a range of subwindow sizes from $N = 32$ to $N = 256$ finding that a reduction of subwindow size from the traditional $N = 256$ to $N = 128$ resulted in no appreciable degradation of accuracy. Westerweel[17] explored in detail the requisite subwindow size to obtain the desired information,

namely the location of the correlation peak to sub-pixel accuracy.

Westerweel[17] sought the effective sampling rate needed to yield a good representation of an original QI image. A signal is band limited if it can be constructed from only signals with frequency components less than some maximum, i.e., its Fourier transform is non-zero over only a finite portion of the frequency domain. Therefore,

$$\mathcal{F}(F)(s,t) = 0, \text{ for } |s| > S, \ |t| > T. \tag{5}$$

The bandwidth is defined as the maximum of S and T. A band limited continuous signal can be reconstructed *exactly* (given infinite samples) provided that the sampling rate is at least *twice* the bandwidth. This rate, known as the Nyquist rate, is $\max(2S, 2T)$.

Goodman[33] has shown that the bandwidth of the image intensity for a thin spherical lens image system (aperture D, focal length f, coherent light with wave length λ) is given by:

$$W = \frac{D}{\lambda z_0} = \frac{D}{\lambda f(M+1)} \tag{6}$$

where $z_0 = f(M+1)$ is the image distance and $M = z_0/Z_0$ is the magnification (Z_0 is the object distance). Now, if an image is obtained with illumination wavelength $\lambda = 0.5$ μm, $f/D = 8$, and $M = 1$, then it follows that $W = 125$ mm^{-1} and therefore that the Nyquist rate, $2W$, is 250 mm^{-1}, explaining the traditional subwindow size of 256 pixels and a subwindow area of 1mm^2 early in the history of PIV. However, this is a significant hurdle for CCD based image acquisition.

Most CCD's pixels are at least 10 μm $\times$ 10 μm, therefore $2W = 100$ mm^{-1} is the best that can be done. However, the reality is that the goal of PIV interrogation is *not* to reconstruct the image exactly, but only to obtain the position of the displacement covariance peak. What is required is not the exact details of the particle shapes (edges) but just the details of their positions (lower wave number information). The covariance function has a spectral density function that is "nearly band limited" meaning its value vanishes for sufficiently large (s,t) but may not be exactly zero. Therefore $F(s,t) \sim 0$ for $|s| > S$, $|s| > T$. Parzen[34] showed that for a 1-D signal the bandwidth, W_p with a circularly symmetric spectral density function is defined as the width of a cylinder with the same volume, therefore

$$W_p = (\pi\, s(0,0))^{-1/2}. \tag{7}$$

It can be shown that if a square region is sampled instead of a circular region and the covariance of particle images with diameter d are of interest, the

covariance width, σ_h, is $\sigma_h = d\sqrt{2}/(2.44\pi)$ which leads to $W_p = 1/(2\pi\sigma_h)$. Thus, given the same optical parameters as above and $d = 20\mu$m ($\sigma_h = 3.7\ \mu$m), $2W_p = 86\ \text{m}^{-1}$, suggesting that 64×64 pixels over a bit less than 1 mm^2 is sufficient. Note that larger particle image diameters reduce the required sampling rate, so that 30μm images $\rightarrow\sim 32\times 32$ pixels/mm^2. Thus the dominant subwindow sizes in currently used PIV algorithms are $N = 64$ and $N = 32$. For more details on band limited signals see Westerweel[17].

4.4. *Sub-pixel Displacement Estimation*

It can be shown that for any non-zero particle image diameter the width of the correlation peak will be greater than one pixel. By including values adjacent to the maximum in R, the center of the peak can be estimated to sub-pixel accuracy.

We decompose the displacement (s, t) as

$$s = s_0 + \epsilon_s, \qquad t = t_0 + \epsilon_t \tag{8}$$

where s_0, t_0 is the integer displacement and ϵ_s, ϵ_t is the fractional part. Therefore

$$-0.5 < \epsilon_s < 0.5 \qquad \text{and} \qquad -0.5 < \epsilon_t < 0.5. \tag{9}$$

In the absence of an estimate for ϵ_s and ϵ_t the error is $\pm\Delta/2$ where Δ is one pixel. For $N = 32$ a displacement of eight pixels has an uncertainty of $\pm 0.5/8 \sim 6\%$.

If we digitize the same subwindow into 256 pixels ($N = 256$) the uncertainty is now $\pm 0.5/64 \sim 1\%$. Therefore to work at small N, we need sub-pixel estimates of the center of the covariance peak. It can be shown that the covariance $\text{cov}\{\epsilon_r,\ \epsilon_t\} = 0$ (i.e., $\overline{\epsilon_r \epsilon_t} = 0$), therefore, we can work with the 1-D problem without loss of generality. For narrow correlation peaks the covariance width is small enough that only the nearest pixels to the peak contain significant information. Since QI techniques generally satisfy the conditions for narrow covariance peaks, three-point sub-pixel estimators are generally sufficient. As an example we will consider the correlation peak shown in Fig. 2e. Fig. 5a shows a close up of this peak, while Fig. 5b shows the highest value, R_0, of the correlation along with the two nearest pixels in the vertical direction, R_{-1} and R_{+1}. The three most commonly used estimators are Center-of-Mass, Parabolic fit and Gaussian fit.

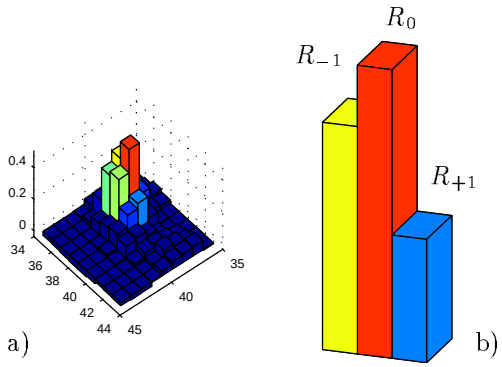

Fig. 5. Close-up of the largest peak in the correlation plane in Fig. 2e.

Center-of-Mass (COM)

The sub-pixel estimation of the Center-of-Mass is calculated from

$$\epsilon_C = \frac{R_{+1} - R_{-1}}{R_{-1} + R_0 + R_{+1}}. \tag{10}$$

This method ignores the need for a peak since R_0 can be less than R_{-1} or R_{+1}. Fig. 5 shows an example of a correlation peak and the positions of R_{-1}, R_{+1} and R_0.

Parabolic

Another option is to assume that the peak has a certain shape, for example a parabola. We can then evaluate

$$\epsilon_P = \frac{R_{-1} - R_{+1}}{2(R_{-1} - 2R_0 + R_{+1})}. \tag{11}$$

This formula actually requires a peak since $|R_0|$ must be greater than $|R_{-1}|$ and $|R_{+1}|$.

Gaussian

The most common assumption in PIV is to use a Gaussian peak-fit. This is mainly due to the fact that spherical particles image as Airy functions and the central lobe of an Airy function is well approximated by a Gaussian curve.

The log of a Gaussian curve is parabolic and therefore

$$\epsilon_G = \frac{\ln R_{-1} - \ln R_{+1}}{2(\ln R_{-1} - 2\ln R_0 + \ln R_{+1})}. \tag{12}$$

Note, again that R_0 must be a peak and $R_i > 0$ for $i = -1, 0, 1$. Clearly all three are a balance of R_{-1} and R_{+1} with a normalizing value in the denominator. Note that ϵ_C will *always* yield a result and ϵ_G has the most restrictions. Therefore we expect ϵ_C to be the most robust and ϵ_G the least. On the other hand, ϵ_C does not even acknowledge that we have a peak while ϵ_G acknowledges and uses shape information. Therefore we may expect ϵ_G to be the best performer and ϵ_C the worst.

What is the expectation of these estimators? We begin with the bias by ignoring the fractional displacement for a moment. Therefore $R_{-1} = R_{+1} = a_1 R_0$ where $0 < a_1 < 1$ and a_1 is proportional to the width of the correlation peak. Westerweel[17] shows

$$E\{\hat{\epsilon}\} = -\frac{2}{N - m_0} \cdot \begin{cases} \dfrac{a_1}{1 + 2a_1} & \text{for } \epsilon_C \\ \dfrac{a_1}{4(1 - a_1)} & \text{for } \epsilon_P \\ \dfrac{1}{4\ln(1/a_1)} & \text{for } \epsilon_G \end{cases} \tag{13}$$

Clearly the terms on the right hand side of the bracket are simply constants, but note that the first term is a function of the displacement m_0 (the integer pixel displacement). This is a negative bias, resulting from the probability of smaller velocities leading to higher pairing. If PIV is to operate in an unbiased manner this bias *must* be corrected. The correction is:

$$R^*(s,t) = \frac{R(s,t)}{F_I(s,t)} \tag{14}$$

where

$$F_I = \left(1 - \frac{|s|}{N}\right)\left(1 - \frac{|t|}{N}\right). \tag{15}$$

We note that this could have been predicted looking at equation (3).

We now consider the fractional part of the displacement. A particle will have a small imbalance in R_{-1} and R_{+1}; therefore we will model the particle image as:

$$R_{-1} = a_1^{1+\epsilon}, \quad R_0 = a_1^{\epsilon} \text{ and } R_{+1} = a_1^{1-\epsilon}. \tag{16}$$

Note that this is *not* Gaussian. Westerweel[17] shows that in this case

$$\epsilon_C = -\frac{\sinh(\epsilon \ln a_1)}{\cosh(\epsilon \ln a_1) + \frac{1}{2}a_1^{\epsilon-1}}, \tag{17}$$

$$\epsilon_P = \frac{\sinh(\epsilon \ln a_1)/2}{\cosh(\epsilon \ln a_1) - a_1^{\epsilon-1}}, \tag{18}$$

$$\epsilon_G = \frac{\epsilon}{2 - 2\epsilon}. \tag{19}$$

The Gaussian sub-pixel fit is independent of a_1, but COM and parabolic are functions of a_1. If we look at the RMS error, $\int_{-1/2}^{1/2}(\text{error})^2\, d\epsilon$, we find that the Gaussian performs the best under *all* particle diameters.

We note that other peak fitting functions have been applied in the literature, for example various versions of spline-curve fits[23]. These, however, have the drawback that they are considerably more computationally intensive without significant gain in accuracy. As will be discussed in the next section, there are alternative methods to improve the sub-pixel fit.

4.5. *Peak-Locking and Solutions*

Westerweel[17] has argued that the requisite bandwidth for a PIV image is constrained only by the need to locate the correlation peak. Hence, if this bandwidth criteria is met, we should be able to locate the correlation peak to subpixel accuracy in an unbiased manner. However, as demonstrated in section 4.4, we find in general a strong bias of subpixel fit estimators toward integer pixel locations. Looking at figure 6 we note that for actual subpixel displacements of $-0.5 < \epsilon < 0$ the bias error is positive – hence the determined displacement is biased *towards* a subpixel displacement of zero. We see the same effect for $0 < \epsilon < 0.5$, namely the bias error is negative resulting again in a bias toward zero pixel displacements. This effect is often referred to as *peak-locking* as correlation peaks have a tendency to 'lock' onto integer-pixel displacements.

The easiest way to identify the presence of peak-locking in a data set is to plot the histogram of the un-calibrated displacements (i.e., the displacements in pixels) as shown in figure 7. Depending on the quantity of interest peak-locking may not be a significant problem. For example, if one is interested in the mean velocity or even the variance, peak-locking contributes little to the error as long as the histogram spans at least two integer pixels of displacement. If the probability density function (PDF) of the underlying velocity field is of known shape then the histogram could be corrected by redistributing the displacements to have the assumed form of the PDF.

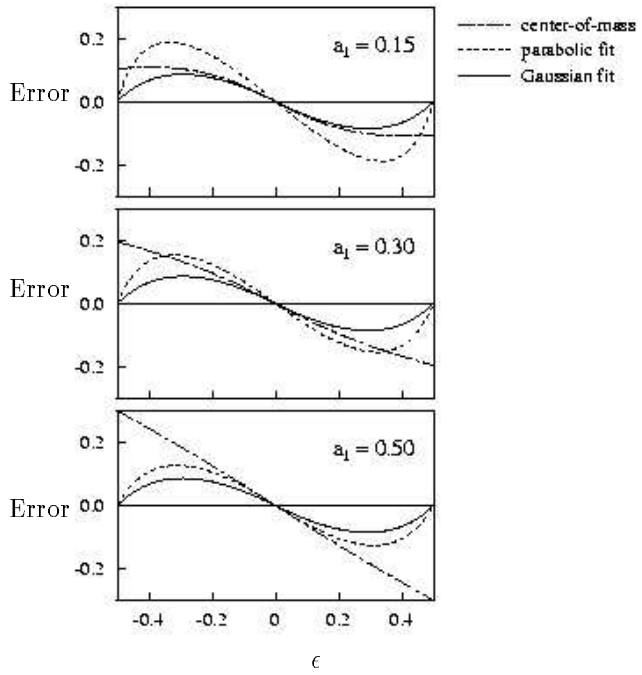

Fig. 6. Sub-pixel bias. Comparison of different sub-pixel peak fits. Figure reprinted with permission from Westerweel[17].

However, if the displacements range over less than two integer pixels the possibility of a strong bias in all calculated statistics exists. Further, if one is interested in derivative quantities (e.g., vorticity, acceleration) or other quantities involving the difference of two displacement calculations (e.g., turbulent structural functions) then errors induced by peak-locking effects may be considerable and must be monitored.

There has been research on minimizing or even eliminating the effects of peak-locking. As there is no peak-locking effect if the subpixel displacement is truly zero the majority of efforts have focused on iterative solutions where the image subwindows are displaced dynamically in fractional steps (a process that is known as continuous window shifting as opposed to the more conventional discrete window shifting where subwindows are shifted an integer number of pixels). Continuous window shifting is generally achieved by interpolating the original image at subpixel locations. The idea is to dynamically shift the subwindow(s) until the subpixel displacement is identically zero. Gui and Wereley[35] have found that continuous

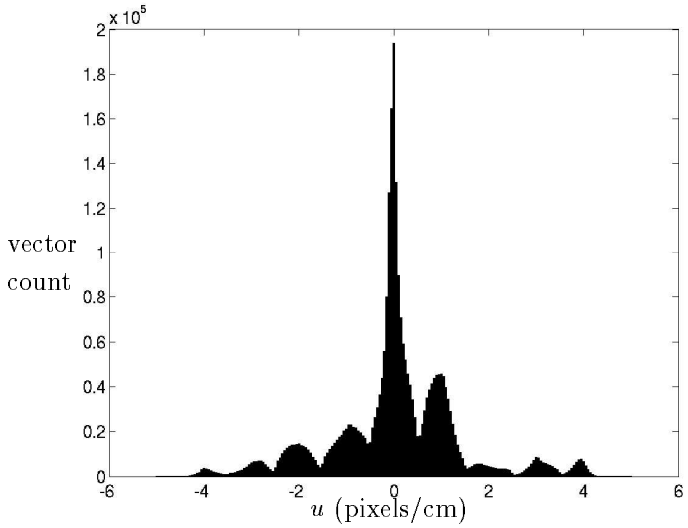

Fig. 7. Example of velocity histogram showing the peaklocking effect.

window shifting using simple bilinear interpolation and as few as four iterations provides a significant minimization of peak-locking effects. Fincham and Delerce[29] employ a slightly different approach to continuous window shifting as they use a spline-thin-shell interpolator to analytically model the first image subwindow and then utilize a PIV result already calculated for the image pair to translate and deform this analytic function based on the velocity field estimate. This analytic function is then compared to the subwindow in the second image in a least-squares sense and iteratively adjusted to minimize the difference. They demonstrate that the final result reduces peak locking and gradient bias while enhancing the robustness of PIV in the presence of locally large shear. It must be pointed out that their approach is computational expensive without any obvious improvement relative to simple bilinear interpolation. As an alternative to interpolation in the spatial domain, Liao and Cowen[36] have shown that interpolations for continuous window shifting can be performed more efficiently in the spectral domain taking advantage of the shift property of the Fourier transform. They demonstrate the elimination of peak locking effects at minimal computational cost, given the Nyquist particle image sampling criteria is met. We also note that bi-linear subpixel interpolation has been used earlier in other contexts of pattern matching which we will return to in section 7.2.

Finally we note that the solutions to peak-locking problems are closely

connected to the evaluation methods in question, and readers may find several relevant publications cited in section 4.7.3. In one of the more recent works, Scarano and Riethmuller[30] give a general review and comparison of several iterative PIV methods.

4.6. *Spurious Vector Detection*

There are three fundamental approaches to test the validity of an individual vector: signal quality, spatial consistency (smoothness in space), and temporal consistency (smoothness in time). Signal quality based tests look at a specific property of the images, such as the magnitude of the correlation peak relative to the magnitude of the background noise in the correlation plane, to test the validity of a vector. Spatial consistency tests compare each vector with certain properties calculated from the local neighborhood of vectors (local filtering), or the entire vector field (global filtering). Temporal consistency tests are similar to spatial consistency tests with the exception that the vector whose quality is being tested is compared to either a statistic developed locally in time (e.g., acceleration) or globally in time (e.g., its distance from the mean relative to the standard deviation[22]).

Filtering using spatial statistics is the standard approach, however, and has the added advantage that it works on individual instantaneous vector fields. That said, the recent advances in camera speed have lead to temporally resolved data. As single-point measurement techniques can only rely on the temporal history of data to remove spurious points there is a well developed literature on the removal of spurious data points from temporal records (e.g., spike detection in acoustic Doppler velocimetry, which is essentially a local gradient (the acceleration) threshold in the time domain). The time domain is a relatively unexploited method for spurious vector detection in PIV which should be developed.

The detection of spurious vectors is an essential part of any PIV measurement, and the topic is well covered in the literature[17,37,38]. Below is a review of the fundamental approaches currently in use.

4.6.1. *Detectability*

It would be nice if we could simply look at a correlation-peak (PIV) and decide, based on some criteria, whether or not we think it is valid. Keane and Adrian[39] proposed the detectability, D_0, as such a measure. They define the detectability as the ratio of the peak value of the first correlation peak to the second and suggest that if D_0 is greater than $1.2 - 1.5$ that the vector is valid. The detectability criterion is often referred to as signal-to-

noise ratio (snr) filtering. Westerweel[37] demonstrates that any of the three methods based on the statistics of the measured vectors in a field perform better. These methods are known as global mean, local mean and local median filters.

4.6.2. *The Global Mean*

Often during PIV measurements we observe that the data contains a few vectors that are very different from the whole ensemble of vectors. These vectors are typically the result of either R_{-1} or R_{+1} being close to zero, see equation (12).

If $v_{i,j}$ is the observed vector at i, j then the global mean $\bar{v} = \frac{1}{N}\sum_{r,s} v_{r,s}$. The variance of this statistics will be $\sigma_v^2 + \sigma_\epsilon^2$ where σ_v^2 is the variance of the actual velocity field and σ_ϵ^2 is the variance of the error. So, one possibility is to try to select the allowable variance of the signal and remove outliers (invalid vectors). However, $\sigma_\epsilon^2 \ll \sigma_v^2$ in practice so this is *hard*, hence, a user determined threshold is generally applied.

4.6.3. *The Local Mean*

Let us now extend this idea by looking at individual vectors and comparing them with their closest neighbors. As opposed to looking at individual vectors compared to the global properties, this filter should be able to remove local "large" vector to vector variations. To do this we evaluate

$$\bar{v}_{ij} = \frac{1}{N_M} \sum_{k,\ell \in M} (v_{i+k,\, j+\ell} - v_{i,j})$$

typically using a 3×3 region ($N_M = 8$) around each vector, where M is the set of points in the neighborhood.

The common approach is now to say that a vector is invalid if it is "very" different from the mean of its neighbors (for example twice as large). We immediately notice that in the cases where the neighborhood actually contains one or more outliers, our criterion for filtering is very difficult to determine, and actually spurious vectors are smoothed out with this method. It turns out we need to know *a priori* which vectors are bad to make reasonable decision criteria. If we can, perhaps using global mean, this performs fairly well.

4.6.4. *The Local Median*

The obvious extension of the local filter is to use the local median value instead of the mean. The median value is the n^{th} value in a $2n+1$ long sequence (data sorted in ascending order) and is well known to be robust to outliers for estimating the mean of data. It should be relatively obvious that spurious data normally will be sorted to either side in this sequence. The advantage is therefore that the median is much more robust to outliers and eliminates the need for *a priori* information on which vectors are bad. Westerweel[37] finds that this performs the best.

4.6.5. *Adaptive Gaussian Filtering (temporal filtering)*

Given the underlying stochastic nature of turbulence it is possible to exploit knowledge of the PDF of the velocity component statistics to remove spurious vectors. Cowen and Monismith[22]demonstrate such a technique for both PIV and PIV stray vector removal. The idea is to choose a robust estimator of the center of the distribution (accumulated in time as the temporal ensemble at a point), such as the median, and then iteratively test the data for lying within the assumed bounds of the PDF (in the case of Cowen and Monismith, a Gaussian) given the number of data points in the ensemble. The data that is retained is used to recalculate the center of the data, where the mean may be used after the first iteration in general, and the retained data is again tested for lying within the bounds of the assumed PDF. This processes is iterated until no data is removed and hence is an adaptive filter.

4.6.6. *Filtering Example*

Fig. 8 shows a velocity field as measured by conventional FFT-based PIV and filtered using a global filter, a local median filter and a signal-to-noise ratio filter. The images in question are identical to the images shown in Fig. 2 and Fig. 11. Fig. 8a–c shows a part of the velocity field as measured from the example images. The three sub-panels depict the velocity field after first a snr-filter has been applied, followed subsequently by a global filter and a local filter.

First the vector field has been filtered by a signal-to-noise ratio filter, discarding all vectors where the ratio of the tallest peak in the correlation plane to the second tallest is less than 1.3.

Secondly the remaining vectors are filtered using a global filter with the criterion defined as

$$v_{i,j} \gtrless \mathrm{mean}(v_{r,s}) \pm k \cdot \mathrm{std}(v_{r,s}), \tag{20}$$

where k is a constant defined by the user (in the present example $k = 3$) and $r = 1, \ldots, R$ and $s = 1, \ldots, S$ and (R, S) is the number of velocity vectors.

Finally a local median filter is applied, and in this case the criterion used for validation is

$$v(i,j) \gtrless \text{median}[v(r,s)] \pm k \cdot \text{std}[v(r,s)], \tag{21}$$

where again k is a user defined threshold (in the present example $k = 3$) and $r = i-1, i, i+1$ and $s = j-1, j, j+1$ are the indexes of the 3×3 neighboring vectors.

4.6.7. *Correlation Based Correction*

Another approach to vector validation was suggested by Hart[40], and in fact his proposal is more a general processing technique than a filtering technique. The idea is to take a correlation plane and compare with one or more adjacent correlation planes. The comparison is performed via an element by element multiplication, and thus represents a zero-dimensional correlation (therefore the name "Correlation Based Correction" or CBC). The point here is that each correlation plane typically will consist of a peak that signifies the displacement plus several additional peaks which are due to noise. The noise-peaks are randomly distributed in the correlation plane, but the displacement peaks are not. Therefore the multiplication of two or more correlation planes cancels the noise peaks, while the displacement peak remains. The principle is shown in Fig. 9. In this way we are able to calculate the displacement of a vector centered between the adjacent regions used in the calculation.

There are basically two ways to use CBC processing. The first is to use the signal to noise ratio in the combined correlation plane as an indicator, discarding vectors where this ratio is smaller than a certain threshold. The second method, suggested by Hart[40], is to compare the peak in the combined correlation plane with the peaks in each of the correlation planes that produced it. If the peak in the combined plane exists as a peak in at least one of the other planes, it is assumed to represent the local displacement. Hart[40] concludes that the method improves sub-pixel accuracy and eliminates spurious vectors, reduces bias errors and improves vector yields. However, he provides no evidence of the reduction in bias errors with regard to peak locking and, in fact, tests performed in the preparation of this manuscript suggest that CBC processing may be sensitive to peak locking.

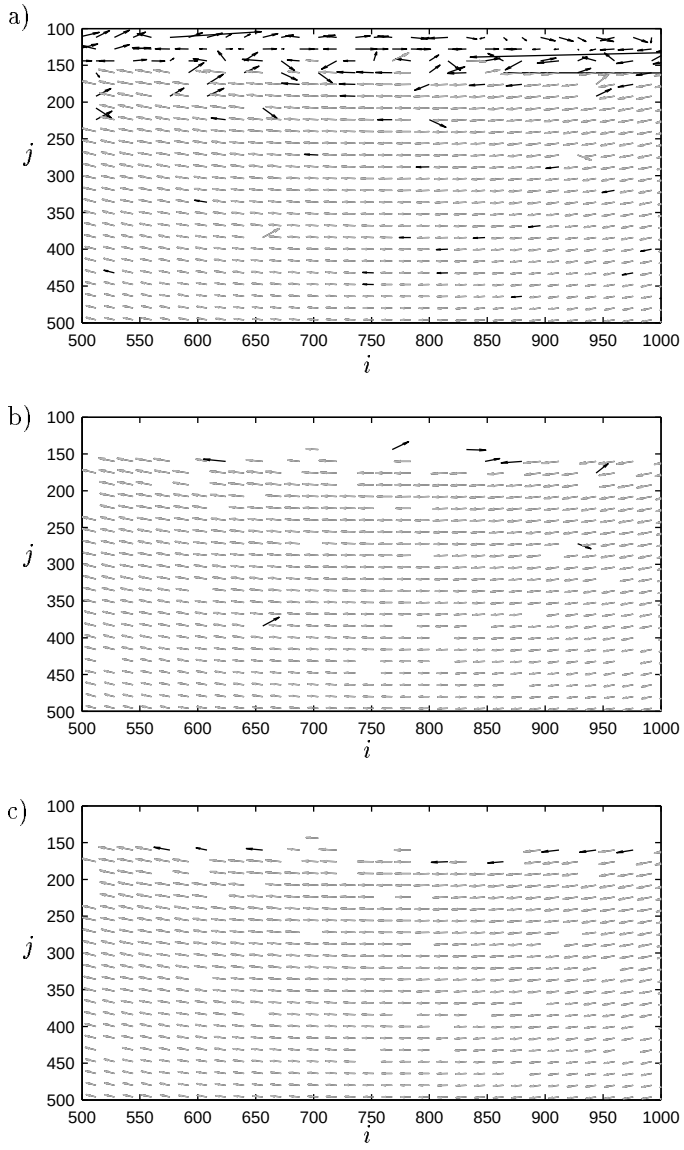

Fig. 8. Example of PIV velocity filtering. a) Velocity-field - black vectors removed by the snr-filter. Gray vectors are the remaining. b) shows the results after the global filter - gray, valid, black, outliers. c) black vectors, outliers identified by the local median filter.

4.7. *Alternative Displacement Estimation Approaches*

We have until now only discussed the "original" approach to PIV, namely by evaluation of the cross-correlation function via the use of FFTs. Further-

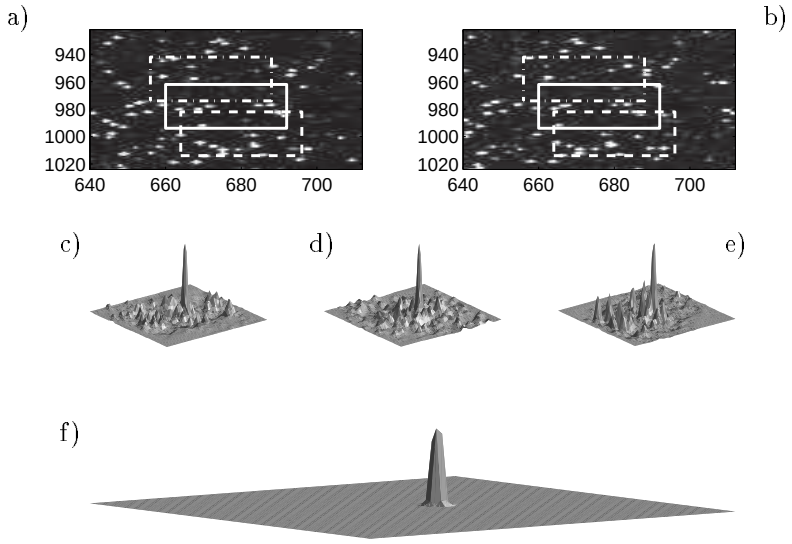

Fig. 9. a) and b), solid square corresponds to panel c), dash-dotted square to panel d) and dashed square to panel e). Multiplication of correlation tables c), d) and e) produces the combined correlation plane f).

more we have seen that in order to normalize the correlation function, we must divide it by the covariance function. The latter is hard to implement using FFTs and therefore a first order approximation is applied by dividing the correlation plane with the standard deviations of the original images.

In the following sections we will review a few different approaches that avoid the use of FFTs. Typically their downside is the additional computational time needed, but on the other hand often at improved accuracy.

The cross-correlation function is used in many applications for pattern matching. By definition it is a measure of the "likeness" of two images (or subwindows). In the cases where there is a certain degree of match between the images, it will produce a plane which will contain a peak at the position of the highest match. In contrast we could have chosen to look for the position where the images are the least unlike. This can be done by looking at the minimum of the squared difference between the images.

4.7.1. *MQD*

The use of the minimum squared distance (MQD) for PIV evaluations was first proposed by Gui and Merzkirch[41] and subsequently expanded in Gui

and Merzkirch[42]. The MQD method evaluates the function

$$R(s,t) = \frac{1}{MN} \sum_{i=0}^{M-1} \sum_{j=0}^{N-1} [F'(i,j) - F''(i+s,\ j+t)]^2,$$

where the location of the *minimum* value of R indicates the displacement between the windows. The method is reported to increase the accuracy compared to the cross-correlation function, and it is claimed that the reason for this is that "... the MQD method contains a term which accounts for a non-uniform distribution of the particle images and for a non-uniformly distributed illumination intensity...". If we turn to the basic mathematics of the method[43], it should be fairly obvious that

$$\sum_{i=0}^{M-1} \sum_{j=0}^{N-1} [F'(i,j) - F''(i+s,\ j+t)]^2 =$$
$$\sum_{i=0}^{M-1} \sum_{j=0}^{N-1} [F'(i,j)^2 - 2F'(i,j)F''(i+s,\ j+t) + F''(i+s,\ j+t)^2], \quad (22)$$

which in turn means that the MQD method equals the (zeroth order) auto-correlation of the first (sub-)image, plus the autocorrelation of the second (sub-)image (or a part of it, that is), minus two times the cross-correlation of image one and two. If the images now contain a nonuniform seeding or illumination, this will show up in both the auto-correlations and in the cross-correlation, but will cancel each other out in the MQD function, roughly speaking. If, on the other hand, the images are uniformly seeded and the illumination is uniform, there is no need to evaluate the full MQD-function.

The evaluation of the auto-correlation function can be understood as a measure of the non-uniformities within the overlapping regions. As an example we may consider each interrogation region as a wave-field, where the particles show up as small waves. A non-uniformity in the seeding density, or a background illumination gradient, will appear as a lower frequency wave in this field. Depending on the amplitude of this lower frequency, our PIV measurement may in some cases be overshadowed by it.

We may argue that the comparison between MQD and cross-correlation is a little unfair since the MQD plane is not normalized. A proper normalization of the cross-correlation plane can only be achieved in the spatial domain, and this is the basis of the next evaluation method we shall look at.

4.7.2. *CIV*

The cross-correlation function, as defined in equation (1), may be used for pattern matching, but the reader may notice that the product will be sensitive to changes in the amplitude in the image matrices. If we double the values in one of the images, the correlation values will double. For this reason the function known as the correlation coefficient (also termed the covariance function earlier in this paper), defined as

$$R(s,t) = \frac{1}{N^2} \frac{\sum_i \sum_j [F'(i,j) - \bar{F}'][F''(i+s,\ j+t) - \bar{F}'']}{\sum_i \sum_j ([F'(i,j) - \bar{F}']^2 [F''(i+s,\ j+t) - \bar{F}'']^2)^{1/2}} \tag{23}$$

is often used in pattern matching[12] (with $M = N$). Here $\bar{F}''$ is the average value of image 2, which is evaluated only once, and $\bar{F}'$ is the mean value of image 1 which needs to be evaluated for every (s,t). As we already have seen, the denominator may be approximated by $N^2 \mathrm{std}(F')\mathrm{std}(F'')$ in the cases where we have uniform illumination or seeding. If this is not the case, the standard deviation of the region in image 1, which overlaps image 2, will actually change as the values of (s,t) change.

The method of direct calculation of correlations in PIV, has been suggested by a few authors, most notably Huang *et al.*[27], but also Fincham and Spedding[23]. The latter authors suggest that there should be two main reasons for calculating the correlations in the spatial domain. Firstly they point out the argument about normalization that we have already presented. Secondly they claim that the use of FFTs place a limitation on the size of the interrogation windows. Traditionally interrogation windows needed to be of size 2^n, but recent developments in FFTs have made this demand less stringent[44].

We note that the benefits of calculating the full correlation coefficient is well known within the image processing community[12,43] and this fact can be found in many introductory level books on image processing.

4.7.3. *PIP-matching and PID*

The concept of Particle Image Pattern (PIP) was first introduced by Huang *et al.*[27] and is basically a direct calculation of the full correlation coefficient, see equation (23). The authors vary the interrogation window sizes in order to maintain all particles in both frames.

The works of Huang *et al.*[27,28] also contain an approach to dealing with large gradients in flows, known as particle image deformation, or PID. Large gradients may lead to interrogation regions with a spatial variation in velocities, again leading to loss of correlation. To adjust for this they

distort the interrogation windows in an iterative manner where the rotation and shear is calculated from the velocity field and subsequently used to distort the images. These distorted windows are then used to calculate a new velocity, and the process is iterated upon until some predefined criteria is met. The window distortion is calculated from the shear components $(\partial u/\partial y,\ -\partial v/\partial x)$ and strain components $(\partial u/\partial x,\ \partial v/\partial y)$, neglecting the translation. A schematic drawing is shown in Fig. 10.

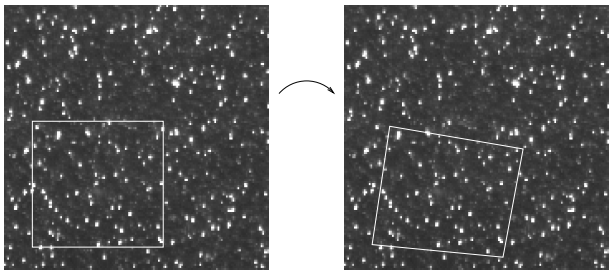

Fig. 10. Schematic drawing of the interrogation window distortion in the PID method.

Scarano and Riethmuller[30] extended the PID method to also use interrogation window offsets, as well as distortion (Originally the PID method used larger interrogation regions in the second window to maintain a high correlation). Additionally they apply progressive grid refinements to maximize the spatial resolution, and in this respect they combine the results of Huang *et al.*[27,28] with those of Westerweel *et al.*[25]. They report a decreased sensitivity with respect to peak-locking effects.

Several authors[46,47,48,49] have applied the PID concept or similar approaches and readers are advised to study these works individually. For the present review, the results by Lin and Perlin[47] are of particular interest since it deals with PIV measurements in water waves. More specifically they combine the PID methods with a particle tracking approach[50] in order to measure velocities in the boundary layer of gravity-capillary waves.

4.7.4. *Other Methods for Displacement Estimation*

Some researchers have suggested the use of the phase correlation functions for pattern matching. This method is detailed in Kuglin and Hines[51], and it is based on using the phase information in the images rather than the amplitude information.

Finally we would like to mention that within the field of Optic (or Opti-

cal) Flow, another approach to motion estimation is often used. Under the assumption that the image intensity, E, is conserved and that the displacements between frames are small, one may solve

$$E_t(x, y, t) + \mathbf{v} \cdot \nabla E(x, y, t) = 0, \tag{24}$$

where subscript denotes partial differentiation with respect to time, $\mathbf{v}$ is the velocity vector and ∇ is the gradient operator. Interested readers may find a good starting point in the paper by Beghdadi *et al.*[52].

4.8. *Comments on the Normalization of the Cross Correlation*

In the cases where the interrogation windows contain an unevenly distributed pattern or where we perhaps have non-uniformities in our illumination, the cross correlation may contain errors. In these cases the calculation of the full correlation coefficient may be more accurate. Since this is a time consuming task to calculate, an approximation is often used for the denominator. In PIV we normally use the standard deviations of the two sub-windows. For example in wavy flows, we often find a thick band of light at the fluid surface, where the light sheet leaves the fluid. For demonstrational purposes we will consider the work by Jensen *et al.*[2], where errors in the PIV calculations are observed close to the free surface. Two images from those experiments are investigated (figure 4 b in that paper), and we focus our attention at this problematic region. Fig. 11a–b shows the original images taken with a time separation of $\Delta t = 0.0012$s. Fig. 11c–d shows two corresponding 64×64 pixel subregions, close to the free surface. The reader may notice that there is a gradient in the background illumination in these two sub-windows. Correlation via FFT produces the correlation plane shown in Fig. 11e and, although difficult to spot visually, the central peak is actually higher than the true displacement peak. The central peak corresponds to a correlation of the gradient in the background illumination within the sub-windows, and since this does not move between exposures, it shows up as a zero-displacement peak. If we, on the other hand, use direct calculation of the cross-correlation function, we get a correlation plane as shown in Fig. 11e. Here the true correlation peak is slightly taller than the false (zero-displacement) peak. We observe that the full correlation calculation is less sensitive to the background illumination gradient. This brief example has an interesting aspect if we compare with the results from the MQD algorithm. Fig. 12 shows the corresponding plane as calculated using the MQD method. Interestingly we can observe that due to the nature of

the algorithm, the false correlation peak has opposite sign of the true displacement peak and the true displacement may easily be identified as the minimum value (black color) in this figure. This indicates that the MQD method may be favorable in many cases where the image quality is not ideal.

In other areas of pattern matching several approaches for calculating the normalized correlation plane are used. For example one method was developed for the motion picture *Forest Gump*[43], which relies on using precomputed integrals of the image and image energy over the interrogation region.

5. PTV

Particle Tracking Velocimetry is probably the oldest of the particle based QI techniques. In its simplest form it can be thought of as low particle image density PIV, where low particle image density means that the concentration of tracers is sufficiently low that the maximum displacement of a particle in some time, Δt, will always be less than the mean nearest neighbour spacing between particle images. In PTV's early incarnations particle images were manually digitized and tracked, but with the aid of computers the process is now automated. However, due to the difficulty in determining a particle's pair when the particle's mean spacing is on the order of, or less than the maximum expected displacement, the technique has remained, until the 1990's[22,50,53,54], a low seeding density technique. That decade saw multiple researchers working with a variety of computer algorithms for particle pairing – the process of identifying which particle image, in any particular image, resulted from the identical particle that produced a particle image in one or more other images (singly exposed multiple images) or a different particle image in the same image (multiply exposed single images). During the 1990's the restriction on displacements relative to the mean particle separation distance was relaxed. The specific seeding density limit depends on the physical parameters of the experiment and the procedure used for particle pairing, and ultimately Cowen and Monismith[22] demonstrated that the limit on seeding density was essentially set only by the requirement of having the majority of the particles produce non-overlapping images.

While the low seeding density restriction on PTV was an early criticism of the technique, it was overcome. A second significant criticism of PTV is the resultant randomly distributed velocity vectors, which is still an important area of research. The traditional solution to dealing with randomly located PTV data is to interpolate it onto a regular grid. This interpola-

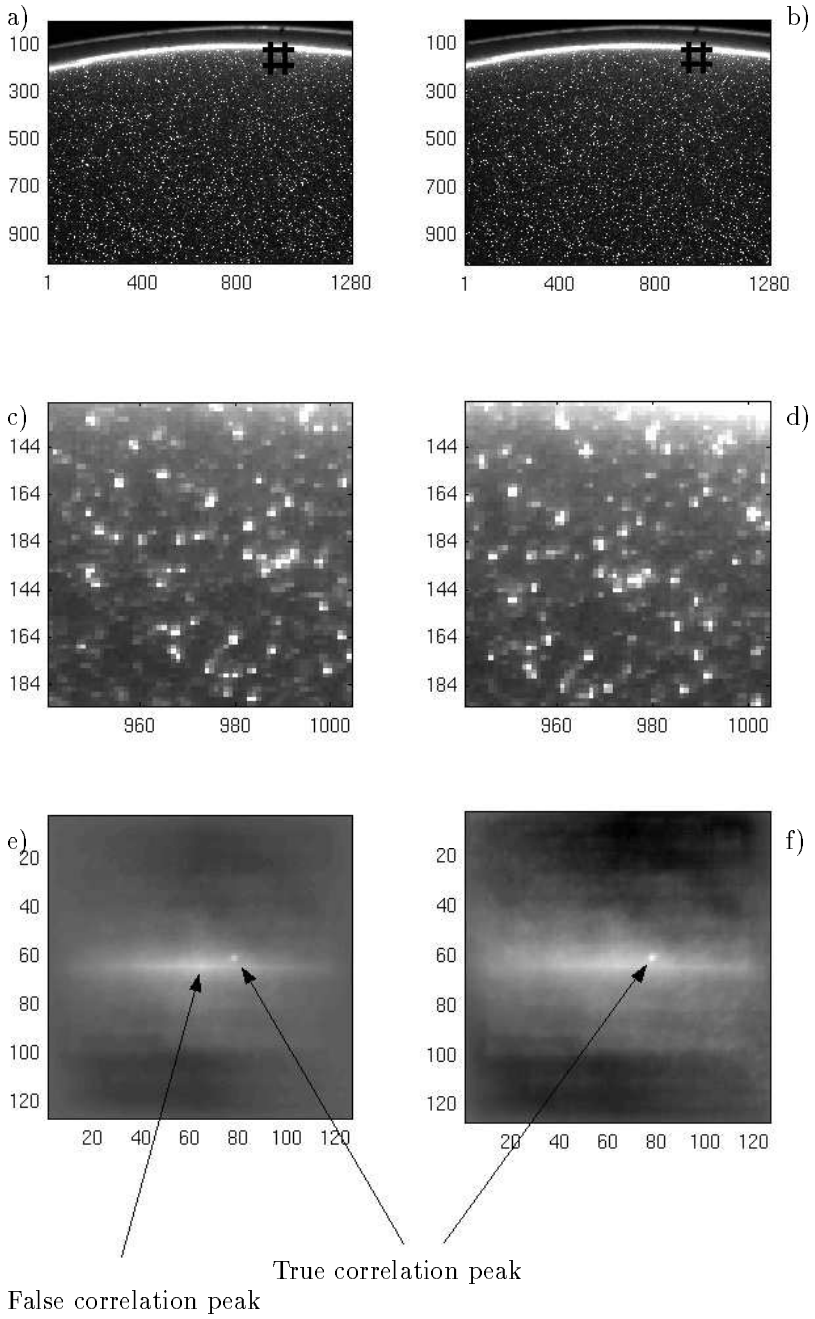

Fig. 11. a)–b) original images, black squares indicate the position of subwindows shown in c)–d). e) FFT-based correlation plane. f) Direct calculation of cross correlation.

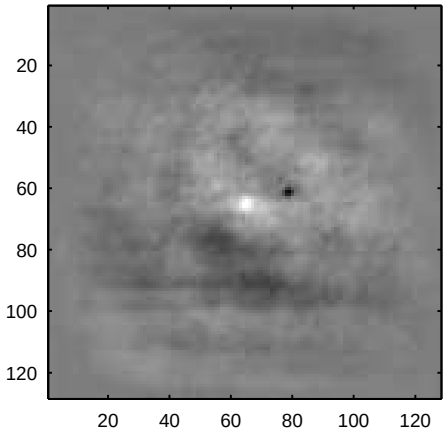

Fig. 12. Figure showing result of MQD-interrogation of the subwindows in Fig. 11 c)–d). Note that the displacement is found as the minimum value (black color) in this plane.

tion has an associated error, that while not a major problem relative to the mean velocity field, is problematic when considering the variance and the co-variance as well as the calculation of any velocity gradients. Several research efforts at determining optimal interpolation methods exist (e.g., Agui and Jiminez[57], Spedding and Rignot[58]) and Cowen and Monismith[22] demonstrate the success of a zeroth order interpolation scheme – simple data binning over small spatial regions.

While the above criticism is generally restricted to PTV, the reality is that the majority of PIV techniques produce quasi-randomly distributed velocity vectors due to the perturbation to the position of the PIV interrogation grid as a result of using first-order displacement differences to estimate the velocity. As discussed, Wereley and Meinhart[26] have resolved this issue through the use of second-order accurate displacement differences that are centered on the interrogation grid node. It should be noted that in the limit of small subwindow size PIV converges to PTV which points to a remaining perturbation to the location of the second-order accurate displacement differences proposed by Wereley and Meinhart – the true location of the estimated average displacement between two subwindows must account for the location of particles in each subwindow (e.g., a particle weighting of the information in each subwindow) and as the subwindow gets small the probability that the center-of-mass of the particle location information is not at the center of the subwindow, as implicitly assumed in all subwindow

techniques, increases.

The fundamental PTV approach to the determination of the velocity is to measure the displacement of an individual particle, located at a random point in space and time, either in doubly-exposed single images or successive singly-exposed images. The velocity is calculated as:

$$\mathbf{u}(x, y, t) = \frac{\mathbf{\Delta x}(x, y, t)}{\Delta t} \tag{25}$$

Clearly in the limit, when $\Delta t \to 0$, we recover the fundamental definition of velocity. However, we must keep in mind that we are approximating the Eulerian velocity as the second-order difference (or in some cases even the first-order difference) of the Lagrangian track of a particle. In fact, it is important to remember that PTV essentially is a Lagrangian measurement technique. Dalziel[53,54] developed an effective and relatively successful particle tracking algorithm, which was commercialized through the code **Dig***Image* and applied in many publications[59,60,61]. In the **Dig***Image* code, for example, the Lagrangian velocity is calculated by a least squares fit to more than two points in the particle path. By fitting a quadratic function we can obtain the Lagrangian acceleration directly, a concept also utilized by Chang *et al.*[62] to obtain acceleration from a single-camera configuration.

5.1. *Particle Detection Algorithms*

All PTV based techniques employ some form of particle detection algorithm. This usually involves thresholding the image, leading to a binary image. The threshold may be set manually or determined automatically. In either case the ideal image will have a histogram that looks like the one shown in Fig. 13. If in fact two separate peaks exist (one for the background and one for the particles) it is a simple process to pick a threshold value between the two peaks. However, as discussed shortly, it is more likely that the histogram will decay monotonically (a negative exponential probability density function), making the threshold choice somewhat unclear. In the latter case a user defined threshold is often the simplest solution.

Preconditioning of the images may, generally speaking, serve to enhance the particle images from the background. Such preprocessing may be accomplished by convolving the images with a top-hat function, which will serve to amplify the particle images. Subsequently thresholding is applied, followed by binarization.

Once an image has been binarized, particles are generally defined in one of two ways, either as horizontally and vertically adjacent "hot" pixels or as horizontally, vertically and diagonally adjacent "hot" pixels.

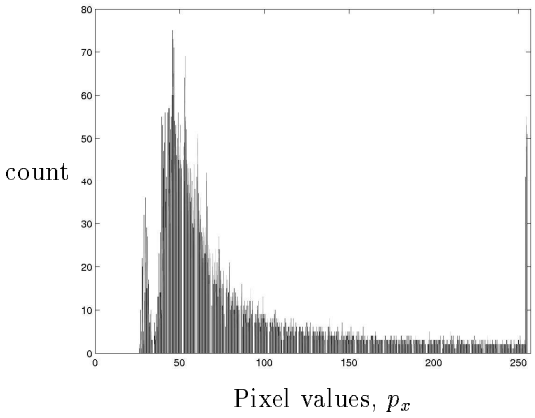

Fig. 13. Histogram of image intensities. Notice the small peak between $20 < p_x < 40$.

Particle tracking may occur over multiple exposed or sequential singly pulsed images. Multiply exposed particle image based techniques can lead to higher accuracy since the probability of incorrectly drawing a vector decreases with each additional exposure. However, these techniques suffer from lower vector density as the probability of tracers remaining in an illuminated region is reduced the longer they are tracked. Multiply exposed image based techniques can also be used to make higher order estimates of the velocity field (by retaining higher order terms in the Taylor Series expansion implied by equation 25) and the estimate of the acceleration field, which will be discussed further in section 6.4.

5.2. *Locating a Particle's Center*

Spherical particles tend to scatter light with an intensity pattern given by a rotated Airy wave. The central lobe of an Airy wave is well approximated as Gaussian. As discussed in section 4.4 as long as the particle image diameter is small (2 - 4 pixels) then a three-point Gaussian sub-pixel estimate of the particle's location is optimal. Unlike correlation peaks, however, we have no guarantee that the particle image diameter will be greater than one pixel. It is important in implementing the image acquisition and particle seeding that proper choices are made to insure that particle images will have diameters greater than one pixel and preferably less than four pixels[22]. Further, the same considerations about peak-locking as discussed in section 4.5 exist. In many cases, however, particle images may overlap (or nearly so) and in these cases the Gaussian fit may not work as well. In this case

the particle image intensity centroid may be used as an estimate of the particle's position.

5.3. *Particle Pairing*

After the individual particles are located at two different time steps, usually the next step will be to perform a pairing or matching of them. That is, we want to know which particle image in the second frame results from a particular particle that was imaged in the first frame.

Hybrid PTV techniques have been developed that employ PIV based algorithms as predictors of the velocity field allowing particle pairing searches to be conducted in small search regions relative to the region to where a particle might have moved[22,50]. This predictor approach overcomes the low seeding density limitation on PTV allowing particles to be tracked in flows where the mean spacing between particles is considerably less than the expected maximum displacements. Ultimately it must be remembered that these PIV predictor techniques suffer from the limitations of the PIV algorithms used to develop the predictive velocity field.

The algorithm in **Dig***Image*[53,54] is based on Operations Research. Particle pairing is accomplished via the use of a cost function. This cost function is dependent on individual particle factors, such as size, intensity, velocity history and shape. Particles are paired after a minimization process, where the cost of pairing particles in one frame with the particles in the next is minimized. The algorithm in **Dig***Image* is similar to the so called transportation algorithm, although not quite the same since particles are allowed to leave and re-enter the light-sheet.

There are many other algorithms that have been proposed for particle tracking that include fuzzy logic[55] and simulated annealing[56]. An intriguing capability of several of these techniques is their ability to extract higher-order flow quantities, such as the local strain rate and rotation rate, directly from the analysis.

6. Higher Order Measurements from Velocity Fields Obtained by QI-techniques

In a great number of applications of QI-techniques the velocity field information is only part of a complete description of the physical processes under study. For example many flows of interest are turbulent and derivative quantities of the velocity field may be of fundamental importance to describe the turbulence. In both turbulent and viscous flows the vorticity field is often an important descriptor of the flow physics. In stratified flows

it may be necessary to simultaneously measure the density field. In wave-structure interaction flows knowledge of the acceleration or pressure field may be of vital interest.

In the present chapter we will turn our attention to the use of the velocity fields, obtained from QI-imaging techniques, to determine higher order quantities. Such quantities may include shear, strain, vorticity, streamlines, acceleration, pressure and estimates of turbulence characteristics (e.g., the Reynolds stresses, turbulence dissipation, and spectra).

6.1. *Vorticity, Strain Rate and Divergence*

The calculation of vorticity, strain rate and divergence is often referred to as the calculation of differential quantities. These quantities are defined as

$$\omega = \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y}, \tag{26}$$

$$s = \frac{1}{2}\left(\frac{\partial v}{\partial x} + \frac{\partial u}{\partial y}\right), \tag{27}$$

$$-\frac{\partial w}{\partial z} = \frac{\partial v}{\partial y} + \frac{\partial u}{\partial x}, \tag{28}$$

where ω, s and $-\partial w/\partial z$ denote the (in-plane) vorticity, strain rate and two-dimensional divergence, respectively.

The differentiation of the velocity is a straightforward **calculation**, using for example a second-order accurate central difference numerical scheme. However, QI velocity measurements contain errors which can lead to strong noise in the differential quantity. Furthermore the data is discrete, which also may introduce noise. In these cases many scientists apply higher order numerical schemes[63], others apply smoothing[64] functions to deal with this noise, and Cowen and Monismith[22] argue for the careful choice of the differentiation length scale to maximize resolution and minimize noise amplification. Readers are advised to consult Raffel *et al.*[19], Nogueira *et al.*[38] and Lourenço and Krothapalli[65] for further details. We also note that it is possible to measure differential properties directly from the images using image distortion techniques[66,67] and other non-correlation based displacement extraction techniques[56].

6.2. *Streamlines and Pressure Estimation*

Defining the streamfunction, ψ, can be very useful in two-dimensional applications. The streamfunction is constant along a streamline, where the streamline is tangent to the velocity field at every point.

Consider a two-dimensional incompressible flow, thus $\nabla \cdot \mathbf{v} = 0$. The streamfunction can be obtained from $d\psi = u\,dy - v\,dx$ where $u = \partial\psi/\partial y$ and $v = -\partial\psi/\partial x$. This implies that the unknown scalar function ψ at the location P is defined as

$$\psi_P = \psi_O + \int_O^P (udy - vdx), \tag{29}$$

where ψ_O is the streamfunction at the origin O (typically set to zero[68]).

The path of the numerical integration can be freely chosen, but in order to minimize noise accumulation, the integrations are often performed on a staggered grid starting from the center of the velocity field and proceeding out towards the edges[68].

In order to estimate the pressure gradient we have to know the acceleration field (see section 6.4). If we make the assumptions that the flow is two-dimensional, steady and incompressible, it is possible to estimate the pressure field from the 2D Navier-Stokes equations[68] or similar approaches[69] using the measured velocity field.

6.3. *Turbulence*

As many fluid flows of interest are turbulent it is not surprising that QI techniques have a development history that reflects a desire to extract information on the turbulent statistics and turbulent structures of the flows. The application of QI techniques to turbulent flows presents unique challenges due to the range in space and time scales, which leads not only to a large requisit dynamic range in velocity to characterize the flow, but also the potential for strong strain rates, which historically are problematic for traditional PIV type techniques (e.g.,[22]).

6.3.1. *Dissipation*

The determination of the dissipation of turbulence, ε, is an important but challenging experimental goal. In theory it should be particularly straightforward using QI techniques based on its definition

$$\epsilon = 2\nu\langle s_{ij}s_{ij}\rangle, \tag{30}$$

where s_{ij} is the fluctuating strain rate tensor defined as

$$s_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right). \tag{31}$$

Hence, the in-plane fluctuating strain rate components can be directly calculated by second-order central differencing and the out-of-plane fluctuating strain rate components can be modeled by an understanding of the anisotropy ratio to the in-plane components[22]. However, presumably due to the noise level present in many velocity field calculations, relatively few researchers have reported dissipation results. Cowen and Monismith[22] report a dynamic algorithm that adjusts the length scale over which the velocity data is differentiated, arguing that the signal-to-noise ratio can be maximized by optimally choosing the length scale of the gradient calculation. Doron *et al.*[74] discuss optimal assumptions about the anisotropy of the out-of-plane gradients impact on the total dissipation. Cowen *et al.*[71] demonstrate several impacts of understating the dissipation structure in the swash zone. They show that the dissipation is in balance with the turbulent kinetic energy and that the turbulence is decaying as free turbulence during the swash uprush. They use the dissipation to estimate the friction velocity and the friction coefficient of the bed.

The single point measurement fields with sufficient temporal resolution to capture temporal frequency based spectra (whose QI-based calculation is described in the next section) have a well developed history of using the inertial subrange spectral energy measurements to estimate the turbulence dissipation. For temporally resolved QI measurements this is another approach to estimating the dissipation. For an example see Liao and Cowen[75] and for a discussion of the errors with respect to direct measurements see Doron *et al.*[74].

6.3.2. *Temporal and Spatial Spectra*

In unsteady flows in general and turbulent flows in particular, it is often of interest to look at flow statistics in the spectral domain. For single-point measurements this is only possible in the temporal frequency based sense. However, a significant advantage of QI techniques is the ability to determine instantaneous spatial information and hence spatial frequency based spectra. With typical QI hardware, say 30 frame per second cameras with 1024×1024 pixel imagers, we often find ourselves in the situation where the high spectral frequency content can only be determined in the spatial domain while the low spectral frequency content must be determined in the temporal domain. Typically a characteristic advective velocity is used to make the Galilean transformation from time to space and the temporal and spatial spectra are combined. In turbulent flows this transformation is known as Taylor's frozen turbulence hypothesis[70] and the characteristic

velocity scale is simply the mean velocity. However, in turbulent flows one must be careful as Taylor's hypothesis requires $(\overline{u'^2})^{1/2} \ll \overline{u}$ and it is not uncommon for this requirement to be violated. Further, for unsteady turbulent flows such as wavy flows, the choice of advective velocity scale is not obvious. Cowen *et al.*[71], working in the swash zone, found that the rms orbital velocity scale can be a reasonable choice of advective velocity scale but not at all wave phases.

6.3.3. *Turbulence versus Wave-Induced Variance*

An important question in considering turbulent flows by any measurment technique is the relative effects of flow unsteadiness and turbulence. For periodic flows in general and wavy flows in particular we often decompose the velocity field as

$$\mathbf{u}(\mathbf{x},t) = \overline{\mathbf{u}}(\mathbf{x}) + \tilde{\mathbf{u}}(\mathbf{x},t) + \mathbf{u}'(\mathbf{x},t) \tag{32}$$

where $\mathbf{u}$ is the instantaneous velocity vector at spatial location $\mathbf{x}$ and time t, $\overline{\mathbf{u}}$ is the temporal mean velocity vector, $\tilde{\mathbf{u}}$ is the phase averaged velocity vector, and $\mathbf{u}'$ is the turbulent velocity perturbation vector. The turbulent perturbations are generally found by subtracting the temporal mean and phase averaged velocity vectors from the instantaneous velocity vector and hence all perturbations with respect to the temporal mean and phase averaged quantities are passed through to $\mathbf{u}'$.

From energy considerations gravity waves are known to propagate as groups and often contain reflection components. Hence at minimum one expects amplitude variations even in a monochromatic wave field. These amplitude variations, even if small, can lead to perturbations from the mean and phase averaged velocity fields that are on the order of the turbulent perturbations, yet are not the result of turbulence (and do not effect the flow in the same manner as turbulence, e.g., they do not enhance the mixing). To account for the amplitude and phase induced variability we can apply a wave-turbulence decomposition technique, originally proposed by Thais and Magnudet[72] for single-point temporal data, to spatially resolved velocity fields. A second-order Stokes wave can be fit to the velocity field in a least squares sense where the wave phase, wave length, and wave amplitude can all be determined by the least squares fit. The velocity field decomposition can be written

$$\mathbf{u}(\mathbf{x},t) = \overline{\mathbf{u}}(\mathbf{x}) + \tilde{\mathbf{u}}_P(\mathbf{x},t) + \tilde{\mathbf{u}}_R(\mathbf{x},t) + \mathbf{u}'(\mathbf{x},t), \tag{33}$$

where $\tilde{\mathbf{u}}_P$ is the irrotational phased-dependent wave-induced velocity vector, $\tilde{\mathbf{u}}_R$ is the rotational phase-dependent wave-induced velocity vector, and

$\mathbf{u}'$ is the turbulent velocity perturbation vector. Importantly, $\tilde{\mathbf{u}}_P$ is used to separate the effects of amplitude and phase-induced variability in the wave-induced motions from the turbulence.

An advantage of the decomposition is that it works independently of the choice of coordinate system. For example, a free-surface following coordinate system may be utilized to explore the free surface boundary layer structure (e.g., Cowen[73]), or an Eulerian grid may be used, referenced to the bed, to look at the ensemble averaged phase dependent turbulence beneath the wave.

6.4. *Acceleration*

Measurements of acceleration have not received the same attention in the literature as PIV itself. These measurements are important, for example, to determine pressure which is needed for calculating forces on structures that are subject to fluid forces. The acceleration of a fluid particle is defined as,

$$\mathbf{a} = \frac{\partial \mathbf{u}}{\partial t} + \left(u\frac{\partial \mathbf{u}}{\partial x} + v\frac{\partial \mathbf{u}}{\partial y} \right) = \frac{D\,\mathbf{u}}{Dt}, \tag{34}$$

where D/Dt is the material derivative and consists of a temporally unsteady and an advective spatially unsteady term.

The acceleration field arising from the temporal unsteadiness of the flow may be calculated from the difference between two consecutive velocity fields, as produced by conventional PIV,

$$\mathbf{a} = \frac{\mathbf{u}_2 - \mathbf{u}_1}{t_2 - t_1}.$$

Here t_1 and t_2 denote the time of measurement of the first and second velocity fields, respectively. However, to calculate the material derivative of the velocity, in addition to the temporal gradients, we also need to evaluate the spatial derivatives. In each calculation small differences in velocity, either in space or time, are normalized by small differences in either position or time, and may produce large errors.

To enable proper control and accuracy of acceleration measurements typically between 2 and 4 consecutive images are required, depending on the choice of PIV-evaluation method (auto or cross-correlation, or a combination of the two). Furthermore we need to accurately control the timing of each exposure.

Jensen *et al.*[2] reported on acceleration measurements in periodic surface waves in a wave tank. They measured the temporal derivative, $(\mathbf{u}(\mathbf{x}, t_2) - \mathbf{u}(\mathbf{x}, t_1))/(t_2 - t_1)$ and used a stereo CCD-camera setup where each camera

takes two single-exposed images. The exposure was controlled by the use of an Acousto-Optical Modulator. In this case careful camera alignment is crucial to avoid introduction of additional errors. Jensen *et al.*[2] reported that their camera calibration typically introduced errors on the order of $1.9 - 2.6\%$ compared to the displacement. Measurements were performed on accelerations due to water waves in a laboratory wave tank and the results were compared to the theoretical description of Stokes waves. In the best case they found a relative standard deviation of about 0.6% for the velocities and about 2% for the accelerations.

Jensen *et al.*[76] extended the acceleration measurements to also include the advective terms. They calculated the material derivative by tracing pseudo particles through a number of iterations. Their method thus uses an interpolated velocity in the second velocity field. They report the need for smoothing of the velocity fields to reduce the effects of spatial noise. Measurements were performed on runup of strongly non-linear waves on a steep beach. Additional results using this method are also published in Jensen *et al.*[77]

Jacobsen *et al.*[78] reported on a 4-CCD-camera looking at the same field of view through the use of beamsplitters. The temporal component of the acceleration in standing waves and trains of irregular waves was measured, with errors being on the order of $20 - 100\%$.

Other investigations[79,80] use similar two CCD-camera setups, employing off-the-shelf PIV cameras that image approximately the same field of view.

Dong *et al.*[79] report a two camera system for measuring the temporal component of accelerations in a rotating flow using pairs of doubly exposed images. They report a relative uncertainty of less than 25%. A similar approach was used by Chang and Liu[81]

Christensen and Adrian[80] extended the acceleration calculation to also include the advective terms. They report that these terms are an order of magnitude smaller than the temporal part in their application, which is acceleration measurements in a turbulent channel flow. They calculate the convective derivative by shifting the evaluation region in the second velocity field by the horizontally averaged mean flow. We note that the work by Christensen and Adrian uses polarization of the laser illumination to avoid a multiple exposure on the second image at the first camera. Such a double exposure was a feature found in the system of Jensen *et al.*[2]

Chang *et al.*[62] proposed an interesting 2D PTV based single-camera technique where they used one doubly exposed and one singly exposed image and PTV to calculate the substantial derivative in the Lagrangian reference frame.

La Porta *et al.*[82] measured fluid particle accelerations in fully developed turbulence, using a 3D PTV method[83] with relatively advanced experimental techniques (recall from section 5 that PTV may give Lagrangian accelerations directly).

7. Conclusion

In this work we have aimed at reviewing the foundations and fundamentals of Quantitative Imaging (QI) techniques in general with a particular focus on Particle Image Velocimetry (PIV) and specific details of its application to wavy flows. The QI techniques have been presented in the context of general pattern matching techniques with the aim of pointing out similarities and differences between the relatively large number of techniques available. With this in mind we conclude by tieing the specifics of QI techniques in fluid flows thus far presented to the more general field of pattern matching techniques. The thoughts below are only a minimal scratch at the surface of the variety of pattern matching techniques in use in quantitative research fields but we hope they give the reader a sense of the range of applications of these techniques.

7.1. *Measurements of Deformation*

The use of PIV techniques has received increasing attention within the field of solid mechanics where it is applied to measure deformation of various substances such as clay, soil or sand subject to pressure or forces[84,85]. In these cases the addition of tracer particles is not as common as it is in fluid mechanics. Sand or soil, for example, will often provide a pattern suitable for matching, whereas a substance like clay can be (very gently) spray-painted on the clay surface to create a suitable pattern.

7.2. *Measurements of Density*

A light-ray that passes through a fluid may experience deflection due to changes in refractive index in the fluid. The method known as Schlieren[11], is a well known quantitative method to determine gradients in a density field. By using a background pattern, any refractive index change will displace this pattern with a distance proportional to the density gradient. This is basically the idea behind the image processing versions of Schlieren, which are known as Synthetic Schlieren[86,87] and Background Orientated Schlieren (BOS)[88,89]. The latter method is based on cross-correlation using FFT (i.e. equation 1) to locate the displacement, while the former method

primarily calculates the full correlation coefficient (equation 23) but also the absolute difference, i.e. $R(s,t) = |F'(i,j) - F''(i+s, j+t)|$, and the minimum quadratic difference methods are used. In the sense that both methods rely on principles of pattern matching, we can say that they are identical.

We should also emphasize that subpixel interpolation of the subwindows were implemented in these techniques. In fact, Dalziel *et al.*[87] applied bilinear subpixel interpolations of their images, similar to that applied by Gui and Wereley[35] in the context of PIV.

7.3. *Measurements of Volume*

In a vast number of industrial and scientific applications it is desirable to measure the volume of an object. As an example we can consider the inspection of manufactured parts where a lower limit tolerance exists with respect to accuracy. Manual inspection is in most cases impossible and the use of various 3D scanning systems are currently in use[90]. There are several, slightly different methods available, some of which are known as structured light or (computer) tomography. Applications include scanning of the human body[91], automatic inspection of car parts[92] and fluid surface measurements[93]. The concept behind this class of methods is basically to illuminate a 3D object with a pattern. The pattern will thus be displaced compared to the scene without the object present. The displacement is directly proportional to the height of the object in the plane we are viewing. By rotating the object or using more than one camera, the 3-dimensional shape may be reconstructed.

Acknowledgements

J. Kristian Sveen gratefully acknowledges the financial support provided by the Research Council of Norway through the Strategic University Program "General Analysis of Realistic Ocean Waves" (GROW). Edwin A. Cowen gratefully acknowledges the financial support of the U.S. Office of Naval Research (grant No. N00014-98-1-0774, Dr. Keith Ward Program Officer) and the U.S. National Science Foundation (Grant No. CTS-0093794). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Office of Naval Research or the National Science Foundation.

References

1. J. K. Sveen, http://www.math.uio.no/~jks/matpiv/, (1998-2003).

2. A. Jensen, J. K. Sveen, J. Grue, J.-B. Richon and C. Gray, *Exp. Fluids* **30**-5, 500 (2001).
3. H. Meng and F. Hussain, *Fluid Dyn. Res.* **8**, 33 (1991).
4. A. K. Prasad and R. J. Adrian, *Exp. Fluids* **15**, 49 (1993).
5. C. E. Willert and M. Gharib, *Exp. Fluids* **12**, 353 (1992).
6. C. Brücker, *Exp. Fluids* **19**, 255 (1995).
7. R. J. Adrian, TAM report, in cooperation with TSI, St. Paul, MN, (1996).
8. L. Prandtl and O. G. Tietjens, *Applied Hydro and Aeromechanics*, (Dover Publ. 1934).
9. G. W. Faris and R. L. Byer, *Science* **238**, 1700 (1987).
10. P. J. W. Roberts, X. Tian and Y. Jung, *Hydraulic Measurements and Experimental Methods 2002 - Proceedings of the Specialty Conference*, EWRI and IAHR, July 30 - August 2, 2002, Estes Park, CO, (2002)
11. W. Merzkirch, *Flow Visualization*, (Academic Press, 1987).
12. R. C. Gonzales and R. E. Woods, *Digital Image Processing* (Addison Wesley 1992).
13. http://mpeg.telecomitalialab.com/
14. H. Foroosh and S. Hoge, in *Video registration*, eds. M. Shah and R. Kumar, (Kluwer Academic, 2003).
15. L. Hesselink, *Ann. Rev. Fluid Mech.* **20**, 421 (1988).
16. R. J. Adrian, *Ann. Rev. Fluid Mech.* **23**, 261 (1991).
17. J. Westerweel, Ph.D. Dissertation, University Press, Delft, (1993).
18. J. Westerweel, *Meas. Sci. Tech.* **8**, 1379 (1997).
19. M. Raffel, C. E. Willert and J. Kompenhans, *Particle Image Velocimetry*, (Springer, Berlin, 1998).
20. R. J. Adrian, *Appl. Optics* **25**, 3855 (1986).
21. R. D. Keane and R. J. Adrian, *Appl. Sci. Res.* **6**, 754 (1992).
22. E. A. Cowen and S. G. Monismith, *Exp. Fluids* **22**, 199 (1997).
23. A. M. Fincham and G. R. Spedding, *Exp. Fluids* **23**, 449 (1997).
24. R. J. Adrian, in *Laser Anemometry in Fluid Mechanics - III*, eds. R. J. Adrian *et al.*, 115 (1988).
25. J. Westerweel, D. Dabiri and M. Gharib, *Exp. Fluids* **23**, 20 (1997).
26. S. T. Wereley and C. D. Meinhart, *Exp. Fluids* **31**, 258 (2001).
27. H. Huang, H. Fiedler and J. Wang, *Exp. Fluids* **15**, 168 (1993).
28. H. Huang, H. Fiedler and J. Wang, *Exp. Fluids* **15**, 263 (1993).
29. A. Fincham and G. Delerce *Exp. Fluids* **29**, 13 (2000).
30. F. Scarano and M. L. Riethmuller, *Exp. Fluids* **26**, 513 (1999).
31. F. Scarano and M. L. Riethmuller, *Exp. Fluids* **Suppl.**, S51, (2000).
32. A. K. Prasad, R. J. Adrian, C. C. Landreth and P. W. Offutt, *Exp. Fluids* **13**, 105 (1992).
33. J. W. Goodman, *Laser Speckle and Related Phenomena*, ed. J. C. Dainty, (Berlin, Springer, 1984).
34. E. Parzen, *Ann. Math. Statist.* **28**, 921 (1957).
35. L. Gui and S. T. Wereley, *Exp. Fluids* **32**, 506 (2002).
36. Q. Liao and E. A. Cowen, *Exp. Fluids* submitted, (2004).
37. J. Westerweel *Exp. Fluids* **16**, 236 (1994).
38. J. Nogueira, A. Lecuona and P. A. Rodriguez, *Meas. Sci. Technol.* **8**, 1493

(1997).
39. R. D. Keane and R. J. Adrian, *Meas. Sci. Technol.* **1**, 1202 (1990).
40. D. P. Hart, *Exp. Fluids* **29**, 13 (2000).
41. L. Gui and W. Merzkirch, *Exp. Fluids* **21**, 465 (1996).
42. L. Gui and W. Merzkirch, *Exp. Fluids* **28**, 36 (2000).
43. J. P. Lewis, Proceedings of *Vision Interface*, 120 (1995).
44. M. Frigo and S. G. Johnson, **ICASSP conf. proc.** *vol. 3*, 1381 (1998).
45. C. Willert and M. Gharib, *Exp. Fluids* **10**, 181 (1991).
46. P. T. Tokumaru and P. E. Dimotakis, *Exp. Fluids* **19**-1, 1 (1995).
47. H. J. Lin and M. Perlin, *Exp. Fluids* **25**, 431 (1998).
48. J. Nogueira, A. Lecuona and P. A. Rodrguez, *Meas. Sci. Technol.* **12**, 1911 (2001).
49. G. M. Quenot, J. Pakleza and T. A. Kowalewski, *Exp. Fluids* **25**-3, 177 (1998).
50. R. D. Keane, R. J. Adrian and Y. Zhang, *Meas. Sci. Technol.* **6**, 754 (1995).
51. C. Kuglin and D. Hines, Proc. 1975 Int'l Conf. on Cybernetics and Society, 163 (1975).
52. A. Beghdadi, M. Meshbah and J. Monteil, *Image Vision Comp.* **21**, 383 (2003).
53. S. B. Dalziel, *Appl. Scien. Res.* **49**, 217 (1992).
54. S. B. Dalziel, *Dyn. Atmos. Oceans* **20** 127 (1993).
55. M. P. Wernet, *Exp. Fluids* **30**, 434 (2001).
56. M Stellmacher and K Obermayer, *Exp. Fluids* **28**, 506 (2000).
57. J. C. Agui and J. Jimenez, *J. Fluid Mech.* **185**, 447 (1987).
58. G. R. Spedding and E. J. M. Rignot, *Exp. Fluids* **31**, 417 (2001).
59. J. Grue, A. Jensen, P-O. Rusås and J. K. Sveen, *J. Fluid Mech.* **380**, 257 (1999).
60. R. E. Hewitt, P. A. Davies, P. W. Duck and M. R. Foster, *J. Fluid Mech.* **389**, 169 (1999).
61. V. Maderich, G. J. van Heijst and A. Brandt, *J. Fluid Mech.* **432**, 285 (2001).
62. K.-A.Chang, E. A. Cowen and P. F.-L. Liu, *3rd International Symposium on Particle Image Velocimetry, Santa Barbara, USA, 16-18 Sept.* (1999).
63. J. Grue, A. Jensen, P-O. Rusås and J.K. Sveen, *J. Fluid Mech.* **413**, 181 (2000).
64. J. D. Luff, T. Drouillard, A. M. Rompage, M. A. Linne and J. R. Hertzberg, *Exp. Fluids* **26**, 36 (1999).
65. L. Lourenço and A. Krothapalli, *Exp. Fluids* **18**, 421 (1995).
66. X. Ruan, X. Song and F. Yamamoto, *Exp. Fluids* **30**-6, 696 (2001).
67. S. Mayer, *Exp. Fluids* **33**-3, 443 (2002).
68. K. Imaichi and K. Ohmi, *J. Fluid. Mech.* **129**, 283 (1983).
69. R. Gurka, A. Liberzon, D. Hefetz, D. Rubinstein and U. Shavit, Proceedings to the 3rd International Workshop on Particle Image Velocimetry, Santa Barbara, California, September 16-18, (1999).
70. G. I. Taylor, *Proc. Roy. Soc. Lond. (A)* **164**, 476 (1938).
71. E. A. Cowen, I. M. Sou, P. L.-F. Liu and B. Raubenheimer, *J. Eng. Mech.* in press, (2003).

72. L. Thais and J. Magnaudet, *J. Geophys. Res.* **100**-1, 741 (1995).
73. E. A. Cowen, Ph.D. Dissertation, Dept. Civil Engineering, Stanford Univ. (1996).
74. P. Doron and L. Bertuccioli and J. Katz and T. R. Osborn, *J. Phys. Ocean.* **31**, 2108 (2001).
75. Q. Liao and E. A. Cowen, *Env. Fluid Mech.* **2**, 9 (2002).
76. A. Jensen, G. K. Pedersen and D. J. Wood, *J. Fluid Mech.* **486**, 161 (2003).
77. A. Jensen, M. Huseby, D. Clamond, G. K. Pedersen and J. Grue, This volume, Chapter 7.4, (2004).
78. M. L. Jacobsen, T. P. Dewhirst and C. A. Greated, *Meas. Sci. Technol.* **8**, 1502 (1997).
79. P. Dong, T.-Y. Hsu, P. Atsavapranee and T. Wei, *Exp. Fluids* **30**, 626 (2001).
80. K. T. Christensen, R. J. Adrian, *Exp. Fluids* **33**, 759 (2002).
81. K.-A.Chang and P. F.-L. Liu, *Phys. Fluids* **10**, 3390 (1998).
82. A. La Porta, G. A. Voth, A. M. Crawford, J. Alexander and E. Bodenschatz, *Nature* **409**, 1017 (2001).
83. M. Virant and T. Dracos, *Meas. Sci. Technol.* **8**, 1539 (1997).
84. D. J. White, W. A. Take and M. D. Bolton, *Geotechnique* **53**-7, 619 (2003).
85. H. Wolf, D. König and T. Triantafyllidis, *J. Struct. Geol.* **25**-8, 1229 (2003).
86. B. R. Sutherland, S. B. Dalziel, G. O. Hughes, and P. F. Linden, *J. Fluid Mech.* **390**, 93 (1999).
87. S. B. Dalziel, G. O. Hughes and B. R. Sutherland, *Exp. Fluids* **28**, 322 (2000).
88. G. E. A. Meier, *Deutsche Patentanmeldung*, DE 199 42 856 A1, (1999).
89. H. Richard and M Raffel, *Meas. Sci. Technol.* **12**, 1576 (2001).
90. F. Arman and J. K. Aggarwal, *ACM Comp. Surv.* **25**-1, 5 (1993).
91. H. Tapp and Z. Dunford, *New Scientist* **6**, 14 (2003).
92. F. Couweleers, Ø. Skotheim, H. Schulerud and K. Kaspersen, *Proc. of SPIE* **5144**, Optical Measurement Systems for Industrial Inspection III, 697 (2003).
93. T. Roesgen, A. Lang and M. Gharib, *Exp. Fluids* **25**, 126 (1998).