

An efficient anti-aliasing spectral continuous window shifting technique for PIV

Qian Liao, Edwin A. Cowen

Abstract A new sub-pixel correlation peak locating algorithm for PIV analysis is introduced. The method is theoretically consistent with the method of continuously shifting interrogation sub-windows by fractional displacements, which has proven to be an effective way to reduce the bias error associated with integer pixel aliasing, or ‘peak-locking’. However the proposed algorithm performs continuous window shifting in the spatial frequency domain using the ‘shift’ property of the Fourier transform, thus it is equivalent to interpolating the original digital image with the Fourier transform reconstruction. Synthetic and real PIV images are used to test the new algorithm’s performance relative to that of traditional (non-iterative) peak-finding methods and other peak-locking reduction algorithms, such as the continuous window shifting technique. The resultant bias error of the proposed algorithm is smaller (by an order of magnitude in some cases), and importantly, the periodic nature of the bias error, the characteristic signature of ‘peak-locking’, is eliminated as long as the discrete particle images have been sampled at a rate greater than the Nyquist sampling frequency. Moreover, this new algorithm is shown to be computationally efficient and it converges faster than the competing algorithms.

1

Introduction

Particle image velocimetry (PIV), with its continued improvement, has become the dominant tool for determining fluid velocities in the laboratory, principally due to its ability to resolve instantaneous 2-D and even 3-D velocity field structures. Although the accuracy of correlation-based PIV has improved significantly in recent years (Hart 2000; Gui et al. 2000; Wereley and Meinhart 2001; Gui and Wereley 2002) this has often been achieved

at relatively high computational cost (Fincham and Delerce 2000).

One of the most important sources of measurement error comes from the inherent discrete sampling of particle images by CCD (and more recently CMOS) based cameras. The mean displacement of the particle image pattern (PIP), which is commonly determined by cross-correlation, is quantized yielding discrete displacement components of integer pixel value. These components are converted to discrete physical displacement upon image calibration. Since PIV was originally proposed in its digital form (Willert and Gharib 1991), this discrete nature has been accounted for and measures have been taken to find sub-pixel accurate estimates of the location of the correlation peak. This can minimize the uncertainty in the PIPs displacement estimate to a level that is an order of magnitude smaller than a pixel. Discrete interrogation sub-window offsets along with iterative schemes (Cowen and Monismith 1997; Wereley and Meinhart 2001; Westerweel et al. 1997) have been introduced to minimize the in-plane loss-of-pair effect, thus a smaller interrogation window can be used to enhance the measurement resolution. Not only are the accuracy and resolution increased, but also the dynamic range of the measured fluctuating velocity field can be increased significantly, without sacrificing sub-pixel precision, by extending the exposure time delay (Δt) (Raffel and Kompenhans 1994).

At first glance it appears that extending Δt leads to the reduced importance of sub-pixel displacement uncertainty, since the increased dynamic range renders the non-dimensional uncertainty (pixels/pixel) smaller. However, two important issues emerge: the wide range of scales in even moderate Reynolds number (Re) turbulent flows and the fundamental assumption of all particle image based quantitative imaging techniques – that displacements can be accurately determined from discrete samples along a Lagrangian track. For even moderate Re turbulence, the separation between the large scale energy containing eddies and smallest fluctuations at the dissipation scale is always more than a decade. This provides a great challenge to PIV interrogation algorithms as both wide dynamic range and excellent sub-pixel precision are desirable if both the largest and smallest structures are to be resolved simultaneously with acceptable accuracy.

In such flows, extending the exposure time delay can analytically be shown to introduce serious measurement errors. Lourenco and Krothapalli (2000) demonstrated that the determined velocity vector by a PIV algorithm is a filtered version of the actual velocity field, $\vec{U}(\vec{x}, t)$, since it

Received: 8 April 2004 / Accepted: 7 October 2004
Published online: 13 January 2005
© Springer-Verlag 2005

Q. Liao (✉), E. A. Cowen
DeFrees Hydraulics Laboratory,
School of Civil and Environmental Engineering,
Cornell University, Ithaca, NY 14853-3501, USA

The authors gratefully acknowledge the financial support of the Office of Naval Research (Grant Nos. N00014-98-1-0774 and N00014-99-1-0591, Dr. Keith Ward, program manager).

is approximated by a finite time difference of the PIP displacement, i.e.

$$\begin{aligned}\vec{U}(\vec{x}, t) &\approx \frac{\vec{X}_{t+\frac{\Delta t}{2}} - \vec{X}_{t-\frac{\Delta t}{2}}}{\Delta t} \\ &\approx \frac{\int_{t-\frac{\Delta t}{2}}^{t+\frac{\Delta t}{2}} \vec{U} dt}{\Delta t} \approx \frac{\int_{\vec{x}-\frac{\Delta \vec{x}}{2}}^{\vec{x}+\frac{\Delta \vec{x}}{2}} \vec{U} d\vec{x}}{\Delta \vec{x}},\end{aligned}\quad (1)$$

where \vec{X} is the Lagrangian location of the moving PIP which is coincident with \vec{x} (the Eulerian location of PIV measurement) at time t ; Δt is the exposure time delay and $\Delta \vec{x}$ is the PIP displacement. In a 1-D sense, the transfer function (spatial frequency response) of this filter is

$$F(\kappa) = \frac{\sin(\pi\kappa|\Delta \vec{x}|)}{\kappa|\Delta \vec{x}|}, \quad (2)$$

where $\kappa \equiv \frac{1}{|\vec{x}|}$, the spatial frequency. Thus if $|\Delta \vec{x}|$ is allowed to be too long (by extending Δt) and it becomes much larger than the length scale of the small eddies in the turbulence field, the filtering can cause serious errors to the calculated velocity spectra or estimates of the fluctuating strain rates and hence the turbulent dissipation. Also extending Δt will increase the probability of the loss of particle pairs in the PIP due to the out-of-plane motion, which will decrease the reliability of PIV correlations.

The limitation on the magnitude of the PIP displacement reemphasizes the importance of sub-pixel precision. In the PIV literature, the sub-pixel measurement error is generally decomposed into bias and random error components

$$E = E_b + E_r \quad (3)$$

The above error components arise, in general, from: discrete sampling, the optical transfer system, the deviation of the particle image shape from an assumed symmetric form, the size of interrogation sub-window, the sub-pixel peak finding algorithm, the in-plane shear, etc. In this paper, we will focus on the sources of perhaps the dominant bias error, the well-known peak-locking phenomenon, or the tendency of determined displacements to alias toward integer pixel displacements.

The peak-locking tendency of various types of PIV algorithms can most easily be identified by the histogram of the measured velocity field, where spurious peaks are found at velocity values that correspond to integer pixel displacements (or ± 0.5 pixel displacements). Peak-locking effects are relatively well studied, often involving error analysis based on synthetic particle image pairs of a uniformly translating flow. These analysis procedures show that if the discrete sub-window offsets are applied, the bias error versus the fractional displacement has a periodic distribution with a period of one pixel. The bias error is always minimal for integer pixel displacements. While its maximum location varies according to different sub-pixel peak location finding algorithms. The error is anti-symmetric with respect to the $1/2$ pixel displacements. Raffel and Kompenhans (1994) showed that the peak-locking phenomenon in the velocity histogram is directly related to the profile of the bias error distribution plus the random error amplitude variation.

Nogueira et al. (2001a) determined that there are three major sources of peak-locking bias errors:

- The under-sampling of particle images, i.e., the image resolution is insufficient to describe a particle (Prasad et al. 1992; Westerweel 1998).
- The sub-pixel peak finding algorithm. Earliest PIV algorithms relied on center-of-mass estimators to determine the sub-pixel correlation peak location. This involves a proper selection of a cut-off threshold, which is in general rather arbitrary. Other estimators have recently proved more popular, and most of them are ad-hoc 1-D or 2-D interpolations or curve-fittings over the local neighborhood of the peak value on the discrete correlation plane. Examples include: three-point exponential curve fits (Willert and Gharib 1991), three-point Gaussian curve fits (Cowen and Monismith 1997), fitting by the nine-point 2-D Taylor expansion of a Gaussian surface (Sholl and Savas 1997), etc. Westerweel (1993) provided a simple analytic model to investigate the performance of several such interpolators and found all of them to be unavoidably biased, yielding bias error distributions that agree well with simulations from synthetic images. Although Gaussian fits are in many cases shown to be superior, the choice of an optimal interpolator is still an open question since other factors such as the particle diameter and in-plane shear are important as well. It should also be noted that, while all of these interpolators themselves are low pass filters (that contribute to the errors in peak location as well), using a Fourier transform reconstruction (i.e., Whitaker's reconstruction) on the discrete correlation plane may decrease the peak-locking effect. Lourenco and Krothapalli (1995) actually took such an approach to find the sub-pixel correlation peak location. However, the reconstruction they used was based on a localized 5×5 cell region, which was arbitrary, also a detailed analysis of the performance was not given.
- The truncation of particle images by the borders of the interrogation sub-window. For Fourier transform based correlation algorithms, this truncation effect introduces oscillations in the Fourier domain over all frequencies, due to the periodic assumption of the discrete Fourier transform (DFT). The spurious oscillations will contaminate the cross(auto)-correlation if the true displacement is not an integer pixel, resulting in peak-locking errors.

Despite all of the above factors, there is no bias error when the true displacement is, in fact, an integer pixel. Thus a solution to reducing peak-locking is to apply the continuous window shifting technique (Gui and Wereley 2002; Nogueira et al. 2001a), i.e., iteratively shift the interrogation sub-window by fractional displacements until the final determined sub-pixel displacement is driven to zero. It should also be mentioned that for PIV algorithms which include particle image distortion (Huang et al. 1993; Scarano and Riethmuller 2000; Fincham and Delerce 2000), the same peak-locking reduction (anti-aliasing) approach is implemented de facto when an iterative PID scheme is applied. Generalizing, for non-correlation

based approaches (e.g., the minimum quadratic difference, or MQD), iterative approaches have also been shown to reduce bias errors associated with integer pixel aliasing (Gui et al. 2000).

Both continuous image shifting and particle image distortion require re-sampling of the original particle images. Interpolations of the original particle images are again found to be unavoidable. Bi-linear and bi-parabolic interpolation, 2-D spline interpolation and Whittaker reconstruction are all reported to have satisfactory results at reducing peak-locking, although no theoretical justification has been given.

Below we propose a new approach to continuous window shifting. Our approach differs from those described above in that the shifting is achieved in the spatial frequency domain, rather than in the image space. This new algorithm takes advantage of the ‘shift’ property of the Fourier transform and thus no image re-sampling is required, making it computationally efficient. The new algorithm avoids the low-pass filtration effect that is inherent to many ad-hoc image interpolators, thus resulting in complete elimination of the peak-locking effect when the appropriate sampling criteria are met.

2

An optimal iterative sub-pixel peak locating algorithm

2.1

Basic algorithm

The proposed peak-locking removal algorithm is compatible with any conventional PIV algorithm that utilizes a spectral based DFT correlation technique. Without loss of generality, in this paper we will work with the dynamic sub-window offset technique of Cowen and Monismith (1997), upgraded to the second-order accuracy in space (Wereley and Meinhart 2001). Namely, if the velocity components are to be determined at a given image location (i, j) , interrogation sub-windows with size $N \times N$, which may be separated by an initially estimated displacement (S_i, S_j) are used to sample local regions in an image pair (referred to as images a and b , respectively), where N may be selected to be a power of 2 (or the sub-images can be padded with zeros to make their size a power of 2) if a standard 2-D FFT algorithm is to be used. The centers of the two sub-window are located at the points $(i - \frac{S_i}{2}, j - \frac{S_j}{2})$ and $(i + \frac{S_i}{2}, j + \frac{S_j}{2})$ on images a and b , respectively. The scheme is seen to be a central finite difference in space and hence it is second-order accurate in space. The PIPs on the interrogation sub-windows are normalized by subtracting their mean and dividing by their standard deviations, such that the resultant cross-correlation values are not larger than 1.

The cross-correlation of the two sub-window images are calculated and the integer displacement of the PIP is determined based on the pixel index of the maximum value on the correlation plane. The initial sub-windows offset, (S_i, S_j) , is then updated. This procedure is iterated until the maximum in the correlation plane is located at $(0,0)$. As a result of the above analysis, the sub-pixel displacement should be confined to the range of ± 0.5 pixel.

Consider that after the final step of a conventional PIV algorithm, the two interrogation sub-windows are stored in arrays f_{ij} and g_{ij} , where $i, j=0, 1, \dots, N-1$. As a result of the spectral correlation determination, the DFT's of f_{ij} and g_{ij} , defined as

$$F_{ij} = \sum_{m,n=0}^{N-1} f_{mn} \exp \left[-i2\pi \left(\frac{mi}{N} + \frac{nj}{N} \right) \right], \quad (4)$$

$$G_{ij} = \sum_{m,n=0}^{N-1} g_{mn} \exp \left[-i2\pi \left(\frac{mi}{N} + \frac{nj}{N} \right) \right], \quad (5)$$

where the $l = \sqrt{-1}$, are also stored. If the cross-correlation of f_{ij} and g_{ij} is defined as

$$h_{ij} = \sum_{m,n=0}^{N-1} f_{mn} g_{m-i, n-j} \quad (6)$$

which makes the assumption that sub-window g_{ij} is periodic, then its DFT is

$$\begin{aligned} H_{ij} &= \sum_{m,n=0}^{N-1} h_{mn} \exp \left[-i2\pi \left(\frac{mi}{N} + \frac{nj}{N} \right) \right] \\ &= F_{ij} G_{ij}^*, \end{aligned} \quad (7)$$

where the superscript ** means the ‘the complex conjugate of’. Hence the cross-correlation plane h_{ij} is calculated as the inverse Fourier transform of H_{ij} .

After the final step of conventional dynamic sub-window PIV is completed, the neighborhood of values to the maximal correlation value $h_{0,0}$ may be used to estimate the sub-pixel location of the correlation plane peak. Assuming the correlation in the i and j component directions are uncorrelated, two 1-D sub-pixel estimating functions can be applied to $(h_{0, N-1}, h_{0,0}, h_{0,1})$ versus $(-1,0,1)$ in the j direction, and $(h_{N-1,0}, h_{0,0}, h_{1,0})$ versus $(-1,0,1)$ in the i direction to estimate the fractional part of the PIP displacement. Without loss of generality, Gaussian interpolation is used as the estimating function in this paper, specifically

$$s_x = - \frac{\log(h_{0,0}/h_{0,1}) + \log(h_{0,N-1}/h_{0,0})}{2[\log(h_{0,0}/h_{0,1}) - \log(h_{0,N-1}/h_{0,0})]}, \quad (8)$$

$$s_y = - \frac{\log(h_{0,0}/h_{1,0}) + \log(h_{N-1,0}/h_{0,0})}{2[\log(h_{0,0}/h_{1,0}) - \log(h_{N-1,0}/h_{0,0})]}, \quad (9)$$

where s_x, s_y are the fractional displacements in the x and y coordinate directions (the j and i image plane directions), respectively. The fractional displacements are next used as initial estimates of the distance by which the interrogation sub-windows are continuously shifted.

According to the ‘shift property’ of the Fourier transform (Bracewell 1998)

$$F[f(x - a)] = \exp(i2\pi\kappa a)F[f(x)], \quad (10)$$

where $F[\cdot]$ is the Fourier transform operator and $\kappa=1/x$ is the spatial frequency. Shifting the sub-windows $f_{i,j}$ and $g_{i,j}$ in the image space by $(-s_y/2, -s_x/2)$ and $(s_y/2, s_x/2)$ to get new images f'_{ij} and g'_{ij} can be achieved in the spatial

frequency domain. The discrete Fourier transforms of the shifted images are

$$F'_{ij} = c_{ij}(-s_y/2, -s_x/2)F_{ij}, \quad (11)$$

$$G'_{ij} = c_{ij}(s_y/2, s_x/2)G_{ij}, \quad (12)$$

where

$$c_{i,j}(y, x) = \exp\left[i2\pi\left(\frac{I}{N}y + \frac{J}{N}x\right)\right], \quad (13)$$

is the correction (shifting) matrix in Fourier space, and

$$I, J = \begin{cases} i, j & \text{when } i, j \leq N/2; \\ i - N, j - N & \text{when } i, j > N/2. \end{cases} \quad (14)$$

Equation (14) reflects the fact that the discrete version of Eq. (10) must ensure that F'_{ij} and G'_{ij} are the DFTs of the shifted images, i.e., F'_{ij} and $F'_{(N-i)(N-j)}$ are complex conjugates.

The Fourier transform of the shifted cross-correlation is then calculated again by Eq. (7) and the cross-correlation plane (h_{ij}) is updated as well, using the inverse FFT. In practice, it is not necessary to re-calculate the entire correlation plane, only the five points around the maximum (i.e., $h_{N-1,0}$, $h_{0,0}$, $h_{1,0}$, $h_{0,N-1}$, $h_{0,1}$), need to be updated since the sub-pixel displacements are confined to the range ± 0.5 pixels. This will reduce the computational cost appreciably. Gaussian interpolations (Eqs. 8, 9) are applied again to calculate the new fractional displacements s'_x and s'_y . Thus the fractional displacements are updated as $s_x = s_x + s'_x$ and $s_y = s_y + s'_y$. The above process is carried out iteratively until the newly calculated displacements satisfy the conditions $|s'_x| < \varepsilon$ and $|s'_y| < \varepsilon$, where ε is the pre-defined precision.

2.2

Fourier series reconstruction of PIV images

A further examination indicates that updating H_{ij} after shifting in Fourier space gives

$$H'_{ij} = F'_{ij}G'^{\star}_{ij} = c_{ij}(-s_y, -s_x)H_{ij}. \quad (15)$$

Taking the inverse Fourier transform of Eq. (15), we get the updated correlation plane

$$h'_{ij} = \sum_{m,n=0}^{N-1} H_{mn} \exp i2\pi \left[\frac{m(I + s_y)}{N} + \frac{n(J + s_x)}{N} \right], \quad (16)$$

where I, J are defined by Eq. (14). It is evident from Eq. (16) that the new algorithm is equivalent to using a Fourier series reconstruction (or the equivalent Whittaker's reconstruction) on the discrete correlation plane to find correlation peak locations to sub-pixel precision. It also shows that if the inverse Fourier transform were applied to Eqs. (11) and (12), it would result in re-sampled particle images shifted by fractional displacements, using the Fourier series reconstruction.

The proposed sub-pixel correlation peak location estimator is different from other traditional peak fitting techniques in that the peak location is a fitting parameter in the latter techniques while the proposed technique

determines the peak location iteratively. This new algorithm is theoretically similar to that of (Lourenco and Krothapalli 1995) but differs in that the prediction of the sub-pixel peak location is used to accelerate the reconstruction process thus an increased convergence rate is expected.

In order to illustrate the algorithm a simple 1-D Gaussian distributed particle image is considered, i.e.,

$$p(x) = \exp \frac{-(x - c)^2}{2\sigma^2}, \quad (17)$$

where c is the location of the particle center and σ is the characteristic size of the particle such that the e^{-2} diameter of the particle is $d=4\sigma$.

The modulus of the continuous Fourier transform of the Gaussian particle is

$$|P(\kappa)| = \sqrt{2\pi}\sigma \exp \frac{-(2\pi\kappa\sigma)^2}{2}. \quad (18)$$

If the resolution (inverse of the spatial sampling frequency) of the discretely sampled particle image satisfies the Nyquist sampling criterion, the original particle image can be fully reconstructed from the discrete samples. Since the spectrum of the Gaussian distribution is Gaussian, the modulus does not go to zero at infinitely high frequencies. Defining σ_k , ($\sigma_k = (2\pi\sigma)^{-1}$ for particle diameter $d=4\sigma$, as indicated by Eq. (18)) as the standard deviation of the Gaussian distribution in the spectral domain, energy beyond $K > 2\sigma_k$ accounts for only 4.5% of the total energy (same order of magnitude as the noise level in many real particle image recordings), thus the Nyquist sampling frequency of Gaussian particle images can be arbitrarily taken as

$$\kappa_N = 4\sigma_k = \frac{2}{\pi\sigma}. \quad (19)$$

If we assume that the transfer function of our CCD (or CMOS) sensor for a Gaussian particle is a delta-function (neglecting the filtering effect of real sensor geometries), and the pixel size is Δ , then the diameter of the particle image $d=4\sigma$ must be at least $\frac{8}{\pi}\Delta \approx 2.6\Delta$ in order not to lose information on the particle intensity distribution. For correlation based algorithms, the requirement is actually a bit less restrictive since the characteristic width of the correlation peak (resulting from the correlation of two or more Gaussian particles) is $\sqrt{2}$ times that of the particles, i.e., the critical particle diameter for correlation peak location determination is about 1.8Δ in order to resolve an unbiased fractional displacement. In the following analysis, we can assume $\Delta=1$ (unit length/pixel), such that $d/\Delta \equiv d$ (pixels), for simplicity.

Figure 1a shows the discrete sampling, reconstruction and fractional shifting of Gaussian particle images with different diameters, where the interrogation sub-window length is $N=16$ pixels, the particle location is set arbitrarily at $c=6.4$ pixels and is intended to be shifted by 3.3 pixels. For $d=4$ pixels, both its Fourier series reconstruction and the shifted profile agree very well with the expected values. For particle sizes smaller than this, under-sampling occurs, resulting in deviated reconstructions. Figure 1c

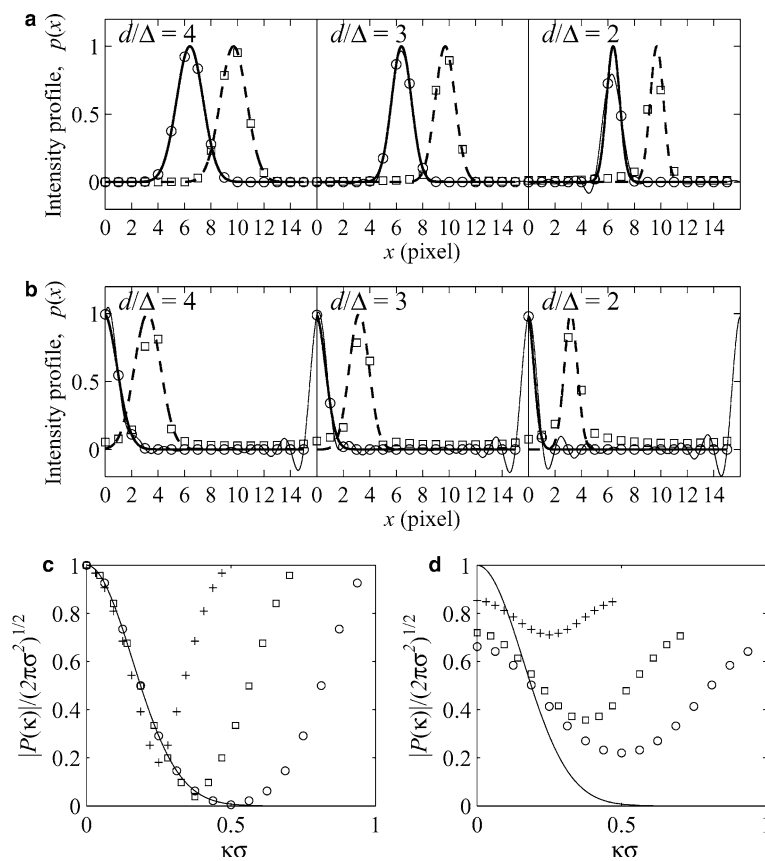


Fig. 1. Single particle image (1-D) with Gaussian intensity profile and its Fourier transforms, where $N=16$, d is the diameter of the particle and image intensities are normalized by the maximum particle image intensity. **a** *bold solid line* real particle image profile, centered at $c=6.4$ pixels; *circles* discrete sample of the particle image; *thin solid line* Fourier reconstruction by the discretely sampled particle image; *bold dashed line*, real particle image shifted by $s=3.3$ pixels; *squares* discrete particle image intensity reconstructed from the spectral shifting technique. **b** Same scheme as in **a**, while the particle image is truncated by the border ($c=-0.1$ pixel). **c**, **d** Modulus of the DFT of the sampled particle images in **a**, **b**, respectively, *solid line* continuous Fourier transform of Gaussian distribution; *circles* DFT of particle image profile with $d/\Delta=4$; *squares* $d/\Delta=3$, *crosses* $d/\Delta=2$

shows the modulus of the corresponding Fourier transforms of particle profiles and the aliasing effect due to under-sampling. In Fig. 1b, the particle is truncated by the sampling window (i.e., $c=-0.1$ pixel). In this case aliasing always occurs, even when the particle diameter satisfies the Nyquist sampling criterion (cf. Fig. 1d), since the truncation introduces spurious energy across the whole spectrum.

2.3 Truncation of particle images by the interrogation sub-window boundary

For the proposed spectral continuous window shifting algorithm, it is not possible to filter out the reconstruction errors due to particle truncations, while for other continuous window shifting techniques mentioned in the introduction, the problem is solved by interpolating the raw image using information outside the interrogation sub-window. We propose two methods to alleviate the particle truncation problem in the spectral continuous window shifting algorithm.

The first option is to ‘remove’ the particles that lie on the edge of the interrogation window. In order to do this, each particle image is identified. Particle identification sub-routines are common components of particle tracking velocimetry (PTV) algorithms. As an example, we employ a binary particle image identification technique with a threshold value based on the local ‘mode’ (most frequent value) of the image (Cowen and Monismith 1997) to generate a binary ‘mask’ image. The mask image consists of ones at locations occupied by a particle and zeros

elsewhere. For each interrogation sub-window, image pixels whose ‘mask’ values are one and also lie on the border of the interrogation sub-window and their adjacent pixels with ‘mask’ values of ones are considered as particle images truncated by the border. These particles are removed by setting their intensity values to the local ‘mode’ of the raw image. Lourenco and Krothapalli (2000) use a similar approach for the purpose of increasing the accuracy of PIV, but with a different particle image detection algorithm. It should be noted that these particle-finding methods are relatively simple and require little computational time compared to the more expensive correlation calculations.

Alternatively, a second approach to reducing the impact of sub-window truncated particle images is to apply a weighting function. In the present case the weight function is chosen as

$$W_{ij} = \left[1 - \left(\cos \frac{\pi i}{N-1} \right)^{20} \right] \left[1 - \left(\cos \frac{\pi j}{N-1} \right)^{20} \right]. \quad (20)$$

The profile of this weighting function with $N=32$ is shown in Fig. 2. The purpose of weighting is to reduce the contribution of the particles near the edges to the cross-correlation, thus reducing the effect of particle image truncation. Equation (20) is good for this purpose since the weight and the first order derivatives are identically zero on the edges. Other weighting functions, such as triangular, cosine, cosine squared or Gaussian functions, are also chosen to test. They perform as well as Eq. (20) does in terms of reducing peak-locking. However,




Fig. 2. Profile of the weight function defined by Eq. (20)

weighting Eq. (20) is better than others in that more information of the image is preserved (about 75% of the area on the sub-image is essentially unchanged by the weighting). Thus it results in a higher percentage of valid interrogations and lower random errors. Weighting functions bring additional benefits such as reducing the bias error that arises from correlation algorithms (Gui et al. 2000) and increasing measurement resolution in flows with large velocity gradients and small scale structures (Nogueira et al. 2001b). Figure 3 shows an example of the two methods applied to a sub-window containing synthetic particle images with imposed background noise level of 10% and $N=32$ pixels.

3 Simulation tests of the proposed algorithm with synthetic particle images

3.1 Generation of synthetic PIV recordings

In order to test the performance of the proposed algorithm and compare it to other established algorithms, synthetic particle image pairs were used to simulate a uniformly translating flow with displacements that vary from $S_x = -1.0$ to 1.0 pixel, and $S_y = 0$. Following Cowen and Monismith (1997), the particles in the simulations are modeled as 2-D Gaussian profiles and are distributed randomly (uniform probability distribution) within the synthetic images, which are 1024×1024 pixels in size. The particles in each image are identical in size and intensity.




Fig. 3. a Original particle images in a 32×32 interrogation sub-window. b Particles truncated by sub-window edges removed from sub-image a. c Sub-image a multiplied with the weighting function, Eq. (20)

The maximum intensity of each Gaussian particle profile is set to 255 counts.

To simulate the spatial integration nature of a digital camera, the intensity value of each pixel in the synthetic image is determined by the integration of the Gaussian surface of a particle over that pixel area and divided by the area of each pixel, thus the synthetic image is a box-filtered version of the true particle intensity distributions, i.e., a unit fill ratio, according to (Westerweel 1998). If at a given pixel position, two or more particle profiles overlap, the intensity value is chosen as the highest value calculated from the individual particle. The particle image density is on average 10 particles per 32×32 pixel sub-window. The most important simulation parameter is the particle size, which is chosen as $d=1, 2, 4$ and 8 pixels ($\sigma=0.25, 0.5, 1$ and 2 , correspondingly), reflecting the case of under-sampling, critical-sampling and over-sampling of the particle images. A second simulation parameter that is varied is the background noise level. The synthetic images either have no noise or are superimposed with a Gaussian noise with mean value 10% of the maximum particle image intensity (i.e., the distribution of noise is Gaussian with mean of 24 counts and RMS of 12 counts), representing a very noisy image acquisition system. The size of interrogation sub-window used in the simulation is either 32×32 or 64×64 pixels.

3.2

Selection of algorithms for simulation tests

The peak-locking effect can be analyzed by evaluating the distribution of measurement errors that have a bias part and a random part (Gui and Wereley 2002). In this paper, the bias error is defined as

$$E_b = \overline{M_x} - S_x, \quad (21)$$

where M_x is the measured horizontal displacement while S_x is the real displacement. The overline indicates an ensemble average, which is calculated over $30 \times 30 = 90$ measurements with 32×32 pixel sub-window size without overlap or with 64×64 pixel sub-window size with 50% overlap. The random error is taken as the root mean square (RMS) of the measured displacement

$$E_r = \sqrt{\overline{(M_x - \overline{M_x})^2}} \quad (22)$$

and the total error

$$E_t = \sqrt{E_b^2 + E_r^2}. \quad (23)$$

Six different sub-pixel peak location finding schemes are tested and compared in the simulations, they are

- A: Three-point 1-D Gaussian curve-fitting (traditional method without peak-locking reduction)
- B: Continuous image shifting with bi-linear particle image interpolation
- C: Continuous image shifting with bi-cubic spline particle image interpolation, with zero second derivatives on the boundaries (i.e., natural splines)
- D: The proposed algorithm (shifting image in Fourier space) without 'particle-truncation-by-border' correction

- E: The proposed algorithm with image corrected by weighting function, Eq. (20)
- F: The proposed algorithm with particles on image borders removed

3.3 Analysis of error distributions

Figures 4, 5, 6, 7, 8 give the resultant bias and random error distributions with respect to the specified fractional displacement for different particle sizes and background noise levels, using 32×32 pixel interrogation sub-windows. The pre-defined precision used to terminate the iterative scheme is $\varepsilon = 10^{-4}$ pixels.

As expected, all of the distributions are found to be anti-symmetric with respect to $S_x = 0.0$, thus only the non-negative part ($S_x = 0.0 \sim 1.0$) of the distributions are shown on these figures. For schemes A, D and E, the bias error distribution was expected to be anti-symmetric with respect to $S_x = 0.5$, where the random error should be maximum. However, the results show that the distributions are skewed. The deviation increases as particle size becomes larger. Further analysis indicates that the deviations are real and are due to particle truncations at image borders. In the presented simulations, the initial estimation of PIP displacement is set to be 0. When S_x is slightly higher than 0.5, particle truncation effect, which always tend to bias the measurement towards zero displacement, will possibly cause the calculated displacement to be less than 0.5. In such a case, further discrete sub-window offset is prevented, while it should happen if there were no truncation effects. For the same reason, if the initial estimation is set to be $S_x = 1$, the error distributions should also be skewed but in the opposite direction. This is shown in Fig. 6 (as comparison between A_1 and A_2). If scheme A is modified such that particles truncated by image borders are removed (as in scheme F, and denoted as scheme A'), the anti-symmetric/symmetric properties should be regained.

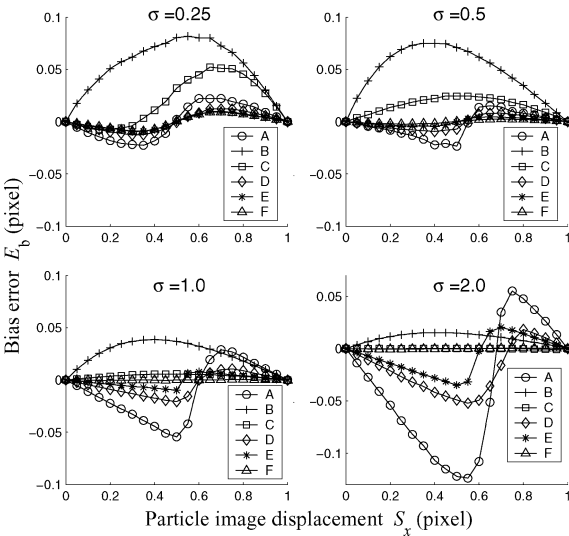


Fig. 4. Bias error distributions for different sub-pixel correlation peak location finding schemes applied to synthetic images without background noise

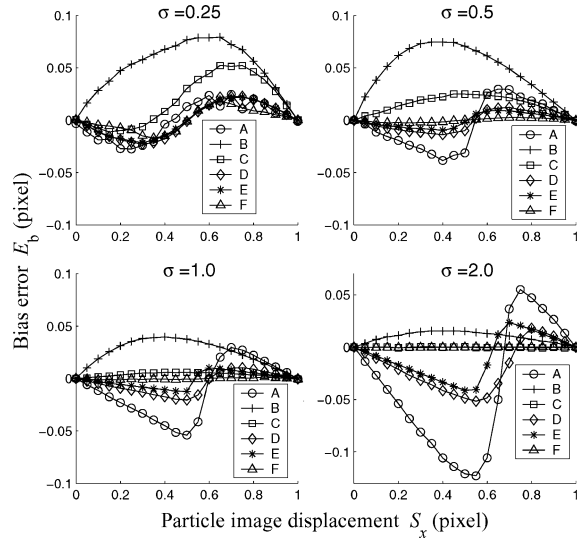


Fig. 5. Bias error distributions for different sub-pixel correlation peak location finding schemes applied to synthetic images with 10% Gaussian background noise

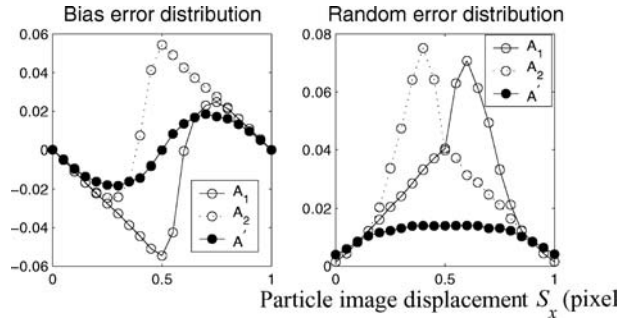


Fig. 6. Bias and random error distributions for the traditional method, scheme A, and the same scheme with particles on image borders removed, scheme A' ; Particle image size $\sigma = 1.0$ pixel, no background noise. A_1 : initial estimation $S_x = 0$, A_2 : initial estimation $S_x = 1$

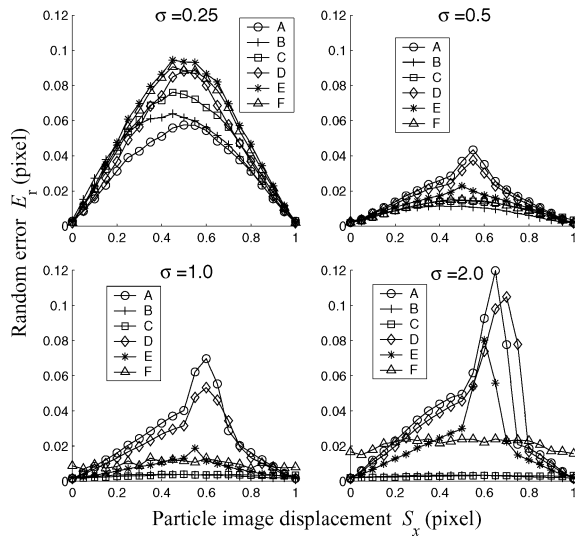


Fig. 7. Random error distributions for different sub-pixel correlation peak location finding schemes applied to synthetic images without background noise

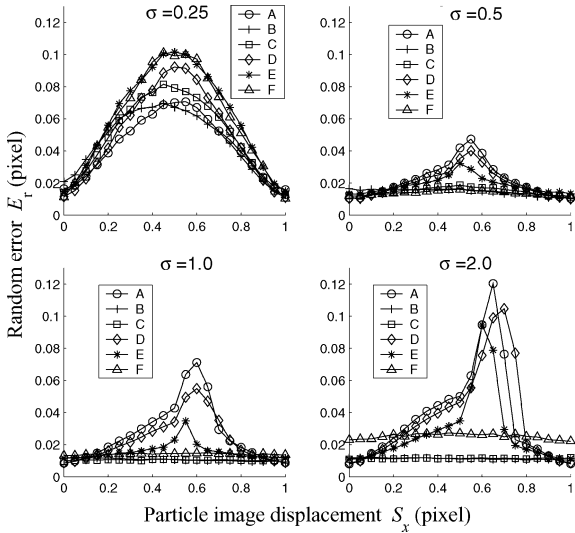


Fig. 8. Random error distributions for different sub-pixel correlation peak location finding schemes applied to synthetic images with 10% Gaussian background noise

This is also shown in Fig. 6. Additionally, both bias and random errors are reduced significantly.

For scheme B and C, the bias error distributions are neither symmetric nor anti-symmetric with respect to $S_x = 0.5$. This is a result of the central spatial finite difference scheme used in the present study. For example, if estimated displacement, $S_x = 0.2$ pixel (S_x is the sum of the discrete displacement and the sub-pixel displacement), the sub-windows in the first and second images are shifted by -0.1 and 0.1 pixels, respectively; however, if $S_x = 0.8$ pixel, sub-windows should be shifted by -0.4 and 0.4 pixels, respectively. Therefore, the magnitudes of errors for $S_x = 0.2$ and 0.8 are not necessarily the same. In other studies where continuous sub-window shifting scheme is adopted, such as that by Gui and Wereley (2002), a forward difference scheme is generally used which results in an anti-symmetry with respect to $S_x = 0.5$. It should also be noted that the gradient of the bias error distribution at $S_x = 0.0$ in schemes B and C is positive while for the other schemes it is negative (as long as there is a periodic variation in the bias distribution). According to Gui and Wereley (2002), this indicates that in most cases, correlation ‘peaks’ in the measured histogram of velocity will be ‘locked’ on mid-pixel displacement for schemes B and C and on integer pixel displacements for other schemes. This argument will be verified by the test on real PIV images shown later in section 4.

A comparison between Figs. 4, 5 indicates that the background noise level is not important as far as bias error is concerned. However, the existence of background noise slightly increases the random errors for all schemes, see Figs. 7, 8.

For the under-sampling condition ($\sigma = 0.25$), all methods suffer from peak-locking as indicated by the periodic variation of the bias error distributions in Figs. 4, 5. However, the bias error of the proposed algorithms (with or without particle truncation correction, schemes D, E and F) are much smaller than that of continuous window

shifting techniques (schemes B and C), whose bias errors are even larger than that of traditional peak-fitting techniques (scheme A).

The continuous window shifting schemes work better when particle size is relatively larger ($\sigma > 10$) and bi-cubic spline interpolation outperforms bi-linear interpolations since it has higher order accuracy. Unless particle size is very large ($\sigma \geq 0.2$), the bi-linear interpolation scheme does not bring any significant improvement compared to the traditional peak-fitting method.

Among the three proposed spectral domain continuous window shifting techniques, scheme F (particles truncated by borders removed) is the best as expected, while the weighting method (scheme E) performs only slightly better than scheme D (no particle truncation corrections). For the cases that the particle size $\sigma = 0.5$ and $\sigma = 1.0$ pixel, scheme F has the smallest bias error relative to all the other methods. Figure 9, in which absolute values of bias errors are plotted on a log scale, indicates that for the particle size $\sigma = 1.0$, scheme F is at least an order of magnitude more accurate than any other scheme. Similar results can also be observed for the case $\sigma = 0.5$. More importantly, the periodic bias error structure essentially disappears. At minimum the magnitude of this periodic variation is much smaller than the statistical convergence error due to limited number of ensembles and hence it is reasonable to conclude that the peak-locking effect has been eliminated.

When the particle size $\sigma = 2.0$ pixels, scheme C performs as well as scheme F at reducing the bias error. For the case of no background noise, the mean magnitude of its bias error distribution is 2.1×10^{-4} pixel, even lower than that of scheme F, which is 3.5×10^{-4} pixel. Moreover, it should be noted that for $\sigma = 2.0$, scheme F has trouble converging: only 43 and 55% of measurements converge successfully in the simulations without and with background noise, respectively. (For comparison, the interrogation success rate for the other simulation cases are above 98%.) The poor convergence is due to a significant fraction of particles being truncated by the interrogation sub-window boundary, which is not surprising as the ratio σ/N grows. The removal of these particles removed the majority of the particles in the sub-window, which increases the

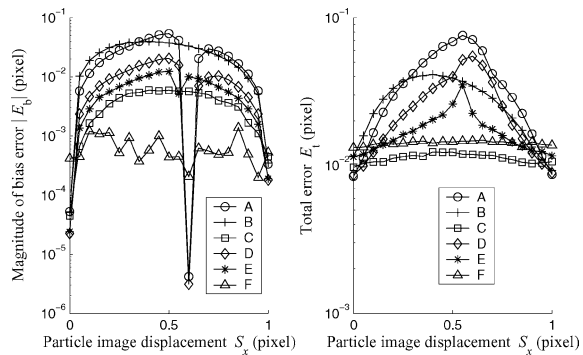


Fig. 9. Bias error and total error distributions of different sub-pixel peak locating schemes applied to synthetic images with 10% Gaussian noise. Particle size $\sigma = 1.0$ pixels, interrogation sub-window 32×32 pixels

uncertainty of the measurement. Obviously, this problem can be alleviated by using larger interrogation sub-windows. For example, by increasing sub-window size, N , to 64 pixels, the success rate of scheme F is increased to 94%, and the mean magnitude of the bias error is 5.0×10^{-4} (for scheme C, it is 3.0×10^{-4} pixel).

For all the schemes, random errors have minimum level when particle size is about $\sigma=0.5 \sim 1.0$, as shown in Figs. 7, 8. This result is consistent with that of (Westerweel 1998). In most cases, the random errors from schemes B and C are smaller than those from the proposed schemes (D, E and F), especially for larger particle sizes. As an example, see Fig. 9, which shows that scheme C outperforms scheme F in that the total error E_t of C is a factor of two smaller than that of F, although the bias error of C is an order of magnitude higher than that of F. This can be explained by the low-pass filtering nature of the image interpolation in schemes B and C. In the future, the combination of the proposed scheme F with an optimally selected low-pass filter should allow the random errors to be decreased accordingly.

3.4

Summary and discussions

Through the simulation tests on synthetic particle images, it has been demonstrated that how different factors affect the uncertainties of sub-pixel estimation of the PIP displacement. As shown in Fig. 4, when the particle size increases from $\sigma=0.25$ to $\sigma=0.5$, the magnitude of both the bias error and the random error decrease for the traditional peak finding method (scheme A). This variation reflects the effects of particle image under-sampling. Then, with the further increases of particle size, both of the errors for scheme A rise, indicating the increasing dominance of the particle truncation errors with the decrease of N/σ .

For small particle size (for instance, $\sigma=0.25$ or 0.5 in the simulations), the continuous window offset with image interpolation (schemes B and C) perform even worse than scheme A. This is explained by the under- and critical-sampling of the particle images, as image interpolation that includes low-pass filtering effect (such as the bi-linear and bi-cubic spline interpolation) is not appropriate for reducing the bias error. For over-sampled particle images, their ‘cut-off’ spatial frequency is lower (with respect to the sampling spatial frequency, i.e., the inverse of the image resolution), hence the low-pass filter does not negatively impact the accuracy of peak estimation significantly. Moreover, continuous window shifting with interpolation is very effective at reducing the bias errors due to the cropping of larger particles by the edges of the sub-windows, since the information outside of the sub-window is used for continuous offset.

As long as the particle images are sampled with frequencies above the Nyquist criterion, sub-window shifting in the spectral domain is the best method of reconstructing the sub-pixel particle image profiles; hence the best in terms of locating the sub-pixel PIP displacement. However, the particle truncation problem introduces strong aliasing into the spectra of the particle

images. Without removing truncated particles, this method (scheme D) is not as good as the image-space window offset technique (especially, scheme C) for small N/σ .

Applying weighting functions to sub-windows helps to reduce the truncation effect at sub-window borders, but truncation effects can not be completely eliminated. Although a weighting function can be used to set the image intensity at borders to zero, it necessarily increases the gradients of the profiles of particles near the borders, or wherever the signal attenuation is occurring. This has a similar effect to that of particle truncations, albeit not quite as severe. Additionally, in the present case, weighting functions are applied to both sub-windows prior to sub-pixel offset, which enhances bias towards integer pixels displacement.

In theory, the ideal solution to sub-pixel accuracy is scheme F. However, its limitations are also obvious: for PIV recordings with high particle density and larger particle size, the scheme causes higher probability of failed interrogations. Increasing the sub-window size can alleviate this problem, but the ‘sample volume’ is larger as well, which is not favorable for turbulent flow measurements. Also, the effectiveness of scheme F depends heavily on the robustness of particle identification, removal and refilling algorithms, especially in unfavorable imaging conditions such as non-uniform background noise and low signal-to-noise ratio.

3.5

Efficiency of the algorithm

Since all peak-locking reduction methods are iterative in nature, it is important to evaluate their computational load and rate of convergence, which are obviously related to the required precision. Generally, the calculation converges when the absolute value of the difference between results from two successive evaluations is smaller than a given precision. However, in peak-locking reduction algorithms, the ultimate goal is to set the image displacement as close as possible to zero. Hence the convergence condition is that the absolute values of sub-pixel displacements in both the x and y directions are smaller than the given precision value, ε , as described in section 1.

In the above simulations on synthetic images, the CPU time for each PIV image pair analysis was also recorded. Figure 10 shows the resultant computation times for the different schemes on the 10% noise added synthetic image pairs with particle size $\sigma=1.0$ pixel and particle displacements $S_x=0.5$ pixels. For this particular case, all schemes converge to a relatively high precision (ε ranges from 7×10^{-2} to 3×10^{-5} pixels). Computation time for all schemes increases linearly with the log of the required precision. However, the rate of increased cost for the proposed algorithms (scheme D, E and F) is much smaller than that of continuous window shifting techniques (schemes B and C). Schemes D and E increase computation time by less than 50%, compared to conventional PIV. Scheme F takes slightly more time. The increased overhead is largely due to computations in particle identification, and it is insignificant considering the great increase of accuracy that it brings.




Fig. 10. CPU time for 900 PIV interrogations with different peak-locking reduction methods normalized by the CPU time of the conventional PIV algorithm (scheme A), where interrogation sub-window length scale is $N=32$ pixels, particle size is $\sigma=1.0$ pixel, displacements are $S_x=0.5$ pixels, $S_y=0.0$ and images have 10% background noise

Similarly, the number of iterations needed increases linearly with $\log \epsilon$ (cf. Fig. 11). The proposed algorithms converge faster than image space based continuous window shifting schemes, by a factor of approximately 50% for this particular case. It should also be noted that although the CPU time needed for scheme C (bi-cubic interpolation) is more than scheme B (bi-linear interpolation), it converges faster and is much more accurate in all cases shown in this paper.

4

Algorithm test with real particle images

For experimentally collected PIV images it is more challenging to evaluate peak-locking reduction algorithms relative to the case of synthetic images, where the velocity field is known a priori. Usually the anti-aliasing abilities are evaluated through the probability distribution function (PDF) of the measured velocities, namely, to see if the PDF is smooth or if it is characterized by ‘peaks’ at the equivalent velocities of integer pixel displacements. This




Fig. 11. Average number of iterations for 900 PIV interrogations with different peak-locking reduction methods, where interrogation sub-window length scale is $N=32$ pixels, particle size is $\sigma=1.0$ pixel, displacements are $S_x=0$ pixels, $S_y=0.0$ and images have 10% background noise

method only gives qualitative verification as the true PDF is not known.

In the present paper, PIV images recorded in two high Re grid turbulence flows are selected as test cases. This canonical flow type is chosen as it is characterized by homogeneous (at a given downstream distance from the grid) and isotropic turbulence. The PDFs of all velocity components are usually observed to be Gaussian distributed (Batchelor 1953). Both experiments were conducted in the DeFrees Hydraulics Laboratory at Cornell University. Case one was carried out in an open channel turbulent water flow in a 2.0-m-wide circulating type flume running with 0.3 m water depth. The flow is driven by two centrifugal pumps and it passes through a series of passive grids before entering the test section. The detailed information about this facility and its flows can be found in (Liao and Cowen 2002). PIV images were collected on a horizontal plane 0.15 m above the bed in the free stream portion of the flow (where the thickness of the slowly growing boundary layer is about $\delta=9.0$ cm), with mean streamwise velocity $U_0=22.89$ cm/s and turbulence intensity is $\sqrt{u'^2}=1.08$ cm/s. The PIV images have spatial resolution of 1024×1024 pixels with a field-of-view of 52.3×52.3 cm. Spectral analysis from the PIV measurements indicate that the turbulence at this station is homogeneous and isotropic with Taylor’s micro-scale Reynolds number $R_\lambda=139$.

Case two is a measurement of grid turbulence in a wind tunnel with cross sectional area of 1.0×1.0 m. In this case, air is drawn in at the upstream end of a 20-m-long test section by a downstream exhaust fan. Air flows through an active grid consisting of an array of moving flaps that can enhance the turbulence level significantly (Mydlarski and Warhaft 1996). The mean streamwise speed at the measurement station, centered 5.5 m downstream from the grid, 0.50 m above the facility floor, and at the lateral midpoint, is $U_0=2.09$ m/s, the turbulence intensity is $\sqrt{u'^2}=0.19$ m/s and $R_\lambda=750$. The PIV images again have spatial resolution of 1024×1024 pixels with a field-of-view of 0.13×0.13 m.

The major purpose of using the two different flows is that they have different particle image size distributions. Spherical hollow glass spheres (manufactured by Potters Industries Inc.) with mean diameter $11.7 \mu\text{m}$ were used as the tracer particles in the first case while in case two the tracers were small water droplets produced by an air pressure driven water jet spray just downstream of the active grid. Examples of the resultant particle images are shown in Fig. 12, which shows that particle images in case one are relatively small and more uniformly distributed than in case two. Note that very large particles can be found in case two, while they are absent from case 1.

Using the particle detection algorithm described in section 3, individual particles and their intensity distributions can be identified thus more quantitative estimates of size distribution can be made. The diameter of a particle image is defined as

$$d_{PI} = 4\sigma = \frac{4(\sigma_x + \sigma_y)}{2}, \quad (24)$$




Fig. 12. Typical interrogation sub-window, with size $N=64$ pixels, taken from a PIV image from each of the two test cases

where σ_x and σ_y are the second central moments of the particle image intensity distributions in the x and y directions, respectively. Figure 13 shows the histograms of particle image size distribution (over 1,000 particles) for the two test cases. The average particle size in case one is about $\langle \sigma \rangle = 0.6$ pixel, which is just at the critical-sampling condition. In case two, $\langle \sigma \rangle = 0.9$ pixel and the distribution is much wider. In particular, there exists a significant number of very large particles with size $\sigma = 2 \sim 2.5$ pixels.

The aforementioned six schemes are applied to the two test cases. For both cases, the interrogation sub-window size is $N=64$ pixels with a 50% overlap. 200 image pairs are recorded in a period when the flow can be considered as statistically steady and homogeneous over the measurement area. Thus there are a total of $30 \times 30 \times 200 = 180,000$ PIV interrogations for each case, which are used to calculate the PDF of the velocity components.

Figures 14 and 15 are the resultant PDFs of the streamwise velocity component, u , for test cases one and two, respectively. A Gaussian PDF, with the same second moment as that of the measured PDF, is also plotted in the figures for comparison. The PDFs are plotted on a semi-log scale as it is easier to identify ‘peaks’ relative to a linear ordinate.

For both test cases, the traditional PIV peak locating algorithm (scheme A) results in serious peak-locking errors. Strong ‘peaks’ are observed at integer pixel displacements. In case one, the average particle size just satisfies the critical-sampling criteria, thus all three algorithms using spectral continuous window shifting (schemes D, E and F) result in ‘good’ PDFs in that the ‘peaks’ are visually gone and they agree well with the expected Gaussian distribution. Although scheme F is




Fig. 13. Histogram of the particle image length scale from each of the two test cases




Fig. 14. PDFs of the PIV measured streamwise velocity, u , with different anti-aliasing schemes for test case one, the wide open channel flow. Bold lines, measured PDF; thin lines, Gaussian distributions

expected to be superior to schemes D and E, it is hard to discern a difference in this case. Continuous window shifting in image space (schemes B and C) reduces the aliasing appreciably in this case, but ‘peaks’ are still




Fig. 15. PDFs of the PIV measured streamwise velocity, u , with different anti-aliasing schemes for test case two, the wind tunnel air flow. Bold lines measured PDF; thin lines Gaussian distributions

obvious in scheme B and apparent in scheme C under careful inspection. ‘Peaks’ are locked at mid-pixel displacement positions for the two cases. This observation agrees with the results from the simulations in section 3 where the gradient of the bias error distribution is positive at integer pixel displacements. Results from case one are comparable with the error distributions found in the simulation of particle size $\sigma=0.5$ shown in Fig. 5. For this particle size, the continuous window shifting in image space (especially with bi-linear interpolation) technique is less effective than the proposed algorithms at reducing peak-locking.

In case two, the average particle size becomes larger. Continuous window shifting in image space based algorithms perform as well as the spectral domain continuous window shifting methods. In this test case, all the anti-aliasing methods work effectively such that no ‘peaks’ are apparent in the PDFs. This result can also be associated with the results from the simulations where the particle size is $\sigma=1.0$ in Fig. 5.

5

Conclusions

Error sources that lead to peak-locking, or integer pixel aliasing effects, in PIV analysis have been reviewed. A new spectral domain image shifting technique is proposed to reduce the bias error in the determination of the fractional particle image pattern (PIP) displacement. The proposed new algorithm is based on the theory of discrete signal sampling thus it utilizes all of the information available to reach the highest sub-pixel accuracy. It is equivalent to a continuous window shifting technique in image space with interpolation based on Fourier series reconstruction of the discretely sampled particle images, but is more computationally efficient. Combined with the technique of removing particle images that are truncated by image borders, the proposed technique is demonstrated to eliminate peak-locking errors through tests with both synthetic and real PIV images, as long as the Nyquist particle image sampling criteria is satisfied. Compared with the image space continuous window shifting techniques, the new algorithm is much more computationally efficient and converges more rapidly. Also the proposed algorithm is more robust to smaller particle image sizes. For large particles the new algorithm still performs as well as the image space continuous window offset in minimizing the bias error, but the random error is larger. This should be solvable if the new algorithm is modified by applying an optimally chosen low pass filter on PIV images or the correlation plane.

References

- Batchelor GK (1953) *The theory of homogeneous turbulence*. Cambridge University Press
- Bracewell R (1998) *The Fourier transform and its applications*, 3rd edn. McGraw-Hill
- Cowen EA, Monismith SG (1997) A hybrid digital particle tracking velocimetry technique. *Exp Fluids* 22:199–211
- Fincham AM, Delerce G (2000) Advanced optimization of correlation imaging velocimetry algorithms. *Exp Fluids* 29:S13–S22
- Gui L, Merzkirch W (2000) A comparative study of the MQD method and several correlation-based PIV evaluation algorithms. *Exp Fluids* 28:36–44
- Gui L, Merzkirch W, Fei R (2000) A digital mask technique for reducing the bias error of the correlation-based PIV interrogation algorithm. *Exp Fluids* 29:30–35
- Gui L, Wereley ST (2002) A correlation-based continuous window-shift technique to reduce the peak-locking effect in digital PIV image evaluation. *Exp Fluids* 32:506–517
- Hart DP (2000) PIV error correction. *Exp Fluids* 29:13–22
- Huang HT, Fiedler HE, Wang JJ (1993) Limitation and improvement of PIV. ii. particle image distortion, a novel technique. *Exp Fluids* 15:263–273
- Liao Q, Cowen EA (2002) The information content of a scalar plume—a plume tracing perspective. *Environ Fluid Mech* 2:9–34
- Lourenco L, Krothapalli A (1995) On the accuracy of velocity and vorticity measurements with PIV. *Exp Fluids* 18:421–428
- Lourenco L, Krothapalli A (2000) True resolution PIV: A mesh-free second order accurate algorithm. In: *Proceedings of the 10th international symposium on applications of laser techniques to fluid mechanics*, Lisbon, Portugal, July 2000.
- Mydlarski L, Warhaft Z (1996) On the onset of high-Reynolds-number grid-generated wind tunnel turbulence. *J Fluid Mech* 320:331–368
- Nogueira J, Lecuona A, Rodriguez PA (2001a) Identification of a new source of peak locking, analysis and its removal in conventional and super-resolution PIV techniques. *Exp Fluids* 30:309–316
- Nogueira J, Lecuona A, Rodriguez PA (2001b) Local field correction PIV, implemented by means of simple algorithms, and multigrid versions. *Meas Sci Technol* 12:1911–1921
- Prasad AK, Adrian RJ, Landreth CC, Offutt PW (1992) Effect of resolution on the speed and accuracy of particle image velocimetry interrogation. *Exp Fluids* 13:105–116
- Raffel M, Kompenhans J (1994) Error analysis for PIV recording utilizing image shifting. In: *Proc 7th Int Symp Appl Laser Tech Fluid Mech*, page 35.5, Lisbon, Portugal, 14–16 July 1994
- Scarano F, Riethmuller ML (2000) Advances in iterative multigrid PIV image processing. *Exp Fluids*, Suppl: S51–S60
- Sholl MJ, Savas O (1997) A fast Lagrangian PIV method for the study of general high-gradient flows. In: *Aerospace Science Meeting & Exhibit*, 35th, Reno, NV, Jan 6–9, 1997. AIAA, pp 97–493
- Wereley ST, Meinhart CD (2001) Second-order accurate particle image velocimetry. *Exp Fluids* 31:258–268
- Westerweel J (1993) *Digital particle image velocimetry: theory and application*. PhD Thesis, Delft University
- Westerweel J (1998) Effect of sensor geometry on the performance of PIV interrogation. In: *The International Symposium on Applications of Laser Technique to Fluid Mechanics*, Lisbon, Portugal, 13–16 July 1998
- Westerweel J, Dabiri D, Gharib M (1997) The effect of a discrete window offset on the accuracy of cross-correlation analysis of digital PIV recording. *Exp Fluids* 23:20–28
- Willert CE, Gharib M (1991) Digital particle image velocimetry. *Exp Fluids*, 10:181–193