Experiments in Fluids 30 (2001) 434–447 © Springer-Verlag 2001

New insights into particle image velocimetry data using fuzzy-logic-based correlation/particle tracking processing

M. P. Wernet

Abstract Digital particle image velocimetry (DPIV) data processing has been developed to the point where DPIV image data are processed via auto- or cross-correlation techniques in near real time and the results are displayed on screen as they are processed. Correlation techniques are highly desirable, since they provide velocity measurements on a regular grid, which are readily comparable to CFD predictions of the flow field. In high-speed flows, particle lag effects are always of concern; however, the correlation operation does not provide any means for minimization or elimination of systematic errors in the recorded particle image data. In this paper, we present a combined correlation processing/particle tracking technique providing "super-resolution" velocity measurements. Fuzzy-logic principles are employed to maximize the information recovery in the correlation operation and to determine the correct particle pairings in the tracking operation. The combined correlation/particle tracking technique is applied to DPIV data obtained in the diffuser region of a high-speed centrifugal compressor producing velocity vector maps with an average density of 6 vectors/mm^2. Inspection of the particle tracking results revealed large particles that were not following the flow. Using preknowledge of the flow field, the biased velocity estimates arising from large particles in the flow were removed, thereby improving the accuracy of the measurements.

1 Introduction

Particle image velocimetry (PIV) is a planar velocity measurement technique wherein a pulsed laser light sheet is used to illuminate a flow field seeded with tracer particles small enough to accurately follow the flow. The light sheet is pulsed at two very closely spaced instants in time, and the positions of the particles are recorded by a camera. The recorded images of the particle positions are then processed to determine the flow velocity over the illuminated plane. PIV has gone through several changes over the last 10 years, most of them driven by CCD camera and computer technology. Originally implemented as a photographic technique, PIV has evolved into a digital technique (DPIV), where image data are acquired and processed in near real time. The ability to acquire single-exposure image frame pairs has driven the trend whereby nearly all DPIV data are processed via correlation processing. In correlation processing, spatially averaged estimates of the flow velocity across a uniform grid are obtained (Willert and Gharib 1991). Computer processing speeds have supported this trend since software-based correlation processing has proved to be a reasonable approach for DPIV data processing.

Particle tracking velocimetry (PTV) has been a longtime rival of correlation processing. As the name implies, the PTV technique requires tracking of the individual particle images in order to estimate the flow velocity. PTV is normally restricted to low seed particle concentration flows so that the individual particle trajectories can be unambiguously tracked. The velocity vectors obtained using PTV are randomly distributed based on the seed particle locations within the flow. Combined correlation processing/particle tracking techniques have been demonstrated by Keane et al. (1995) and by Wernet (1995). In Keane et al. (1995), double exposure imagery of a liquid flow were processed via autocorrelation, followed by particle tracking. The particle tracking results were noisy, requiring filtering and spatial averaging before being compared with the correlation results. Wernet (1995) presented a combined correlation tracking technique utilizing fuzzy logic for correlation peak validation and for the particle tracking operation. Single-exposure image data of an underexpanded supersonic nozzle flow were processed via cross-correlation, followed by particle tracking. The normal shock in the flow was readily observable in the tracking results along with the repeated acceleration/deceleration in the flow downstream of the shock. The tracking results covered a wider dynamic range than the spatially averaged correlation results. The combined technique enabled particle tracking to be applied in high-speed particle concentration image data ($\sim$10 particles/mm^2), which were previously processed exclusively using correlation techniques. The particle tracking stage produces "super-resolution" velocity measurements that accurately record velocity gradients and/or shocks so that structures in the flow that may have been smeared by the spatially averaged correlation operation are now resolved down to the fidelity of the particle flow dynamics.

DPIV data-processing techniques employing recursive correlation have also been used to achieve "super-resolution" velocity measurements (Hart 1998); however,

Received: 21 October 1999/Accepted: 19 August 2000

M. P. Wernet
NASA Glenn Research Center, MS-77-1
21000 Brookpark Road, Cleveland, OH 44135, USA

variations in seed particle density across the image and local velocity gradients impedes the applicability of the recursive technique. In this recursive approach, subregion image shifting is applied in each successive pass, using the estimated displacement from the previous pass to enable successively smaller subregions. The recursive correlation technique is purported to obtain velocity measurements with a spatial resolution on the order of the particle size. It is not possible to obtain independent velocity estimates on the scale of the individual particles. The limiting spatial resolution for obtaining independent velocity estimates from DPIV image data is physically bounded by either the average separation of the particles in the fluid or the particle displacement between exposures, whichever is smaller. Velocity estimates more closely spaced than the smaller of the average particle separation or the particle displacements between exposures are interpolated data. Therefore, the combined correlation processing/particle tracking approach, wherein the individual particle displacements are determined, yields the highest spatial resolution independent velocity measurements from seeded flow fields.

Velocity biasing is a common ailment of all particulate-based velocity-measurement techniques. Even if the flow tracers perfectly follow the fluid fluctuations, the measurements can be biased. Velocity bias in LDV measurements arises from the flow turbulence, which affects the arrival rate of particles through the probe volume. More particles are convected through the probe volume per unit time during periods of high flow velocity than during periods of low flow velocity. Hence, the ensemble-averaged velocity estimates are biased towards higher velocities (McLaughlin and Tiederman 1973). The velocity bias is roughly proportional to the square of the turbulence intensity. Several processing strategies were proposed for correcting the velocity bias, but they were only implemented in cases of high flow turbulence, due to the larger bias introduced under these conditions.

The LDV velocity bias is one example of a measurement bias that can occur when measuring fluid velocities in particulate seeded flows. All particulate seeding techniques are sensitive to particle biases of various types. LDV and DPIV are both sensitive to the particle lag that can occur when large particles do not accurately follow the flow. In DPIV correlation processing, large particles are the dominant contributors to the size and shape of the correlation peak; therefore the resulting correlation peak is biased towards the flow properties of the large particles. There are many other factors that affect the size and shape of the correlation peak: shear, flow turbulence and particle dynamics are all potential contributors. Keane and Adrian (1990) showed that velocity gradients across the interrogation subregion resulted in a reduction in the amplitude and a broadening of the displacement peak. In addition to the bias caused by these non-uniformities in the flow across the interrogation region, they also showed that there exists a statistical bias in DPIV velocity estimates. The statistical bias stems from the fact that lower-velocity particles are more likely to remain in the plane of the light sheet and be captured during the two image exposures than higher-velocity particles, resulting in a bias towards lower velocities. The effect of the statistical bias can be minimized in cases of high seed particle concentration, where the statistical influence of the loss of particle pairs is reduced. The spatial averaging that occurs in correlation processing yields a correlation peak that is an amalgam of all of these factors, and although the contributions of these various effects can be described, they cannot be removed from the velocity estimates.

In particle tracking, the process of determining the correct particle pairings is very tedious, but easily accomplished by manual identification and pairing by a human operator. Attempts to automate the particle pairing process are difficult to implement efficiently and prone to error. Neural nets can be used to perform the particle tracking, but require training for each type of flow field to be measured (Cenedese et al. 1992). Alternatively, fuzzy-logic principles can be used to perform the particle tracking. The control of complex processes has been aided by the development of fuzzy control systems, which have been used to control traffic flow, appliances, and even subways for optimal energy efficiency and passenger comfort. Fuzzy logic employs a tolerance for imprecision to achieve system control. An exact model for the system inputs and outputs is not required. Fuzzy inference control utilizes membership functions and a rule base developed by the user to process information. The physical mechanisms underlying the process are impertinent to the controller. The process of identifying and tracking particles in a flow is a good candidate for fuzzy control, since the procedure is not clear cut, but involves some gray area decisions.

The power of the particle tracking technique is realized not only in its ability to produce high spatial resolution velocity measurements but also in its ability to separate and identify the individual particle behaviors. All of the flow non-uniformities, small-scale structures and particle lag effects can be investigated using particle tracking, which is not possible using correlation processing. The particle tracking results can be spatially averaged to produce a uniform grid of velocity estimates, similar to those obtained using correlation processing. However, before the spatial averages are computed, corrections to the individual velocity estimates can be made, i.e., biases can be corrected or removed. Particle tracking is not traditionally used, since it is unable to compete with correlation processing in highly seeded flows. However, with the adaptation of correlation and particle tracking into a combined technique, particle tracking is a viable option. The combined correlation/particle tracking technique has been implemented in a software package by Wernet (1999), which is now available from NASA Glenn Research Center.

This paper will discuss the basic principles used in the combined correlation/particle tracking processing used in the PIVPROC software developed at NASA Glenn. The implementation of a fuzzy processor to obtain the correct particle pairings from a pair of single-exposure PIV images will be discussed. DPIV image data obtained in the diffuser region of a high-speed centrifugal compressor will be used to examine the particle dynamics. Even though attempts are made to ensure that the flow is seeded with nearly monodispersed seed particles, the particle tracking results reveal that large particles are present that are not accurately following the flow, which leads to a velocity bias

towards lower velocities in the correlation processed data. Averaging the particle tracking data over a uniform grid of points (using the identical grid and subregion size used in the correlation processing operation), confirms that the data obtained from particle tracking contains nearly identical information to that obtained by direct correlation processing, including the biases. Systematic errors can be manifested in many ways that are not always easily discernible. However, in the compressor measurements presented here, large particle bias is readily identified in the flow exiting the impeller. Knowledge of the true flow field properties enables the determination of a filtering strategy by which large, slow-moving particles can be removed from the particle tracking results, yielding spatially averaged velocity vector maps of a more superior quality than are possible by using straight correlation processing. The bias corrected instantaneous velocity vector map shows that the particle lag velocity bias levels are most significant where the large particle concentration in the flow is high and ranges from 7 to 11%. Application of the bias correction prior to computing the time average of a 50-frame series of velocity vector maps shows a uniform bias throughout the flow of approximately 4%.

2
Correlation processing vs particle tracking

The purpose of this paper is not to reiterate what is already known about correlation processing of DPIV image data. Here we will only be concerned with direct digitally acquired pairs of single-exposure image data. Single-exposure image data enables cross-correlation data processing to be utilized, which provides the maximum dynamic range and highest signal to noise from DPIV image data. Single-exposure image frame pairs are also readily amenable to particle tracking processing. In either processing technique, the accuracy of the velocity vector estimates is determined by the accuracy to which the spatial position of the correlation peak or individual particle centroids can be determined, as will be discussed next.

In correlation processing, the DPIV image frame is divided up into a grid of small interrogation regions, and the cross correlations of the subregions from the first and second exposure are computed. The output correlation plane contains a correlation peak corresponding to the average displacement of the particles across the subregion. The particle displacements are restricted to be one-quarter of the subregion size for correlation processing (Keane and Adrian 1990). Each subregion must contain at least five to ten particles for a good correlation result. Hence, the seed particle concentration in the recorded images will ultimately limit the selected correlation subregion size and the concomitant velocity vector grid density. Flow features smaller than the correlation subregion size will be low-pass filtered, due to the spatial averaging nature of the technique. Larger-sized subregions containing more particles will yield higher reliability, but even more spatially smoothed velocity vector maps. For a subregion size of 64×64 pixels, a maximum displacement of 16 pixels can be achieved. The correlation peak locations are nominally estimated to an accuracy of 0.1–0.2 pixels, yielding a full-scale accuracy of 0.6% (Westerweel 1997).

Particle tracking yields the individual particle displacements between exposures. The first step in the particle tracking operation is the determination of the individual particle centroids in the two single-exposure image frames. A search region around each initial exposure particle is used to identify candidate second exposure particles. These candidate particle displacements are then analyzed in the tracking operation to determine the correct particle pairings. The accuracy in the individual particle displacement estimates is given by $\sqrt{2}$ times the uncertainty in the particle centroid estimates. For particle images spanning on the order of 1–2 pixels, the particle centroid estimates can be determined to nominally 0.2 pixels (Wernet and Pline 1993). For single-pixel particles, the error reaches the maximum value of 0.5 pixels. Assuming a mean particle centroid estimate accuracy of 0.35 pixels and a particle displacement search region size of 16 pixels, a nominal single particle displacement accuracy of 3% of full scale is achieved. Hence, the individual particle displacements are not as accurate as the spatially averaged estimates obtained from the correlation operation. However, if the particle tracking results were used to compute a spatially averaged velocity vector map, with the same sized subregions that are used in the correlation operation, then comparable accuracy velocity estimates would be obtained. Assuming the image contains sufficient seed particle concentration for there to be ten particles per correlation subregion, then if particle tracking were used, ten individual particle displacements could be averaged together to obtain a spatially averaged velocity estimate. Hence, the PTV technique is capable of providing spatially averaged velocity vector maps on regular grids with accuracies on the order of 1% of full scale.

In addition to the measurement precision errors discussed above, the spatially averaged velocity estimates contain other noise sources: flow turbulence σ_{TI}, velocity gradients across the subregion σ_{VG}, systematic errors σ_B resulting from particle bias effects and particle dropout errors σ_{PD} producing a loss of correlation. For correlation processing, the error in the estimated velocity, assuming negligible timing error, is given by:

$$\sigma^2_{\Delta X-\mathrm{Correl}} = \sigma^2_{CP} + \sigma^2_{TI} + \sigma^2_{VG} + \sigma^2_{B} + \sigma_{PD} \qquad (1)$$

where σ_{CP} is the error in estimating the correlation peak position. When computing the cross correlation, these additional error terms are all combined into the correlation peak, causing a broadening and reduction in amplitude of the peak. Once combined, these error sources cannot be deconvolved. The error in estimating the correlation peak position increases linearly with the width of the correlation peak (Wernet and Pline 1993), therefore as these error sources broaden the correlation peak, the uncertainty in the velocity estimates increases.

The results obtained from the particle tracking operation can be compared with the correlation results by computing the spatially averaged velocity field. The subregion for computing the spatial average of the particle tracking results is assumed to be the same size as the correlation processing subregion. The particle drop out error term is now replaced by a term representing the spurious vectors detected in the tracking process, σ_{SV}. The

error in the spatially averaged particle tracking estimated velocity is given by:

$$\sigma^2_{\Delta X-\mathrm{Track}} = \frac{2\sigma^2_{\mathrm{PT}}}{N} + \sigma^2_{\mathrm{TI}} + \sigma^2_{\mathrm{VG}} + \sigma^2_{\mathrm{B}} + \sigma^2_{\mathrm{SV}} \qquad (2)$$

where σ_{PT} is the error in estimating a single particle position, and N is the number of particles per subregion. The measurement precision error is inversely related to the number of vectors in the subregion. The other error sources are typically larger than the measurement precision error. Computing the spatial average of the particle tracking results is a two-step process, first the particle tracking is performed and then the spatial average is computed. Hence, the error sources are not pre-convolved into the final velocity estimate; instead, they are still separate entities. The need for error correction is first determined by inspection of the particle tracking results. If the error sources can be identified and separated from the flow information, then the data quality can be improved.

Any preknowledge about the flow field can be applied via a filtering process before computing the spatial average. Preknowledge may include maximum and minimum velocities, seed particle size distribution, bias errors and spurious vectors. Spurious vectors typically result from the particle tracking operation due to loss of particles or incorrect matching, and can be considered a noise source. Spurious vectors can be removed or minimized either through the use of a robust particle tracking estimator, or through post-processing filtering. The presence of bias errors must be determined by inspection of the data. Errors arising from velocity gradients across the subregion, once considered the bane of correlation processing can now be corrected. Before computing the spatial average, or in the process of computing the spatial average, checks can be made for the existence and magnitude of velocity gradients across the subregion. The size and shape of the averaging subregion can be modified to minimize the velocity gradient effect, thereby improving the spatial resolution of the estimated velocity field. Although the measurement precision may not be as high as correlation processing, the power of the combined correlation/particle tracking technique is realized via the prefiltering capabilities afforded before the computation of the spatially averaged velocity field, producing velocity field estimates of higher accuracy (free of systematic errors) than possible using correlation techniques.

3
Fuzzy logic

The control of complex processes has been aided by the development of fuzzy control systems. Fuzzy logic employs a tolerance for imprecision to achieve system control (Zadeh 1965). An exact model for the system inputs and outputs is not required. Fuzzy inference control utilizes membership functions and a rule base developed by the user to process information. The physical mechanisms underlying the process are irrelevant to the controller. The process of identifying and tracking particles in a flow is a good candidate for fuzzy control, since the procedure is not clear cut, but involves some gray area decisions.

Fuzzy-logic principles have been used in the three separate DPIV data-processing scenarios: to directly track particles in low to moderate seed particle concentration flows, to validate correlation peaks in correlation processing, and to perform particle tracking in the combined correlation processing/particle tracking approach for high seed particle concentration flows. Each of these applications will be discussed in the following sections. First, the use of fuzzy logic to track particles in a low seed particle concentration case will be described. This application provides a foundation for the use of fuzzy logic in DPIV/PTV data processing. In the fuzzy-logic particle tracking technique, local particle displacement information is used to identify candidate particle tracks. The fuzzy inference processor is used to determine the most probable particle trajectories based on common-sense rules that an observer would use to identify particle tracks. Next, fuzzy-logic principles are used to improve the data validity in normal correlation processing. Finally, the use of fuzzy logic in the combined correlation/particle tracking processing is demonstrated.

3.1
Fuzzy-logic processor applied to PTV

The DPIV experiment is set up such that two single-exposure image frames are acquired. The particle centroids on frame #1 are used as starting points for possible particle displacements. The user specifies a search region radius, R_s, typically 10–20 pixels, to search for frame #2 particles. Each frame #2 particle within a radius R_s from the initial particle centroid is a candidate displacement vector. All possible displacements of the initial particle to the second particle locations within the search region are recorded and stored as lists of candidate displacement vectors for each initial particle. The number of candidate velocity vectors is given by $\rho\pi R_s^2$, where ρ is the average particle concentration across the image. Hence, for high data density areas or for large search regions, many initial particles may be competing for the same second exposure particle centroids. At this stage in the processing, the vector field is very convoluted and noisy, as indicated by the competing vectors shown in Fig. 1. The fuzzy inference processor operates on these lists of candidate displacement vectors to determine the most likely displacement vector for each initial particle centroid location.

The list of vectors for each initial particle is compared with all other initial particle displacement vector lists to determine whether there is any commonality. If two separate initial particles do claim the same second exposure particle, then all possible vector pairings between the candidate lists for each of these initial particles are compared. The main assumption is that, if two initial particles are close enough to interact (claim the same second particle), then the pair of vectors that look the most similar (in direction and magnitude) must be the correct pair of displacement vectors for the two separate initial particles. This assumption also holds for tertiary and higher interactions.

There are four inputs to the fuzzy PTV processor for each vector pair: distance between the vector midpoints in pixels (Sep); average vector magnitude (Mag); difference

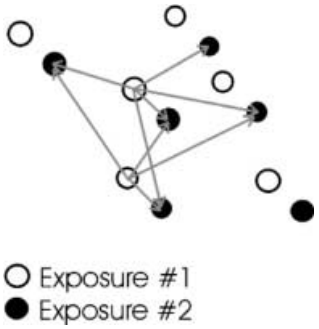

Fig. 1. Particle centroids from a small region of images #1 and #2 have been combined on the same plane. All possible candidate velocity vectors from two of the initial exposure centroids have been drawn to the surrounding second exposure centroids to illustrate the interactions between candidate vectors

in vector magnitudes (MagDif); and the sum of the squares of the differences of the x and y components of the two velocity vectors (Delta). The first (Sep) and third (MagDif) measures, when combined, act as a velocity gradient measure. The last measure operates in the opposite manner to the dot product. This measure is small when the vectors are similar in magnitude and direction, and large when they are different. Each input measure is assigned to a fuzzy set, where the degree of membership for each element in the set varies between 0 and 1. Standard 25–50% overlapping triangular input membership functions are used. The degrees of membership for each input are processed through a rule base of "IF ... THEN" blocks. The rule base defines an output fuzzy set. For a given vector pair, up to 16 rules may fire, depending on the number of unique combinations of membership values. In lieu of the more common centroiding technique, the fuzzy PTV processor output is computed via the singleton technique with a weighted average, which is computationally simpler and more efficient (Brubaker and Sheerer 1992).

The fuzzy PTV rule base is shown in matrix format in Table 1. Each input measure has three elements in its fuzzy set: {Small, Med, Large}. The input measures are computed for each vector pair analyzed. The fuzzy processor output from the rule base is the fuzzy set of confidence (Conf_Out) in the vector pair. The set of confidence singletons and their associated numerical weights are: {Very High = 1.0; High = 0.8; Med = 0.5; Low = 0.1; None = 0}. For example, in Table 1, Sep = Small, Mag = Med, MagDif = Small, and Delta = Med yields the output membership Conf_Out = Med.

The membership functions assign a numeric weight, or degree of membership, for all elements in each input measure's respective fuzzy set. The number of non-zero memberships depends on the amount of overlap of the membership functions. The midpoints of the membership functions are: Sep_Mid = R_s; Mag_Mid = $R_s/2$; MagDif_Mid = 5 pixels; Delta_Mid = 5 pixels, where R_s is the search region size in pixels, which is nominally 20 pixels. Figure 2 shows a sample set of membership functions for three sets: Small, Medium and Large. An input measure of 7.5 pixels is observed to be a member of the set Small to degree $\mu_{\text{Small}} = 0.25$ and also a member of the set Medium to degree $\mu_{\text{Medium}} = 0.75$.

All unique combinations of the degrees of membership for the fuzzy sets for each input measure are combined via a fuzzy "AND", which is the minimum value of the four membership values. The actual confidence level in the vector pair is computed by taking the product of the fired rule output value (Conf_Out) and the fuzzy "AND" of the degree of memberships. A sum of these products is computed over all of the fired rules. The total is normalized by the sum of the minimum degree of membership values for each rule fired. The output of the fuzzy processor represents the crisp estimate of the confidence in the velocity vector pair, ranging from 0 (low confidence) to 1 (high confidence).

A simplified example is shown in Fig. 3 where two exposure #1 particles are competing for two exposure #2 particles. All possible vectors are drawn and their properties are listed in Table 2. From the vectors in Fig. 3, only two candidate vector pairs are possible: V_2V_3 and V_1V_4. The other two pairings V_1V_3 and V_2V_4 are not possible,

Table 1. Rule base for fuzzy PTV processor. Each unique combination of inputs defines an element of Conf_Out (the output set), which are the values in the body of the table

	MagDif Small			MagDif Med			MagDif Large		
	Mag Small	Mag Med	Mag Large	Mag Small	Mag Med	Mag Large	Mag Small	Mag Med	Mag Large
Sep Small									
Delta Small	High	High	High	High	Med	Med	Med	Med	Low
Delta Med	High	Med	Low	Med	Med	Med	Low	Low	Low
Delta Large	Med	Low	Low	Low	Low	Low	Low	Low	Low
Sep Med									
Delta Small	High	Med	Med	Med	Med	Low	Med	Low	Low
Delta Med	Med	Med	Low	Med	Med	Low	Low	Low	Low
Delta Large	Low	Low	Low	Low	Low	Low	Low	Low	Low
Sep Large									
Delta Small	Med	Med	Low	Med	Low	Low	Low	Low	Low
Delta Med	Low	Low	Low	Low	Low	Low	Low	Low	Low
Delta Large	Low	Low	Low	Low	Low	Low	Low	Low	None

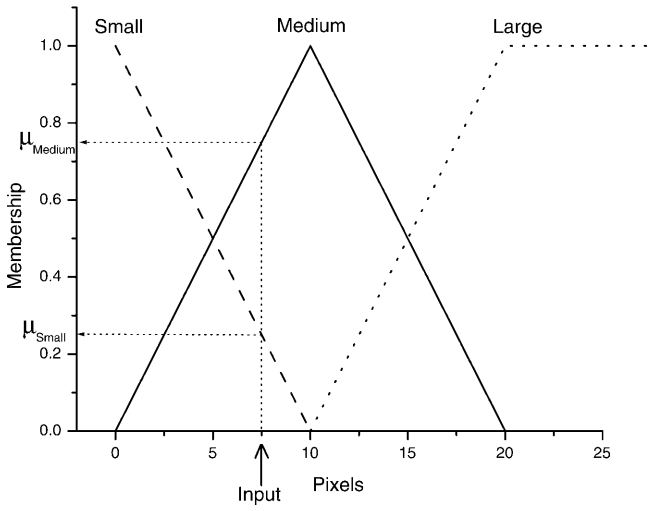

Fig. 2. Input membership functions with 50% overlap for the average vector magnitude fuzzy set

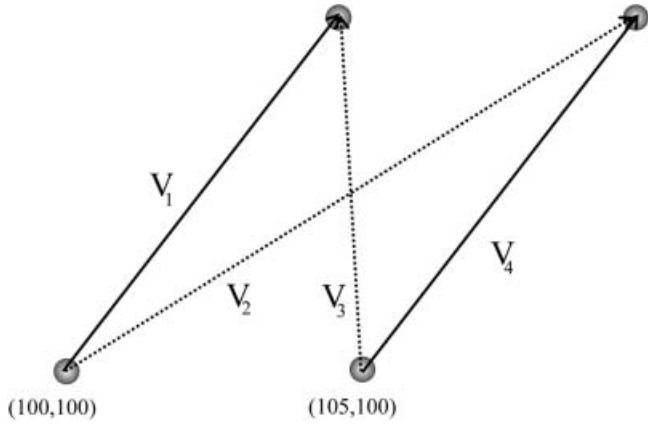

Fig. 3. Two pairs of interacting vectors with separate initial points. *Dashed line* indicates incorrect vectors, *solid line* denotes correct vectors

since the two separate initial particles are competing for the same second exposure particle, which would produce physically impossible vector pairs. The two possible combinations, and their respective input measures (all units in pixels) to the membership functions are given in Table 3.

Looking at the first case, V_2V_3 has the fuzzy set memberships:

Sep $\{\mu_{\text{Small}} = 0.92,\ \mu_{\text{Med}} = 0.08,\ \mu_{\text{Large}} = 0\}$;
Mag $\{\mu_{\text{Small}} = 0.15,\ \mu_{\text{Med}} = 0.08,\ \mu_{\text{Large}} = 0\}$;
MagDif $\{\mu_{\text{Small}} = 0.50,\ \mu_{\text{Med}} = 0.00,\ \mu_{\text{Large}} = 0\}$;
Delta $\{\mu_{\text{Small}} = 0.00,\ \mu_{\text{Med}} = 0.00,\ \mu_{\text{Large}} = 1\}$.

Table 2. Sample data for the interacting vectors shown in Fig. 3

	X_i	Y_i	$\bar{V}$ [pixels]	$\theta°$
V_1	100	100	8.5	45
V_2	100	100	11	30
V_3	105	100	6	81
V_4	105	100	7.1	51

Table 3. Fuzzy processor input measures and outputs for sample vector pair combinations

	Sep	Mag	MagDif	Delta	Confidence
V_2V_3	1.7	8.5	5	71.6	0.17
V_1V_4	4.3	7.8	1.4	2.5	0.65

With these fuzzy sets, the following four rules are fired, which are all possible combinations of the fuzzy memberships for each input set:

- IF(Sep = Small AND Mag = Small AND MagDif = Small AND Delta = Large) THEN
  - Conf_Out_1 = Med
  - $\mu\text{\_out}_1 = \text{MIN}(\mu_{\text{Sep}} = \text{Small},\ \mu_{\text{Mag}} = \text{Small},\ \mu_{\text{MagDif}} = \text{Small},\ \mu_{\text{Delta}} = \text{Large}) = 0.15$
- END IF
- IF(Sep = Small AND Mag = Med AND MagDif = Small AND Delta = Large) THEN
  - Conf_Out_2 = Low
  - $\mu\text{\_out}_2 = \text{MIN}(\mu_{\text{Sep}} = \text{Small},\ \mu_{\text{Mag}} = \text{Med},\ \mu_{\text{MagDif}} = \text{Small},\ \mu_{\text{Delta}} = \text{Large}) = 0.5$
- END IF
- IF(Sep = Med AND Mag = Small AND MagDif = Small AND Delta = Large) THEN
  - Conf_Out_3 = Low
  - $\mu\text{\_out}_3 = \text{MIN}(\mu_{\text{Sep}} = \text{Med},\ \mu_{\text{Mag}} = \text{Small},\ \mu_{\text{MagDif}} = \text{Small},\ \mu_{\text{Delta}} = \text{Large}) = 0.08$
- END IF
- IF(Sep = Med AND Mag = Med AND MagDif = Small AND Delta = Large) THEN
  - Conf_Out_4 = Low
  - $\mu\text{\_out}_4 = \text{MIN}(\mu_{\text{Sep}} = \text{Med},\ \mu_{\text{Mag}} = \text{Med},\ \mu_{\text{MagDif}} = \text{Small},\ \mu_{\text{Delta}} = \text{Large}) = 0.08$
- END IF

where the MIN( ) function is the the fuzzy "AND" operation. The fuzzy processor output confidence for this vector pair is given by:

$$\text{Confidence} = \frac{\sum_{i=1}^{N_R} \text{Conf\_Out}_i \cdot \mu\text{\_out}_i}{\sum_{i=1}^{N_R} \mu\text{\_out}_i} \tag{3}$$

$$\text{Confidence} = \frac{0.5 \times 0.15 + 0.1 \times 0.5 + 0.1 \times 0.08 + 0.1 \times 0.08}{0.15 + 0.5 + 0.08 + 0.08} = 0.17 \tag{4}$$

where N_R is the number of rules fired, and the numerical weights assigned to the Conf_Out set elements have been defined above. The fuzzy processor output confidence levels for the two-vector pairings are shown in Table 3, of which the V_1V_4 pair has the higher confidence level. This is a very simplified case; in most instances each exposure #1 particle will have interactions with several other exposure #1 particles. Initially, all of the vectors are assigned a confidence level of 0. The computed confidence level for the vector pair is compared against the current confidence level for each vector in the pair, and the maximum confidence value is stored for each vector. When all the in-

teracting vector pairs have been analyzed, the list of candidate vectors for each initial point will have confidence levels relative to one another. The vector in the list with the highest confidence level is assumed to be the most probable vector for the current initial point, and moved to the top of the list.

A second pass through the data is performed to analyze non-interacting velocity vectors lists. A non-interacting vector is defined as a vector with multiple candidate vectors for a single initial point, but does not share a common second-particle centroid with any other vector. The overall flow field pattern is used to determine the correctness of these outlying velocity vectors. The fuzzy processor is used to compare each vector in the list with all of the other previously analyzed velocity vectors in the flow field. Again, the vector with the highest confidence value in the list is assumed to be the correct vector.

A final pass through the data is performed in which the highest confidence velocity vectors for each initial particle centroid are compared. All the confidence levels are reset to 0. Each remaining valid vector is compared with all other valid vectors in the flow field. The resulting confidence level for each vector is then compared with a fixed threshold level. Only vectors with confidence levels above the set threshold level are considered valid.

The fuzzy weights in the rule base and membership functions are generic and applicable to a wide class of flow fields. For the direct particle tracking, all the particle displacements are initially considered; hence there is a lot of interaction between the adjacent particle images. For low seed particle density flows, the number of interacting vectors is processed in a reasonable amount of time. However, as the total number of particles N_P in the image increases, the computational load on the fuzzy processor grows dramatically $\approx N_P(\rho\pi R_s^2)^2$, due to the increased interaction between neighboring particles. The tracking operation is still successful, but is no longer computationally efficient. The combined correlation/particle tracking operation described below turns out to be a more efficient approach for highly seeded flows. The correlation vector map provides an efficient reference for determining the individual particle displacements, instead of determining the particle displacement from just the random particle positions between exposures.

3.2
Fuzzy-logic correlation peak detection

The correlation processing technique requires identification of the correlation peak on the correlation plane corresponding to the average displacement of particles across the subregion. Noise on the images and particle dropout contribute to spurious peaks on the correlation plane, leading to misidentification of the true correlation peak. The subsequent velocity vector maps contain spurious vectors where the displacement peaks have been improperly identified. Typically, these spurious vectors are replaced in a post-processing step by a weighted average of the neighboring vectors, thereby decreasing the independence of the measurements. In the PIVPROC program, fuzzy-logic techniques are used to determine the true correlation displacement peak even when it is not the maximum peak on the correlation plane, hence maximizing the information recovery from the correlation operation and minimizing the number of spurious velocity vectors.

The particle tracking fuzzy inference engine has been used in the PIVPROC program to detect the correct auto- and cross-correlation plane displacement peaks. Ideally, when the image data are of high quality and high seed density, the highest amplitude peak on the correlation plane represents the average displacement of particles across the subregion being processed. An example of a high signal-to-noise case is shown in Fig. 4. However, particle out-of-plane motion, velocity gradients, image noise, and low particle concentration are all contributing sources for a noise peak to be misidentified as the average particle displacement across the subregion. In these cases, the peak corresponding to the average motion of the particles across the subregion between exposures is not the highest amplitude peak, and possibly not even the second highest amplitude peak on the correlation plane. Figure 5 shows a pair of noisy input subregions and the resulting correlation plane output. Vectors have been drawn on the correlation plane, indicating the possible displacement vectors that could be derived from this correlation result. The correct average particle displacement for the subregions is the same as the case shown in Fig. 4, down and to the right. However, as is observed on the correlation plane, the brightest peak is not the one corresponding to the average displacement across the subregion. The brightest peak is up and to the right, yielding an incorrect estimate of the local flow velocity.

In the actual DPIV processing, each correlation plane is scanned for the five highest amplitude peaks, which are then stored. After all subregions in the image have been processed, the fuzzy inference operation is applied. The five highest amplitude peaks detected on each subregion correlation plane are treated as candidate velocity vectors for that subregion. The fuzzy-logic processor uses flow continuity to determine the appropriate correlation peak.

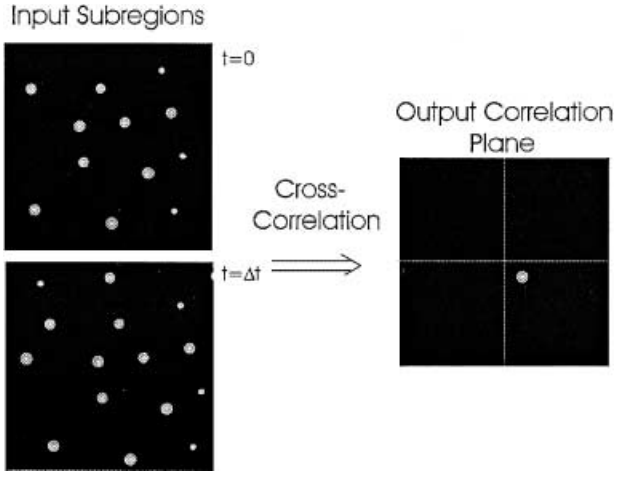

Fig. 4. Two single-exposure input subregions and the corresponding output cross-correlation plane. The location of the single bright correlation peak from the origin is the average displacement across the subregion

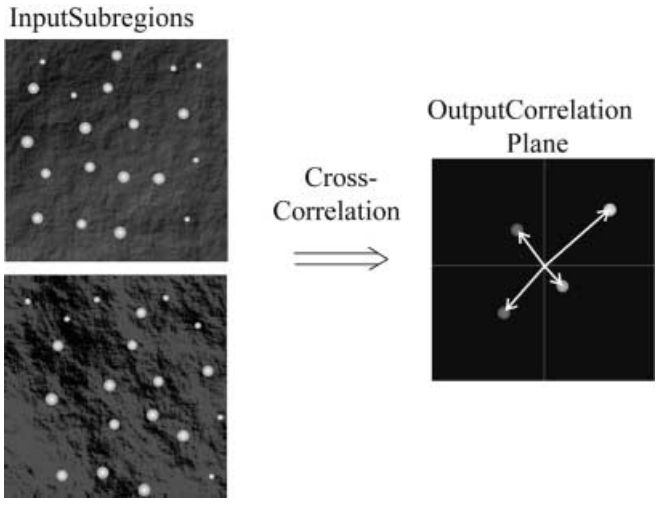

Fig. 5. Noisy input subregions result in spurious peaks on the correlation plane. The brightest peak is not always the correct displacement peak

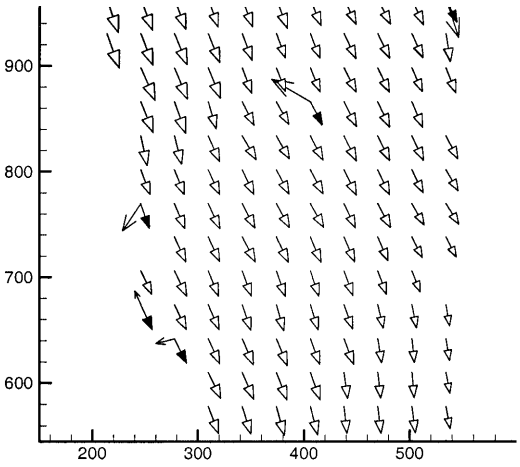

Fig. 6. A vector field containing spurious vectors, which have been correctly identified using the fuzzy processor. The main flow direction is down and to the right. The vectors with *hollow heads* represent displacements that were correctly identified by both processing techniques. The original spurious vectors are shown with *open line vector heads*. The vectors with *solid filled heads* represent displacements that were initially incorrectly identified, but have been subsequently correctly identified by the fuzzy processor

The stored correlation peaks from each subregion are compared on a pairwise basis with the results from the surrounding four subregions. The displacement peaks resulting in velocity vectors with the most similar qualities are given the highest confidence weighting. Spurious vectors will not correlate across the surrounding subregions and will thereby be discriminated. The displacement with the highest confidence weighting for each processed subregion is taken as the correct correlation displacement peak. Hence, the fuzzy inference technique is very similar to the weighted average replacement technique, except that the surrounding velocity vectors are used to identify the correct displacement peak from the correlation information, instead of merely replacing the spurious vector. Flows with sharp accelerations/decelerations (including shocks) have been processed using the fuzzy correlation peak detection technique with no adverse affects (Wernet 1995, 1999). The positions of the shocks were accurately determined and the fuzzy processor did not reduce our ability to resolve the flow structures.

A comparison of actual flow field image data processed with and without the fuzzy peak detection is shown in Fig. 6. The main flow direction is down and to the right. The vectors with hollow heads represent displacements that were correctly identified by both processing techniques. The original spurious vectors are shown with open line vector heads. The vectors with solid filled heads represent displacements that were initially incorrectly identified, but have been subsequently correctly identified by the fuzzy processor.

3.3
Correlation processing combined with fuzzy-logic particle tracking

The availability of a high-quality velocity vector map obtained from the cross-correlation operation offers the opportunity to perform particle tracking on length scales smaller than the correlation subregion size. Instead of using the interaction of nearest neighbors to determine the correct particle pairings, the correlation velocity vector map can be used as a guide for the particle tracking. Using the same pair of single-exposure image frames used to compute the cross-correlation vector map, the first exposure image is scanned for particle centroids. All second exposure particles located within a user-specified search region around each first exposure particle are detected and stored. Next, the fuzzy inference engine is employed to determine which detected particle pairings are correct. The first-exposure particles and their associated list of candidate second-exposure particles are now individually examined. The four nearest neighboring velocity vectors from the cross-correlation vector map to the initial particle location are found and used to compute a local spatially averaged velocity vector, called a "benchmark vector". The benchmark vector is then used in a pairwise fuzzy comparison with all of the candidate vectors in the list for this initial particle. The candidate vector in the list most similar to the benchmark vector is assigned the highest confidence weighting. Benchmark vectors are computed for all remaining initial particle locations and used to identify the most probable velocity vector for each initial particle. Proceeding in this manner, the correct particle pairs for all of the initial exposure particles are determined. In practice, the particle tracking operation has a success rate (properly matched particle pairs) of approximately 30–60%.

The PTV data are randomly distributed according to the seed particle distribution within the flow at the time the image data are acquired. The PTV data can be converted to a regular grid of points by computing spatial averages. Computing spatial averages of the PTV data on subregions of the same size and spacing used in the correlation operation, a direct comparison between the correlation and particle tracking results can be made. For each subregion,

both the vector tail and head coordinates are verified to be within the current subregion before they are included in the average estimate. This ensures that the same set of particle images used in each subregion to obtain the correlation results is identical to the set of tracked particles used to compute the spatially averaged tracking result.

4 Discussion

The data shown in the following figures is from the application of DPIV to the diffuser region of a high-speed centrifugal compressor. For a more thorough discussion of the centrifugal compressor facility, the DPIV installation and the measurement results obtained from the facility, see Wernet (1998). The impeller tip speed is 490 m/s at 21,000 rpm. Previous measurements of the flow exiting the impeller using LDV were in general agreement with CFD predictions of the flow (Skoch et al. 1997). Hence, the flow exiting the impeller is well characterized, with a large circumferential velocity component and a small radial component. The nominal flow velocity at the impeller exit is 375 m/s due to the 50° backsweep of the impeller vanes from radial at the exit. As the flow migrates out across the vaneless space into the diffuser passage, the radial component is reduced as the flow turns to match the diffuser vane surface angle.

A 1,000 × 1,000 pixel cross-correlation CCD camera with 9-μm pixels was used to acquire the image data. Optimally imaged particles are obtained from the 0.16 magnification optical system by setting the recording camera lens to *f*/5.6 so that the diffracted images of the seed particles span 1–2 pixels on the CCD camera sensor (Wernet and Pline 1993). The time between laser pulses is 1.8 μs and the imaged field of view is approximately 65 × 65 mm. A once-per-rev signal generated from the impeller shaft is used to trigger the laser firing in order to capture DPIV image pairs at a particular circumferential orientation of the impeller. The compressor flow field is globally seeded with alumina particles which have a size distribution mean diameter and standard deviation of 0.7 ± 0.2 μm. The particles are dispersed in a pH-stabilized ethanol dispersion which is atomized in the plenum section of the compressor facility (Wernet and Wernet 1994). The ethanol evaporates, leaving the dry alumina seed particles in the flow stream. Using the particle dynamics results of Melling (1997), these alumina particles should exhibit a frequency response of 3 kHz. The light sheet is inserted into the flow using a periscope-type probe. The location of the illuminated plane is at 12% span (0% being at the diffuser hub and 100% being at the casing wall). For the measurements shown here, the light sheet propagates upstream through the diffuser passage. The diffuser vanes are overdrawn in the raw DPIV image and drawn in the processed vector plots.

4.1 Instantaneous vector maps

The centrifugal compressor data were selected since the flow has previously been characterized and due to the readily observable particle dynamic effects. A sample raw DPIV image file is shown in Fig. 7, illustrating the seed particle concentration variation across the image. A high concentration of seed particles is observed at the top of the image, which corresponds to the through-flow fluid exiting the impeller. The impeller blades are 2 mm thick at the top and 17 mm thick at the base. At this deep immersion in the blade passage, the impeller blades are near their maximum thickness (15 mm), and evidence of the viscous blade wake is observed in the region just downstream from the impeller blade. The viscous blade wake region has a markedly lower seed particle concentration. Just below the blade is the through-flow fluid packet (again moderate seed particle concentration) from the previous blade passage, and at the bottom of the image we see the remnants of the viscous blade wake from the previous blade, which is now out of the field of view. Again, this lower blade wake fluid has a low concentration of seed particles.

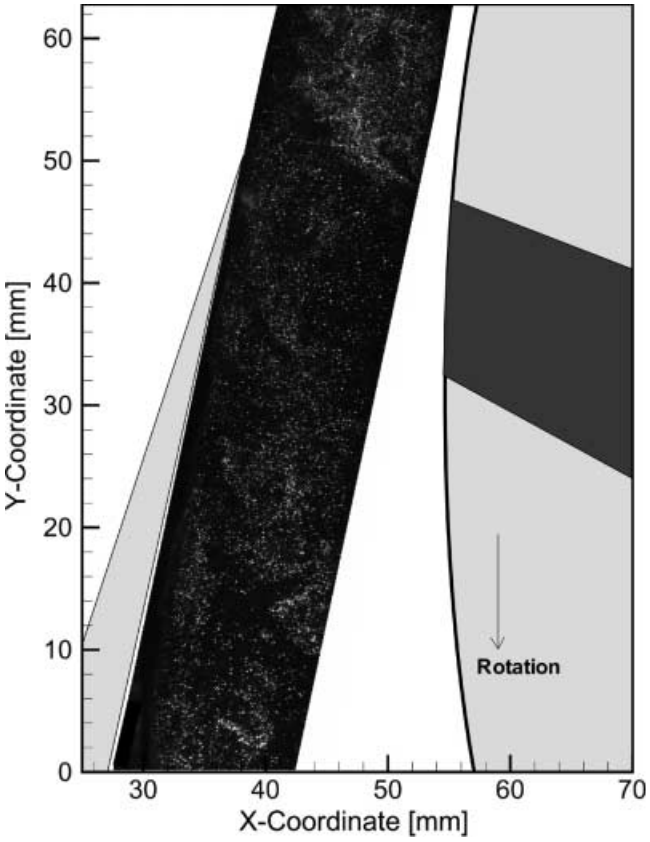

Fig. 7. Section of raw PIV image file overdrawn on a schematic layout of the compressor facility. The impeller is seen on the *right* along with the impeller blade position. The illuminated plane shown here is at 12% span. The light sheet is propagating *from the bottom to the top* of the figure

The correlation results obtained by using a 64 × 64 pixel subregion size with 75% overlap are shown in Fig. 8a. The velocity vectors represent the angle and magnitude of the measured velocity. In addition, the velocity vectors are color coded by their magnitude, in order to accentuate the changes in the flow. The velocity vector map reveals no evidence of the underlying large variations in seed particle concentration as observed in the raw image file shown in Fig. 7. The viscous blade wake region is characterized by a lower-velocity magnitude and more radial flow angle

than the higher speed flow observed above and below the blade.

The particle tracking operation uses the correlation processed data shown in Fig. 8a as an input. The tracking operation is tolerant of some spurious vectors in the correlation vector map; however, if too many spurious vectors are located in a concentrated region, then the tracking operation may not perform optimally in that region. The first stage in the particle tracking operation is the particle centroid detection. A global image threshold level of 13 gray levels was used to detect over 6,400 particles in each image. The results of the tracking operation are shown in Fig. 8b, where over 4,000 velocity vectors have been tracked, yielding a tracking success rate of over 60%. The spatial resolution of the velocity measurements was 6 vectors/mm^2 across the illuminated region of the flow, and spatial resolutions as high as 8 vectors/mm^2 were realized in some regions. The complete particle detection and particle tracking operation completes in less time than required to complete the correlation operation. The particle tracking vector map clearly indicates the underlying seed particle concentration in the image. A high concentration of vectors is observed in the topmost portion of the image, where the through-flow fluid exits the impeller. A small blank region in the high-speed flow at the top of the image is caused by an excessively high particle concentration, which precludes the software from resolving the individual particle images due to the use of a global threshold level for particle centroid detection. The viscous blade wake region has a lower concentration of vectors, as predicted by the low seed particle concentration in the raw image file. The more sharply turning velocity vectors are readily observed in this region. At the bottom of the image, we observe the most sparsely seeded region of the flow.

The randomly sampled particle tracking data were then processed to compute a spatially averaged velocity field, corresponding to 64 $\times$ 64 pixel subregions with 75% overlap. Figure 8c shows the spatially averaged velocity field, which agrees very closely with the 64 $\times$ 64 pixel subregion correlation processed data shown in Fig. 8a. The small blank region at the top of Fig. 8b results in a low velocity region in the spatially averaged vector map in Fig. 8c. Otherwise, the good agreement between the spatially averaged results and the correlation results confirms that the particle tracking operation is recovering essentially the same information as the correlation operation.

At this point, we can see that the particle tracking operation has revealed information on the underlying seed particle concentration, a fact that is usually masked by correlation processing. A closer inspection of the particle tracking data in Fig. 8b reveals that there are a significant number of low-magnitude velocity vectors that have a large radial flow angle. The flow should exit the impeller and turn to follow the diffuser vane surface. These vectors represent large particles that are exiting the impeller with a large radial velocity component and not following the flow. These particles must be larger than the mean particle size of 0.7 μm diameter, and are obviously biasing the spatially averaged velocity estimates. As discussed above, the particle images recorded on the CCD array result entirely from diffraction, therefore, no determination of the true particle size can be made from the raw DPIV image data. However, since the individual particles have been tracked, the velocity bias caused by these low-velocity particles can easily be filtered out.

The criteria used to discriminate the biased vectors will ultimately determine the quality of the data. Two filtering techniques were investigated: hard velocity cutoff and an automated spurious vector removal technique. Preferably, an automated statistical technique could be applied which would rely on the local flow properties to remove spurious or biased velocity vectors. The automated procedure for removing outliers used here is based on Chauvenet's criterion (Taylor 1982), in which the probability of occurrence of a given point deviating from the mean is computed. The main assumption here is that the parent velocity distribution is Gaussian. For each subregion, the mean and standard deviation are computed. Then the number of standard deviations that each measurement in the subregion lies from the mean is computed. The probability that a given measurement could deviate from the mean by this many standard deviations is computed from the Normal Error Integral and multiplied by the number of points in the distribution. If the computed probability is less than a preset level, then the point is removed. This technique has been used with good success for spurious vector removal when computing the time-averaged mean velocity from a series of correlation-processed DPIV measurements. When applied to the collection of particle tracking measurements in each spatially averaged subregion, the automated procedure did not perform as well as a simple velocity cutoff filter.

Application of a hard velocity cutoff filter was justified, based on inspection of the particle tracking measurements and what was known about the flow exiting the compressor. The mean impeller exit velocity was known to be 375 m/s due to the rotational speed of the impeller, yet velocities less than 225 m/s were observed with flow angles markedly different from the mean flow. Velocity fluctuations due to flow turbulence in the impeller exit region are not expected to exceed 15% (which will be verified via the time-average measurements below), yet these low-velocity vectors deviated from the mean by nearly 45%. The low velocity coupled with the large radial flow angle identified these measurements as originating from large particles that were not accurately following the flow. A low-velocity cutoff limit of 225 m/s was thus imposed on the data.

Figure 9a shows the 294 velocity vectors with magnitudes of 225 m/s and below that have been removed from Fig. 8b. The recomputed spatially averaged velocity vector map from the filtered data is shown in Fig. 9b. The velocity vector map is now more uniform and shows smaller regions of low-velocity flow than the biased velocity vector maps shown in Fig. 8a and c. The average percentage change in the velocity vector magnitudes in the filtered vs unfiltered spatially averaged particle tracking data is shown as a color surface plot in Fig. 10a. Most of the vector map is little or unaffected, as indicated by the 0% change. The velocity bias is most significant in regions where the low-velocity vector populations were highest, as shown in Fig. 9a. In these regions, the average change in

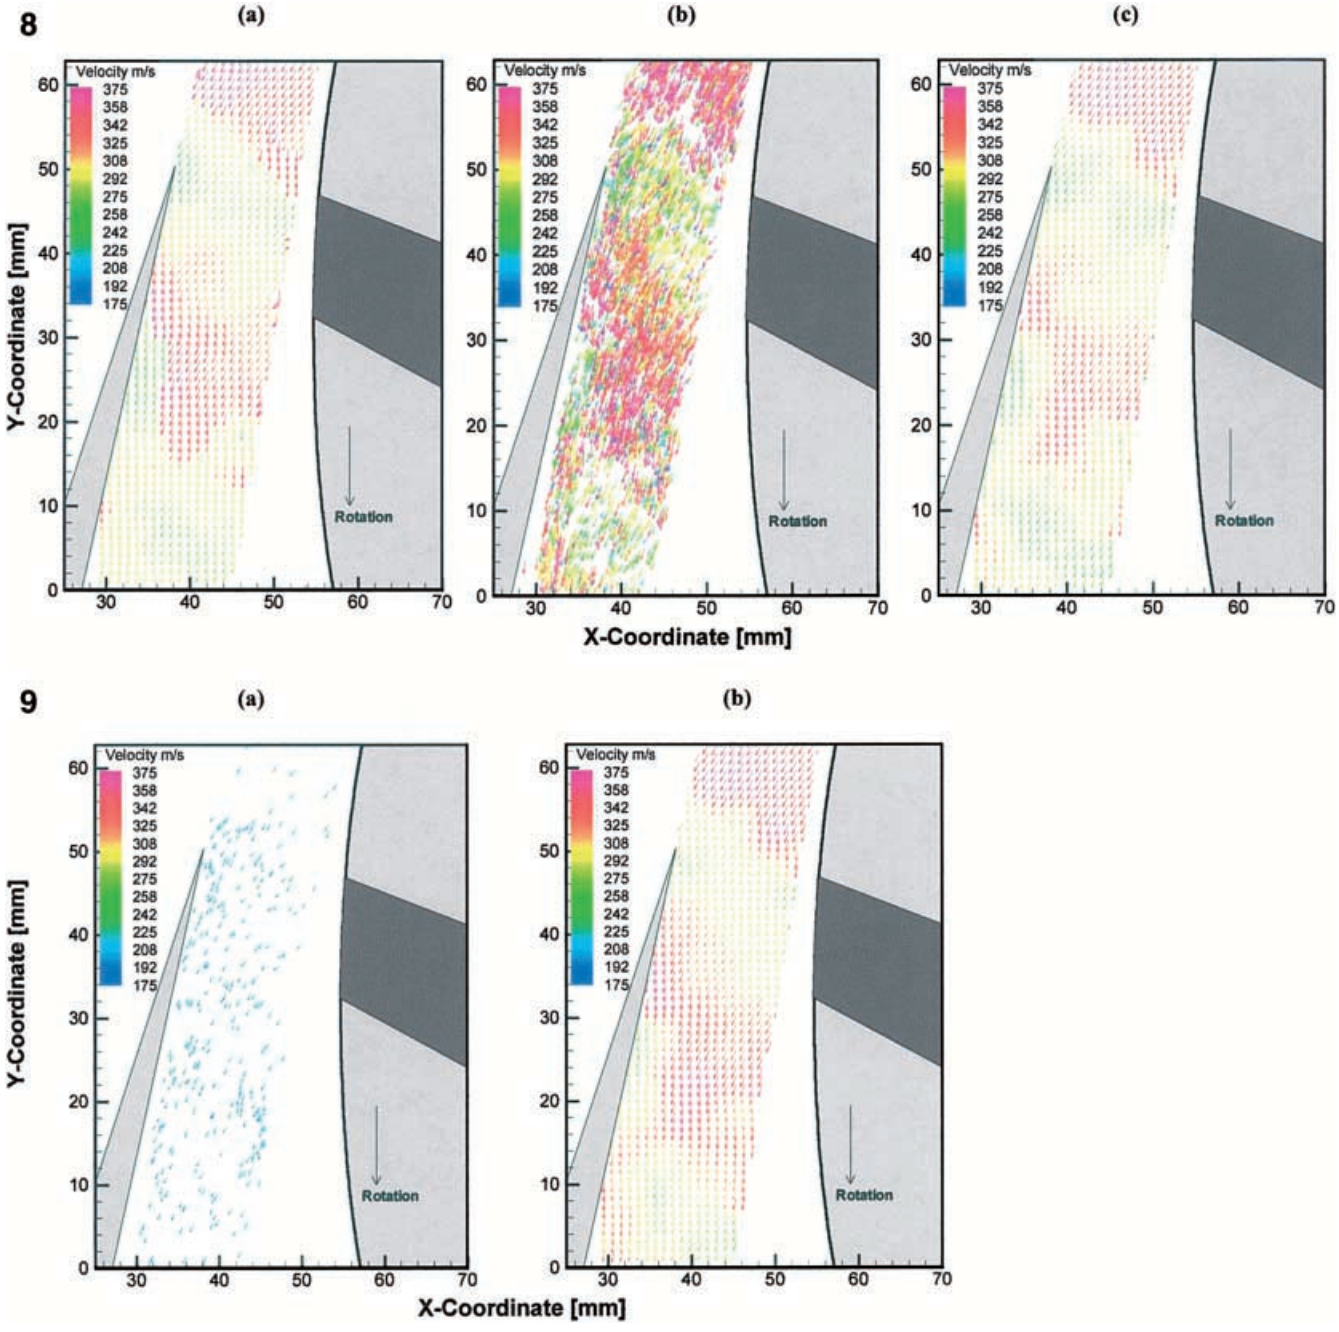

Fig. 8a–c. Processed vector fields: **a** cross-correlation vector map using 64 × 64 pixel subregions with 75% overlap; **b** particle tracking result using correlation map in **a** as a guide, over 4,000 vectors have been tracked; **c** spatial average of particle tracking results using same grid spacing and subregion size used in the correlation operation in **a**

Fig. 9. **a** Vector map showing the spatial distribution of the velocity vectors with magnitude <225 m/s that have been removed from Fig. 8b; **b** Spatial average of the filtered particle tracking data. The velocity magnitudes are now larger than they were in the spatial average of the unfiltered tracking data shown in Fig. 8c

the velocity ranges from 7 to 11%. These rather significant increases in the mean velocity demonstrate the detrimental effect that a relatively small population of low-velocity particles can have on the velocity vector maps.

Since the individual particle velocities have been measured prior to computing the spatially averaged velocity vector maps, we can also calculate the rms velocity at each subregion. The rms variation in the unfiltered, spatially averaged velocity data from Fig. 8c is plotted in Fig. 10b. The rms variations are shown as a color-coded surface plot and the values range from 12 to 30%. Figure 10c shows the rms velocities for the low-velocity cutoff filtered data from Fig. 9b, where the magnitude of the variations has been significantly reduced. The most significant reductions are observed corresponding to the spatial regions containing the largest populations of low-velocity

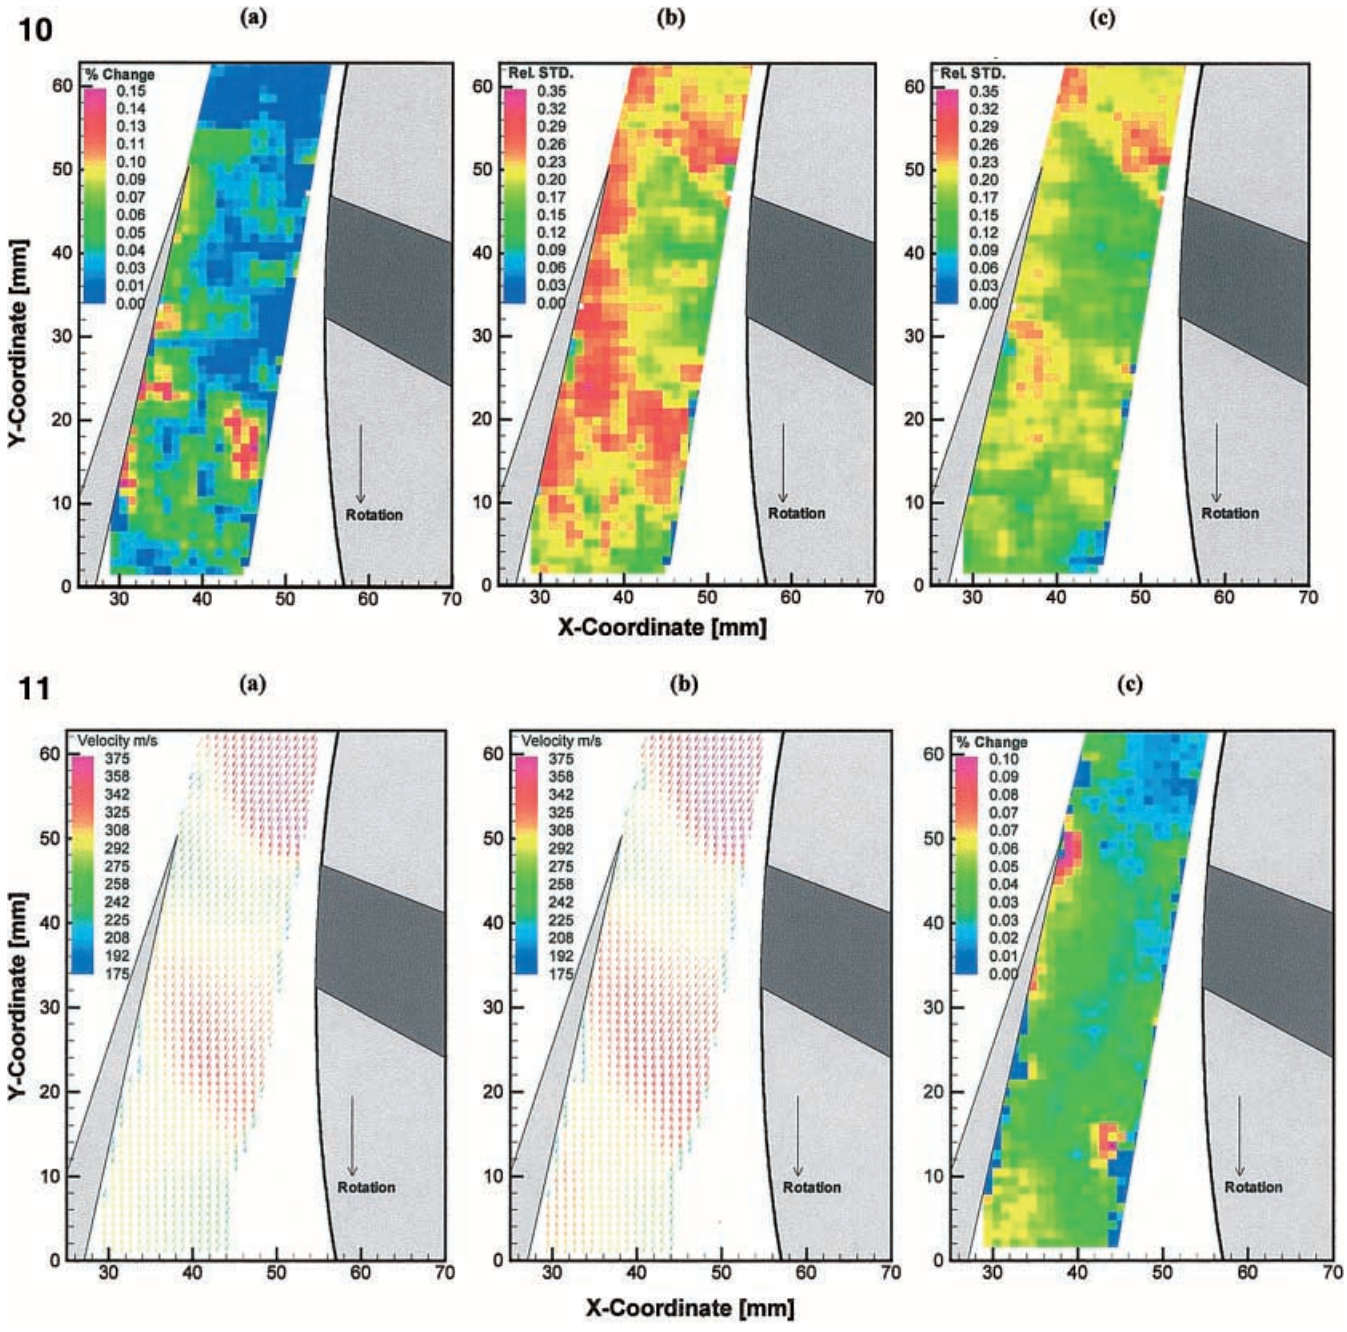

Fig. 10. **a** Change in velocity between the filtered and unfiltered spatially averaged velocity vector maps shown in Figs. 8c and 9b respectively; **b** Rms variation in the unfiltered spatially averaged particle tracking result. The largest variations are observed where the low-velocity vector concentrations are highest; **c** Rms variation in the filtered, spatially averaged velocity vector map. The variations caused by the low-velocity particles have been reduced, yielding a lower magnitude, more uniform rms map
Fig. 11a–c. Time-average vector fields using 50 frames: **a** spatially averaged PTV data using 64 × 64 pixel subregions with 75% overlap, no filtering; **b** time average as in **a**, 225 m/s low-velocity cutoff filter applied to PTV data; the high-velocity regions have increased in size; **c** percentage change in filtered relative to unfiltered data in Fig. 11a and b, illustrating an almost constant 4% increase in the velocity across the entire measurement region

magnitude vectors that were filtered out before computing the spatial average.

4.2
Time-averaged vector maps

The discussion above concerned large particle bias correction with regard to a single instantaneous velocity vector map. Time-average measurements of the compressor flow are required to reduce the measurement uncertainty and to ensure that high-quality velocity vector maps are obtained for comparison with CFD predictions. Therefore, DPIV image data were collected in 100-frame image sequences, from which 50 velocity vector maps could be obtained by processing the frame pairs. The once-per-rev synchronization electronics ensured that the impeller was in the same location each time an image pair

was acquired. These 50 velocity vector maps were then used to compute the time-averaged velocity vector map of the compressor flow.

A sequence of 50 frame pairs was acquired at the same compressor span and operating conditions used to acquire the instantaneous velocity vector map discussed above. The 50 frame pairs were processed via the same correlation/particle tracking procedure. Spatially averaged vector maps were generated from the PTV vector maps. The 50 spatially averaged PTV maps were then used to compute the time-average flow field and the relative standard deviation of the time-average flow. Figure 11a shows the time-average flow of the spatially averaged PTV data. The nominal flow features are similar to the instantaneous velocity fields shown in Fig. 8a and c. The computed relative standard deviations for this time-average flow range from 5% up to 13%. Assuming the DPIV measurement error is on the order of 1% and that there are no other significant error sources in the measurements (for now the large particle bias error will be ignored), the relative standard deviation is an approximate indicator of the flow turbulence intensity. These measurements confirm the flow turbulence intensity in the diffuser region of the centrifugal compressor does not exceed 15%, and validates the rationale for applying the low-velocity cutoff filter to the data. In Fig. 11b, the 225 m/s low-velocity cutoff filter has been applied to the PTV data prior to computing the spatial and time averages. The filtered flow field shows larger regions of high-velocity fluid than the unfiltered data in Fig. 11a. The spatial extent of the low-velocity viscous blade wake region has also been noticeably reduced. The ratio of the filtered to unfiltered data, illustrating the percentage change in the flow between the two cases, is shown in Fig. 11c. Recall that the percentage change in the filtered to unfiltered instantaneous velocity data shown in Fig. 10a showed isolated regions of significant increases in velocity, which were coupled with the local distribution of large particles in the flow. In contrast, the ratio of the filtered to unfiltered time-averaged data shown in Fig. 11c is very uniform. There is a region of large percentage change observed in the lower right region of the flow, which is caused by either generally low overall particle concentrations in this region or an excessive concentration of low-velocity particles. The change in the flow magnitude is approximately constant at 4% across the measured region. The uniformity arises from the statistical averaging over the 50 frames. Over the course of the 50 frames, large particles will have been evenly distributed throughout the flow field; hence, their effect will be present throughout the measurement area. The effect of the large particles in the time-average flow (4%) is not as large as the effect observed in the instantaneous case (7–11%). However, at 4%, the effect on the mean flow is significant and correction is warranted.

5 Conclusions

Direct correlation processing of DPIV image data masks the underlying seed particle uniformity in the flow and is susceptible to velocity bias from particles not accurately following the flow. The combined correlation/particle tracking operation has demonstrated that "super-resolution" velocity measurements are possible which can reveal small-scale structures in the flow. Over 4,000 velocity vectors were measured with an average density of 6 vectors/mm^2. The time required to perform the particle tracking operation is approximately the same as that required to perform correlation processing of the image. The particle tracking results were then spatially averaged on a regular grid using subregions of identical size and spacing as used in the correlation operation. The spatially averaged tracking results agreed very closely with the correlation results, indicating that the particle tracking operation is extracting the same information from the DPIV images as the correlation operation.

Particle lag effects are always of concern in high-speed flows. When possible, seed particles are selected that will accurately follow the flow. In the present study, refractory seed particles with a high specific gravity were required to endure the high flow temperatures. Although efforts were made to ensure that the particles were small enough to follow the flow and dispersed so that they did not agglomerate, some large particles were observed in the particle tracking results. These large, low-velocity particles, which were biasing both the correlation and spatially averaged tracking results, were easily removed from the particle tracking data using a low-velocity cutoff filter. Not all of the flow regions were biased. In regions where the large particle concentrations were high, the velocity bias was determined to be on the order of 7–11% of the velocity vector magnitude for the instantaneous velocity vector maps. The spatially averaged, filtered tracking results yielded a vector map with larger regions of high-velocity flow. In addition, inspection of the rms velocities obtained from the spatially averaged PTV results showed significant reductions in regions where the low-velocity particles were removed. In contrast, the time-averaged PTV velocity vector maps showed a more spatially uniform effect of the large particle bias removal, producing a uniform increase in the velocity of approximately 4% across the measured region of the flow.

The combined correlation/particle tracking processing technique yields excellent insight into the underlying particle dynamics in DPIV image data, which is only possible using particle tracking techniques. Particle tracking techniques provide the only means possible for removing large particle velocity bias effects from DPIV image data. In general, we expect the velocity biases resulting from larger particles to occur in regions of high acceleration/deceleration, or sharply turning flows. In the present work, the velocity bias resulted from particle lag of large particles. In other flows, there may different kinds of biases that arise (velocity gradients) that must be removed or minimized to improve the quality of the velocity estimates. In each flow measurement case, the appropriate filtering technique that minimizes the bias must be determined without compromising the integrity of the data. The combined correlation/particle tracking processing technique enables the identification of these biases in the DPIV data and provides the opportunity for their removal, without imposing a severe processing time penalty.

References

Brubaker DI; Sheerer C (1992) Fuzzy logic system solves control problem. Electronic Design News, 18 June, pp 121–127

Cenedese A; Romano GP; Paglialunga A; Terlizzi M (1992) Neural net for trajectories recognition in a flow. In: Proceedings of the 6th International Symposium on Applications of Laser Techniques to Fluid Mechanics, Instituto Superior Técnico, Lisbon, Portugal, pp 27.1.1–27.1.5

Hart DP (1998) Super-resolution PIV by recursive local-correlation. In: Proceedings of the VSJ–SPIE98, Yokohama, Japan, 6–9 December

Keane R; Adrian RJ (1990) Optimization of particle image velocimeters. I. Double pulsed systems. Meas Sci Technol 1: 1202–1215

Keane R; Adrian RJ; Zhang Y (1995) Super-resolution particle image velocimetry. Meas Sci Technol 6: 754–768

McLaughlin DK; Tiederman WG (1973) Biasing correction for individual realization of anemometer measurements in turbulent flows. Phys Fluids 16: 2082–2088

Melling A (1997) Tracer particles and seeding for particle image velocimetry. Meas Sci Technol 8: 1406–1416

Skoch GJ; Prahst PS; Wernet MP; Wood JR; Strazisar AJ (1997) Laser anemometer measurements of the flow field in a 4: 1 pressure ratio centrifugal impeller. ASME 97-GT-342, June

Taylor JR (1982) An introduction to error analysis. University Science Books, Oxford University Press, Mill Valley, Calif.

Wernet JH; Wernet MP (1994) Stabilized alumina/ethanol colloidal dispersion for seeding high temperature air flows. In: Proceedings of the ASME Symposium on Laser Anemometry: Advances and Applications, Lake Tahoe, Nev., 19–23 June

Wernet MP (1993) Fuzzy logic particle tracking velocimetry. In: Proceedings of the SPIE Conference on Optical Diagnostics in Fluid and Thermal Flow, San Diego, Calif., 11–16 July

Wernet MP (1995) Fuzzy inference enhanced information recovery from digital PIV using cross-correlation combined with particle tracking. SPIE Conference on Optical Diagnostics in Fluid and Thermal Flow, 2546, 9–14 July

Wernet MP (1998) Digital PIV measurements in the diffuser of a high speed centrifugal compressor. AIAA-98-2777, 20th Advanced Measurement and Ground Testing Conference, Albuquerque, N.M., 15–18 June

Wernet MP (1999) Fuzzy logic enhanced digital PIV processing software. 18th International Congress on Instrumentation for Aerospace Simulation Facilities (ICIASF), Toulouse, France, 14–17 June

Wernet MP; Pline A (1993) Particle displacement tracking technique and Cramer–Rao lower bound error in centroid estimates from CCD imagery. Exp Fluids 15: 295–307

Westerweel J (1997) Fundamentals of digital particle image velocimetry. Meas Sci Technol 8: 1379–1392

Willert CE; Gharib M (1991) Digital particle image velocimetry. Exp Fluids 10: 181–193

Zadeh LA (1965) Fuzzy sets. Inf Control 8: 338–352