

Acceleration of Tomo-PIV by estimating the initial volume intensity distribution

N. A. Worth · T. B. Nickels

Received: 9 January 2008 / Revised: 15 March 2008 / Accepted: 28 March 2008 / Published online: 8 June 2008
© Springer-Verlag 2008

Abstract Tomographic particle image velocimetry (Tomo-PIV) is a promising new PIV technique. However, its high computational costs often make time-resolved measurements impractical. In this paper, a new pre-processing method is proposed to estimate the initial volume intensity distribution. This relatively inexpensive “first guess” procedure significantly reduces the computational costs, accelerates solution convergence, and can be used directly to obtain results up to 35 times faster than an iterative reconstruction algorithm (with only a slight accuracy penalty). Reconstruction accuracy is also assessed by examining the errors in recovering velocity fields from artificial data (rather than errors in the particle reconstructions themselves).

1 Introduction

One of the characteristic features of turbulent flow is its three-dimensionality. Recent advances in particle image velocimetry (PIV) have enabled resolution of full three-dimensional (3D) velocity fields [Holographic PIV, HPIV (Hinsch 2002); scanning PIV (Hori and Sakakibara 2004); particle tracking velocimetry, PTV (Maas et al. 2004); and defocussing PIV (Willert and Gharib 1992)]. However, despite their ability to capture the 3D flow field, each of the aforementioned 3D PIV/PTV techniques has a number of limitations. These are discussed by Elsinga et al. (2006), who recently proposed a new PIV technique called

tomographic-PIV (Tomo-PIV). In Tomo-PIV, illuminated particles within a predefined volume are imaged from a number of angles (typically employing between 3 and 5 cameras), and through application of a computational tomographic reconstruction algorithm, the light intensity through the volume can be reconstructed. These light intensity volumes can then be directly cross-correlated, to obtain 3D velocity vector fields.

Tomographic-PIV is capable of producing high quality, fully three-dimensional velocity fields, using a completely digital approach. Advantages over other 3D techniques include—simpler setup and higher speed sampling rates than HPIV and scanning-PIV, truly instantaneous measurements compared to quasi-instantaneous in scanning-PIV, and higher possible seeding densities in comparison with PTV and defocussing-PIV. However, Elsinga et al. (2006) suggest that seeding density is limited in Tomo-PIV, with both scanning and HPIV offering higher spatial resolution. It is also noted that technique accuracy is limited by imperfect reconstructions caused by optical defects (Wieneke and Taylor 2006), and relies heavily on high precision calibration (Elsinga et al. 2006). Furthermore, the reconstruction problem is computationally intensive, and obtaining time-resolved properties from hundreds of vector fields requires significant resources (Schröder et al. 2006).

To understand the source of these costs, the reconstruction problem should be considered. An algebraic series expansion approach is commonly used, permitting basis functions to be combined linearly to approximate the object, reducing the reconstruction problem to the form, $Wf = p$ (Herman 1980), where p is some measurable projected property; f represents the object; and W is a two-dimensional (2D) weighting matrix, which provides a relationship between every discretised projected data point and every discretised point in the object volume. Iterative

N. A. Worth (✉) · T. B. Nickels
Department of Engineering, Cambridge University,
Trumpington Street, Cambridge CB2 1PZ, UK
e-mail: naw30@cam.ac.uk

methods are often the tools of choice for image reconstruction due to their simple, versatile nature, ability to handle constraints and noise, and process large data sets efficiently (Natterer 1999). In Tomo-PIV therefore, the high computational costs stem from processing huge numbers of voxels (volume elements) in the discretised volume using an iterative algorithm (see Sect. 4.1).

Elsinga et al. (2006) performed parametric studies using a 2D computational model, distributing Gaussian particles randomly in a 2D area, around which a number of 1D camera planes are positioned (Fig. 1). Artificial images, created by projecting the object onto each camera plane, can be used with camera positional information to reconstruct the object tomographically using the multiplicative algebraic reconstruction technique (MART) algorithm (Eq. 1). This algorithm updates each voxel sequentially for every intersecting image pixel line of sight, with the magnitude of the update dependent on the pixel intensity, p_i , the current projection, W_{ij}^k , and the current voxel intensity, W_{ij} .

$$f_j^{k+1} = f_j^k \left(\frac{p_i}{W_{ij}^k} \right)^{\mu W_{ij}} \quad (1)$$

In the study by Elsinga et al. (2006) accuracy was assessed using a correlation coefficient, $Q = f_a \cdot f_r / \|f_a\| \|f_r\|$, between vectorised artificial, f_a , and reconstructed, f_r , 2D areas, hereafter referred to as the reconstruction correlation coefficient (RCC) method. Although Q is a measure of correlation between reconstructed and artificial objects, the later use of these intensity fields is not considered. It is obvious that close agreement between the fields (high Q values) will give better PIV results after cross-correlation, but it is more difficult to interpret lower Q values. For example, Elsinga et al. (2006) discuss the reduction in accuracy due to *Ghost Particles*; regions of intensity formed by the coincidental convergence of camera lines of sight where no real particle exists. Although these will obviously lower the Q value, if *Ghost Particles* are uncorrelated, as expected, then their effect after cross-correlation may be insignificant. Additionally, the variation of the Q value with camera angle and calibration error

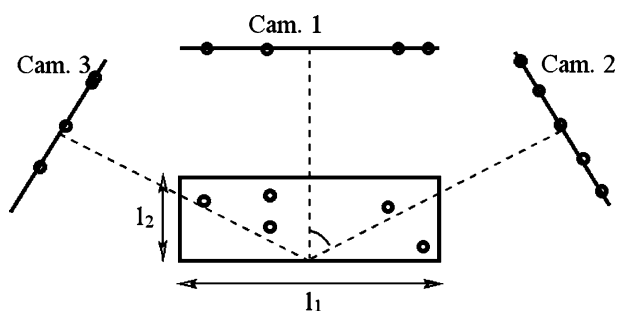


Fig. 1 Schematic of numerical 2D Tomo-PIV setup

(Elsinga et al. 2006) may be simply a result of stretched or misshapen particles, and although these may deviate from the original artificial distribution, their effect on cross-correlation is yet to be determined.

This paper covers the following two areas: firstly a new first guess method will be introduced, in order to reduce the computational costs of this technique; secondly this method will be assessed in conjunction with a number of other setup parameters using a cross-correlation based analysis, hereafter referred to as the cross correlation velocity (CCV) method. By taking account of the error in recovering the velocity field from the reconstructed volumes, this is perhaps a more applicable measure for this application. Wieneke and Taylor (2006) used a similar CCV approach to assess Tomo-PIV for a “fat light sheet”, using real and synthetic data to quantify the accuracy of the out-of-plane gradient at several seeding densities. The present investigation extends these findings by considering a number of other setup parameters.

2 A new method for initialising the volume intensity distribution

Using an iterative algebraic approach requires an initial volume intensity distribution to be set, which can then be updated using a reconstruction algorithm. The most basic first guess is uniform intensity throughout the volume; a solution commonly employed in other investigations (Elsinga et al. 2006, 2007). As in CFD, a poor first guess will slow solution convergence, with a uniform intensity distribution forcing the algorithm to search every voxel during the initial iteration. A better approach would be to set the intensity field so that only regions that will contain blobs of intensity after the reconstruction are given an initial value. In other words, find regions within the volume to which non-zero intensity contributions are made by all cameras, and set any voxels outside these regions to zero intensity. This is justified by considering the update of a voxel within the volume: irrespective of updates from other high intensity pixels, if just one strongly correlated pixel contains zero intensity this will dominate the voxel update, and its intensity will tend towards zero.

Several new first guess methods were created by calculating the current projection $p^T W$ of each camera through the volume (Fig. 2). This extends intensity from each pixel along its line of sight, creating a series of constant intensity streaks through the volume for each camera. Overlapping regions were either set to a uniform value, summed, or multiplied (Fig. 3), creating different initial fields. This places intensity only where all lines of sight converge, and in the latter two cases the magnitude of that intensity becomes dependent on all camera contributions. If any

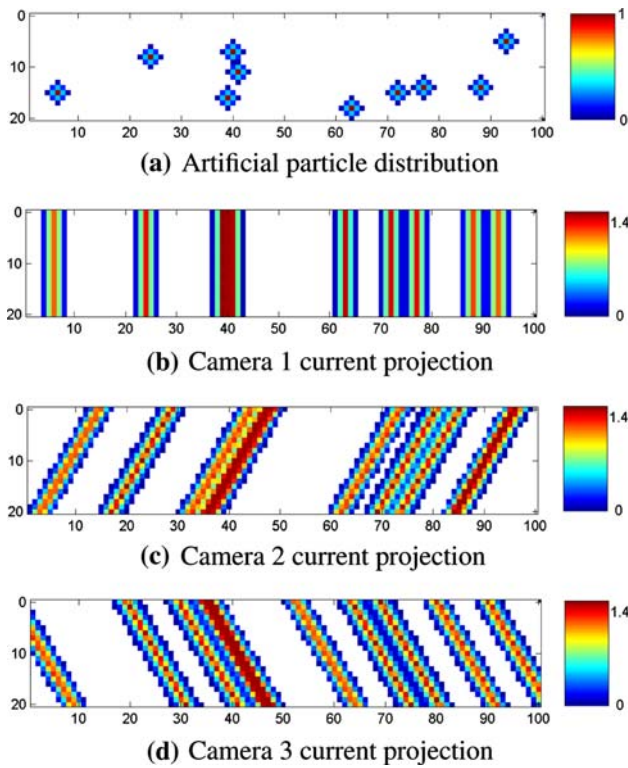


Fig. 2 2D projections (3 cams at 0° and $\pm 30^\circ$)

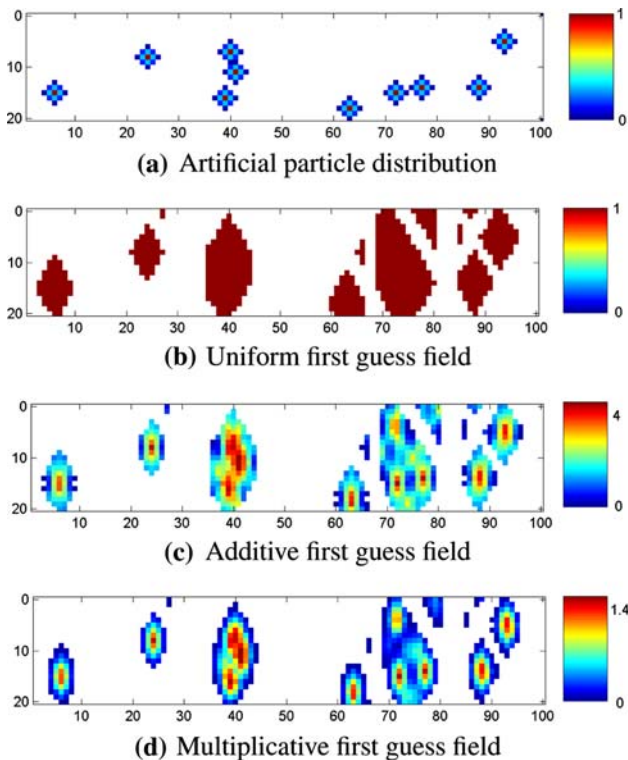


Fig. 3 First guess scheme comparison (3 cams at 0° and $\pm 30^\circ$)

contribution to a particular voxel is zero, then the voxel intensity will be zeroed; accelerating the inevitable fate of the voxel under the MART scheme. In the case of the multiplicative first guess (MFG) scheme, the intensity magnitude is also normalised by raising it to the power $1/N_{\text{cam}}$, where N_{cam} is the number of cameras, to ensure relative field magnitude continuity. Thus, the elimination of a huge number of voxels from the iterative calculation is performed through simple matrix multiplication, without requiring the expense of the MART algorithm (Eq. 1), thereby reducing computational costs significantly (detailed in Sect. 4.1).

The currently proposed MFG scheme appears to offer the most accurate estimate of the solution, and is therefore recommended over the other previously mentioned schemes. The effect of adding cameras to the system can also be shown to further increase the accuracy of this first guess (see Fig. 4), offering more well-defined intensity distribution in local regions of high particle density, and more effective elimination of *Ghost Particles*.

Although not considered in this study, an additional PIV-based first guess mechanism could be used. If the maximum possible particle displacement is well understood, the probable region in which a particle may be found in a subsequent frame of a pair of images, or indeed subsequent pairs of images should be determinable to within certain limits. Therefore a spherical radius of possible

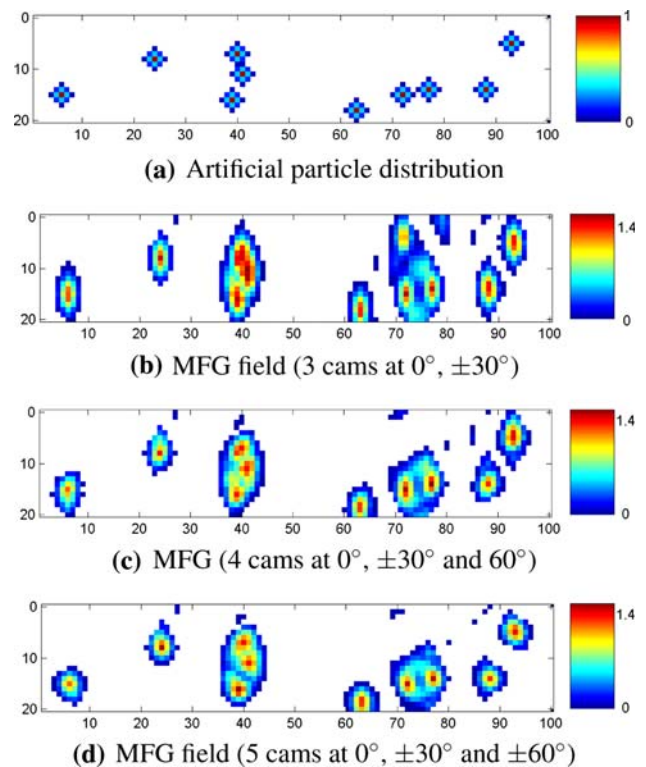


Fig. 4 Effect of additional cameras on the MFG intensity field

intensity surrounding the likely particle location could be created, which could also take account of the mean flow velocity (shown in Fig. 5). Taking account of local flow velocity could permit a tighter radius of probable locations, although this would be more involved. Combined with the previous criterion, further computational savings may be possible. However, highly intermittent flow may require the use of large radii, and the implementation costs of this scheme may render it unworkable.

3 Computational setup

The MART reconstruction algorithm (Herman and Lent 1976) has been implemented in FORTRAN77 to create a 3D Tomo-PIV program. This follows a similar methodology to that of Elsinga et al. (2006), using similar basis functions, pixel/voxel intersection relationships, and default setup parameters. Differences include using a finite focal length pinhole camera model to move between image and global coordinates, and an optional “first guess” process to reduce computational workload (see Sect. 2).

Unless otherwise stated the following default setup values should be assumed: three 200×200 pixel cameras with a pixel size of $4.65 \times 4.65 \mu\text{m}$, fitted with ideal 50 mm lenses, placed approximately 1,200 mm from the volume at yaw angles 0° and $\pm 30^\circ$; a $20 \times 20 \times 6$ mm volume, discretised at 10 voxels/mm, seeded with 3 voxel diameter Gaussian particles at a seeding density of $N_{\text{ppp}} = 0.05$; zero calibration error and image noise; artificial images assume perfect linear summation of projected light intensity through the volume. These default setup parameters were selected in relation to a current experimental project (not detailed within this paper). When using additional cameras in the parametric tests these are positioned at $\pm 60^\circ$. During camera angle tests, one camera remains at 0° , with the other two cameras at yaw angles $\pm \theta^\circ$. The calibration error refers to the magnitude of an artificial translation applied to each image. This simulates a

constant misalignment error with the other camera images, with each camera shifted in a different direction to create maximum misalignment ($[+, +]$; $[+, -]$; $[-, -]$ for the three cameras along the principle directions, u_{im} and v_{im} , in image coordinates).

Seeding density is given as the number of particles per pixel (N_{ppp}), which can be related to the particles per volume through knowledge of the camera resolution and volume of interest. Using this definition removes the volume size dependence of the results, because changes in volume depth will be balanced by the volumetric seeding density, with the number of particles in each image remaining constant. Therefore, coupled with an ideal pinhole camera model, the results presented here should be applicable to any volume size, although the relative percentage of the volume close to an edge may cause minor discrepancies. To demonstrate this a larger volume ($50 \times 50 \times 15$ mm volume, discretised at 10 voxels/mm, and imaged with 500×500 pixel cameras) has also been tested.

The flow field around a simple vortex line was chosen as a test case. A vortex line with circulation, $K = 30 \text{ mm}^2/\text{s}$ ($K = 60 \text{ mm}^2/\text{s}$ for the large volume), passes through the centre of the volume, and is set at 45° to induce a 3D velocity field. A random distribution of particles is created within the volume, and each particle is individually displaced according to this velocity field, for a timestep of $\Delta t = 1$ s, giving rise to voxel displacements of $\delta u \approx 0.3$ – 2 voxels. The artificial particle fields are projected to create artificial images for each camera position, which are then used in conjunction with the pinhole camera model parameters to reconstruct the light intensity distribution through the volume.

To obtain the velocity field, each volume is first pre-processed using an intensity magnitude threshold criterion (to remove some low intensity reconstruction artifacts), after which a simple single-pass FFT-based cross-correlation of small interrogation volumes is employed. The number of particles per volume is kept approximately constant at ten particles, and therefore as seeding density is increased, the interrogation volume size is reduced accordingly, increasing the measurement resolution (see Table 1). This basic scheme may be unable to adequately handle the high-velocity gradient region close to the vortex centre line. Therefore, to ensure testing focuses on the accuracy of the tomographic technique, and not the cross-correlation scheme accuracy, these high-velocity gradient regions are excluded from the results. In the current flow, regions of high-velocity and high-velocity gradient are coincident, and therefore the latter can be excluded simply by removing vectors from regions where the theoretical maximum voxel displacement exceeds a threshold value. Using a maximum limit of $\delta u \approx 2$ voxels excludes

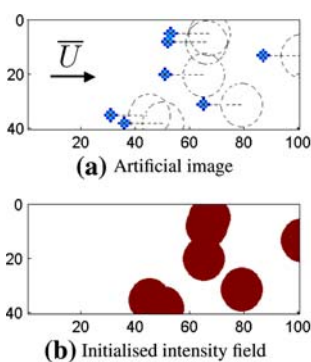


Fig. 5 Predictive regions first guess procedure

Table 1 Interrogation volume size and vector yield

Seeding (N_{ppp})	Small vol.		Large vol.	
	IV size	Vectors	IV size	Vectors
0.050	24	900	32	7200
0.075	20	1805	28	10404
0.100	20	1805	26	13690
0.125	18	2205	24	17600
0.150	16	3456	22	23232

approximately 6% of vectors (depending on IV size, and the associated vector spacing). Smaller interrogation volumes will result in lower velocity gradients, which may introduce an accuracy bias for the higher seeding density results, given that these cross-correlation schemes will have to resolve the velocity for lower gradients. However, through the elimination of the high gradient vortex core region, these differences are slight and therefore are not expected to significantly affect the results. It should be noted that more sophisticated cross-correlation algorithms exist that can resolve these high-velocity gradient regions, however, for our purpose of comparing reconstruction effects on the velocity field it is not necessary to include these regions, and a single-pass method is suitable. After obtaining the vector fields, spuriously high magnitude vectors are removed from the results using a simple threshold magnitude criterion.

Vector fields from artificial and reconstructed volumes are then compared with the ideal velocity distribution. The local error in voxels is given by Eq. 2, where u_i is the ideal velocity, and u_c is the velocity obtained from cross-correlation. The mean error is then taken using Eq. 3. The cross-correlation of the artificial volume gives a baseline measure of accuracy for the CCV process, allowing the accuracy of the tomographic technique to be assessed independently.

$$E = \sqrt{(u_i - u_c)^2 + (v_i - v_c)^2 + (w_i - w_c)^2} \tag{2}$$

$$\bar{E} = \sum_{i=1}^n \frac{E}{n} \tag{3}$$

4 Results

4.1 Computational costs

The computational cost of the MFG method can be compared to an iterative calculation using the MART algorithm (Eq. 1) by considering the number of floating point operations required. In 3D, assuming a pixel to voxel ratio of 1, each voxel should receive an intensity contribution from

approximately 4 pixels in each image. The MFG method updates every voxel with each correlated pixel, requiring two numerical operations: first to multiply the current pixel–voxel weighting value, W_{ij} , by the current pixel intensity, p_i ; second to multiply this product with the previous voxel intensity. A further operation is required to normalise the voxels, raising them to the power $1/N_{cam}$. The number of required calculations can be expressed as, $F_{MFG} = N_{vox}(8N_{cams} + 1)$, where N_{vox} is the total number of voxels in the volume.

The iterative calculation is more involved, and crucially requires calculation of the current object projection, Wf , for each voxel update. This calculation alone involves two numerical operations for each voxel in the current pixel’s line of sight. The number of voxels can be estimated by assuming approximately 4 voxels will contribute to the pixel in each slice of the volume, resulting in $4L$ voxels, where L is the average length of the line of sight through the volume in voxels. Elimination of low or zero intensity voxels is employed in the current program to accelerate the algorithm, and due to its significant impact on the total cost, this must be considered. As such, the value of N_{vox} will change after each iteration and L will change during each iteration. For cost estimation the former should be averaged over the entire iterative calculation, as should the latter in addition to being averaged over each iteration, with these averaged values denoted by \bar{N}_{vox} and \bar{L} . Therefore, combined with an additional four numerical operations in the MART algorithm, the iterative cost can be expressed as, $F_{MART} = 4k_t N_{cams} \bar{N}_{vox} (8\bar{L} + 4)$, where k_t is the total number of iterations.

The following example uses the large volume setup geometry (see Sect. 3). Estimates of the \bar{N}_{vox} variable are made by examining the number of zero and low intensity voxels after successive iterations. Dividing these by the total number of camera pixels permits estimation of L . A uniform field at $k = 0$ results in initial values of $N_{vox} = 37.5 \times 10^6$ and $L \approx 150$. N_{vox} reduces accordingly for successive iterations in comparison to the initial value: 18, 11, 9, 8, and 8%. After averaging appropriately, values of $\bar{N}_{vox} \approx 11 \times 10^6$ and $\bar{L} \approx 30$ are obtained. Therefore, in this example, the MFG method alone performs approximately 175 times fewer calculations than the uniform guess solution after 5 iterations.

If used before the MART algorithm, the MFG method dramatically reduces the initial value of N_{vox} to around 40% of the uniform field method, with the following reductions after successive iterations: 15, 10, 9, 8, and 8%. After calculating average values of $\bar{N}_{vox} \approx 6 \times 10^6$ and $\bar{L} \approx 20$, the MFG method prior to 5 MART algorithm iterations is shown to reduce the number of calculations by a factor of approximately 2.5. It is possible that further cost reductions could be made by using higher threshold values after the

MFG to exclude more voxels from the MART iterations (for example, $k = 0$ solutions in Sect. 4.3 are obtained by cross-correlating only around 10% of the total voxels).

Actual program run-times are shown in Fig. 6. These run times relate to a single large volume reconstruction, using a 64 bit, 3.4 GHz *Intel Xeon* workstation. For a seeding density of $N_{ppp} = 0.05$, the MFG method more than halves the computational cost, for a given number of iterations (see Fig. 6a). As seeding density is increased, the benefit of the MFG method becomes less pronounced, as fewer voxels are excluded from the calculation as a result of higher image densities (see Fig. 6b). Using the MFG method without any MART algorithm iterations (the $k = 0$ solution) is shown to be significantly faster than the iterated solutions. The large magnitude of the speed increase possible with this new method, albeit offset by a slight decrease in accuracy (see Sect. 4.3), may be particularly

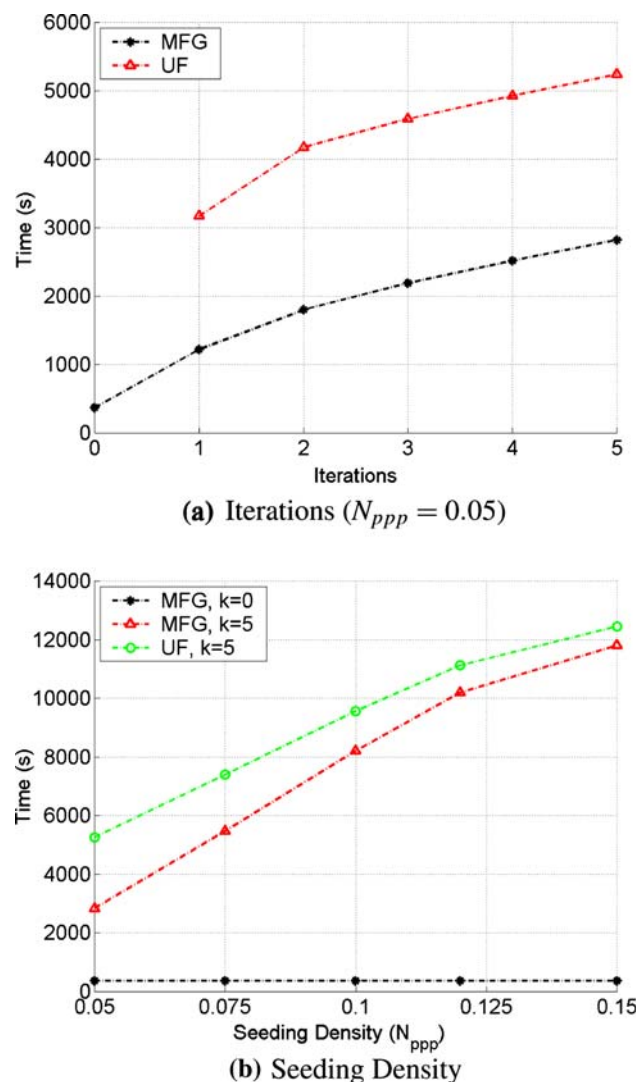


Fig. 6 Computational costs of the large volume reconstruction

useful for time-resolved measurements, requiring hundreds of velocity fields.

It should be noted that the cost of this method is constant in Fig. 6 due to the current programming practice of performing the first guess using all of the information in each image (including the zero intensity lines of sight). If the sparsity of the images is used, the cost of this method would become dependent on the seeding density, and therefore should decrease considerably, in particular for the lower seeding densities. Further differences between the theoretical estimates of cost and actual measured run-times arise in part through programming overheads, but most significantly from the inability to store the complete W matrix, which for a large 3D domain will require a huge amount of RAM (3.6 GB for the current large volume example). Therefore, in the absence of sufficient computational resources these weighting values will have to be recalculated every time they are needed, which will significantly increase the cost of the MFG method, but have a smaller relative impact on the MART algorithm. Thus, in order to obtain the maximum speed benefit from the MFG method, a workstation with a large amount of RAM is recommended, allowing the W matrix to be stored.

In the practical application of Tomo-PIV, images are likely to contain a certain level of background noise, without a significant number of zero-intensity pixels. However, image preprocessing is standard practice in Tomo-PIV (Elsinga et al. 2006), and therefore good quality experimental PIV images, which have been preprocessed will allow similar computational savings to be made.

It should also be noted that this new method only accelerates the reconstruction step, and the cross-correlation step remains unchanged. However, implementing the cross-correlation step using an efficient FFT subroutine is significantly less costly than the reconstruction step, and therefore the time savings obtained have a large impact on the total solution time.

4.2 Reconstruction visualisation

Particle reconstructions are presented in Fig. 7, using red and blue colour maps to represent the particles before and after the small displacement. To make it easier to see, the vortex core was aligned with the z -axis, and a smaller $100 \times 100 \times 50$ voxel domain employed. Thresholding is applied to exclude the weaker intensity voxels, and due to differences in the intensities magnitudes after MFG and UF methods, different threshold values are required. The good qualitative agreement between Fig. 7b and c demonstrates that irrespective of the initial field, the MART algorithm tends to iterate to a single solution. Slight differences between the zero iteration solution (Fig. 7) and other solutions can be seen, with the former containing more

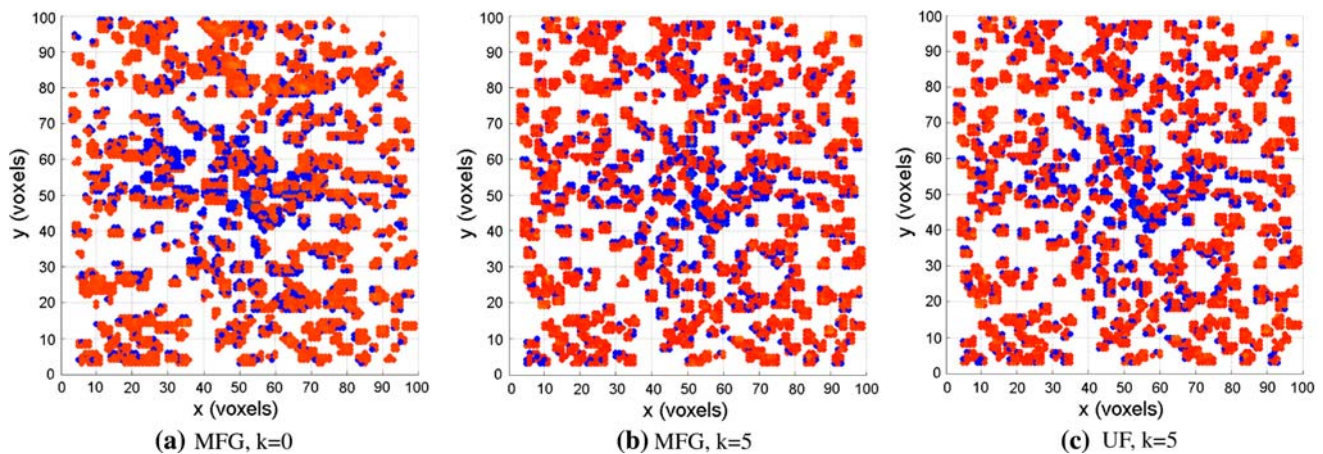


Fig. 7 Comparison of reconstructed displaced particle fields X–Y plane view (*MFG* multiplicative first guess, *UF* uniform field, *k* number of iterations)

misshapen particle reconstructions, although thresholding issues prevent more definitive conclusions. However, in practice, as previously stated, the accuracy of the particle reconstructions is not our primary concern, but rather the velocity fields which will be obtained from them. Therefore, this is discussed in the next section.

4.3 Cross-correlation PIV-based accuracy analysis (CCV approach)

The number of iterations, seeding density, camera angle, image noise, calibration error, and effect of the first guess scheme have all been investigated, and the results are shown in Fig. 8.

Figure 8a shows the effect of successive iterations on the velocity error (CCV approach), using 3, 4 or 5 cameras and different initial fields. For all camera and initial field setups, increasing the number of iterations decreases the error magnitude. The largest change occurs after the first iteration, with subsequent iterations resulting in increasingly small gains in accuracy as the solution approaches convergence. Comparing the number of cameras shows that adding a fourth camera significantly increases accuracy, although the effect of a fifth camera is less dramatic, giving only a slight increase in accuracy.

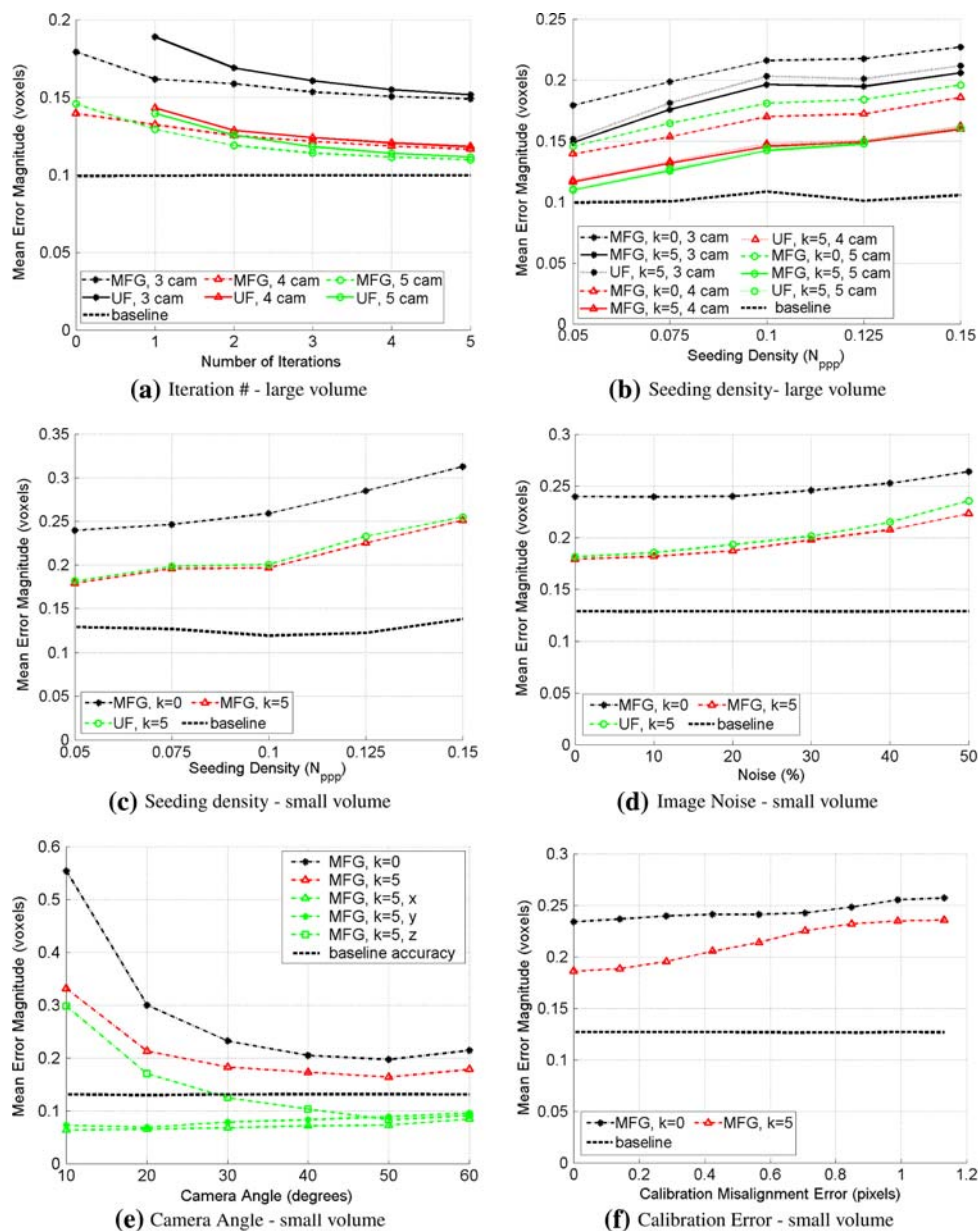
The multiplicative first guess method is shown to enhance the rate of convergence in comparison with the uniform initial field, especially during the early iterations. The MFG alone appears to give remarkably accurate solutions, within approximately 0.03 voxels of the solution after 5 iterations. It should also be noted that the MFG with zero iterations for the 4 and 5 camera solutions gives more accurate results than the 3 camera results after 5 iterations. Therefore, depending on user preference, it may be desirable to employ the first guess method without any MART iterations, significantly reducing run-time at the expense of a little accuracy.

Comparing the errors with the CCV approach baseline accuracy shows that the latter is the main source of error, and particularly in the case of the 4 and 5 camera reconstructions, the difference between the ideal solution and the reconstructions is small. Although the previous study (Elsinga et al. 2006) recommended at least five iterations, the present results seem to suggest reasonable convergence after three iterations. Reducing the number of iterations will allow for a reduction in computational costs, although it is noted that the first iteration is significantly more expensive than subsequent iterations, due to the removal of insignificant low intensity voxels and filtering of insignificant updates by the program.

Figure 8b and c shows an approximately linear increase of error magnitude with seeding density. Varying the seeding density from $N_{ppp} = 0.05–0.15$ results in an accuracy decrease of approximately 0.05 voxels. Fluctuations in the curves are attributed to different particle distributions at different seeding densities, and also through the decreasing interrogation volume (IV) size, which results in varying average particle concentrations in each interrogation volume. Fluctuations in baseline results are similarly caused through differences in the particle distributions, and also by the exclusion of different spurious vectors throughout the volume (only non-spurious vector locations from both artificial and reconstructed volumes were used to set the baseline accuracy).

The previously observed effects of adding cameras and altering the number of iterations appears relatively consistent, although the difference between the iterated 4 and 5 camera solutions is shown to decrease for higher seeding densities. The same trends are observed for both volume sizes, supporting the idea that these results are applicable to any volume size. Slight differences between the magnitude of the error for the large and small volumes can be attributed to use of different vortex strengths and separate

Fig. 8 Error variation with number of iterations and seeding density



random particle distributions. The zero iteration, 4 camera solution is more accurate than that of the 5 camera solution, and this trend persists irrespective of seeding density. This counter-intuitive behaviour may be caused by excessive intensity removal by the first guess procedure, although the difference is slight, and instantly recovered after iteration.

Figure 8d shows that, in contrast to the linear variation previously suggested (Elsinga et al. 2006), the velocity field obtained from the reconstructed volumes appear to be resistant to modest amounts of random noise ($\sim 30\%$). This may be due to the uncorrelated nature of the noise, to which the cross-correlation technique is somewhat impervious.

Figure 8e shows that camera angles of around 50° give optimum accuracy, although there is little variation

between 30° and 60° . This variation of accuracy with camera angle is slightly different to that shown by Elsinga et al. (2006), although an optimum angle of around 50° for this camera arrangement is predicted by other studies (Worth and Nickels 2007) using a RCC approach. As previously identified, camera configurations using narrow angles are associated with large errors, although the interpretation of this error in this instance is different. RCC-based methods return low accuracy due to particle *smearing*, where intensity is placed erroneously. However, the reduction in accuracy in the present investigation demonstrates the increased sensitivity of the cross-correlation to these elongated reconstructed particles. The cross-correlation peaks become wider and flatter in the direction of the elongation (usually the out of plane direction (z -axis

in the present study) due to camera depth of field considerations), becoming more sensitive to reconstruction imperfections or noise, and resulting in larger errors in that direction. This can be seen by examining the individual error components in Fig. 8e. The x and y error components are low for narrow configurations, increasing only slightly with camera angle. The z component is strongly dependent on camera angle, and dominates the total error in the low camera angle configurations. For practical application, the depth of field will also require consideration, although these results suggest that narrow configurations should be avoided as far as possible.

The effect of camera misalignment on error magnitude is shown in Fig. 8f. For the iterated solution, as the misalignment is increased from zero, the error rises slowly at first, increasing more rapidly for more moderate values, before showing a slower increase for high values. This curve shape may be caused by the size of the reconstructed particles. The artificial particles are 3 voxels in diameter, resulting in particles of 4 pixels diameter on the images (for a pixel to voxel ratio of approximately 1). As the misalignment is increased, the volume in which all camera lines sight converge to reconstruct a single particle will decrease. The relatively flat areas of the graph may represent a consistent reconstructed particle size, for example the first plateau may be caused by particles of typically $3 \times 3 \times 3$ voxels, and the upper plateau by particles of typically $2 \times 2 \times 2$ voxels, with the mid-range values showing the transition between these two typical states. If the particle image diameter is increased, the misalignment error may become insignificant, as the reconstructed particles would still be large enough to cross-correlate accurately. However, accuracy might still be affected through more crowded images, and this measure may therefore only be applicable to low seeding density regimes. A typical planar-PIV rule of thumb is that particles of approximately 3 pixels in diameter should contain sufficient information to perform accurate cross-correlation. However, given inherent camera misalignment and calibration mapping errors between three or more cameras, recorded particle image diameters may have to be increased to ensure that this rule of thumb is still applicable to the reconstructed particles.

The zero iteration solution appears to show remarkable robustness to moderate misalignments. This may be due to creation of larger more intense particles (see Fig. 3), with the high intensity particle centres on each image still contributing significantly to each particle reconstruction. This could occur as high intensity from one projection intersects with weaker intensity from the edge of the particle from the other projections, resulting in moderate intensity voxels. However, after each successive iteration, the high intensity update, would be followed by more than

one (depending on the number of other misaligned camera contributions) low intensity updates, resulting in a much less intense overall particle reconstruction, and a corresponding drop in accuracy. The limit of this appears to be reached after a misalignment of around 0.7 pixels, when high intensity particle centres are so misaligned that they begin to intersect with either extremely low or zero intensity projections from the other cameras, and the error begins to increase correspondingly.

5 Conclusion

In order to reduce the computational workload associated with the Tomo-PIV technique a new “multiplicative first guess” technique has been proposed. When using this method as a precursor to the standard MART approach the reductions in cost are realised primarily through the elimination of a large number of zero intensity voxels. A study of tomographic reconstruction accuracy has been performed using a CCV analysis of a simple angled vortex line flow field, to demonstrate the potential of this new method and to determine the effect of other setup parameters.

Using the MFG method accelerates convergence, and even provides a reasonably accurate solution before implementing the MART algorithm. This introduces the possibility of extremely rapid volume reconstructions (with only a slight accuracy penalty), which may be particularly useful for obtaining time-resolved measurements. This method is shown to be 175 and 35 times faster theoretically and during practical application respectively. The MFG method also appears enhance the robustness of the solution to degradation through camera misalignment, although after iteration this advantage appears not to hold. The results of the present study also indicate that fewer iterations could be used to reduce computational expense with only a fractional reduction in accuracy.

The parametric study results for seeding density and camera number appear to largely support the parametric study of Elsinga et al. (2006), although the present study suggests that the benefit a fifth camera is less clear, especially at high seeding densities. The present study also suggests that the solution may not be as sensitive to image noise as previously indicated, and that slightly wider camera angles give optimum accuracy, although only yaw has been considered. Moreover, the results of the present study quantify these parametric relationships using a more representative measure of attainable PIV accuracy as opposed to a RCC-based approach.

These findings have important implications for reducing the large computational costs associated with Tomo-PIV, and may be useful for setup parameter selection and accuracy assessment. The sparsity of the reconstruction

problem in Tomo-PIV sets it apart from other reconstruction problems, and furthermore the use of cross-correlation releases some of the restrictions on reconstruction quality and artifacts. Taking this into consideration has led to these novel findings. Further investigation into thresholding and smoothing, and more advanced first guess procedures may further increase the predictive ability of the initial estimate, further reducing the cost of this promising new technique.

Acknowledgments The first author wishes to acknowledge funding from the Engineering and Physical Sciences Research Council, through a Cambridge University Doctoral Training Award.

References

- Elsinga G, Scarano F, Wieneke B, van Oudheusden B (2006) Tomographic particle image velocimetry. *Exp Fluids* 41:933–947
- Elsinga G, Kuik D, van Oudheusden B, Scarano F (2007) Investigation of the three-dimensional coherent structures in a turbulent boundary layer with Tomographic-PIV. In: 45th AIAA Aerospace Sciences Meeting and Exhibit, Reno, Nevada
- Herman G (1980) Image reconstruction from projections: the fundamentals of computerized tomography. Academic Press, London
- Herman G, Lent A (1976) Iterative reconstruction algorithms. *Comput Biol Med* 6:273–294
- Hirsch K (2002) Holographic particle image velocimetry. *Meas Sci Technol* 13:R61–R72
- Hori T, Sakakibara J (2004) High-speed scanning stereoscopic PIV for 3D vorticity measurement in liquids. *Meas Sci Technol* 15:1067–1078
- Maas H, Gruen A, Papantoniou D (2004) Particle tracking velocimetry in three-dimensional flows. *Exp Fluids* 15:133–146
- Natterer F (1999) Numerical methods in tomography. *Acta Numer* 8:107–142
- Schröder A, Geisler R, Elsinga G, Scarano F, Dierksheide U (2006) Investigation of a turbulent spot using time-resolved tomographic PIV. In: 13th international symposium on applications of laser techniques to fluid mechanics, Lisbon, Portugal
- Wieneke B, Taylor S (2006) Fat-sheet PIV with computation of full 3D-strain tensor using tomographic reconstruction. In: 13th international symposium on applications of laser techniques to fluid mechanics, Lisbon, Portugal
- Willert CE, Gharib M (1992) Three-dimensional particle imaging with a single camera. *Exp Fluids* 12:353–358
- Worth N, Nickels T (2007) A computational study of tomographic reconstruction accuracy and the effects of particle blocking. In: 5th joint ASME/JSME fluids engineering conference, San Diego