

# Performance of digital image velocimetry processing techniques

S. P. McKenna, W. R. McGillis

106

**Abstract** Digital particle image velocimetry (DPIV)-processing techniques have become increasingly more sophisticated in recent years. However, much work is still done using standard traditional methods of analysis. This paper investigates several traditionally based techniques for cross-correlation image processing in terms of computational efficiency and measurement accuracy. Direct spatial domain correlation, standard fast Fourier transform (FFT) correlation, a dynamic FFT correlation technique, and a new hybrid correlation method are discussed and evaluated. In addition, a particle-tracking velocimetry scheme based on that of Cowen and Monismith (1997) is examined in the same context as the DPIV methods. A detailed examination of the behaviors of each correlation method reveals that direct spatial domain correlation is more accurate than FFT-based methods, with the standard FFT correlation showing the weakest performance. Using the more robust methods (dynamic FFT and hybrid correlation), accuracy can be improved significantly over the standard FFT method in many cases, while still remaining computationally efficient. The particle-tracking algorithm studied was found to yield comparable accuracy to the DPIV routines and can provide much higher spatial- resolution possibilities.

## 1 Introduction

This paper examines the performance of a number of digital particle image and particle-tracking velocimetry

(DPIV and PTV)-processing techniques. Focus is restricted to the more common approaches to DPIV and PTV processing found in practice today. One reason for this particular focus is that many commercial image velocimetry packages available today are based upon traditional forms of analysis (e.g., fast Fourier transform techniques). Understanding the limitations and expected performance of these types of processing codes is necessary. A second reason for this focus is the relative simplicity of these methods and the lack of intensive algorithm development required for implementation.

In addition to focusing on basic processing schemes, this study also puts emphasis on practicality. Theoretical considerations are not used, but rather synthetic image flow fields are processed to determine expected algorithm performance in terms of accuracy and computational efficiency. For this investigation, no a priori knowledge of the flows is used to guide the processing, thereby providing a truly comprehensive evaluation of the processing methods. Finally, the material presented in this paper can be found in greater detail in McKenna (2000).

## 2 Processing digital image data for velocimetry

Traditional forms of DPIV-processing algorithms have relied on auto- or cross-correlation schemes. Cross-correlation implementations are typically favored over auto-correlation methods since the former yield directionally unambiguous displacements and have been found to provide superior aggregate performance, e.g., lower correlation noise, greater dynamic range, and less gradient bias (Keane and Adrian 1992). Early DPIV cross-correlation analyses were performed in the Fourier domain, as this provided the most efficient means of processing the data (e.g., Willert and Gharib 1991). However, with the continual development of faster CPUs, the need for the transformation into frequency space is becoming less imperative and some methods of correlation are performed 'brute force' in the spatial domain directly (Huang et al. 1993, Roesgen and Totaro 1995). Here, four cross-correlation algorithms are studied: (1) direct spatial domain correlation, (2) standard FFT (fast Fourier transform) correlation, (3) dynamic FFT correlation, and (4) a new hybrid dynamic FFT/direct spatial domain correlation. In addition, a particle-tracking technique that is based closely on the hybrid method of Cowen and Monismith (1997) is examined. Each scheme was investigated for accuracy and efficiency for a number of prescribed displacement fields.

Received: 3 September 1999 / Accepted: 21 June 2001  
Published online: 29 November 2001

S. P. McKenna (✉)  
Department of Ocean Engineering  
Massachusetts Institute of Technology  
Cambridge, MA 02319, USA  
e-mail: smckenna@whoi.edu

W. R. McGillis  
Department of Applied Ocean Physics and Engineering  
Woods Hole Oceanographic Institution  
Woods Hole, MA 02543, USA  
e-mail: wmcgillis@whoi.edu

This work was sponsored by the Andrew W. Mellon Foundation Endowed Fund for Innovative Research and the National Science Foundation grant no. OCE-9729972. During the course of this work, S. P. McKenna was supported as an Office of Naval Research Graduate Fellow and this assistance is gratefully acknowledged. Woods Hole Oceanographic Institution contribution no. 9962.

## 2.1

### Cross-correlation algorithms for DPIV

Extracting velocity information from a pair of single-exposed DPIV images involves subdividing the images (hereafter, A and B) into smaller subimages and quantifying the average spatial shift of particles from subimages in A to subimages in B. The spatial offset between A-B subimage pairs can be determined using the statistical technique of cross-correlation. In two dimensions, the discrete cross-correlation between two real-valued functions,  $f(x, y)$  and  $g(x, y)$ , is given by Gonzales and Wintz (1987) as

$$\Phi_{fg}(x, y) = f \circ g = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} f(m, n) g(x+m, y+n) \quad (1)$$

for  $x = 0, 1, \dots, M-1$  and  $y = 0, 1, \dots, N-1$ . Equation (1) can be used to express the correlation between subimages  $f$  and  $g$ , of dimensions  $A \times B$  and  $C \times D$ , respectively. In order to avoid wraparound errors, these samples should be extended (e.g., through zero-padding) and made periodic with periods  $M$  and  $N$  in the  $x$ - and  $y$ -directions such that  $M = A + C - 1$  and  $N = B + D - 1$ ; however, in practice, this precaution may or may not be taken. The cross-correlation can be computed in either the spatial domain directly or in the Fourier domain (with certain potential restrictions), the latter being preferred when computational efficiency is the primary goal. Before addressing issues of implementation, Eq. (1) should be slightly modified to yield a de-measured, variance-normalized cross-correlation coefficient:

$$C_{fg}(x, y) = \frac{[f(x, y) - \bar{f}] \circ [g(x, y) - \bar{g}]}{\sqrt{\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} [f(m, n) - \bar{f}]^2 [g(m, n) - \bar{g}]^2}} \quad (2)$$

The quantities  $\bar{f}$  and  $\bar{g}$  represent the respective means of subimages  $f$  and  $g$ . A variance-normalized coefficient is a more robust measure of correlation between two subimages since it accounts for intensity variation (both local and global) between images A and B. This measure of correlation has been shown to be more accurate than alternative non-normalized measures by Burt et al. (1982) and Huang et al. (1997). In each of the algorithms examined here, correlation peak locations are found to subpixel accuracy using two three-point Gaussian curve fits – one in the  $x$ -direction, one in the  $y$ -direction. Thus, in this investigation, Eq. (2) was used as the correlation measure, leaving only the implementation of the cross-correlation calculation (spatial domain versus frequency domain) and advanced subimage matching strategies as the differences between the DPIV algorithms considered.

#### 2.1.1

##### Direct spatial domain correlation (template matching)

Direct spatial domain correlation is perhaps the most intuitively straightforward manner in which the cross-correlation can be computed, in that it can be made to mimic a simple template-matching operation. However, this approach is also the most numerically intensive. Specifically,




Fig. 1. Schematic of direct spatial domain cross-correlation calculation. The shaded subimages ( $B_{ij}$ ) show a sampling of all possible image B subimages used to generate the complete correlation surface. The remainder A-B pairings are found by tiling the  $B_{ij}$  subimages horizontally and vertically within the search region as indicated by the arrows

it uses the coefficient given by Eq. (2) with a direct multiplication operation like that given by Eq. (1). In this work, the direct correlation has been implemented as depicted by Fig. 1. A subimage in A of size  $M \times N$  is identified and a search region in B of size  $R \times S$  is centered about the subimage in A. The subimage in A (the template) is correlated with each possible  $M \times N$  subimage in the search region of B, building up a correlation surface of size  $(R - M + 1) \times (S - N + 1)$ . In this way, issues of subimage extension and periodicity are avoided and the actual unique content from image B is utilized.

There are a number of advantages to direct correlation in the spatial domain. Most importantly, this approach minimizes the out-of-pattern (or, in-plane loss-of-pairs) effect (Keane and Adrian 1992, Huang et al. 1993) since, under ideal conditions, a properly defined search region allows for all of the particles in the subimage from A to be found in a subimage from B. This has been referred to as particle image pattern matching, or PIPM (Huang et al. 1993). By allowing for PIPM, direct correlation is a powerful method of cross-correlation, and this was found to be the case as will be shown in Sect. 3. A further advantage of the direct spatial domain approach is its flexibility. Unlike FFT-based schemes that can be restricted by certain Fourier transform algorithms (e.g., square, power-of-two, equal-size subimages), direct correlation has no limitations on subimage/search region shapes or sizes. This can be helpful when flows are quasi-unidirectional (e.g., channel flows). In such instances, the search region can be stretched in the flow direction and reduced in the cross-flow direction, reducing the number of unnecessary

computations. Eliminating unnecessary calculations is particularly desirable for the direct correlation since this method of correlation is an order  $N^4$  operation. To ameliorate this issue of computational load, possible speed enhancements have been implemented for direct spatial domain correlation schemes (Roesgen and Totaro 1995; the present hybrid technique).

### 2.1.2 FFT correlation

Because the direct spatial domain correlation operation is so computationally demanding, use of Fourier transform methods is popular in DPIV processing. In frequency space, the double summation in Eq. (1) can be replaced by a complex conjugate pointwise multiplication of the two-dimensional Fourier-transformed subimages. This is expressed by the correlation theorem (Gonzales and Wintz 1987),

$$f(x, y) \circ g(x, y) \Leftrightarrow F(\xi, \eta)G^*(\xi, \eta) , \quad (3)$$

where  $F$  denotes the Fourier transform of the function  $f$ , and  $G^*$  represents the complex conjugate of the Fourier transform of function  $g$ . In practice, the FFT implementation involves taking an  $N \times N$  subimage from image A and a co-located  $N \times N$  subimage in B, and generating the correlation surface in a single step using Eq. (3) with Eq. (2).

When working with an FFT calculation of the correlation, it must be understood that the correlation theorem assumes the samples being correlated are periodic. For DPIV, this means that the two raw subimages being correlated, which are not periodic, are considered periodic with periods equal to the subimage dimensions. This condition reveals under which circumstances the FFT computation of the correlation and the direct spatial domain calculation of the correlation will be equivalent. The two will be equal only when the direct correlation is executed using subimages that are made periodic with periods equal to the subimage dimensions. However, as Fig. 1 illustrates, the direct-correlation approach discussed here does not impose this periodicity, but instead uses the actual content of image B to calculate the correlation in a template-matching sense – the information is available and the direct correlation method exploits it, avoiding errors due to in-plane loss of pairs. In the case of the FFT approach, windowing methods like those used in spectral analysis to deal with the artifacts of non-periodic data are undesirable in DPIV since they can introduce systematic errors and/or degrade the correlation signal-to-noise ratio (Raffel et al. 1998). Furthermore, Westerweel (1993) has argued that the use of window functions is unnecessary in practice. The method of zero-padding the subimage data in order to remove wraparound errors is also a possible topic of debate. Wraparound errors emerge when the domain for the computed correlation is not large enough and part of the unresolved correlation is folded back (aliased) onto the correlation space. For a common  $N \times N$  subimage, this would typically require padding the sample with zeros to arrive at a sample of size  $M \times M$ , with  $M = 2N$ . Westerweel notes that in practice, for particle displacements less than about 1/4 of the subimage size ( $N/4$ ), the cross-

correlation effectively vanishes beyond displacement offsets of  $N/4$ . This being the case, the choice of  $M = N$  is acceptable. However, calculating the cross-correlation via FFTs without zero-padding can have consequences for larger displacements (Westerweel 1993) and for multiple-exposure images (Westerweel 1997). In the spirit of simplicity and based on the practical findings of Westerweel, neither zero-padding nor windowing is used in the FFT-based methods studied here.

The FFT correlation method, while very efficient, only being of the order of  $N \log N$  in complexity, can suffer dramatically from the out-of-pattern effect since the subimage in B typically will fail to contain all of the particles found in the subimage from A. As the mean particle displacement increases (e.g.,  $>1$  pixel) this effect becomes more pronounced and errors can become significant.

### 2.1.3 Dynamic FFT correlation (DFFT)

Westerweel et al. (1997) have shown that out-of-pattern effect difficulties can be removed by using an image-shifting technique. The authors find that by shifting the location of the subimage in image B by an amount equal to the integer part of the displacement (in pixels, px), the precision of the cross-correlation estimates can be enhanced. The improved precision stems from results showing that the root-mean-square (RMS) error in estimating displacements is proportional to the displacement itself for displacements less than 1/2 px (Adrian 1991, Willert and Gharib 1991). By shifting the location of the subimage in B by the integer value closest to the total displacement and performing the cross-correlation, the residual displacement will always be less than 1/2 px, thus yielding less variation in the measurement error. The Westerweel et al. (1997) approach was shown to improve the displacement estimate accuracy for cases where a universal mean flow component existed (simulated isotropic turbulence and experimental grid turbulence) and also for the more general case of a pipe flow where the shifting was implemented on a subimage-by-subimage basis using an initial correlation pass, followed by a second pass using individually offset subimages.

In their hybrid particle-tracking algorithm, Cowen and Monismith (1997) present a dynamic subimage-shifting procedure for FFT-based correlation methods. This strategy was adopted for the dynamic FFT (DFFT) algorithm used here. The scheme begins with a standard FFT pass through the entire image. This pass provides an initial estimated displacement vector for each node location. For each node, the  $(dx, dy)$  displacement estimates are rounded to their nearest integer values and the results are used as the amounts by which to dynamically shift the image B subimage in the next pass. This process is repeated iteratively for all nodes on a node-by-node basis, stopping when the rounded integer-pixel displacements for both components of the displacement are zero. Five dynamic iterations are performed in all. Scarano and Riethmuller (1999) report a similar DFFT method in which, along with dynamic shifting, iterative subimage size refinement takes place, allowing for increased spatial resolution.

### 2.1.4

#### Hybrid correlation

To achieve both high accuracy and efficiency, a hybrid cross-correlation scheme was developed in this work that makes use of the dynamic FFT approach and the direct spatial domain calculation of the correlation. It is a straightforward extension of both that exhibits the favorable accuracy found in the direct correlation calculation along with the superior computational efficiency of the FFT methods. The first phase of the hybrid scheme performs a full DFFT correlation operation for the image pair. The second phase further refines the cross-correlation with a limited direct spatial domain correlation operation. An example serves to illustrate the methodology. Consider a node located at  $(x, y)$  that has yielded a converged DFFT result of  $(-1.017 \text{ px}, 6.230 \text{ px})$  displacement using  $N \times N$  subimages. The next step locates an  $(N + 6) \times (N + 6)$  search region in image B centered at coordinates  $(x - 1, y + 6)$ . This yields a  $7 \times 7$  directly computed correlation surface that is dominated by the correlation peak. The choice of a  $7 \times 7$  correlation surface is somewhat arbitrary, but it allows for refinements to the DFFT displacement up to 2.5 px based on the operation of the present subpixel peak locator, with an acceptable increase in computation time. This refinement pass is done once for all nodes in the image.

### 2.2

#### Particle tracking

Spatial resolution is a limitation in all correlation-based DPIV approaches. Flow scales smaller than the subimage dimensions cannot be captured since the result from an auto- or cross-correlation operation represents the mean displacement of all particles within the subimage region. Therefore, subimage size plays an important role in determining the achievable spatial resolution of the flow. Reducing the size of the subimage increases the resolution, but there is a limit on how small the subimage can be based on the statistical requirement for an adequate number of particle images, or unique image texture, in the subimage. A further limitation of typical correlation procedures, which also derives from the integrating nature of the correlation operation, is their difficulty in handling velocity gradients. The main problem encountered when velocity gradients are present, particularly for standard static FFT correlation schemes, is gradient biasing. For instance, when an in-plane gradient exists, faster moving particles leave the subimage region and the remaining, slower moving particles act to bias the displacement estimate toward slower velocities. Using dynamic correlation methods like the ones described here reduces this error. However, even with dynamic correlations that achieve PIPM, small-scale gradients are still undetectable because of the finite size of the correlation subimages. An alternative technique that can overcome resolution and gradient-biasing effects is particle tracking. Particle tracking, by definition, tracks individual flow tracer particles, and assuming ideal particles that exactly follow the flow, can provide a high-accuracy measurement of the flow field. The challenge is to successfully track large numbers of particles to achieve high spatial resolution.

In this work, a particle-tracking scheme based on the ideas of Cowen and Monismith (1997) was implemented. The scheme is also similar to those of Keane et al. (1995) and Rehm and Clemens (1999). It is a hybrid technique that uses dynamic FFT DPIV results to guide the particle-matching algorithm. This approach allows for much higher particle-seeding densities than have been used in previous particle-tracking methods. Details of the method can be found in Cowen and Monismith (1997), therefore it is only summarized here. The routine begins by performing a coarse dynamic FFT DPIV pass over the entire image domain. Next, the original images are binarized using a user-adjustable threshold and individual particles in both images are identified and labeled using a region-growing algorithm. Particles between image A and image B are matched in a multi-step process that utilizes the coarse DPIV pass results as initial displacement estimates that are increasingly refined and subsequently used to direct the particle-matching algorithm. Each particle displacement is found by differencing the two matched particle centers, which are found from either an intensity-weighted centroiding method or, as in the Cowen and Monismith (1997) approach, two three-point Gaussian fits. A final quality control check is imposed: the particle displacement,  $dp$ , is compared to the quantities  $(1 \pm \alpha)U_{\text{est}}$  where  $U_{\text{est}}$  is the final estimated displacement and  $\alpha$  is a selectable factor. If  $(1 - \alpha)U_{\text{est}} < dp < (1 + \alpha)U_{\text{est}}$ , the result is retained, otherwise, the particle match is discarded. This check effectively operates like an outlier removal routine and is successful at removing wildly incorrect particle matches. All particles are tracked in the same fashion, and the resulting irregularly spaced data are put to a regular grid using a standard triangulation-based interpolation routine.

### 3

#### Algorithm accuracy and efficiency

To ascertain the accuracy and computational efficiency of the five processing schemes described above, numerically generated PIV images were used with known prescribed displacements. Images were generated by populating a black background with randomly located particle images. A random number generator was used to determine the horizontal ( $x$ ) and vertical ( $y$ ) coordinates of each particle image in the field. Particle images were assumed to possess normally distributed intensity profiles of the form

$$I(x, y) = I_o \exp \left[ \frac{(x - x_o)^2 + (y - y_o)^2}{-2\sigma^2} \right], \quad (4)$$

where  $I_o$  is the central intensity (also randomly generated) at the particle center  $(x_o, y_o)$ , and  $\sigma$  is related to the particle diameter as  $D_p = 4\sigma$  (based on the  $e^{-2}$  intensity value). A constant particle diameter of  $D_p = 2.8 \text{ px}$  was used. The digital-image array was filled by integrating Eq. (4) over each member pixel for each particle image. The particle image density was set at 20 particles per  $32 \times 32$ -px image region. This density is realistic in practice and is comparable to values used in other simulated image tests. Once particles were distributed in image A, a displacement function was used to shift the particle locations in image B

and the pixel intensity integrations repeated. Two displacement functions were investigated: uniform translation and uniform solid-body rotation. Out-of-plane particle motions were not simulated.

Errors associated with DPIV algorithms often are decomposed into two components: mean bias error and RMS error. Given the actual displacement,  $d_a$ , and the estimated DPIV/PTV displacements,  $d_i$  ( $i = 1, 2, \dots, N$ ), the mean bias error can be defined as

$$E_{\text{bias}} = d_m - d_a, \quad (5)$$

where  $d_m$  represents the spatial mean of the  $N$  estimated displacements. Mean bias errors arise when the actual correlation peak is not well represented by the peak-fitting curve or procedure (Huang et al. 1997). This is often the case when out-of-pattern effects are prominent, or when strong gradients within the subimage act to broaden and/or distort the symmetry of the correlation peak. The RMS error was defined as

$$E_{\text{rms}} = \sqrt{\frac{1}{N} \sum_{i=1}^N (d_i - d_m)^2}, \quad (6)$$

and reflects the deviation of the displacement estimates about their mean. RMS errors are due to any number of influences: poor particle seeding, velocity gradients, out-of-plane particle motions, imaging imperfections (e.g., non-uniform illumination, camera non-linearities, parallax), and data acquisition noise. As in Huang et al. (1997), the errors associated with the methods of processing the digital-image data are the main focus of this investigation. For all results shown here, the only post-processing performed on the displacement fields was a simple outlier removal operation. Outliers were determined on a node-by-node basis by differencing each component of displacement with its eight neighbors and comparing the results to a user-selected threshold. If the differences exceeded this threshold for more than four neighbors, the value was flagged an outlier and replaced with a bilinearly interpolated value. With slight modification to the above scheme, edge and corner values in the data arrays were also processed with the outlier removal operation. Use of outlier removal was consistent with the theme of practical application of these methods.

### 3.1

#### Uniform translation

Following the approach of Huang et al. (1997), a series of subpixel uniform displacements ranging from 0 to 1 px was considered first. The image size was  $640 \times 640$  px. The situations for  $16 \times 16$ -px and  $32 \times 32$ -px subimage size were studied separately with zero overlap (all measurements independent). For each case, the processing strategy was held consistent for all processing methods. In these uniform translation tests, outlier removal was not necessary. The DPIV processing results for the RMS and mean bias errors are shown in Fig. 2 for the  $32 \times 32$ -px case. Because the direct spatial domain correlation incorporates the most image information and does not inherently suffer from out-of-pattern effects, this method produces the most accurate results, as expected. This figure also begins to

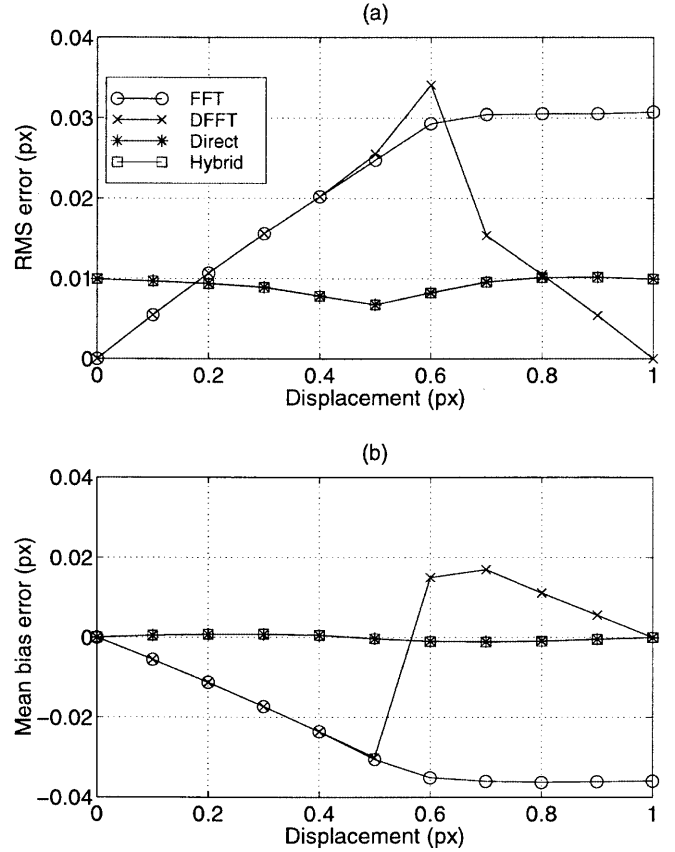


Fig. 2a, b. Variation of a RMS error, and b mean bias error, for all four DPIV methods over the 0-1 px displacement range with  $32 \times 32$  subimage processing

demonstrate the favorable performance of the hybrid scheme and the reasoning behind it. The hybrid technique is able to retain the same accuracy as the full direct correlation and as will be shown in Sect. 3.2, does so with an increase in speed of at least an order of magnitude.

Considering the RMS errors, the variability of the error for the direct and hybrid techniques remains essentially constant over the complete subpixel range. The reason for this finite variability is the statistical variability in the content of the individual subimages themselves. With direct calculation, the peak is comprised of correlations between several different A-B subimage combinations and consequently, the peak can assume an asymmetric shape due to the particular spatial distributions of particle images and their intensities in these subimages. This asymmetry should be a random process, assuming randomly located particles with randomly generated intensities. Thus, subpixel Gaussian curve fitting of the peak, as was done here, will yield random variability in the computed peak location. The FFT-based methods produce less correlation peak asymmetry near zero displacements because of the assumed periodicity of the calculation. This is observed for the FFT and DFFT methods near displacements of zero and for the DFFT method near 1 px, where it has shifted to produce a zero residual displacement. The standard FFT correlation shows a proportional increase in RMS error as the pixel displacement increases for displacements less than 1/2 px, confirming what was noted

earlier. Beyond  $1/2$  px, but less than 1 px, the RMS error is nearly independent of the displacement. The behavior of the DFFT scheme is somewhat more subtle. For displacements less than  $1/2$  px, the DFFT results are identical to that of the FFT method since no shifting takes place. At  $0.7$  px and beyond, the DFFT results are symmetric (because of shifting) and at  $0.5$  px and  $0.6$  px there are discrepancies. The explanation lies with the dynamic nature of the DFFT method. When the displacement is  $1/2$  px, most DFFT results will converge to rounded values of  $0$  px because of the typical underestimation of displacements computed using FFTs (see Fig. 2b). A small number of nodes may converge to  $1$  px, and for these nodes, the DFFT will shift the second subimage by  $1$  px. Consequently, the mean result at  $1/2$  px will be different for the DFFT compared to the FFT, which does not shift. Similarly, at  $0.6$  px most DFFT nodes will converge to  $1$  px, but a small number may round to  $0$  px because of underestimation. This results in the increased variability observed at this displacement.

Turning to the mean bias errors, for the direct and hybrid methods, the mean bias incurred is extremely small ( $0.0005$  px mean magnitude over the complete  $0$ – $1$ -px range). For the two FFT-based methods, the mean bias errors are comparable in magnitude to the RMS errors, a result also described by Huang et al. (1997). The non-symmetric behavior of the DFFT mean bias errors over the  $0$ – $1$ -px range follows the same logic as described above for the RMS errors. Evident is the consistent underestimation bias of the displacement using the static FFT scheme. This is a consequence of in-plane loss of pairs. Westerweel (1993) has shown that this bias is directly proportional to the width of the correlation peak relative to the subimage dimension. In general, this width is determined by not only the particle image size but also by the variation of the displacement field over the subimage. To further amplify this behavior, Fig. 3 shows the displacement frequency distributions due to a  $1.0$ -px shift for the FFT and direct spatial domain correlation methods. This result clearly shows both the significant negative mean bias error and




Fig. 3a, b. Displacement frequency distributions for a FFT and b direct spatial domain correlation methods. Actual displacement is  $1.0$  px

the substantial variability of the displacement estimates obtained from the FFT calculation. Westerweel (1993) derived an analytical result that enables elimination of this mean bias error for FFT methods (also, Raffel et al. 1998). The FFT implementation studied here does not incorporate this correction since it is unclear as to how prevalent the use of this correction is in practice. Essentially, the correction involves a compensation factor in the correlation calculation to account for the in-plane loss of pairs. Interestingly, the present hybrid technique achieves the same effect in a different manner: the FFT implementation of the correlation is employed and the effect of in-plane loss of pairs is eliminated via the local, direct spatial domain calculation of the correlation peak.

Processing with the  $16 \times 16$ -px subimages yielded similar qualitative behavior in both components of the error over the  $0$ – $1$ -px range. Summarizing, for the  $16 \times 16$ -px case, RMS errors were higher by about a factor of 4 and mean bias errors were higher by about a factor of 2.5 over that shown for the  $32 \times 32$ -px result. This implies (and was further validated with tests using  $64 \times 64$ -px subimages) that, for these displacements, the error in the displacement estimates decreases with increasing subimage size. This would be expected since larger subimage sizes incorporate a greater statistical sample, or alternatively, more unique subimage templates; the drawback, of course, is reduced spatial resolution.

Investigating the errors for pixel displacements in the  $0$ – $1$ -px range is useful in ascertaining the detailed behavior of the correlation schemes. Here, it has revealed certain shortcomings of FFT, as well as DFFT, methods of DPIV processing that are undesirable in certain flow situations. However, rare is it that DPIV images yield such small displacements in general, making processing schemes that use offset subimages very advantageous (e.g., Westerweel et al. 1997). To examine the behavior of the four processing methods for pixel displacements larger than  $1$  px, the same procedure as above was performed for displacements of  $1.25, 2.25, \dots, 10.25$  px (Fig. 4). First to note is the relatively poor performance of the standard FFT method, in terms of consistent negative mean bias error (underestimation) and larger RMS variability in the error. This is a consequence of its static nature and the resulting out-of-pattern effects incurred. For the other three methods, patterns very similar to the results for the  $0$ – $1$ -px range were found to repeat for subsequent single-pixel intervals (see Fig. 6) and the results at  $dx.25$  px are observed. Thus, increasing displacements have little effect on the behavior of the accuracy of these three methods. This is noteworthy because subimage shifting is taking place automatically, thereby handling general flow patterns without any special knowledge of the flow. Results using  $16 \times 16$ -px subimages showed both larger RMS and mean bias errors, and somewhat less systematic behavior over the range of displacements. Considering just the FFT results as an indicator, the maximum RMS error was nearly an order of magnitude higher (at a displacement of  $5.25$  px – beyond this point, the correlation is nearing the Nyquist limit and the estimates become poor) compared to the  $32 \times 32$ -px case. The largest FFT mean bias error was higher by roughly a factor of 3 for the  $16 \times 16$ -px case.

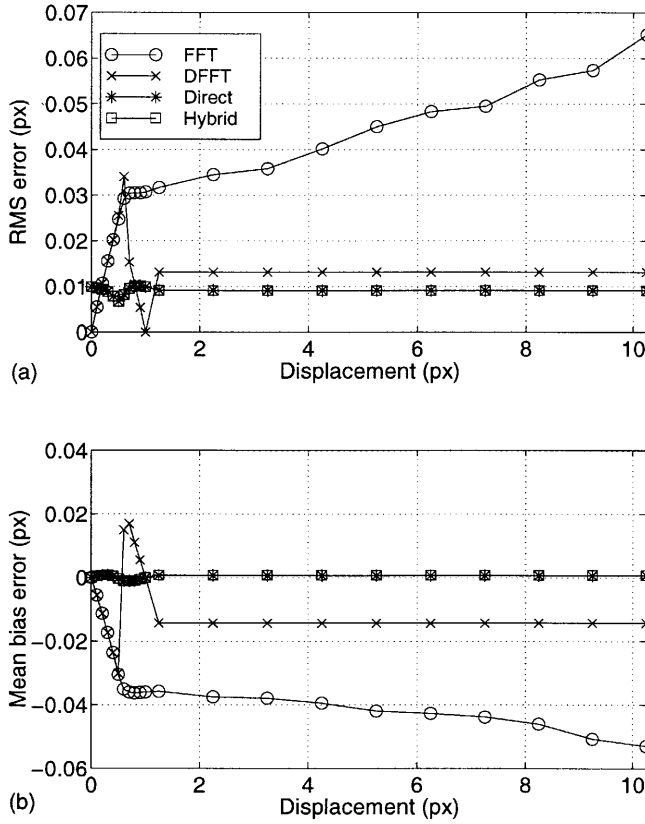


Fig. 4a, b. Variation of a RMS error, and b mean bias error, for all four DPIV methods over the 0–10.25-px displacement range with  $32 \times 32$ -px subimage processing

The mean, relative displacement error,

$$E_{\text{rel}} = \frac{1}{N} \sum_{i=1}^N \left| \frac{d_i - d_a}{d_a} \right| \times 100\% , \quad (7)$$

for the  $32 \times 32$ -px subimage size case over the detailed 0–1-px range and the extended displacement range is shown in Fig. 5. For this case, as well as for the  $16 \times 16$ -px case (not shown), the direct and hybrid methods show the smallest errors, with the DFFT method yielding similar results. For all methods, the errors beyond about 2 px are relatively constant – less than 2% for the  $32 \times 32$ -px case, and less than 5% for the  $16 \times 16$ -px case. These results show that the directly calculated correlation methods (direct and hybrid) are noticeably superior to the traditional static FFT method over the complete displacement range studied and are slightly better than the DFFT method over this range.

To further understand the detailed behavior of the DPIV errors at larger displacements, a more detailed analysis was performed for the 5–6-px displacement range and the  $32 \times 32$ -px results are given in Fig. 6. The errors for the DFFT, direct, and hybrid methods appear to be cyclical over the displacement range studied. Over this range, both the RMS and mean bias errors show qualitative similarity to the 0–1-px range results, indicating that PIPM is achieved. This is not the case for the FFT method and consistent errors are observed. Unlike the result for the 0–1-px range, the relative error also exhibits a cyclic




Fig. 5a, b. Variation of relative error over the a 0–1-px displacement range, and b 0–10.25-px displacement range, for all four DPIV methods with  $32 \times 32$ -px subimage processing

pattern over the 5–6-px displacement range. This is because the relative error for the 0–1-px range is based on scaling the absolute error by a small displacement value ( $< 1$  px), whereas for the 5–6-px range, the absolute error is similar but the scaling value is much larger.

At this point, the results of PTV processing of the uniform displacement tests are included. Figures 7 and 8 show the  $32 \times 32$ -px DPIV error results with the raw, irregularly gridded results of PTV included. Although exhibiting favorable mean bias errors of less than 0.012 px, the RMS errors are quite large. The primary reason behind this is particle pair mismatching. Although these events were rare, the errors incurred were relatively significant, leading to the large RMS errors seen in Fig. 7a and the large relative errors for the smallest displacements in Fig. 8a. Figure 8c shows the behavior of the PTV error at larger displacements in more detail. The PTV error also appears to exhibit a cyclic structure for displacements greater than 1 px.

### 3.2 Uniform rotation

To simulate rotational flow, synthetic image pairs were used to achieve constant vorticity, solid-body rotations. Two rotations were investigated. The rotation used for the  $16 \times 16$ -px subimage case had a maximum pixel displacement of 5.6 px and the rotation used for the

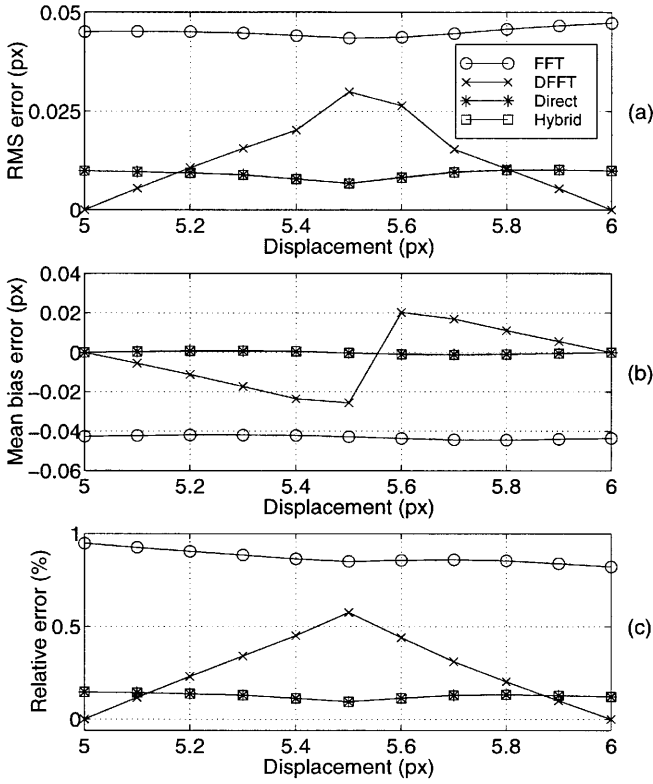


Fig. 6a-c. Variation of a RMS error, b mean bias error, and c relative error, for all four DPIV methods over the 5-6-pixel displacement range with  $32 \times 32$ -pixel subimages

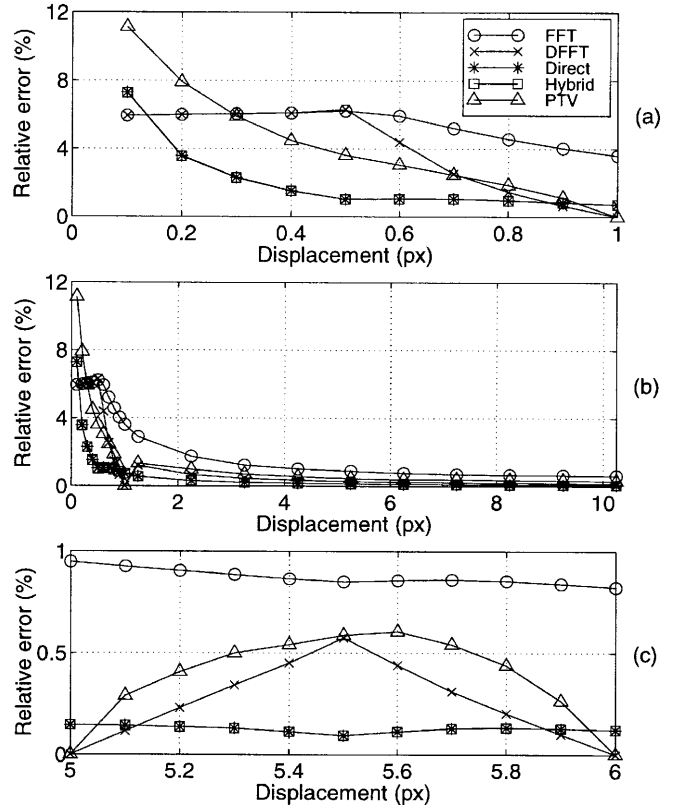


Fig. 8a-c. Variation of relative error over the a 0-1-pixel displacement range, b 0-10.25-pixel displacement range, and c 5-6-pixel displacement range, showing raw, ungridded PTV performance ( $32 \times 32$ -pixel subimage size used for all methods)




Fig. 7a, b. Same as Fig. 2, but with raw, ungridded PTV results included

$32 \times 32$ -pixel subimage case had a maximum of 11.3 pixels. The choice of these displacements was guided by the often used rule of thumb that the maximum displacement within a given subimage should be less than about 1/3 of the subimage dimension (e.g., Willert and Gharib 1991). In addition to assessing accuracy, these tests also were used as the basis for the algorithm efficiency comparison. Computational performance is presented in Table 1. The speed of the standard FFT method is readily apparent; it is at least twice as fast as the other DPIV schemes, being 1-2 orders of magnitude faster than the direct spatial domain calcula-

Table 1. Computational performance of DPIV/PTV processing algorithms. Computational times are the result of processing with a 400-MHz Intel PC running Linux. For the PTV results, subimage size reflects the size used during the coarse DPIV pass

Scheme	Time (s)	Vector count	Vectors/s
<b>16 x 16-pixel subimages</b>			
FFT	0.71	1369	1928
DFFT	1.43	1369	957
Direct	9.73	1369	141
Hybrid	3.18	1369	431
PTV	4.62	5376	1164
<b>32 x 32-pixel subimages</b>			
FFT	0.81	361	446
DFFT	1.63	361	221
Direct	35.11	361	10
Hybrid	3.29	361	110
PTV	4.62	5101	1104



tion. While not as efficient as the FFT-based algorithms, the hybrid technique is found to be anywhere from about three to ten times the speed of the direct correlation. The speed of the PTV scheme is impressive considering the large number of vectors the method is able to yield.

The outcome of the accuracy tests is summarized in Table 2. Outlier removal became a factor in the rotation tests and is reported in this table. In this case, where a single mean displacement is of little use, the definition of the RMS error was modified slightly,

$$E_{\text{rms}} = \sqrt{\frac{1}{N} \sum_{i=1}^N (d_i - d_a)^2} , \quad (8)$$

with the actual local displacement being used instead. The mean (relative) vorticity error was defined as

$$E_{\text{vort}} = \frac{1}{N} \sum_{i=1}^N \left| \frac{\omega_i - \omega_a}{\omega_a} \right| \times 100\% , \quad (9)$$

where  $\omega_a = 2\theta$ ,  $\theta$  being the angle of rotation. The estimated vorticities,  $\omega_i$ , were computed from the local eight-point circulation calculation around each node. This approach has been shown to be preferable over a first-order central difference method (Westerweel 1993). Table 2 reveals how the performance of the algorithms is affected by the subimage size used. In terms of displacements, the  $16 \times 16$ -px subimage size DPIV processing shows much more error and variability in accuracy than does the case with  $32 \times 32$ -px subimages. This is largely the effect of the reduced statistical sample size increasing the probability of spurious correlation peaks. Typically, at least four or five particles per subimage are required to obtain an unambiguous measurement of the displacement – at the particle seeding density used here, the average number of particles per  $16 \times 16$ -px subimage is just slightly greater than five.

**Table 2.** Simulated uniform rotation results. *Boldface* indicates best performance. DPIV processing parameters are identical to those used in the uniform translation tests except for outlier detection. For the  $16 \times 16$ -px case, an outlier threshold of 2 px was used; for the  $32 \times 32$ -px case, a threshold of 4 px was used. DPIV outlier fraction represents the fraction of the total number

Examining the algorithms, at the  $16 \times 16$ -px subimage size, the FFT algorithm shows significant RMS error and large error in the derived vorticity. The dynamic FFT scheme reduces these errors by approximately a factor of 2. The direct and hybrid methods further reduce this error by another factor of 2. Thus, at this fine spatial resolution, the direct and hybrid techniques are still viable, whereas the FFT-based routines are being strained. This would indicate that in addition to extending the useable valid dynamic range, the direct and hybrid techniques also can offer greater spatial resolution without sacrificing accuracy. The  $32 \times 32$ -px subimage results show much improved accuracy and less performance variability. All DPIV methods show favorable results, the differences in errors, both absolute and relative, being small. Note, however, that the direct and hybrid schemes reduce the displacement errors by nearly a factor of 2 over that of standard FFT processing. Based on these results, the errors of the DPIV methods would be expected to become less disparate as the subimage size became larger, at which point choice of a processing method can be done based on efficiency concerns alone.

Three forms of particle-tracking results are also included in Table 2. The first, PTV, represents the accuracy based on the irregularly spaced individual particle vectors. The second, PTV<sub>DPIV-grid</sub>, represents the results after transforming the raw PTV vectors to a regular grid identical to the DPIV grid formed by the other four processing methods. The third, PTV<sub>fine-grid</sub>, represents the raw particle-tracking results gridded to a fine-resolution regular array. The array density was chosen such that approximately the same number of vectors appear in the gridded result as in the raw result, thereby mimicking the actual spatial resolution.

In both subimage size cases, the PTV result, which successfully tracked 80–85% of the detected particles, shows the smallest relative displacement error. Examining

of computed vectors that were determined to be outliers. PTV outlier fraction represents the fraction of the total number of particles that were eliminated from tracking due to poor local correlation or considerable deviation from the local mean displacement

Scheme	Outlier fraction (%)	RMS error (px)	Relative error (%)	Vorticity error (%)
<b>16 × 16-px subimages</b>				
FFT	7.23	0.326	6.93	28.85
DFFT	2.41	0.165	3.08	16.84
Direct	0.73	0.088	2.03	7.21
Hybrid	<b>0</b>	0.060	1.98	6.86
PTV	11.08	0.089	<b>1.15</b>	–
PTV <sub>DPIV-grid</sub>	–	0.062	1.19	<b>6.77</b>
PTV <sub>fine-grid</sub>	–	<b>0.059</b>	1.23	14.29
<b>32 × 32-px subimages</b>				
FFT	<b>0</b>	0.095	1.93	2.39
DFFT	<b>0</b>	0.060	1.28	2.07
Direct	<b>0</b>	<b>0.052</b>	1.18	<b>1.70</b>
Hybrid	<b>0</b>	<b>0.052</b>	1.18	<b>1.70</b>
PTV	8.15	0.169	<b>0.75</b>	–
PTV <sub>DPIV-grid</sub>	–	0.108	1.25	2.97
PTV <sub>fine-grid</sub>	–	0.117	1.14	11.13

the  $16 \times 16$ -px subimage case, the smoothing effect of the gridding procedure is observed, yielding a reduction in the particle-tracking RMS errors. This smoothing also has the effect of marginally increasing the relative error, as would be expected for a vortical flow with displacement gradients. The incurred vorticity error as calculated from the  $PTV_{DPIV-grid}$  result is essentially the same as that for the two spatial domain correlation DPIV methods. For the fine-grid result, the vorticity error is doubled ( $2.11 \times$ ). This is not surprising since the finer grid is close to twice ( $1.97 \times$ ) the resolution of the DPIV grid and the vorticity is still computed using eight-point circulation contours. The vorticity calculation as implemented here involves a division by the node spacing and will show greater sensitivity as the node spacing is reduced. Thus, the errors in displacement, which essentially remain the same for both PTV gridded resolutions, are amplified during the vorticity calculation for the finer grid resolution. Similar PTV results are found for the  $32 \times 32$ -px subimage case; however, the RMS errors for the raw data are nearing 0.2 px. It should be kept in mind, though, that this flow field has significant displacements – this fact is reflected by the excellent PTV relative error results. It appears that the gridding operations in the  $32 \times 32$ -px case have the same effect of reducing the RMS scatter and increasing the relative errors of the PTV results. The behavior of the relative vorticity error follows the earlier discussion. In this case, the  $PTV_{fine-grid}$  results are at a spatial resolution 3.94 times the DPIV grid results and, as expected, the vorticity error increases by a similar factor of 3.74 since the displacement errors do not change significantly between the two resolutions. Rehm and Clemens (1999) report similar behavior of derived quantities at higher resolutions. Finally, it was found that except for the FFT method with  $16 \times 16$ -px subimages, all other processing schemes obtained mean values for the estimated vorticity extremely close to the actual values (less than 1% error). The FFT results for the  $16 \times 16$ -px subimage case underestimated the actual vorticity by about 5%.

#### 4

##### Summary

These numerical experiments have shown that, on a detailed level, the generation of the cross-correlation using the direct spatial domain calculation provides high-accuracy results from digital-image data for a variety of flow fields. The method accomplishes this by incorporating more of the available actual image content when computing the correlation surface, rather than relying on subimage extension and periodicity assumptions. The result is a correlation that avoids the negative effects of in-plane loss of pairs and minimizes the displacement estimate error. The dynamic FFT method is a viable processing scheme, but the standard FFT method has to be considered the least accurate of the group studied and its use would not be suggested unless speed is a primary concern. Both the DFFT and the hybrid methods can improve upon the accuracy of the FFT processing at only minimal computational cost. In fact, the hybrid method, which is an original approach to PIV data processing, was found to mimic the accuracy of the direct spatial

domain method, while only requiring a fraction of the computational effort. The particle-tracking routine studied has shown its ability to provide accurate displacement results at resolutions several times greater than standard DPIV techniques. Note, DPIV data often are oversampled (overlapping subimages), which increases the number of measurement points. However, this does nothing to reduce the size of the scales resolvable; the smallest scales measurable are still dictated by the size of the subimage. The PTV scheme used here would be better served by a more sophisticated gridding procedure than the one employed for these results (e.g., the thin-shell spline method of Spedding and Rignot 1993), but this subject is beyond the scope of this paper. Additional improvements that reduce particle pair mismatching would also be helpful (see, for example, Rehm and Clemens 1999).

##### References

- Adrian RJ (1991) Particle-imaging techniques for experimental fluid mechanics. *Annu Rev Fluid Mech* 23: 261–304
- Burt PJ; Yen C; Xu X (1982) Local correlation measures for motion analysis, a comparative study. In: Proceedings of the Conference on Pattern Recognition and Image Processing, Las Vegas, Nev., 14–17 June 1982, pp 269–274
- Cowen EA; Monismith SG (1997) A hybrid digital particle tracking velocimetry technique. *Exp Fluids* 22: 199–211
- Gonzales RC; Wintz P (1987) Digital image processing. Addison-Wesley, Reading, Mass.
- Huang HT; Fiedler HE; Wang JJ (1993) Limitation and improvement of PIV. Part I: limitation of conventional techniques due to deformation of particle image patterns. *Exp Fluids* 15: 168–174
- Huang H; Dabiri D; Gharib M (1997) On errors of digital particle image velocimetry. *Meas Sci Tech* 8: 1427–1440
- Keane RD; Adrian RJ (1992) Theory of cross-correlation analysis of PIV images. *Appl Sci Res* 49: 191–215
- Keane RD; Adrian RJ; Zhang Y (1995) Super-resolution particle imaging velocimetry. *Meas Sci Tech* 6: 754–768
- McKenna SP (2000) Free-surface turbulence and air-water gas exchange. Massachusetts Institute of Technology, Cambridge, Mass.
- Raffel M; Willert C; Kompenhaus J (1998) Particle image velocimetry: A practical guide. Springer, Berlin Heidelberg New York
- Rehm JE; Clemens NT (1999) An improved method for enhancing the resolution of conventional double-exposure single-frame particle image velocimetry. *Exp Fluids* 26: 497–504
- Roesgen T; Totaro R (1995) Two-dimensional on-line particle image velocimetry. *Exp Fluids* 19: 188–193
- Scarano F; Riethmuller ML (1999) Iterative multigrid approach in PIV image processing with discrete window offset. *Exp Fluids* 26: 513–523
- Spedding GR; Rignot EJM (1993) Performance analysis and application of grid interpolation techniques for fluid flows. *Exp Fluids* 15: 417–430
- Westerweel J (1993) Digital particle image velocimetry: Theory and application. Delft University Press, Delft, The Netherlands
- Westerweel J (1997) Fundamentals of digital particle image velocimetry. *Meas Sci Tech* 8: 1379–1392
- Westerweel J; Dabiri D; Gharib M (1997) The effect of discrete window offset on the accuracy of cross-correlation analysis of digital PIV recordings. *Exp Fluids* 23: 20–28
- Willert CE; Gharib M (1991) Digital particle image velocimetry. *Exp Fluids* 10: 181–193