

Second-order accurate particle image velocimetry

S. T. Wereley, C. D. Meinhart

258

Abstract An adaptive, second-order accurate particle image velocimetry (PIV) technique is presented. The technique uses two singly exposed images that 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. An adaptive interrogation region-shifting algorithm is used to implement the central difference approximation. 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. The second-order accuracy resulting from the central difference approximation can be obtained with relatively little additional computational effort compared to that required for a standard first-order accurate forward difference approximation.

The adaptive central difference interrogation (CDI) algorithm has two main advantages over adaptive forward difference interrogation (FDI) algorithms: it is more accurate, especially at large time delays between camera exposures; and it provides a temporally symmetric view of the flow. By comparing measurements of flow around a single red blood cell made using both algorithms, the CDI technique is shown to perform better than conventional FDI-PIV interrogation algorithms near flow boundaries.

Cylindrical Taylor–Couette flow images, both experimental and simulated, are used to demonstrate that the CDI algorithm is significantly more accurate than conventional PIV algorithms, especially as the time delay between exposures is increased. 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 (PIV), 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 that track the motion of tracer particles, such as PIV, is displaced from the centroid of the first interrogation region by a distance equal to half the displacement vector, and 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). Wereley et al. (1998) and Lourenco and Krothapalli (1998) both implemented second-order accurate PIV techniques without proof of their effectiveness. 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 nonuniformly spaced data, several researchers in fields related to PIV have addressed how to map nonuniformly 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 2000). Regardless of the accuracy, the remapping process will always introduce some error into the measurement.

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

Received: 15 June 2000/Accepted: 2 February 2001

S. T. Wereley
School of Mechanical Engineering, Purdue University
West Lafayette, IN 47907-1288, USA
e-mail: wereley@purdue.edu

C. D. Meinhart
Department of Mechanical and Environmental Engineering
University of California, Santa Barbara, CA 93106, USA

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.

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 that significantly improves spatial resolution and overall accuracy while reducing velocity bias errors in comparison to conventional PIV algorithms. This algorithm was initially developed by Wereley et al. (1998) and presented without demonstrating its second-order accuracy. Here the analysis is expanded and a theoretical, computational and experimental analysis of the algorithm will be given. We will consider two-component PIV in which the images are recorded as pairs of singly exposed images on consecutive frames of some digital medium (CCD camera typically). Most of the ideas considered here are extensible to other forms of two-component and three-component PIV.

1.1 Spatial resolution

The spatial resolution of any physical measuring system can be generally defined as the volume of space required to make the 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

$$\begin{aligned} & \{ \kappa_x, \kappa_y, \kappa_z \} \\ & = \{ \min(W_{x1}, W_{x2}) + |\Delta X|, \min(W_{y1}, W_{y2}) + |\Delta Y|, \delta z_c \} \end{aligned} \quad (1)$$

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 that varies across the region of interest.

The typical autocorrelation constraint that the difference between the particle displacement and the spatial offset between first and second interrogation regions not exceed a third of the interrogation region size can be eliminated by always using a spatial shift equal to the local flow velocity. Then the window size W can be chosen by signal strength or spatial gradient considerations alone. 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 et al. (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 the point that 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 Accuracy

The issue of accuracy is related to that of spatial resolution. Larger interrogation windows and larger spatial shifts will clearly produce more coarse estimations of the 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 $t + \Delta t$, given by

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

The accuracy of the velocity calculation can be greatly improved by using a relatively simple CDI scheme in which the velocity at time t is calculated using particle images at times $t - \Delta t/2$ and $t + \Delta t/2$, given by

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

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 desired 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 2–4 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 0 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 greatly reduce 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 to account for possible diffusion of the tracer particles. 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 to determine ensemble-averaged flow fields is to obtain a series of instantaneous velocity fields and then ensemble averaging 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 et al. (2000) 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 Sect. 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 relatively narrow gap of a cylindrical Couette flow system.

2

Central difference interrogation fundamentals

A central difference interrogation (CDI) scheme has several distinct advantages over the conventional forward difference interrogation (FDI) scheme. The most obvious advantage already mentioned is that the technique is second-order accurate, 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

$$\varepsilon_B = \frac{1}{u} \sqrt{\frac{2D}{\Delta t}} \quad (4)$$

where u is the local velocity, D is the diffusion coefficient of the particles, and Δt is the time delay. Equation (4) demonstrates that when all other factors remain constant the relative error due to Brownian motion can be reduced by increasing the time delay.

Further advantages of the CDI algorithm are that it allows the velocity field to be calculated closer to flow boundaries as well as in a symmetric manner with respect to flow boundaries. These advantages can be demonstrated by examining the following stagnation type flow. Figure 1a and b show velocity fields of a flow around a human red blood cell calculated by an adaptive spatial shift FDI scheme and also by the adaptive spatial shift CDI scheme. These results were taken from a set of experiments conducted by Juan Santiago at the Beckman Institute at the University of Illinois and are described further by Wereley et al. (1998). The blood cell is trapped between a coverslip and a slide glass and surrounded by deionized, particle-seeded water (300 nm diameter polystyrene latex, 10^{-4} concentration by volume). 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 blood cell. The images are interrogated with 28×28 pixel interrogation windows that are overlapped by 75%. While this overlap oversamples the images, it functions as a very accurate (yet computationally inefficient) interpolator of the velocity field. For both algorithms, measurement regions 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 “x” symbol because they will tend to produce velocity measurements with serious errors. The CDI scheme is able to accurately measure velocities closer to the surface of the cell than is the FDI scheme. The CDI scheme has a total of 55 invalid measurement points, equivalent to $59.4 \mu\text{m}^2$ of image area that cannot be interrogated, while the forward difference has 57 invalid measurement points or $61.6 \mu\text{m}^2$. Although this difference of two measurement points may not seem significant, it amounts to the area that the FDI algorithm cannot interrogate being 3.7% larger than the area the CDI algorithm cannot interrogate. Furthermore, the distribu-

tion of the invalid points is significant. By carefully comparing Fig. 1a with Fig. 1b, it is apparent that the FDI scheme has three more invalid measurement points upstream of the blood cell than does the CDI scheme while the CDI scheme has one more invalid point downstream of the blood cell. This difference in distribution of invalid points translates into the centroid of the invalid area being nearly twice as far from the center of the blood cell in the FDI case ($0.66 \mu\text{m}$ or 7.85% of the cell diameter) versus the CDI case ($0.34 \mu\text{m}$ or 4.05% of the cell diameter). This difference means that the FDI measurements are less symmetrically distributed around the blood cell than the CDI measurements are. In fact, they are biased toward the time at which the first image was recorded. Computing the average distance between the invalid measurement points and the surface of the blood cell indicates how closely to the blood cell surface each algorithm will allow the images to be accurately interrogated. On average, the invalid measurement points bordering the red blood cell generated using the FDI scheme are 12% farther from the cell than the invalid measurement points generated with the CDI scheme. Consequently, the adaptive CDI algorithm is more symmetric than the adaptive FDI algorithm and also allows measuring the velocity field nearer the cell surface.

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 half-way between $t - \Delta t/2$ and $t + \Delta t/2$, the spatial shift between interrogation windows at each measurement point in the first and second image must be equal to the distance that the particles move between images, just as in conventional forward difference algorithms. In practice, a uniformly spaced grid of measurement points is defined at the time t and the interrogation windows in the first image




Fig. 1. Comparison of a forward difference adaptive window offset analysis: a vs. 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 the second-order CDI scheme

are shifted by an amount $-V(x, y)\Delta t/2$ while the interrogation windows in the second image are shifted by an amount $V(x, y)\Delta t/2$, netting a spatial shift of $V(x, y)\Delta t$ 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.

This procedure is outlined in flowchart form in Fig. 2. The first step in the 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 Sect. 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. (CleanVec is a PIV vector validation program written by Ron Adrian's group in Theoretical and Applied Mechanics at the University of Illinois. Further information about CleanVec is available at <http://hairpin.tam.uiuc.edu>.) CleanVec compares a candidate's 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 the median of its neighbors by more than an amount specified by the user, it will be replaced by either the second choice or third choice correlation peak, which are found along with the primary correlation peak and stored for error correction purposes. If neither of these secondary peaks matches the neighborhood any better, a vector is calculated by interpolating from the neighbors' values. Gaussian smoothing is used to reduce high frequency noise in the displacement field.

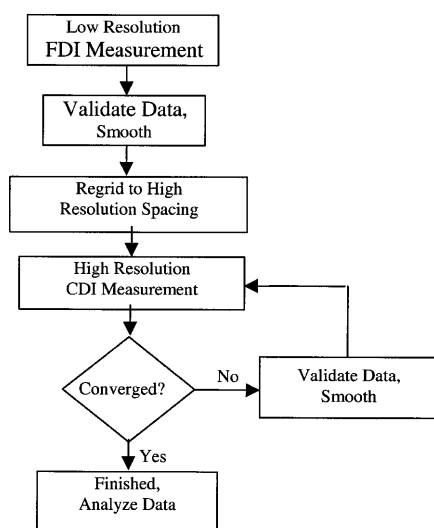


Fig. 2. Flowchart demonstrating iterative procedure necessary adaptive image shifting and second-order accuracy

Using this 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 initial low-resolution velocity field at each new high-resolution measurement point. This step is used purely as an initial estimate of the high-resolution velocity field and is the only interpolation step in the algorithm. 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 the CDI algorithm, the centroid of the first interrogation window is displaced from the location of the measurement point by the integer number of pixels that is nearest $-0.5 \times \Delta X$, i.e. opposite to the flow direction, while the centroid of the second interrogation window is displaced from the first by the integer number of pixels nearest ΔX . Here, ΔX represents the local expected displacement measurement from a previous iteration. 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 sufficient quality that further post-processing is not necessary.

4 Application to cylindrical Couette flow

In order to study the relative accuracy of different PIV 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 more closely match the known velocity field will be the more accurate one. Two adaptive image-shifting techniques will be compared. The first is a conventional FDI scheme and the second is the CDI scheme. These two adaptive shifting techniques both require two images to compute and approximately the same amount of computational effort when compared to each other, but do require additional computational effort when compared with constant shift techniques. Cylindrical Couette flow is used as a test bed to compare the two PIV algorithms because it is a relatively simple flow with a closed-form solution in which experimental velocity measurements can be compared to results from 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

$$v_r = v_z = 0$$

$$v_\theta(r) = r\omega(r) = A_0r + B_0/r \tag{5}$$

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 Eq. (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 Fig. 3a). 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 mathematics. The standard first-order accurate forward difference discretization of the velocity field is

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

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

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

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

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

$$V_{x,\text{act}} = 0$$

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

The difference between these two velocities will tell us how much of an effect the limitations of the PIV algorithm will have on the overall accuracy of the measurements. To simplify the resulting expression, the trigonometric functions will be expanded as power series. After significant algebraic manipulation, the error in the X and Y direction are found to be

$$\Delta V_x = -\frac{1}{2}r_0\Delta t\omega(r_0)^2 + \frac{1}{24}r_0\Delta t^3\omega(r_0)^4 + \dots$$

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

where ΔV_x and ΔV_y represent the bias error in each direction.

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 . In the special case 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

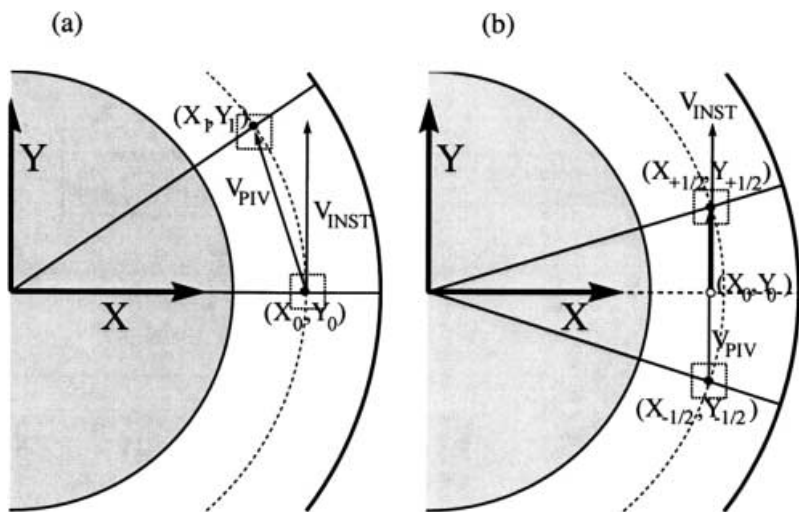


Fig. 3. Schematic views illustrating: a the FDI algorithm; b the CDI algorithm as applied to the cylindrical Couette flow system. The dashed squares represent spatially offset interrogation regions while the dashed circles represent the path that a fluid element would follow

smaller than the azimuthal component, first-order accuracy in Δt can be restrictive.

4.1.2

Central difference error evaluation

Although the CDI technique is ostensibly second-order accurate in Δt , the accuracy of the technique must be carefully examined to prove 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 Fig. 3b). 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

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

and

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

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

$$V_{x,\text{act}} = 0$$

$$V_{y,\text{act}} = r_0 \cos\left\{\omega(r_0) \frac{\Delta t}{2}\right\} \omega\left(r_0 \cos\left\{\omega(r_0) \frac{\Delta t}{2}\right\}\right) \quad (12)$$

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

$$\Delta V_x = 0$$

$$\Delta V_y = r_0 \omega(r_0)^2 \Delta t^2 \left\{ \frac{1}{12} \omega(r_0) + \frac{1}{8} r_0 \omega'(r_0) \right\} + \dots \quad (13)$$

where ΔV_x and ΔV_y represent the bias error in each direction.

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 bias error in the azimuthal component is on the order of Δt^2 as expected.

4.2

Monte Carlo simulations

One method for evaluating the accuracy of PIV algorithms is to process a set of simulated particle images displaced according to a prescribed velocity field. This approach has certain advantages over evaluating algorithms using experimental results, which will naturally have errors and inaccuracies associated with them. Sets of simulated images can easily be constructed 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), Keane and Adrian (1990), and many subsequent studies. Image pairs are simulated by creating a randomly distributed particle image 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 prescribed velocity field; in this case Eq. (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 any 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 gray level 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. Twenty-six sets of image pairs were generated for each set of operating conditions explored.

4.3

Experimental measurements

The experimental flow system, shown in Fig. 4, 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




Fig. 4. Experimental setup. The camera-light sheet orientation indicates that velocity components in a radial-azimuthal plane will be measured

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 $\Gamma = 47.8$. 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 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 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 consecutive time instants. The laser pulses were separated by 7.5, 15, 30, and 60 ms, depending on the run, in order to study the dependence of the measurement error on the time step. The laser sheet had a thickness of approximately 1 mm, with $\pm 0.5\%$ 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 Computer telecentric lens (model TEC-M55) with a 55 mm focal length and an aperture of f2.8 with $2 \times$ extender was used to avoid distortion caused by parallax. Twenty-six sets of image pairs, the same number of image sets generated in the Monte Carlo investigation, were recorded for each set of operating conditions explored.

4.4 Velocity field comparison

The sets of synthetic images produced through the Monte Carlo technique (discussed in Sect. 4.2) and the experimental images acquired (discussed in Sect. 4.3), were each analyzed twice. The first analysis was the FDI, while the second was the CDI. While studying the convergence of the iterative, adaptive algorithms is not the primary purpose of the present work, it is certainly an important part of it. After each iteration is completed, the velocity produced from that iteration is used as the adaptive shifting file for the next iteration without alteration or correction. In general, the velocity data could be validated or smoothed at this point but that was not necessary for these measurements. After repeating this procedure several times, the measurement at a given point will either have reached a final value that does not change with subsequent iterations or it will be oscillating back and forth between two or more possible solutions. This oscillation is caused by the adaptive shift in one iteration leading to a measurement that does not produce the same

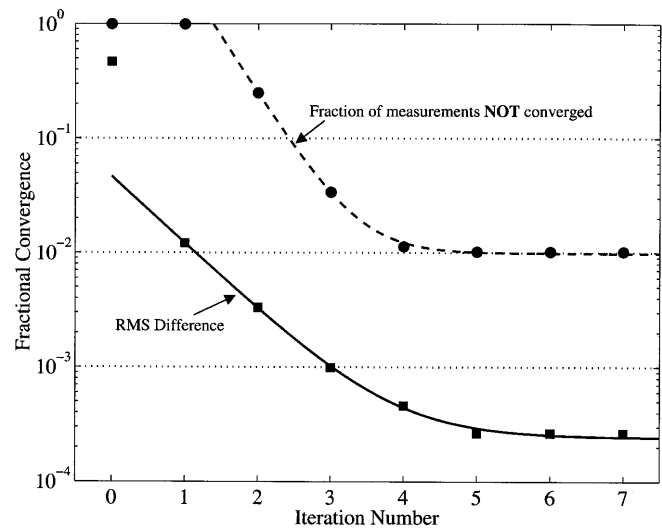


Fig. 5. Convergence characteristics of CDI algorithm. The square symbols (■) and the solid curve show the RMS difference between successive iterations normalized by the RMS of the velocity field. The circular symbols (●) and the dashed curve show the fraction of the velocity measurements that have not converged as a function of iteration number. Both data sets are modeled quite well by a decaying exponential function

shift in the next iteration. While this oscillation might at first seem troubling, it is very easy to explain and quantify. Figure 5 shows two measures of the convergence rate of the CDI algorithm as applied to the Taylor-Couette measurements here. The first measure is how quickly the rms difference between the measured velocity and actual velocity approaches a final value and the second measure is the fraction of measurement points that have not converged. Both of these measures are plotted in Fig. 5 as a function of iteration number. Both are also well modeled by a decaying exponential which is superposed on each set of data point symbols. This figure demonstrates that with only a single iteration, the measurements are accurate to within 1% of the correct value, while after three iterations, the measurements are accurate to within 0.1%, and after five iterations, the measurements are accurate to within 0.02%, which is as accurate as they will ever get in these experiments. Figure 5 also shows that after four iterations, all but 1% of the measurements have reached a steady-state value after which further iterations will produce no additional convergence. The iteration scheme used in this work is a very simple one, merely feeding back the velocity field produced in one iteration to calculate a more accurate velocity field in the next iteration, which worked adequately in this situation. A more robust approach would be to include some damping in each iteration to prevent the measurements from running away from the correct solution due to a bad measurement in one iteration, and to prevent the oscillating measurements mentioned above. In addition, the fundamental assumption of PIV, that the particle images be randomly distributed, makes running away from the correct measurements statistically unlikely. As a safety check, any CDI measurement that differs from the FDI measurement made at the same point by more than some

threshold amount can be replaced with the FDI measurement.

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 data sets. 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 bias errors are expressed as fractional errors relative to the linear speed of the surface of the inner cylinder $v_{\theta,ic}$, 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 δv_r in the forward difference formulation expressed in Eq. (9) can be rewritten as

$$\frac{\delta v_r}{v_{\theta,ic}} = -\frac{1}{2} \frac{v_{\theta}(r)}{v_{\theta}(r_{ic})} \omega(r) \Delta t = -\frac{1}{2} \Omega^*(r) \quad (14)$$

where $\delta v_r/v_{\theta,ic}$ is the fractional error and $\omega(r)$ is the angular rotation rate of the fluid at radius r which is given in Eq. (5). $\Omega^*(r)$ is the modified angular displacement between laser flashes, measured in radians, of a fluid element at a radial position r . This modified angular displacement varies from zero at the outer cylinder to Ω_{ic} at the inner cylinder. In between the two cylinders, the modified angular displacement is scaled by the ratio of the local fluid velocity to the inner cylinder surface speed. From this

relationship it is clear that the fractional error of the forward difference scheme will depend both on the radial position as well as on the time between laser flashes Δt . The parameter Ω^* can have several interpretations. It can be thought of as a dimensionless time step Δt because it is directly proportional to Δt for a fixed radial position. It can also be thought of as a dimensionless measure of radial position because it is a monotonic, albeit nonlinear, function of radial position when the time step Δt is fixed. Consequently it is an ideal parameter for comparing the experimental and computational results because it will allow both the effects of time step size as well as radial position to be explored.

To study the dependence of algorithm accuracy on increasing time delay lengths Δt , a succession of time delays and radial positions 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 is on the order of Δt^2 for both algorithms, it does not provide any insight into which algorithm will be more effective and will not be discussed here. However, the radial error should be proportional to Δt as well as a function of radial position, as Eq. (14) shows for the forward difference algorithm, while the error should not depend on the time delay for the central difference algorithm.

The dependence of the radial velocity error on radial position was explored first. The radial velocity component was divided by the surface speed of the inner cylinder and plotted as a function of both the parameter Ω^* (bottom axis) as well as the radial position (top axis) in Fig. 6. For a fixed inner cylinder speed and the relatively large radius




Fig. 6. Typical distribution of velocity measurements, each marked by a plus (+) sign, about the predicted bias error curve. The left side of the plot corresponds to stationary outer cylinder while the right side of the figure corresponds to the rotating inner cylinder. The scatter of the points indicates the random error (± 0.0048) while the general trend indicates the velocity bias error

ratio of this experimental apparatus ($\eta = 0.81$), the parameter Ω^* is nearly a linear function of radial position with Ω^* equal to zero at the outer cylinder, a radial position equal to 1.00, and Ω^* equal to the angular displacement of the inner cylinder between laser flashes at the inner cylinder, a radial position equal to the radius ratio $\eta = 0.81$. In Fig. 6, each of the “+” symbols represents a single velocity measurement point while the solid line represents the predicted $-1/2$ slope line. The dashed lines represent one standard deviation (± 0.0048) about the $-1/2$ slope line. The concentration of the symbols is greatest at the left side of the figure – the side of the figure that corresponds to the outer cylinder – and smallest at the right side of the figure – the side of the figure that corresponds to the inner cylinder. Note that the symbols are not distributed all the way to the inner or outer cylinders because the finite size of the interrogation region excludes the centroids of the measurement points from reaching either wall. The concentration of measurement points varies significantly because of the area effect, i.e., larger radial positions correspond to larger areas containing more measurement points. The general sloping distribution of the measurement points confirms the prediction about how the bias error depends on radial position while the scatter in the measurement points (± 0.0048) demonstrates the size of the random error in relation to the systematic bias error. Near the outer cylinder the random error dominates, while near the inner cylinder the bias error dominates.

To quantify the bias error, i.e., radial velocity component, as a function of increasing time delays Δt , for each Δt the radial velocity measurements are plotted as a function of radial position (as in Fig. 6) and a second-order polynomial is fitted to the measurements to allow for curvature effects. The radial velocity that would be measured at the surface of the inner cylinder (where the bias error is worst) is calculated by finding the polynomial’s value at the inner

cylinder radial position. This quantity is plotted in Fig. 7 as a function of angular displacement of the inner cylinder Ω^* . Figure 7 shows that for the FDIs 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 CDIs 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 shifting CDI technique has been shown to produce more accurate velocity measurements when compared to a conventional adaptive shifting FDI technique. Section 3 demonstrates that the adaptive CDI algorithm is significantly better than FDI algorithms when measuring near flow boundaries. On average, the invalid measurement points bordering the red blood cell generated using the FDI scheme are 12% further from the cell than the invalid measurement points generated with the CDI scheme. This difference in distribution of invalid points translates into the centroid of the invalid area being nearly twice as far from the center of the blood cell in the FDI case (7.85% of the cell diameter) versus the CDI case (4.05% of the cell diameter). Therefore, the CDI algorithm is able to measure closer to flow boundaries and is more symmetric than the FDI algorithm. The experiments and simulations from Sect. 4 demonstrate the increased accuracy of the CDI algorithm when applied to 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. Since adaptive FDI techniques are becoming standard, even in

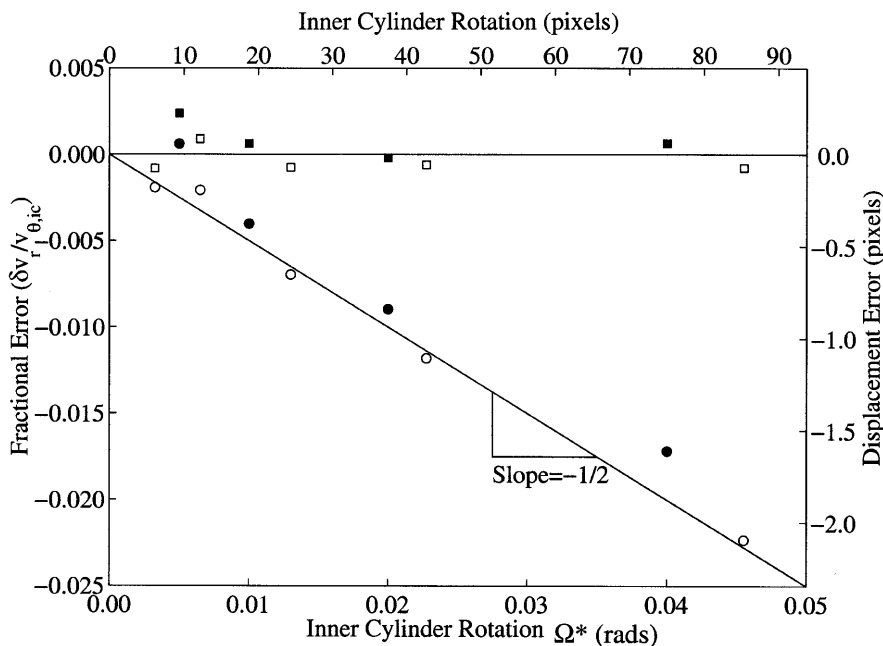


Fig. 7 Comparison of the radial velocity errors in the forward difference and central difference discretizations. Square symbols (\square , \blacksquare) represent CDI while circles (\circ , \bullet) represent FDIs. Filled symbols (\bullet , \blacksquare) represent interrogations of experimental data while the hollow symbols (\circ , \square) 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

commercially written software, the advantages of the CDI algorithm can be attained for little additional computational cost. Consequently, the CDI algorithm should be the algorithm of choice in demanding situations where high accuracy is required.

References

- Agui J; Jimenez J** (1987) On the performance of particle tracking. *J Fluid Mech* 186: 447–468
- Cohn RK; Koochesfahani MM** (2000) The accuracy of remapping irregularly spaced velocity data onto a regular grid and the computation of vorticity. *Exp Fluids* 29(7): S061–S069
- Cowen EA; Monismith SG** (1997) A hybrid digital particle tracking velocimetry technique. *Exp Fluids* 22: 199–211
- Guezennec YG; Kiritsis N** (1990) Statistical investigation of errors in particle image velocimetry. *Exp Fluids* 10: 138–146
- Huang H; Dabiri D; Gharib M** (1997) On errors of digital particle image velocimetry. *Meas Sci Technol* 8: 1427–1440
- Keane RD; Adrian RJ** (1990) Super-resolution particle imaging velocimetry. *Meas Sci Technol* 1: 1202–1215
- Keane RD; Adrian RJ** (1993) theory and simulation of particle image velocimetry. In: *Laser anemometry advances and applications, Fifth International Conference, Veldhoven, Netherlands, 23–27 August 1993, Proc SPIE Series, vol 2052*, pp 477–492
- Lourenco L; Krothapalli A** (1998) Mesh-free second order accurate algorithm for PIV processing. In: *Proceedings of the International Conference on Optical technology and image processing in fluid thermal and combustion flow. Visualization Society of Japan, Yokohama*
- Meinhart CD; Wereley ST; Santiago JG** (2000) A PIV algorithm for estimating time-averaged velocity fields. *J Fluids Eng* 122: 285–289
- Santiago JG; Wereley ST; Meinhart CD; Beebe DJ; Adrian RJ** (1998) A particle image velocimetry system for microfluidics. *Exp Fluids* 25(4): 316–319
- Spedding GR; Rignot EJM** (1993) Performance analysis and application of grid interpolation techniques for fluid flows. *Exp Fluids* 15: 417–430
- Urushihara T; Meinhart CD; Adrian RJ** (1993) Investigation of the logarithmic layer in pipe flow using particle image velocimetry. In: *So R et al (eds) Near-wall turbulent flows. Elsevier, New York*, pp 433–46
- Wereley ST; Lueptow RM** (1994) Azimuthal velocity in supercritical circular Couette flow. *Exp Fluids* 18: 1–9
- Wereley ST; Lueptow RM** (1998) Spatio-temporal character of nonwavy and wavy Taylor Couette flow. *J Fluid Mech* 364: 59–80
- Wereley ST; Meinhart CD; Santiago JG; Adrian RJ** (1998) Velocimetry for MEMS applications. In: *Proceedings of the ASME/DSC Micro-fluidics symposium, Anaheim, CA, November 1998, vol 66. ASME, New York*, pp 453–459
- Westerweel J** (1993) Analysis of PIV interrogation with low-pixel resolution. *Proc. SPIE Series, vol 2005*, pp 624–635
- Westerweel J** (1997) Fundamentals of dynamic particle image velocimetry. *Meas Sci Technol* 8: 1379–1392