Accuracy Improvements in Particle Image Velocimetry Algorithms

Steven T. Wereley
School of Mechanical Engineering
Purdue University
West Lafayette, IN 47907-1288
wereley@purdue.edu

Carl D. Meinhart
Department of Mechanical and
Environmental Engineering
University of California
Santa Barbara, CA 93106
meinhart@engineering.ucsb.edu

ABSTRACT

An adaptive, second-order accurate particle image velocimetry (PIV) technique is presented. The technique uses two singly pulsed images which are interrogated using a modified cross correlation algorithm. Consequently, any of the equipment commonly available for conventional PIV (such as dual head Nd:YAG lasers, interline transfer CCD cameras, etc.) can be used with this more accurate algorithm. At the heart of the algorithm is a central difference approximation to the flow velocity (accurate to order Δt^2) versus the forward difference approximation (accurate to order Δt) common in PIV. In order to use the central difference approximation, an adaptive interrogation region shifting algorithm must be used. Adaptive shifting algorithms have been gaining popularity in recent years because they allow the spatial resolution of the PIV technique to be maximized. Adaptive shifting algorithms also have the virtue of helping to eliminate velocity bias errors. Consequently, since in many cases an image shifting algorithm would be used anyway, the second order accuracy of the central difference approximation can be had for very little additional cost.

The adaptive central difference interrogation (CDI) algorithm has two main advantages over adaptive forward difference interrogation (FDI) algorithms: it is more accurate, especially at larger time delays between exposures; and can maximize the signal-to-noise ratio of the velocity measurements. The technique is shown to perform better than conventional PIV algorithms when measuring near flow boundaries and in areas of low signal-to-noise ratio by comparing measurements of flow around a single red blood cell. Also, the technique is shown to be significantly more accurate than conventional PIV algorithms as the time delay between exposures is increased by interrogating cylindrical Couette flow images acquired both experimentally as well as computationally. The results of the interrogations are shown to agree quite well with analytical predictions and confirm that the CDI algorithm is indeed second order accurate while the conventional FDI algorithm is only first order accurate.

1. INTRODUCTION

Since the early days of particle image velocimetry, researchers have realized that PIV algorithms do not generally make maximal use of the information contained in the images they interrogate. In particular, Spedding and Rignot (1993) observed that the most appropriate spatial location for the displacement vector calculated by measurement techniques which track the motion of tracer particles (like PIV) is displaced from the centroid of the first interrogation region by a distance equal to half the displacement vector, not at the centroid of the first interrogation region (as is commonly implemented). Attributing the displacement measurement to its proper location is equivalent to a central difference approximation (order Δt^2) to the velocity field whereas the conventional implementation of PIV is equivalent to a forward difference approximation (order Δt). Unfortunately, moving the measured displacement vector from the centroid of the first interrogation region destroys the uniform spacing of the measured displacement data that most researchers prefer. Because other velocimetric techniques, such as particle tracking velocimetry and scalar image velocimetry, also produce non-uniformly spaced data, several researchers have already addressed how to map non-uniformly spaced data onto a uniform grid and what are the effects of doing so (Spedding and Rignot, 1993, Agui and Jimenez, 1987, Cohn and Koochesfahani, 1999). No matter how accurately the remapping process is performed, it will introduce some error and take some amount of time.

In this paper we will explore another approach, namely adaptively shifting the locations of the first and second interrogation regions such that when the calculated displacement vectors are located at their most proper locations (displaced from the centroids of the first interrogation regions by half the displacement distance), they will occur on a uniformly spaced grid, preserving the accuracy of the original calculations and eliminating the need for remapping. We will present an adaptive, iterative, second order accurate PIV algorithm which significantly improves spatial resolution and spatial accuracy while reducing velocity bias errors over conventional PIV algorithms. This algorithm was initially developed by Wereley, et al. (1999) and presented without demonstrating its second order accuracy. A theoretical, computational and experimental analysis of the algorithm is given. We will consider here two-dimensional PIV in which the images are recorded as pairs of singly-pulsed images on consecutive frames of some digital medium (CCD camera typically). Most of the ideas considered here are extensible to other forms of two-dimensional and three-dimensional PIV. Because the goal of this work is to design an algorithm which maximizes accuracy, not speed, the length of the interrogation region sides (in number of pixels) will not be limited to powers of two.

1.1. Spatial Resolution

The spatial resolution of any physical measuring system can be generally defined as the volume of space required to make a measurement. Similarly, the temporal resolution of a measurement system can be defined as the total amount of time over which the measurement is averaged. Using these definitions, the spatial resolution of cross-correlation PIV is determined by the size of the interrogation windows used to interrogate the images and the spatial shift between the interrogation windows. The temporal resolution is the time between each image exposure. Generally, the spatial resolution κ is a three-dimensional function of position that can be expressed as

$$\left\{\kappa_x, \kappa_y, \kappa_z\right\} = \left\{\min\left(W_{x1}, W_{x2}\right) + \left|\Delta X\right|, \min\left(W_{y1}, W_{y2}\right) + \left|\Delta Y\right|, \delta z_c\right\} \tag{1}$$

where W_{x1} and W_{y1} represent the size of the first interrogation window and W_{x2} and W_{y2} represent the size of the second interrogation window. The spatial shift between first and second interrogation windows in the x direction is represented by ΔX and in the y direction by ΔY. The quantity δz_c represents the depth of the measurement domain in the out of plane direction. These parameters can be constants and express the overall spatial resolution for the entire image or they can vary from measurement point to measurement point and produce a spatial resolution which varies across the region of interest.

It is possible to eliminate one of these constraints entirely and make the experiment optimization process much simpler. The typical constraint that the difference between the particle displacement and the spatial offset between first and second interrogation regions not exceed 1/3 of the interrogation region size can be eliminated by always using a spatial shift equal to the local flow velocity. Keane and Adrian (1993) were the first to suggest such an approach and performed the theoretical calculations to show that this method would be superior to other methods. Urushihara, Meinhart, and Adrian (1993) first implemented a simple, flexible spatial shifting technique. The second order accurate adaptive technique presented here is similar in nature to the first order accurate technique of Cowen and Monismith

(1997), who implement a multi-pass adaptive spatial shift as part of their super resolution particle tracking routine. At every measurement point they started with an initial estimate of the velocity at that point which was used as the spatial shift. The resulting velocity measurement was then used in the next iteration as the spatial shift. Measurement points at which this procedure did not converge to a constant spatial shift within three iterations were flagged as invalid.

1.2. Spatial Accuracy

The issue of spatial accuracy is related to that of spatial resolution. Larger interrogation windows and larger spatial shifts will clearly produce more coarse estimations of the local or instantaneous velocity field. However, simply minimizing the size of the interrogation windows or the distance the particles move between the two images is not the solution because a certain minimum number of particle images are required to produce reliable measurements in any given experimental situation. In addition, those particles must move a sufficient distance to reduce the relative error in determining the location of the correlation peak to acceptable levels. One solution is to use a higher-order approximation velocity calculation. To take advantage of the existing specialized equipment for acquiring image pairs—double pulse lasers and interline transfer cameras—it is necessary to limit the data acquisition to pairs of images. Most PIV algorithms use a simple forward differencing interrogation (FDI) scheme in which the velocity at time t is calculated using particle images recorded at times t and t+ Δt, given by

$$\left.\frac{d\mathbf{X}}{d\tau}\right|_{\tau=t} = \frac{\mathbf{X}(t+\Delta t)-\mathbf{X}(t)}{\Delta t} + \frac{\Delta t}{2}\left.\frac{d^2\mathbf{X}}{d\tau^2}\right|_{\tau=t} + \ldots \tag{2}$$

The accuracy of the velocity calculation can be greatly improved by using a relatively simple central differencing interrogation (CDI) scheme in which the velocity at time t is calculated using particle images at times t-Δt/2 and t+ Δt /2, given by

$$\left.\frac{d\mathbf{X}}{d\tau}\right|_{\tau=t} = \frac{\mathbf{X}(t+\Delta t/2)-\mathbf{X}(t-\Delta t/2)}{\Delta t} + \frac{(\Delta t)^2}{24}\left.\frac{d^3\mathbf{X}}{d\tau^3}\right|_{\tau=t} + \ldots \tag{3}$$

The velocity approximation is accurate to order Δt for the forward difference formulation while the central difference formulation is accurate to order Δt^2 using the same two images. This improved accuracy can correspond to better accuracy for a given time delay or allow increasing the time delay between images while maintaining a constant accuracy. The relative contribution of other errors, such as Brownian motion or the uncertainty in finding the correlation peak location, to the overall error can then be reduced.

1.3. Velocity Bias

One area where measurement errors can be significantly reduced at little additional computational expense is velocity bias. PIV measurements made with equally-sized interrogation windows in the first and second images are generally biased towards the spatial shift of the interrogation windows. Westerweel (1997) demonstrates that this error is on the order of 0.1 pixels for common PIV parameters—32×32 pixel interrogation windows, particle diameters in the range of 2 to 4 pixels, 10 particle images per interrogation window, and displacements ranging from 0 to 10 pixels. Westerweel also demonstrates that this error decays linearly to zero as the particle displacement decreases from 0.5 pixels to zero pixels. While this error may seem relatively insignificant, it can seriously degrade the calculation of derivative quantities such as vorticity and shear stress. Even though the velocity bias error becomes very small for displacements smaller than 0.5 pixels, Huang, et al. (1997) demonstrate that simply spatial shifting identically-sized interrogation windows is insufficient to fully eliminate velocity bias. The velocity bias can be eliminated by several techniques. Westerweel (1993) demonstrates that dividing the image correlation function by the correlation of the windowing functions, commonly denoted F_I, will completely eliminate the velocity bias. Recently, Huang, et al. (1997) demonstrated an efficient algorithm that eliminates velocity bias errors by renormalizing the values of the correlation function in the vicinity of the peak location prior to calculating the subpixel peak location.

Keane and Adrian (1993) suggest that bias errors can be nearly eliminated by making the second interrogation window sufficiently larger than the first so that all particle images in all first interrogation windows will likely be contained in the second. They also propose that spatially-shifting the second interrogation window by the integer part of the displacement will substantially reduce this error. In a further analysis of Keane and Adrian's theory, Westerweel

(1997) points out that using interrogation windows of differing sizes results in a constant-valued portion at the center of the window correlation function F_I, corresponding to small displacements, which eliminates the velocity bias. Since our adaptive procedure uses spatial shifts that very accurately reflect the velocity of the particles in the interrogation windows, the approach suggested by Keane and Adrian can be implemented with minimal computational expense by making one of the interrogation windows as little as 2 pixels larger than the other in each direction. Because the application of our technique is measuring microscale flows in which Brownian motion of the seed particles can be a serious factor, we typically make the second interrogation window larger than the first. The size of the second interrogation window is increased by the number of pixels that particle diffusion and spatial gradients are expected to affect the flow plus a few additional pixels, typically two, in each direction to produce the constant-valued region at the center of the window correlation function.

1.4. Ensemble-Averaged Correlation Function

For both the experimental and simulated flows presented here, an ensemble of image pairs is acquired or generated which can be interrogated at higher spatial resolution than single pairs of images. Consequently the effect of the spatial averaging inherent in PIV measurements is reduced, allowing the accuracy of the central differencing algorithm to be fully explored. The standard method for determining ensemble-averaged flow fields is to obtain a series of instantaneous velocity fields and then ensemble average them. When the signal-to-noise ratio of the images is low, which can happen for a variety of reasons, such as faint seed particle images, high background noise, and Brownian motion of seed particles, this procedure can produce many spurious vectors, called 'outliers', due to the peak search algorithm mistakenly identifying a noise peak as the signal peak. These outliers can be located anywhere in the correlation plane and are completely uncorrelated with the signal peak location. Consequently they can cause significant noise. Often these vectors can be eliminated from the ensemble by using statistical arguments. Even so, it is preferable for these outliers not to be created in the first place.

Recently, Meinhart, Wereley, and Santiago (1999) demonstrated a novel ensemble averaging technique for improving the signal-to-noise ratio of steady-state velocity measurements. Instead of ensemble averaging after the calculation of the series of velocity fields, at every measurement point they ensemble averaged the correlation function before the peak search algorithm was performed. By ensemble averaging in this way, the signal peaks add constructively while the randomly distributed noise peaks do not, making the signal peak more distinguishable from the noise peaks. This procedure computes an ensemble averaged velocity field without first calculating the instantaneous velocity field. One of the primary advantages of ensemble averaging in this manner is that the effective particle image concentration is increased, which allows for reducing the size of the interrogation windows. In the experiments and simulations described in Section 4, 26 image pairs were acquired or generated which allowed the spatial resolution to be increased to the point at which 20 statistically-independent measurements can be made across the gap of a cylindrical Couette flow system.

2. CENTRAL DIFFERENCE INTERROGATION FUNDAMENTALS

A central difference interrogation (CDI) scheme has many advantages over the conventional forward difference interrogation (FDI) scheme. The most obvious advantage already mentioned is that the technique is second order accurate in space and time, providing increased accuracy over the forward difference scheme for the same time delay or, alternately, the same accuracy as the forward difference scheme for a significantly longer time delay. In many situations it is advantageous to use a larger time delay. For example, in microscale flows the diffusion of submicron seed particles may be a significant source of error for microscale velocity measurements (Santiago, et al., 1998). The error associated with Brownian motion ε_B was shown to be

$$\varepsilon_B = \frac{1}{u}\sqrt{\frac{2D}{\Delta t}} \tag{4}$$

where u is the local velocity, D is the diffusion coefficient of the particles, and Δt is the time delay. Consequently, the relative error due to Brownian motion can be reduced by increasing the time delay.

Another major advantage of the CDI algorithm is that it allows better velocity field calculations in the vicinity of flow boundaries. There are two reasons for this improvement. First, since an adaptive technique is required to implement CDI, the technique will automatically reduce the spatial shift between interrogation windows in areas where the

particle displacements are small, i.e. near boundaries. This increased spatial resolution can be especially important near curved surfaces and stagnation points where conventional algorithms will place interrogation windows inside boundaries and generate spurious results. Another reason why the CDI algorithm provides better results near boundaries is that the measurements are symmetric in space about the boundary instead of biased toward the wall location in the first image. These advantages are demonstrated below.

Figures 1a and 1b show velocity fields of a flow around a human red blood cell calculated by a conventional constant spatial shift FDI scheme and also by the adaptive spatial shift CDI scheme. The blood cell is trapped between a coverslip and a slide glass and surrounded by deionized, particle seeded water (300 nm PSL, 10^{-4} volume concentration). The flow is driven from lower left to upper right around the blood cell using the surface tension of a water droplet placed at the edge of the coverslip to establish a flow past the red blood cell. The images are interrogated with 28×28 pixel interrogation windows that are overlapped 75%. While this overlap actually oversamples the images, it functions as a very good interpolator of the velocity field. For both techniques, measurement points that resulted in more than 20% of the combined area of the first and second interrogation windows inside of the blood cell are eliminated and replaced with an '×' symbol because they will tend to produce velocity measurements with serious errors. The adaptive CDI scheme is able to accurately measure velocities much closer to the surface of the cell than is the FDI scheme. There is essentially a single layer of invalid measurements surrounding the blood cell for the CDI measurements, while the FDI scheme has three rows of invalid measurements upstream of the cell and a single layer around the remainder of the cell. The CDI scheme has 55 invalid measurement points, equivalent to 59.4 μm^2 of image area that cannot be interrogated, while the forward difference has 63 invalid measurement points or 70.9 μm^2. The forward difference scheme has 19% more area where vectors cannot be calculated compared to the CDI scheme. In terms of the blood cell area, which is 55.7 μm^2, the area in which the CDI scheme cannot calculate vectors is only 7% larger than the cell itself while the area in which the FDI scheme cannot calculate vectors is 27% larger than the blood cell. Consequently the CDI algorithm resolves the variations in the velocity field near the cell surface better than the conventional forward difference scheme because the spatial shift between interrogation windows is based on the local velocity and not on the global mean velocity.

Another feature of the CDI scheme is that the invalid measurements are symmetric about the blood cell, whereas the conventional scheme has more invalid measurements upstream from the blood cell than downstream from it. This nonuniform representation of the location of the blood cell makes calculating surface pressure or shear stress difficult

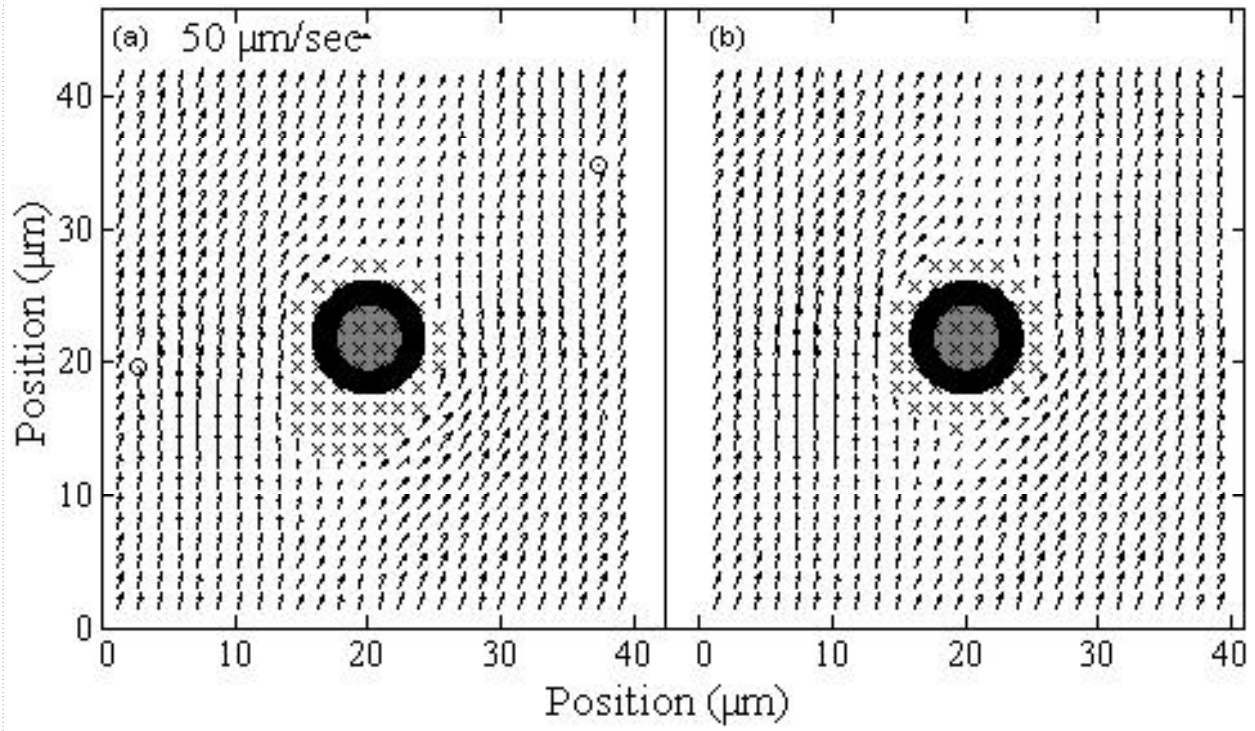

Figure 1. Comparison of a forward difference, fixed window offset analysis (a) versus a central difference, adaptive window offset analysis (b) of the flow around a single red blood cell. This figure demonstrates quantitatively the advantage of using an adaptive, second order analysis scheme.

and inaccurate using conventional forward difference measurements. In addition, the adaptive nature of the CDI algorithm allows the signal-to-noise ratio to be maximized for a fixed interrogation size. For PIV the signal-to-noise ratio is commonly defined as the ratio of the maximum peak amplitude (signal) to the second peak amplitude (largest noise peak). The advantage of the CDI algorithm can be seen by averaging the signal-to-noise ratio for all measurement points outside the blood cell. The signal-to-noise ratio is 2.70 for the CDI algorithm and 2.51 for the conventional FDI algorithm, which is equivalent to an increase in the signal-to-noise ratio of 7.6% for the CDI algorithm. In Figure 1, the circle symbols (o) represent measurement points at which no measurement could be made because of insufficient signal strength. There are two such points for the forward difference algorithm but none for the CDI algorithm.

3. CENTRAL DIFFERENCE INTERROGATION ALGORITHM

The forward difference scheme (conventional PIV) can be easily implemented using a constant spatial shift in a single iteration, while the central difference scheme requires a more complicated algorithm. In general, the goal of PIV is to measure a velocity field that contains spatial gradients. These spatial gradients translate into particles moving at different velocities as a function of spatial position. In order to define the measurement time t as halfway between t-Δt/2 and t+Δt/2, the spatial shift between interrogation windows in the first and second image must be equal to the distance the particles move between images. In practice a uniformly-spaced grid of measurement points is defined at the time t and the interrogation windows in the first image are shifted by an amount $-\vec{\mathbf{V}}(x,y)\cdot\Delta t/2$ while the interrogation windows in the second image are shifted by an amount $+\vec{\mathbf{V}}(x,y)\cdot\Delta t/2$, netting a spatial shift of $\vec{\mathbf{V}}(x,y)\cdot\Delta t$ centered between the particle locations in the first and second images. Clearly, employing this variable spatial shift requires an advanced knowledge of the velocity field. An iterative adaptive approach is required in which increasingly accurate estimates of the velocity field are used to compute the spatial shift, in turn yielding increasingly accurate velocity field estimates.

The first step in this procedure is to obtain a low spatial resolution measurement of the entire velocity field using a constant spatial shift between all first and second interrogation windows at every measurement point. This coarse measurement provides an initial estimate of the velocity field. In the case of the Couette flow measurements presented in Section 4, 64 pixel square interrogation windows with a spatial shift between first and second interrogation windows equal to the estimated mean displacement of the particles in the images was used. No overlap between neighboring windows was used. After the velocity field was calculated, the vectors were validated using CleanVec, a validation program written at University of Illinois by Carl Meinhart and Steve Soloff. CleanVec compares a candidate measurement with the median of its neighbors' values to determine whether it is more likely a valid measurement or an invalid measurement attributable to noise. If a vector differs from its neighbors too significantly, by an amount specified by the user, it will be replaced by either the second choice or third choice correlation peak. If neither of those peaks matches the neighborhood any better, a vector is calculated by interpolating from the neighbors' values. Gaussian smoothing is used to remove noise in the displacement field.

Using the initial validated velocity field measurement, a grid of new measurement points is chosen for subsequent high spatial resolution interrogation of the images. A bicubic interpolation scheme is used to resample the low resolution velocity field at each new high resolution measurement point. When using a forward difference formula to approximate the velocity, each measurement point is located at the centroid of each first interrogation window. In the central difference formulation, however, the measurement point is located halfway between the centroids of the first and second interrogation windows. At each measurement point in our algorithm, the centroid of the first interrogation window is displaced from the location of the measurement point by *round*(-0.5*$\Delta\mathbf{X}$) while the centroid of the second interrogation window is displaced from the first by *round*($\Delta\mathbf{X}$), where *round* determines the nearest integer and $\Delta\mathbf{X}$ is the local displacement. The grid of measurement points is chosen so that the maximum area of the images is interrogated without any of the interrogation windows going beyond the edges of the images. The new high resolution displacement field is then calculated using the interpolated displacement field to determine the spatial shifting of the interrogation windows. After each iteration, the vectors are checked for convergence to see if the velocity field is progressing to a value that does not change between iterations. If another iteration is needed, the data are validated again. If not, the data are usually of high enough quality that further post-processing is not necessary.

4. APPLICATION TO CYLINDRICAL COUETTE FLOW

In order to study the relative merits of different particle image velocimetry algorithms, it is necessary to compare the results from interrogating a known velocity field using each of the different schemes. The scheme producing results that most closely match the known velocity field with the least amount of work will be the best one. Two adaptive image shifting techniques will be compared here. The first is a conventional forward difference scheme and the second is our new central difference scheme. These two adaptive techniques both require only two images to compute but do require additional computational effort when compared with constant shift techniques. Still, in most circumstances the adaptive image shifting technique allows the spatial resolution of the measurements to be maximized while the bias errors are reduced. Cylindrical Couette flow was chosen to compare the two PIV algorithms because it is a relatively simple flow in which velocity measurements can be compared to Monte Carlo simulated images and analytical calculations to determine the relative accuracy of each technique.

Cylindrical Couette flow is produced between two differentially rotating concentric cylinders. The steady state velocity field is nearly one-dimensional far from the ends of the apparatus and obeys the simple and well known solution to the Navier-Stokes equations

$$\begin{aligned} v_r &= v_z = 0 \\ v_\theta(r) &= r \cdot \omega(r) = A_0 \cdot r + B_0 / r \end{aligned} \tag{5}$$

where v_r, v_θ, and v_z are the radial, azimuthal, and axial velocity components, respectively, r is the radial position, ω is the angular velocity, and the no-slip boundary conditions at the inner and outer cylinder uniquely determine the constants A_0 and B_0. Since the velocity field is known, a flow can be simulated by evolving randomly-distributed particles forward in time using the velocity field in Equation (5). Wereley and Lueptow (1994, 1998), among others, have shown that experimental measurements of this flow very closely approach the analytical solution, making it a good candidate for comparing simulated flows to experimental results.

4.1. Analytical Calculations

4.1.1. Forward Difference Error Evaluation

Because PIV is usually performed on a Cartesian grid while cylindrical Couette flow exhibits azimuthal symmetry, notation can become somewhat confusing. To bridge this notational gap, we will consider a fluid element located at the Cartesian coordinates (X_0,Y_0) which coincide with the polar coordinates (r_0, θ_0) (see Figure 2a). The angular velocity of the fluid element at a position r_0 is given by $\omega(r_0)$. This fluid element will be swept through an arc of length $r_0\omega(r_0)\Delta t$ during the time Δt to the new position (X_1,Y_1) which coincides with (r_0,θ_1). Because the velocity field is independent of azimuthal position, we can, without loss of generality, set $\theta_0=0$ for the remainder of the analysis to simplify the arithmetic. The standard first order accurate forward difference discretization of the velocity field is

$$\vec{\mathbf{V}}_{meas} = \frac{\vec{\mathbf{X}}(t+\Delta t) - \vec{\mathbf{X}}(t)}{\Delta t} \tag{6}$$

Since the fluid elements trace out curved paths, and this discretization is a two point, linear approximation to the derivative, there will necessarily be some discrepancy between the measurement and the actual velocity. In order to compare the two PIV algorithms, we will quantify this error for both schemes. The Cartesian components of the velocity field can be written as

$$V_{x,meas} = \frac{X_1 - X_0}{\Delta t} = \frac{r_0 \cos\{\omega(r)\Delta t\} - r_0}{\Delta t} \tag{7a}$$

$$V_{y,meas} = \frac{Y_1 - Y_0}{\Delta t} = \frac{r_0 \sin\{\omega(r)\Delta t\}}{\Delta t} \tag{7b}$$

The actual instantaneous velocity at the point (X_0,Y_0) can be written as

$$V_{x,act} = 0$$
$$V_{y,act} = r_0\omega(r_0) \tag{8}$$

The difference between these two velocities will tell us how much of an effect the limitations of the PIV technique will have on the overall accuracy of the measurements. To simplify the resulting expression, the trigonometric functions will be expanded as power series. We have (after significant algebraic manipulation)

$$\Delta V_x = -\frac{1}{2} r_0 \Delta t \cdot \omega(r_0)^2 + \frac{1}{24} r_0 \Delta t^3 \omega(r_0)^4 + \ldots$$
$$\Delta V_y = -\frac{1}{6} r_0 \Delta t^2 \omega(r_0)^3 + \frac{1}{120} r_0 \Delta t^4 \omega(r_0)^5 + \ldots \tag{9}$$

The error in X (or radial) component is proportional to Δt while the error in the Y (or azimuthal) component is proportional to Δt^2. In a general flow with no particular symmetry, the error in both components should be order Δt, but the azimuthal symmetry in this flow increases the accuracy of the azimuthal component to order Δt^2. For this example of measuring Couette flow, it is not necessary to measure the radial velocity component since it should be zero (unless to verify that the experiment is properly aligned). However in cases where higher-order flow states are being measured, such as Taylor vortex flow and wavy Taylor vortex flow, where the radial velocity component is of interest and is much smaller than the azimuthal component, this first order accuracy in Δt can be restrictive.

4.1.2. Central Difference Error Evaluation

The accuracy of the PIV algorithm can be significantly improved by assuming that the measurement point is halfway between the centroids of the two interrogation windows. This assumption is equivalent to a central difference discretization of the velocity derivative which is second-order accurate. An iterative approach must be used to define a uniformly-spaced grid of measurement points. First, a coarse interrogation of the velocity field is made using a standard forward difference interrogation technique. Then this velocity field is resampled and the locations of the interrogation windows are adjusted so that the measurement points are again distributed on a uniform grid. This procedure can be repeated as often as necessary, but usually two iterations are sufficient.

Although the CDI technique is second-order accurate in Δt, the accuracy of the technique needs to be examined to quantify its advantages. Toward this end, we will consider a fluid element, initially at $(X_{-1/2},Y_{-1/2})$, corresponding to $(r_0,\theta_{-1/2})$, which moves with the fluid surrounding it to $(X_{+1/2},Y_{+1/2})$, corresponding to $(r_0,\theta_{+1/2})$ (see Figure 2b).

(a) (b)

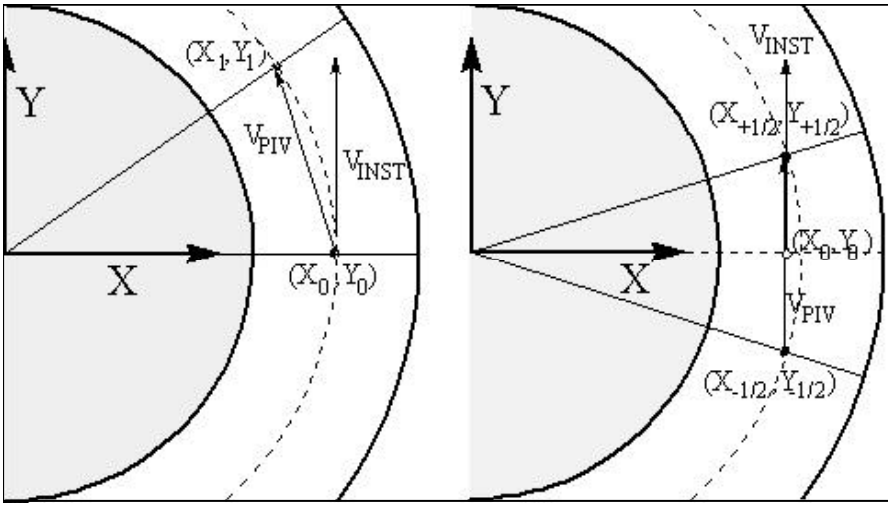

Figure 2. Schematic views illustrating the (a) forward difference interrogation algorithm and the (b) central difference interrogation algorithm as applied to the cylindrical Couette flow system

Consequently, the measurement point at (X_0, Y_0) is no longer on the path followed by the fluid element. The components of the velocity field can be written as

$$V_{x,meas} = \frac{X_{+1/2} - X_{-1/2}}{\Delta t} = 0 \text{ and} \tag{10}$$

$$V_{y,meas} = \frac{Y_{+1/2} - Y_{-1/2}}{\Delta t} = \frac{2r_0 \sin\left\{\omega(r_0)\frac{\Delta t}{2}\right\}}{\Delta t} \tag{11}$$

The actual instantaneous velocity at the point (X_0, Y_0) can be written as

$$\begin{aligned} V_{x,act} &= 0 \\ V_{y,act} &= r_0 \cos\left\{\omega(r_0)\frac{\Delta t}{2}\right\}\omega\left(r_0 \cos\left\{\omega(r_0)\frac{\Delta t}{2}\right\}\right) \end{aligned} \tag{12}$$

The difference between the actual velocity and measured velocity can be better assessed after expanding both solutions as a power series and combining terms of the same order in Δt

$$\begin{aligned} \Delta V_x &= 0 \\ \Delta V_y &= r_0\omega(r_0)^2\Delta t^2\left\{\frac{1}{12}\omega(r_0) + \frac{1}{8}r_0\omega'(r_0)\right\} + \ldots \end{aligned} \tag{13}$$

Using the central difference formulation for the velocity approximation, we would expect the error in both velocity components to be on the order of Δt^2. However, there is *no* error in the radial velocity component due to the azimuthal symmetry of the flow. The error in the azimuthal component is order Δt^2 as expected.

4.2. Monte Carlo Simulations

One method for evaluating the accuracy of different PIV algorithms is to process a set of simulated images containing a known velocity field. A good alternative to performing painstaking experiments (which will naturally have errors and inaccuracies associated with them) is simulating particle-seeded flows for a variety of operating parameters. This approach was used to assess the importance of operating parameters in the PIV process by Guezennec and Kiritsis (1990) and Keane and Adrian (1990) and many subsequent studies. Image pairs can be simulated by creating a randomly distributed particle field (hence the use of the term 'Monte Carlo') in the flow domain and then evolving the positions of those particles forward in time according to the velocity equation (5) above. Based on the observation of many real particle images, the particle image intensity can be modeled as a gaussian curve. Because CCD cameras are composed of many small pixels, the gaussian distribution of image intensities should be averaged over the area of each pixel as demonstrated by Huang, et al. (1997). While this step is very important when investigating the effect of subpixel resolution interpolation routines, it is not important in the current investigation because the imprecision in the particle image shape will be constant among the algorithms being evaluated.

In order to facilitate direct comparisons between the simulations and the experiments, the parameters of the simulations were closely matched to the experimental parameters. Even though the important parameters were matched, the simulated images were idealized in that all the particle images had the same diameter and the same intensity. All of the simulated particle images were set to the same maximal intensity (107) that matched the mean particle brightness above the background level in the experiments. The background level of the simulations was set to be zero since the background only contributes to a DC component in the correlation plane. The mean particle image diameter (3.1 pixels) and the mean particle image concentration (3.3 particle images in a 16×16 pixel interrogation window) were also matched to typical experimental conditions.

4.3. Experimental Measurements

The flow system consisted of two concentric acrylic cylinders, a rotating inner cylinder and a stationary outer cylinder with nominal radii r_i=4.24 cm and r_o=5.23 cm, respectively. The resulting gap width was $d=r_o-r_i$=0.99 cm, and the radius ratio was $\eta= r_i/r_o$=0.81. The inner cylinder was driven by a stepper motor at a rate of 0.106 Hz, below the transition to Taylor vortex flow. The stepper motor allowed the rotational speed to be precisely controlled. The actual rotation rate varied by ± 0.001 Hz which was monitored by an optical encoder having a resolution of 300 pulses per revolution. The two cylinders were held concentric by precision bearings mounted in opposing aluminum endcaps. The ratio of the length of the annulus to the gap width was Γ=47.8. The particles were illuminated by a laser light sheet formed in a plane perpendicular to the common axis of the cylinders and imaged through an acrylic window mounted in the lower endcap. The working fluid was a water-glycerine mixture, 40% by weight, with density 1.1 $g \cdot cm^{-3}$ and viscosity 3.5 cSt at 21.5°C. The flow was seeded with 16 μm silvered hollow glass spheres having a density of 1.6 $g \cdot cm^{-3}$. The particles remained suspended due to their small size and have been shown to accurately follow the flow (Wereley and Lueptow, 1998).

Two 25 mJ per pulse Nd:YAG lasers operating at 20 Hz were used to illuminate the flow at two time instants. The laser pulses were separated by 7.5, 15, 30, and 60 milliseconds, depending on the run, to study the dependence of the measurement error on the time step. The laser sheet had a thickness of approximately 1 mm, with ±0.5% variation in thickness in the measurement field. It was carefully aligned in a radial-azimuthal plane to avoid crosstalk between velocity components. Pairs of images each containing particle images from a single laser flash were obtained using a TSI cross-correlation CCD camera, positioned perpendicular to the laser sheet. A Computar telecentric lens (model TEC-M55) with a 55 mm focal length and an aperture of f2.8 with 2X extender was used to avoid distortion caused by parallax.

4.4. Velocity Field Comparison

The sets of synthetic images produced through the Monte Carlo technique (discussed in Section 4.3) and the experimental images acquired (discussed in Section 4.4), were each analyzed twice. The first analysis was a forward difference interrogation, while the second was a central difference interrogation. They were analyzed as described in Section 3 using the adaptive image shifting technique. To study the dependence of algorithm accuracy on increasing time delay lengths, a variety of time delays were studied. After the velocity fields were calculated for each data set, the Cartesian velocity vectors were decomposed into polar velocity vectors based on the location of the inner and outer cylinder walls. Any radial velocity component signifies an error because the flow should be purely azimuthal. The azimuthal error was found by subtracting the measured azimuthal component from that predicted by Equation (5). Since the azimuthal error should be order Δt^2 for both algorithms, it will not provide any insight into which algorithm will be more effective. However, the radial error should scale as Δt for the forward difference algorithm while the error should not depend on the time delay for the central difference algorithm. To quantify the systematic error in the measurements, a second order polynomial is fitted to the radial measurements. The radial velocity that would be measured at the surface of the inner cylinder (worst case scenario) is calculated by finding the polynomial's value at the inner cylinder radial position. This quantity is plotted in Figure 3 as a function of angular displacement of the inner cylinder.

Because the experimental conditions did not exactly match the Monte-Carlo simulation conditions, the data are expressed as dimensionless quantities to facilitate comparisons between the two. The radius ratio η was matched between the experiments and the simulations while the actual size (in pixels) of the inner and outer cylinders differed somewhat. For the simulations, the outer cylinder radius measured 890.0 pixels while for the experiments it measured 2316.1 pixels. Particle images were similarly larger in the experiments so that the interrogation windows used in analyzing the experiments were square, measuring 16 pixels on a side, while for the simulations they were squares measuring 8 pixels on a side. The errors are expressed as fractional errors relative to the linear speed of the surface of the inner cylinder since that is the only velocity scale available. The time delay between images Δt was expressed nondimensionally as the angular rotation measured in radians between the two laser flashes. Using these two dimensionless quantities, the radial velocity measurement error in the forward difference formulation expressed in Equation (9) can be rewritten as

$$\frac{\delta v_r}{v_\theta} = -\frac{1}{2}\Omega \tag{14}$$

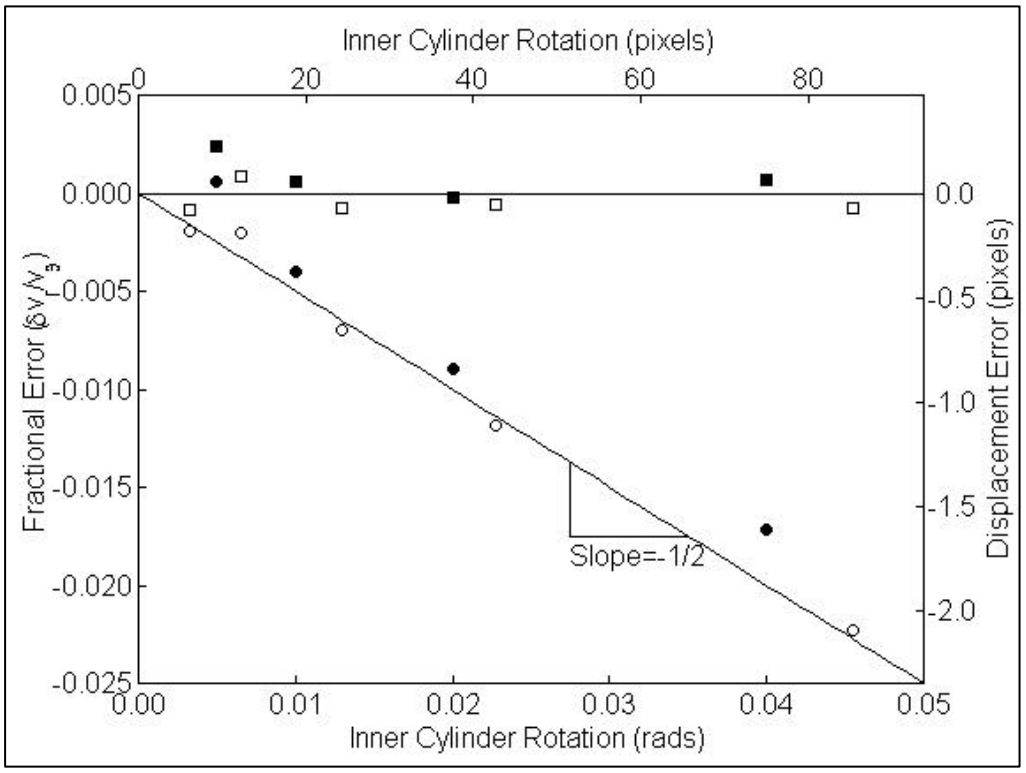

Figure 3. Comparison of the radial velocity errors in the forward difference and central difference discretizations. Square symbols (□,■) represent central difference interrogation while circles (○,●) represent forward difference interrogations. Filled symbols (●,■) represent interrogations of experimental data while the hollow symbols (○,□) represent the Monte Carlo results. The bottom horizontal axis is scaled in radians of rotation of the inner cylinder while the top horizontal axis is scaled in pixels of displacement of the inner cylinder. The left vertical axis is scaled as the fractional velocity error while the right vertical axis is scaled as the displacement error in pixels. The slope of –1/2, predicted by Equation (14), is plotted on the figure for reference.

where $\delta v_r/v_\theta$ is the fractional error, and Ω is the rotation of the inner cylinder between laser flashes measured in radians.

The trends in the data agree with the analytical predictions quite well. For the forward difference interrogations of both the experimental images as well as the simulated images, the fractional errors calculated at each cylinder rotation Ω (dimensionless Δt) fall very near the predicted –1/2 slope line. Similarly, the central difference interrogations for both the experimental and simulated images are scattered about the predicted error of zero. We would expect these trends to continue until other error terms begin to dominate, such as the spatial gradients associated with larger inner cylinder rotations.

5. CONCLUSIONS

The adaptive central difference interrogation (CDI) algorithm has three main advantages over conventional interrogation techniques: it is more accurate; it produces better results near boundaries; and it can maximize the signal-to-noise ratio of the measurements. The adaptive shifting CDI technique has been shown to produce more accurate velocity measurements when compared to an adaptive shifting forward difference interrogation (FDI) technique. The experiments and simulations from Section 4 demonstrate this increase in accuracy for the cylindrical Couette flow system. The fractional errors associated with measuring the radial velocity grow in proportion to Δt for the FDI technique while they remain in the vicinity of zero for time delays examined for the CDI technique.

Section 3 demonstrates that the adaptive CDI algorithm is significantly better than nonadaptive FDI algorithms when measuring near flow boundaries. The CDI algorithm has an area only 7% larger than the blood cell in which it cannot make valid measurements while the FDI algorithm cannot make measurements in an area 27% larger than the blood

cell. Section 3 also demonstrates that the CDI technique produces measurements with consistently higher signal-to-noise ratios. These advantages clearly demonstrate the CDI algorithm's utility in making measurements in demanding situations where high accuracy is required. Consequently the CDI technique surpasses the FDI technique in overall accuracy, performance near boundaries, and in low signal-to-noise ratio areas.

ACKNOWLEDGMENTS

This work was supported by AFOSR/DARPA contract number F49620-97-1-0515, DARPA contract F33615-98-1-2853, and the College of Engineering at UCSB. Special thanks to Alp Akonur and Professor Richard Lueptow in the Mechanical Engineering Department at Northwestern University for providing the experimental images of the cylindrical Couette flow.

REFERENCES

Agui, J. and Jimenez, J., 1987, 'On the Performance of Particle Tracking,' *J. Fluid Mech.*, **Vol. 186**, pp. 447-468.

Cohn, R. K. and Koochesfahani, M. M., 1999, 'The Accuracy of Remapping Irregularly Spaced Velocity Data onto a Regular Grid and the Computation of Vorticity,' pp. 225-230, Third International Workshop on Particle Image Velocimetry, Santa Barbara, CA, Sept. 1999.

Cowen, E. A. and Monismith, S. G., 1997, 'A hybrid digital particle tracking velocimetry technique,' *Exp. Fluids*, **Vol. 22**, pp. 199-211.

Dudgeon and Mersereau, 1984, *Multidimensional Digital Signal Processing*, Prentice-Hall.

Guezennec, Y.G. and Kiritsis, N., 1990, 'Statistical investigation of errors in particle image velocimetry,' *Exp. Fluids*, **Vol. 10**, pp. 138-146.

Huang, H., Dabiri, D., and Gharib, M., 1997, 'On errors of digital particle image velocimetry,' *Meas. Sci. Technol.*, **Vol. 8**, pp. 1427-40.

Keane, R. D., Adrian, R. J., 1990, 'Super-resolution particle imaging velocimetry,' *Meas. Sci. Tech.* **Vol. 1**, pp. 1202-1215.

Keane, R. D., Adrian, R. J., 1993, 'Theory and Simulation of Particle Image Velocimetry,' *Proc. SPIE*, **Vol. 2052** (Laser Anemometry Advances and Applications. Fifth International Conference, Veldhoven, Netherlands, 23-27 Aug. 1993), p.477-92.

Meinhart, C.D., Wereley, S.T., and Santiago, J.G., 1999, 'A PIV Algorithm for Estimating Time-Averaged Velocity Fields,' ASME/JSME Fluids Engineering Conference, San Francisco, CA, July 18-23.

Santiago, J. G., Wereley, S. T., Meinhart, C. D., Beebe, D. J., and Adrian, R. J., 1998, 'A Particle Image Velocimetry System for Microfluidics,' *Exp. Fluids*, **Vol. 25**:4, pp.316-19.

Spedding, G. R. and Rignot, E. J. M., 1993, 'Performance Analysis and Application of Grid Interpolation Techniques for Fluid Flows,' *Exp. Fluids*, **Vol. 15**, pp. 417-430.

Urushihara, T., Meinhart, C. D. & Adrian, R. J., 1993, 'Investigation of the logarithmic layer in pipe flow using particle image velocimetry,' In *Near-Wall Turbulent Flows*, R. So, et al. (Eds.), New York: Elsevier, pp. 433-46.

Westerweel, J., 1997, 'Fundamentals of dynamic particle image velocimetry,' *Meas. Sci. Technol.*, **Vol. 8**, pp. 1379-92.

Westerweel, J., 1993, 'Analysis of PIV interrogation with low-pixel resolution,' *Proc. SPIE*, **Vol. 2005,** pp. 624-635.

Wereley, S.T. and Lueptow, R.M., 1994, "Azimuthal velocity in supercritical circular Couette flow," *Exp. Fluids*, **Vol. 18**, pp. 1-9.

Wereley, S.T. and Lueptow, R.M., 1998, "Spatio-temporal character of nonwavy and wavy Taylor Couette flow," *J. Fluid Mech.*, **Vol. 364**, 59-80.

Wereley S.T., Meinhart, C.D., Santiago, J.G., and Adrian, R.J., 1998, "Velocimetry for MEMS Applications," Proc. ASME/DSC, Vol. 66, 453-459, (Micro-fluidics Symposium, Anaheim, CA, Nov. 1998).