

**UNIVERSITÀ DEGLI STUDI DI NAPOLI “FEDERICO II”
DIPARTIMENTO DI INGEGNERIA INDUSTRIALE**

**SCUOLA DI DOTTORATO IN INGEGNERIA
AEROSPAZIALE, NAVALE E DELLA QUALITÀ**

XXV CICLO



**TESI DI DOTTORATO IN
INGEGNERIA AEROSPAZIALE**

Tomographic Particle Image Velocimetry

Developments and applications to turbulent flows

Stefano Discetti

Tutors:

Prof. Ing. Tommaso Astarita
Prof. Ing. Giovanni Maria Carlomagno
Prof. Ronald J. Adrian

Coordinatore

Prof. Ing. Luigi De Luca

Aprile 2013

Contents

| | |
|---|-----------|
| INTRODUCTION | I |
| CHAPTER 1 - 3D PARTICLE IMAGE VELOCIMETRY | 1 |
| 1.1 Particle image velocimetry | 2 |
| 1.1.1 Fundamentals of PIV | 2 |
| 1.1.2 Particle imaging | 3 |
| 1.1.3 Interrogation strategies | 4 |
| 1.1.4 Advanced interrogation strategies | 6 |
| 1.1.5 Limits of 2D PIV | 7 |
| 1.2 Stereoscopic Particle Image Velocimetry | 8 |
| 1.2.1 Optical calibration | 10 |
| 1.2.2 Reconstruction of the flow field | 10 |
| 1.2.3 Disparity correction and self-calibration | 11 |
| 1.2.4 Multi-Plane Stereoscopic PIV | 12 |
| 1.3 Three-Dimensional Particle Image Velocimetry | 13 |
| 1.3.1 Scanning PIV | 13 |
| 1.3.2 Holographic PIV | 15 |
| 1.3.3 Defocusing PIV | 16 |
| 1.3.4 3D Particle Tracking Velocimetry (PTV) | 17 |
| 1.3.5 Tomographic Particle Image Velocimetry (Tomo-PIV) | 19 |
| CHAPTER 2 - TOMOGRAPHIC PIV | 23 |
| 2.1 Algebraic reconstruction techniques for Tomographic PIV | 24 |
| 2.1.1 The reconstruction problem | 24 |
| 2.1.2 Algebraic reconstruction techniques | 26 |
| 2.1.3 MART performances | 28 |
| 2.1.4 Volume self-calibration | 29 |
| 2.2 The problem of the ghost particles | 31 |
| 2.2.1 Estimate of the number of ghost particles | 32 |
| 2.2.2 The role of the ghost particles in the velocity measurement | 33 |
| 2.3 Improvements and advances on the reconstruction techniques | 34 |
| 2.3.1 Fast reconstruction techniques | 34 |

| | |
|--|------------------|
| 2.3.2 Developments on the accuracy of MART and advanced reconstruction methods | 36 |
| 2.3.3 Multi-exposure methods | 37 |
| 2.4 Three-dimensional motion analysis | 38 |
| <u>CHAPTER 3 - MULTI-RESOLUTION MART</u> | <u>40</u> |
| 3.1 Multi-Resolution approach to Tomo-PIV | 41 |
| 3.2 Parametric study of MR-MART | 43 |
| 3.2.1 Quality of the reconstruction | 44 |
| 3.2.2 Acceleration of the process | 45 |
| 3.2.3 Memory storage reduction | 47 |
| 3.3 Validation via synthetic experiment | 48 |
| 3.3.1 Low seeding density simulations | 49 |
| 3.3.2 High seeding density simulations | 52 |
| <u>CHAPTER 4 - SPATIAL FILTERING IMPROVED TOMOGRAPHIC PIV</u> | <u>54</u> |
| 4.1 Reconstruction quality enhancement by spatial filtering | 55 |
| 4.2 Performances of SFIT-MART | 57 |
| 4.2.1 Effect of the experimental parameters | 57 |
| 4.2.2 Spatial resolution of SFIT-MART | 62 |
| 4.2.3 Experimental application to fractal generated turbulence | 65 |
| <u>CHAPTER 5 - EFFICIENT 3D PIV INTERROGATION ALGORITHMS</u> | <u>71</u> |
| 5.1 Fast correlation computation | 72 |
| 5.1.1 Multi-resolution predictor estimation | 73 |
| 5.1.2 Fast corrector computation with sparse direct correlations | 74 |
| 5.1.3 Efficient corrector computation with overlapping windows | 74 |
| 5.2 Performance assessment | 79 |
| 5.2.1 Test case 1 - Uniform displacement | 80 |
| 5.2.2 Test case 2 - One-dimensional shear displacement | 81 |
| 5.2.3 Test case 3 - Circular jet | 85 |
| 5.2.4 Experimental test case – swirling jet | 88 |
| 5.2.5 Conclusions and guidelines | 92 |
| 5.3 Effects of poor discretization of the weighted cross-correlation | 93 |
| 5.3.1 The theoretical model for the MTF estimation | 93 |
| 5.3.2 Spatial resolution and stability of block weighting windows | 94 |

| | |
|---|------------|
| 5.3.3 Performance assessment by 3D simulations | 98 |
| 5.3.4 Conclusions and guidelines | 101 |
| <u>CHAPTER 6 - LOW-COST TOMO-PIV SYSTEMS</u> | 102 |
| 6.1 Proposed experimental setup | 103 |
| 6.1.1 Suppression of the bias effect of the coherent ghost particles motion | 103 |
| 6.1.2 Accuracy improvement via MTE-MART | 105 |
| 6.2 Numerical simulations | 106 |
| 6.2.1 Spatial resolution | 106 |
| 6.2.2 Measurement noise and MTE | 109 |
| 6.3 Summary and conclusions | 115 |
| <u>CHAPTER 7 - THE DECAY OF FRACTAL-GENERATED TURBULENCE</u> | 116 |
| 7.1 Wake of fractal objects | 118 |
| 7.1.1 Square space-filling fractal grids | 119 |
| 7.1.2 The wake interaction lengthscale | 120 |
| 7.1.3 Quasi-isotropic and homogeneous “high” $Re\lambda$ turbulence | 122 |
| 7.1.4 Power-law or exponential decay? | 123 |
| 7.1.5 The violation of the dissipation anomaly | 125 |
| 7.1.6 Questions raised and role of 3D data | 126 |
| 7.2 Experimental details | 127 |
| 7.2.1 Experimental apparatus and design point | 127 |
| 7.2.2 Calibration, tomographic reconstruction and PIV interrogation | 129 |
| 7.3 Results | 131 |
| 7.3.1 Homogeneity assessment | 131 |
| 7.3.2 Large and small scale isotropy | 133 |
| 7.3.3 Longitudinal integral and Taylor lengthscales | 135 |
| 7.3.4 Dissipation measurements | 139 |
| 7.4 Final remarks | 142 |
| <u>CHAPTER 8 - JET FLOWS PAST A SUDDEN EXPANSION</u> | 144 |
| 8.1 Fluidic precessing jet nozzles | 145 |
| 8.1.1 The phenomenology of precession in axisymmetric jets | 146 |
| 8.1.2 The effects of geometry and Re on the precession probability | 148 |
| 8.1.3 Determination of the Strouhal number of the precession | 150 |

| | |
|---|-------------------|
| 8.2 Experimental details | 152 |
| 8.2.1 Water facility | 152 |
| 8.2.2 Tomographic PIV system | 153 |
| 8.2.3 Calibration correction | 154 |
| 8.2.4 Tomographic reconstruction and CC analysis | 156 |
| 8.3 Results | 158 |
| 8.3.1 Instantaneous flow features | 158 |
| 8.3.2 Mean flow features | 160 |
| 8.4 POD analysis and Low Order Reconstruction | 163 |
| 8.4.1 POD implementation | 163 |
| 8.4.2 POD modes | 165 |
| 8.4.3 Low Order Reconstruction of the circular precessing jet | 170 |
| 8.4.4 Phase computation: precession frequency measurements for the FPJ | 171 |
| 8.5 Final remarks | 173 |
| | |
| <u>CONCLUSIONS AND PERSPECTIVES</u> | <u>175</u> |
| | |
| <u>REFERENCES</u> | <u>179</u> |
| | |
| <u>LIST OF PUBLICATIONS</u> | <u>191</u> |
| | |
| <u>ACKNOWLEDGMENTS</u> | <u>194</u> |

INTRODUCTION

Understanding turbulence is universally considered the Holy Graal of fluid mechanics. The inherent unsteady three-dimensional nature of turbulent flows has challenged developers and users of numerical and experimental techniques for many decades, and most probably it is a war far yet from being won. The underlying cascade process needs investigation instruments with large dynamic and spatial range (i.e. able to measure with satisfactory accuracy and reliability the turbulent fluctuations from the very large scale of the energy-containing range down to the dissipative range); the requirement becomes more stringent as the Reynolds number increases (i.e. the separation between large and small scales widens). Most often, one has to give up information at the small scales due to limited spatial resolution or poor sensitivity of the technique; in numerical simulations the computational cost effectively results in an upper boundary for the Reynolds number in Direct Numerical Simulations (DNS) of the Navier-Stokes equation, or the introduction of sometimes questionable turbulence models to solve the closure problem (as in the Reynolds Averaged Navier-Stokes (RANS) equations). The advancement of the Computational Fluid Dynamics (CFD), in this sense, has been impressive over the years. However, instead of replacing the experimental fluid dynamics, the demands of the CFD community actually have pushed towards significant advancements in the field of development of the experimental techniques, as it is starving for more accurate and detailed real data for turbulence models validation.

In this direction, much has been done since the advent of Particle Image Velocimetry (PIV), which has unquestionably contributed to significant leaps forward in the understanding of many fluid dynamics problems thanks to its ability to provide field measurements on planar domains (see Westerweel et al 2013 for an extensive review). On the other hand, the analysis is physically limited to "slices" of phenomena occurring in a 3D space; *turbulence does require a full three-dimensional (3D) three components (3C) anemometric technique.*

From this point of view Tomographic Particle Image Velocimetry (Tomo-PIV, Elsinga et al 2006a) appears to be very promising. The technique is based on the reconstruction of the volumetric pattern of light intensity scattered by seeding particles illuminated by a pulsed light, simultaneously recorded from several viewing directions. The reconstructed distributions at different time instants are then interrogated by a 3D cross-correlation algorithm to determine the 3C velocity field. The interest in the development of the Tomographic PIV technique is testified by its relevance within the AFDAR (Advanced Flow Diagnostic for Aeronautical

Introduction

Research, funded within the 7th Framework Programme, grant agreement N. 265695) and in which most of the current thesis is framed.

Even though only recently developed, Tomo-PIV is quite well assessed in the measurement of turbulent flows, ranging from the classical wall turbulence (Humble et al 2009, Elsinga et al 2010, Atkinson et al 2011, Gao et al 2011, Schröder et al 2011, Elsinga et al 2012), the wake of streamlined (Ghaemi & Scarano 2011) or bluff bodies (Hain et al 2008, Scarano & Poelma 2009) to more advanced applications in aeroacoustics (Violato et al 2011) and in biomedicine (Buchmann et al 2011). Nevertheless, the introduction of a new piece in the chain of the well-assessed PIV measurement procedure, i.e. the tomographic reconstruction, and the extension to the 3D scenario of the cross-correlation interrogation algorithms raise new questions, open great margins of improvement and introduce new difficulties.

The aim of the first part of the thesis is to address these topics; in particular, the following themes are of crucial importance:

Computational cost of the tomographic reconstruction. The tomographic reconstruction procedure with the widely used Multiplicative Algebraic Reconstruction Technique (Herman & Lent 1976) can be very intensive. Hain et al (2008) quoted a processing time of 30 min on a eight cores computer only for the reconstruction of each volume with a size of $1430 \times 1389 \times 217$ voxels (being the voxels the 3D equivalent of the pixels) using 4 MART iterations, while in the case of adoption of higher-resolution cameras (11 Mpixels) the computational cost increases to 34 hours on a dual-core computer for the reconstruction of a volume with a size of $3512 \times 2562 \times 599$ voxels. Such a large processing time is certainly unacceptable when the number of samples to be analyzed is considerable. A reduction of the computational cost could be obtained either by advanced parallelization (Open Multi-Processing, Graphics Processing Units, etc.) or by developing more efficient reconstruction algorithms. In this thesis a Multi-Resolution version of MART is proposed, leading to a remarkable processing time reduction (up to 20 times) over a wide range of seeding densities.

Accuracy. Due to the limited available number of simultaneous views, the problem of the tomographic reconstruction is underdetermined, i.e. different particles distributions can satisfy the set of projections. The pursuit of a reliable reconstruction algorithm at a fairly high seeding density and with an acceptable computational cost is still in progress. In this thesis a variation of MART is proposed, based on properly oriented artificial diffusion applied on the particles distributions in the iterative procedure. The technique, named SFIT (Spatial Filtering Improved Tomography), allows an improvement of the quality of the reconstruction without increasing the computational cost.

Computational cost of the 3D-PIV. The extension of the well-established high accuracy PIV interrogation algorithms to the third dimension is trivial from a conceptual point of view, but problematic in the practical implementation. Memory

storage issues and large computational cost are the most relevant problems. An alternative efficient algorithm, based on multi-resolution interrogation, sparse direct cross-correlation and reduction of the number of redundant operations is proposed in the present thesis.

Hardware cost. A relevant balance sheet item for a Tomographic PIV system is related to the use of several cameras to achieve accurate tomographic reconstructions. However, differently from the case of planar PIV, the requirement of double-shutter cameras is much easier to be removed, as the objects to be cross-correlated consist of the clouds of particles reconstructed at different time instants instead of the row images. In this thesis a low-cost approach, based on two (or more) independent tomographic systems (composed by cheap cameras working in single-frame mode) is proposed. Of course the price to pay is in the complication of the optical arrangement of a number of cameras twice as large as the standard tomographic PIV systems. On the other hand, the observation of the same set of particles from different viewing directions at different instants provides the advantage of more diversified information on the objects to be reconstructed, thus enabling a significant improvement of the results using multi-exposure reconstruction methods.

In the second part of the thesis the focus is on two possible applicative scenarios of the technique. In the first application the nearly isotropic and homogeneous turbulence generated by square fractal grids (i.e. grids with a square pattern repeated at different scales) is investigated. Such turbulence is characterized by a very unusual decay, apparently at odds with the so-called dissipation anomaly, i.e. the scaling of the turbulent dissipation with the rate of transfer of the energy from the large scales to the small scales. Tomographic PIV is certainly well suited for the assessment of homogeneity and isotropy (at least at large scales); conversely, it is extremely challenged by the small scales measurement, and in particular the dissipation is difficult to be estimated as it requires spatial resolution down to the Kolmogorov scale. Assessing the performances of Tomographic PIV in such framework is of great importance, particularly in order to provide an instrument for validation of numerical codes and turbulence models.

In the second application the Tomographic PIV exhibits its great potential in the field of 3D quantitative visualization of the organization of coherent structures. The evolution of a circular jet (with and without a swirling component) after a sudden expansion in a cylindrical chamber is investigated. The flow organization is extremely complex, as it intermittently switches between two working conditions, i.e. a quasi-axisymmetric expansion within the chamber and an asymmetric configuration with reattachment point precessing in a gyroscopic like motion. The flow is unsteady, intermittent and three-dimensional, thus making the understanding of the flow topology rather difficult. The Proper Orthogonal

Introduction

Decomposition is applied to the 3D data in order to isolate the main features of the outflow mode; furthermore, since the flow field is dominated (energetically speaking) by a significant periodic component, a low order reconstruction, obtained combining the most energetic modes, is implemented to observe the key dynamics of the large structures.

Chapter 1 - 3D Particle Image Velocimetry

Turbulent flows are prevalent in nature and in engineering applications. The intrinsic three-dimensional features of turbulence have challenged the developers and practitioners of measurement techniques for many years. The pursuit of a technique able to combine both large dynamic and spatial range (i.e. large ratio between the maximum and minimum detectable velocity and spatial scales, respectively), and at the same time provide a complete description of the three-dimensional pattern of coherent structures and its time evolution, is still an open field of research with great margins of improvement. Furthermore, the advent of advanced 4D numerical simulations (the 4th dimension being the time one), even though with limited Reynolds Re number, does require the development of measurement techniques capable of providing accurate data sets for comparisons and validation.

In the last decades the improvement of the capabilities of the measurement techniques has been impressive. The limits in the quantitative analysis of flow visualization and the frustration of relying on sometimes questionable assumptions to extrapolate spatial information on coherent structures in a dimensional space larger than that of measurement (see for example the application of the Taylor's frozen turbulence hypothesis in hot-wire anemometry measurements) have been softened by the introduction of field measurements with optical methods. The nomenclature *particle image velocimetry* (PIV) appeared in the literature for the first time in the work by Adrian (1984), where it is stressed that the illumination of particle tracers, within a flowing fluid, by a laser light would produce particle images on a recording medium. Illumination by thin laser sheet pulses with a short pulse provides images of particles in a plane. Individual tracking of particles (in case of sparsely populated images) or interrogation by statistical operators have been widely used in the many different approaches to measure two components of the velocity field in the nearly two dimensional illuminated domain. Again, the researchers had to accept giving up part of the information, since neither the third component of the instantaneous velocity field nor 5 out of 9 terms of the velocity gradient tensor were available. More details are reported in Sec 1.1.

Stereoscopic PIV (Willert 1997, Soloff et al 1997) determined a step forward in this sense by allowing the measurement of the third component of the velocity field, even though in a 2D domain, by using a two-cameras system in stereoscopic arrangement. While adding the complications of an additional camera, a calibration technique to map the spatial coordinate reference system onto the image system and procedures to account for the possible misalignment of the laser sheet with the reference plane for the optical calibration, Stereo PIV has the great advantage of

accounting for the errors due to the perspective effect of the out-of-plane motion, affecting the standard monocular PIV.

The extension of PIV to the fully three-dimensional three components (3D3C) velocity measurements scenario is more complicate, and it is still ongoing. In the first part of this chapter a description of the main features of monocular and stereoscopic PIV is provided. Subsequently, a brief survey of some of the implemented solutions for the 3D extension of PIV is reported, with no pretension of being exhaustive; a more complete review is reported in fundamental books of PIV theory (Raffel et al 2007, Adrian & Westerweel 2011). The chapter ends with a description of the working principle of Tomographic PIV.

1.1 Particle image velocimetry

1.1.1 Fundamentals of PIV

Particle image velocimetry is a non-invasive anemometric technique for field velocity measurements. A sketch of a standard PIV system is provided in Fig. 1.1. A basic PIV system consists of a pulsed light source collimated in a light sheet and illuminating a cloud of seeding particles within a fluid flow. The particles are required to be small enough to accurately track the fluid motion (typically a few microns). The small scattering section of such particles and the relatively short exposure time to capture particle images without blurring led for a long period of time to use almost exclusively high intensity pulsed laser light sources. The light scattered by the particles is recorded on a camera (for example a high resolution digital camera) on a single-frame (double exposure) or on two separate frames (single exposure). The optic axis of the camera is set perpendicular to the laser sheet within alignment tolerances.

The recorded images are divided in interrogation spots (often referred as *windows*). For each interrogation spot a velocity vector is determined using statistical methods: auto-correlation in the case of double-exposure images (Adrian & Yao 1984); cross-correlation in the case of single-exposure images (Willert & Gharib 1991). The velocity components are found from measurements of the particle image displacements ΔX and ΔY in the image plane according to the relation

$$\begin{pmatrix} u \\ v \end{pmatrix} = \frac{M_0}{\Delta t} \begin{pmatrix} \Delta X \\ \Delta Y \end{pmatrix} \quad (1.1)$$

where u and v are the fluid velocities in the object plane, Δt is the time separation between the two pulses, and M_0 is the lateral magnification from the object plane to the image plane, taken to be constant in the absence of aberrations and misalignment of the two conjugate planes. The magnification M_0 is the ratio of the




Fig. 1.1 Sketch of a standard PIV experimental setup (Raffel et al 2007).

image distance z_0 (distance between the lens and the sensor plane) to the object distance Z_0 (distance between the lens and the object plane).

1.1.2 Particle imaging

The particle images are formed on the sensor of the camera (in recent applications predominantly based on CCD = Charge Coupled Device, or CMOS = Complementary Metal-Oxide Semiconductor) via a high quality lens. The particle image diameter depends both on the particles size and on the lens properties. Given that $f_{\#}$ is the ratio of the focal length and the aperture diameter, the diameter of an imaged particle is (Adrian 1991):

$$d_{\tau} \approx \sqrt{d_{geom}^2 + d_{diff}^2} = \sqrt{(M_0 d_p)^2 + (2.44 f_{\#} (1 + M_0) \lambda_L)^2} \quad (1.2)$$

where λ_L is the wavelength of the illuminating source, d_p is the particle physical diameter, d_{geom} is the geometric image diameter due to magnification and d_{diff} is the diffraction-limited diameter, which formula is determined by approximating the Airy disk with a Gaussian function.

The particle image diameter should be minimized in order to reduce the uncertainty in the displacement measurement (see chapter 5 of Raffel et al 2007 for further details) and to ensure strong light intensity (since the collected light scales with the inverse of the particle image area); on the other hand, small particle images can determine a bias effect towards integer pixel values, referred in the literature as *peak-locking* (Westerweel 1993).

Considering that the light sheet has a finite thickness Δz_0 , the imaged particles will be in focus only if Δz_0 is smaller than the depth of field, estimated as follows (Adrian 1991):

$$\delta z \approx 4.88 \left(1 + \frac{1}{M_0}\right)^2 f_{\#}^2 \lambda \quad (1.3)$$

The features of the recorded image depend upon the *source density* N_s , defined as:

$$N_s = C \Delta z_0 \frac{\pi d_{\#}^2}{4M_0^2} \quad (1.4)$$

where C is the particles concentration.

If $N_s < 1$, the probability of particles overlapping is rather small (see Sec 1.3.4), and the image will be composed of randomly located individual spots; otherwise, since the light source is coherent, they will generate randomly oriented fringe patterns.

1.1.3 Interrogation strategies

The solutions to measure the velocity field from PIV images are strongly seeding density dependent. The most intuitive approach is based on tracking of individual particles in subsequent exposures (PTV, *Particle Tracking Velocimetry*). The dramatic drop in percentage of valid matching as the image density increases limits the application to very sparse images; in a simplifying picture, it is reasonable to assume that the particles spacing should be larger than the particles displacement in order to obtain a high rate of valid detections. Actually, more advanced algorithms, based on the spatial and temporal coherence of the displacements of neighbouring particles, can ideally afford image density up to 0.3ppp (see for example Ohmi & Li 2000). On the other hand, the accuracy is limited by the particle identification scheme, which suffers in case of overlapping particles.

In case of larger image density (but still with $N_s < 1$, so that the scattered light of the particles will form particle images more than speckles), statistical methods are more appropriate. Correlation techniques (either auto-correlation or cross-correlation) fit perfectly for this purpose. Keane & Adrian (1992) performed a theoretical and numerical study on the performances of auto- and cross-correlation, achieving the final result of a simple curve (Fig. 1.2), representing the valid detection probability as a function of a dimensionless parameter, on which data from different systems (interrogation spot size, interrogation strategy, number of exposures) collapse. The parameter is the product of the mean number of particle per interrogation spot N_l and the fraction of particles remaining in the interrogation spot despite of out-of-plane and in-plane displacement (F_0 and F_l , respectively); in practice, it represents the mean number of particle pairs in each interrogation spot.




Fig. 1.2 Valid detection probability as a function of the mean number of particle pairs in the interrogation spot (Keane & Adrian 1992).

Correlation analysis can be performed either in the spatial or in the frequency domain via the Wiener-Kinchin theorem (Smith 1999). Considering interrogation spots with intensity $A_{i,j}$ and $B_{i,j}$, the normalized cross-correlation coefficient map in its discretized version is determined by (supposing square interrogation window of linear size D_I for simplicity):

$$\Phi(\Delta X, \Delta Y) = \frac{\sum_{i,j}^{D_I} (A_{i,j} - \mu_A)(B_{i+\Delta X, j+\Delta Y} - \mu_B)}{\sqrt{\sum_{i,j}^{D_I} (A_{i,j} - \mu_A)^2 \sum_{i,j}^{D_I} (B_{i+\Delta X, j+\Delta Y} - \mu_B)^2}} \quad (1.5)$$

where μ_A and μ_B are the mean intensities of the interrogation spots. The main advantage of performing this operation in the frequency space is in terms of computational cost: using the Fast Fourier Transform (FFT) the number of operation required is $o(2D_I^2 \log_2 D_I)$ against $o(D_I^4)$ for the case of straightforward application of (1.5). However, the introduction of artefacts due to the imposed periodicity has to be accounted.

The maximum of the cross-correlation map provides the location of the most likely displacement of the particles within the interrogation spot. Since the diffraction-limited spot is approximately Gaussian, one can assume that reasonably a peak resulting from convolution of Gaussian signals is Gaussian as well. For this reason, among the others, a Gaussian peak interpolator has demonstrated to be the most accurate in case of small particle images (Raffel et al 2007).

1.1.4 Advanced interrogation strategies

Observing Fig. 1.2 it is evident that a high probability of valid detection is achieved when $N_I F_0 F_I > 7$. To some extent, high seeding density is desired (within the limit of $N_S < 1$), while in-plane and out-of plane losses of pairs have to be very limited. The effect of out-of-plane motion can be reduced by properly orienting the laser sheet, or reducing Δt ; however, due to the finite thickness of the laser sheet (typically 0.5-1mm), the out-of-plane motion is a limiting parameter only in presence of a relatively strong mean velocity component orthogonal to the illuminated region. On the other hand, the factor F_I cannot be maximized by reducing the in-plane motion (i.e. setting a smaller Δt), since it is the object of the measurement; considering that the minimum resolvable velocity is approximately 1-10% of the mean particles diameter, a small Δt would determine a small dynamic range and low measurement accuracy.

A strategy to reduce (or in some cases eliminate) the in-plane loss of pairs consists in using interrogating regions with different size (for example, the interrogation spot on the second image can be made larger). This approach has been thoroughly analyzed in several works (Adrian 1991, Westerweel 1997, 2000). Its main advantages are the intrinsic simplicity and the low computational cost, since it is a single-step method. However, a multi-pass cross-correlation with a discrete windows offset (Westerweel et al 1997), applied after the evaluation of a predictor using interrogation windows larger than the final ones, leads to higher signal-to-noise ratios and improved measurement precision, since the algorithm actually tracks the particle patterns. A symmetric windows shift (Wereley & Meinhart 2001) provides even better performances, as it is second-order accurate.

A further improvement can be obtained by applying a continuous (i.e. sub-pixel) shifting, as proposed by Lecordier et al (2001). In this scenario the intensity values at sub-pixel locations are needed, and the interpolation scheme is of fundamental importance in assessing both bias and random errors (Astarita & Cardone 2005, Astarita 2006).

All the presented methods account for in-plane displacement, but not for velocity gradients within the interrogation window, i.e. the velocity dynamic range is improved with respect to the standard one-step interrogation, while the vorticity dynamic range is about the same. Pushing along the path of sub-pixel interpolation, one could imagine of interpolating the velocity field onto each pixel, and deforming the entire interrogation window to maximize the particles pattern matching. Huang et al (1993) proposed the *particle image distortion* technique, based on bilinear interpolation of the predictor displacement field to compensate both for translation and rotation of the interrogation windows using kinematic formulae. Jambunathan et al (1995) developed a different algorithm, based on the interpolation of the predictor displacement field on each pixel of the first image to evaluate the distortion of the second image. Since then image deformation methods (see Scarano

2002 for a review) evolved and spread as the standard high-accuracy interrogation algorithms.

A slightly different approach is based on particle tracking on high density images with the aid of biased search after calculating a predictor displacement field using cross-correlation analysis. The approach is referred as *super-resolution PIV* (Keane et al 1995). With respect to the standard cross-correlation analysis, requiring 10-15 particle images for each interrogation spot, the super-resolution approach can potentially provide a velocity measurement for each image pair, determining a mean spatial resolution of $D_I/M_0 N_I^{1/2}$ (compared to D_I/M_0 for standard PIV). However, the method has 3 main limits: the particle image density cannot be as high as standard PIV, otherwise even the biased search would fail; the data will be distributed on an unstructured grid with non-uniform distribution of the spatial resolution (for example regions of high swirl, where larger resolution is needed, can be poorly seeded due to centrifugal effects); since the results are based only on individual couples, the improved spatial resolution is obtained at the expense of the measurement accuracy, reduced by the factor $d_\tau/N_I^{1/2}$.

1.1.5 Limits of 2D PIV

As already mentioned, the standard PIV provides only 4 out of 9 components of the velocity gradient tensor, thus leading to the need of assumptions to extract 3D information (for example, the turbulent dissipation rate is measured relying on local isotropy hypothesis, Tanaka & Eaton 2007) or possible misleading data interpretation, such as ambiguous coherent structures identification.

Furthermore the out-of-plane displacement introduces two limitations. The first one, discussed in Sec 1.1.3-4, is the out-of-plane loss of pairs, reducing the signal-to-noise ratio, or, in different terms, the spatial resolution, since more particle images are needed to keep $N_I F_0 F_I > 7$. The second one is more subtle, as it comes from the perspective effect. Due to the finite depth of field, the particle images recorded on the cameras will come from the planes within the region $|z - z_0| < \Delta z_0/2$. Assuming that (x_p, y_p, z_0) is the particle location, it is possible to show that (Adrian & Westerweel 2011):

$$\begin{pmatrix} \Delta X_p \\ \Delta Y_p \end{pmatrix} = M_0 \begin{pmatrix} \Delta x_p \\ \Delta y_p \end{pmatrix} + M_0 \begin{pmatrix} x_p/z_0 \\ y_p/z_0 \end{pmatrix} \Delta z_p \quad (1.6)$$

where $\Delta X_p, \Delta Y_p$ are the components of the image displacement vector, and $\Delta x_p, \Delta y_p, \Delta z_p$ are the components of the physical displacement vector. Even a modest out-of-plane motion can determine a significant error. For example, in case of $\Delta z_p/D_I = 0.02$, $z_0 = 400mm$ and a field of view of $100 \times 100 mm^2$ will determine an error of 1% on a displacement equal to $|\Delta x_p|/D_I = 0.25$.

Furthermore, it is often forgotten that in (1.1) the magnification plays the same role of the image displacement in determining the accuracy of the measurement.

The magnification uncertainty can cause quite significant errors if the light sheet location is not determined accurately, or if the light sheet is too thick, or tilted. For example, 1mm error in the location of the light sheet when using a 100mm object distance results in a bias error of 1% of each velocity reading. The magnification can be quite accurately estimated performing an optical calibration (see Sec 1.2.1). However, it must be emphasized that the calibration is based on the assumption that the laser sheet plane coincides with the reference calibration target plane (without loss of generality, from this point on this plane is referred as $z = 0$, since the origin of the physical reference system is usually enclosed within it). In other words, the estimated magnification is accurate (within the calibration uncertainty) in the $z = 0$ plane, but not throughout the thickness of the laser sheet, nor even in the measurement plane. Furthermore, in many applications the target cannot be placed into the measurement domain, thus preventing the procedure described above. In these cases other (usually less reliable) solutions have to be implemented (for example precision markers machined on the walls of the test section to define the scale). Discetti & Adrian (2012) proposed a method to measure the magnification with high accuracy, provided that the optical access for a second camera is available. The procedure is the same addressed in Sec. 1.2.3, as it is based on a self-calibration procedure to identify the laser sheet location into the physical domain. The great advantage is that the method can be applied even with an *ex situ* calibration, implementing the refraction index changes into the mapping functions between spatial coordinates and image coordinates, as proposed by Wieneke (2005).

1.2 Stereoscopic Particle Image Velocimetry

As observed in the previous section, the in-plane displacement is the result of the projection of the 3C velocity field (the third component of it being relevant when the requirements of paraxial viewing $x_p/z_0 \ll 1$, $y_p/z_0 \ll 1$ are not closely met, and/or the out-of-plane motion is significant). As a matter of fact (1.6) is a system of two equations in three unknowns; with planar PIV one has to accept to neglect the second term on the right hand side of (1.6) and measure two of the three velocity components. On the other hand, the third component could be inferred if another simultaneous measurement is available from a different point of view and the second term on the right hand side of (1.6) is stressed, i.e. using two cameras in stereoscopic arrangement.

In Fig. 1.3 two possible stereoscopic arrangements are reported. The linear displacement (also referred as translation system) has the advantages of uniform magnification, uniform quality of the particle images, and the same depth of field requirements of monocular PIV. On the other hand, the common area is generally limited, especially for long focal length lenses.




Fig. 1.3 Basic stereoscopic PIV arrangements: **a)** lateral displacement **b)** angular displacement with Scheimpflug condition fulfilled (Willert 1997).

A possible solution, shown in Fig. 1.3a, consists in off-setting the sensors; however, sometimes this is not possible due to physical constraints of the camera body and due to degradation of the lens performances as the viewing angle increases.

The restriction on the angle subtended between the cameras lead to large errors in the measurement of the out-of-plane component, as documented by Lawson & Wu (1997a, 1997b). This constraint can be eliminated by using a rotational system (angular displacement, Fig. 1.3b), which allows to increase the angle while working with the optimum lens performances. On the other hand, the magnification is no longer uniform, and, moreover, in order to obtain well-focused particle images over the field of view the object plane, the lens plane and the image plane have to be collinear. This requirement, known as Scheimpflug condition (Scheimpflug 1904), further exasperates the non-uniformity in the magnification.

Prasad (2000) classified the stereoscopic reconstruction methods in geometric and calibration-based. Only the calibration-based methods will be discussed, as

many studies have demonstrated their superiority in terms of accuracy on the geometric methods.

1.2.1 Optical calibration

The calibration-based methods require the measurement of a mapping function, determining the correspondence \underline{F} between spatial coordinates \underline{X} and image coordinates \underline{x} :

$$\underline{X} = \underline{F}(\underline{x}) \quad (1.7)$$

Generally, the calibration consists in placing a target with highly-contrasted markers (for example white dots on a dark background) with known physical position. Soloff et al (1997) used 1mm circular dots with 5mm spacing on a black background. The target is placed possibly in the same location of the laser sheet, and translated to sweep a volume in which the illuminated region is enclosed (incidentally, if a 3D target is used there is no need to translate it; see for example the 3D wavy target used by Scarano et al 2005). One of the markers is identified as the origin of the physical reference system, while the other points coordinates are determined by setting the orientation of the axes. The identification of the markers onto the images can be performed by a template-matching technique, such as cross-correlating the target images with a template, and searching for the local maxima of the cross-correlation map. The final result is a discrete correspondence between spatial coordinate and image coordinates. A regression procedure is then used to fit a chosen function for the correspondence \underline{F} .

Generally speaking, the most adopted camera models are:

- Pinhole camera model (Tsai 1987), based on 6 extrinsic parameters (a translation vector and three Euler angles) and 6 intrinsic parameters (pixel aspect ratio; radial distortion factors of first and second order describing the distortion caused by the lenses; focal length; image coordinates intersection of the optical axis with the image plane);
- Generic interpolating functions, as: polynomial functions, 3rd order in x and y , 2nd order in z (Soloff et al 1997); rational functions (Willert 1997); bicubic splines (Lawson & Wu 1997a, 1997b).

1.2.2 Reconstruction of the flow field

The 3C reconstruction of the velocity field can be performed with 3 different approaches: generalized least-square reconstruction (Soloff et al 1997), mapping and warping (the last two methods are variations of the technique proposed by Willert 1997).

The *generalized reconstruction* is based on the interrogation of the raw (*warped*) images with a standard PIV algorithm. The 3C reconstruction of the

physical displacement $\underline{\Delta x}$ in the location \underline{x} is obtained by observing that a first-order Taylor series expansion of:

$$\underline{\Delta X} = \underline{F}(\underline{x} + \underline{\Delta x}) - \underline{F}(\underline{x}) \quad (1.8)$$

leads to the following relation:

$$\underline{\Delta X} \cong \underline{\nabla}(\underline{F}) \cdot \underline{\Delta x} \quad (1.9)$$

The expression (1.9) is a system of 4 equations (one for each component of the image displacement and for each camera) in 3 unknowns (the three components of the physical displacement). Ideally, 2 equations are linearly dependent; in general, due to measurement noise, a least square solution of (1.9) is used to determine $\underline{\Delta x}$.

In the case of the *mapping* technique, the images are de-warped onto a common grid in the physical space, and then interrogated with a standard PIV algorithm. The 3C displacement is obtained by using geometrical relations (Willert 1997). On the other hand, the *warping* approach, similarly to the case of the generalized reconstruction method, consists in interrogating the warped images, and subsequently dewarping the two vector fields to the corresponding physical positions. In general, the mapping technique is affected by potential loss of image quality due to interpolation (unless the dewarping process is included in the image deformation process). The warping approach, on the other hand, suffers for non-uniform and anisotropic spatial resolution due to the local magnification differences between the two cameras.

1.2.3 Disparity correction and self-calibration

The most severe source of error in Stereo-PIV is related to the misalignment between the calibration plate and the laser sheet. Neglecting the misalignment generates two sources of error: position error, i.e. the reconstruction is performed using velocity measurements relative to different positions; reconstruction error, i.e. the local gradient matrix in the method by Soloff et al (1997), or the viewing angles in the mapping/warping approaches, are computed in the wrong position. A sketch of the misalignment error is provided in Fig 1.4.

Willert (1997) proposed a procedure based on cross-correlation of the simultaneous dewarped images of the two cameras of the imaging system in the $z = 0$ plane; the local misalignments (also referred as *disparity* map) are used to correct the position errors. A more complete procedure, accounting also for the reconstruction error, is the self-calibration proposed by Wieneke (2005). In the self-calibration procedure the disparity map is used to estimate the position of the measurement plane; subsequently, the mapping functions are modified so that the measurement plane will be coincident with the $z = 0$ plane. A slightly different version was proposed by Giordano & Astarita (2009), arguing that, since the




Fig. 1.4 Sketch of the misalignment error between calibration and measurement plane (Giordano & Astarita 2009).

relation between the disparity map and the local misalignment involves the viewing angles, the initial misalignment determines an error in calculating the viewing angles themselves. Considering Fig. 1.4 as a reference, geometrical relations lead to:

$$\Delta z = \frac{D_u}{\tan(\alpha_1) + \tan(\alpha_2)} \quad (1.10)$$

Initially, one can compute α_1 and α_2 in P_1 or P_2 , or halfway, introducing an uncertainty. The process is iterated to reduce this source of error.

1.2.4 Multi-Plane Stereoscopic PIV

Stereo-PIV is a 2D3C technique, thus it provides only 6 of the 9 components of the velocity gradient tensor. The missing information can be recovered by using a multi-plane system, as proposed by Kähler & Kompenhans (2000). The imaging system consists of two pairs of cameras, while the illumination is provided by a four-pulse laser system in which the first couple of pulses is orthogonally polarized with respect to the second couple. The polarized light is separated by polarizing beam splitters, as shown in Fig. 1.5. The polarization is preserved if the particle tracers are small enough, i.e. with size comparable to the incident light.

Kähler & Kompenhans (2000) used a versatile four cavity laser system, which allowed them to place the 4 laser sheets in the same position (so that the four frames can be used for multi-frame interrogation, or to provide measurements of




Fig. 1.5 Schematic of the optical configuration for dual plane Stereo-PIV: 1-4 digital cameras; 5 lens; 6 mirror; 7 polarizing beam splitter; 8 light absorbing material; α opening angle (Kähler & Kompenhans 2000).

the fluid acceleration), or with an offset in the direction orthogonal to the laser sheets to accommodate for a large out-of plane displacement. In this second-case, all the three components of the vorticity vector and the nine components of the velocity gradient tensor can be measured.

1.3 Three-Dimensional Particle Image Velocimetry

1.3.1 Scanning PIV

Quasi-simultaneous 3D velocity measurements can be obtained by scanning a volume with a high pulse frequency laser sheet (Brücker 1995). The laser sheets could be either observed by a single camera or two cameras in stereoscopic arrangement (Hori & Sakakibara 2004), providing 3D2C or 3D3C velocity measurement, respectively (actually, in the first case the out-of-plane velocity component can be extracted by applying the continuity equation; however, the measurement is affected by noise due to differentiation of the in-plane velocity components). This approach is known as scanning PIV. Scanning PIV has potentially the same resolution of the standard PIV, and the velocity field computation can be performed using the same interrogation algorithms. On the other hand, due to depth of field requirements (i.e. larger $f_{\#}$ to increase δz , according to (1.3)), in order to achieve the same exposure level of an equivalent 2D PIV experimental




Fig. 1.6 Vorticity-magnitude iso-surfaces in the volume and vorticity magnitude contours in three spanwise planes at three different instants in the periodic motion of the flapping wing (David et al 2012).

setup one should generally increase the laser power or, whenever it is possible, increase the tracers diameter. This could be difficult, since high-frequency lasers provide lower power output than low-frequency lasers. Furthermore, a highly precise scanning mechanism is required for a successful and repeatable measurement (in Stereo-Scanning PIV the disparity map computation can be used either as a correction procedure for misalignments and for proper laser sheets identification).

Perhaps the most critical aspect is related to the scanning frequency. The volume must be swept within a time interval that is smaller than the characteristic timescale of the flow under investigation. This aspect limits the applicability of the technique to low-speed flows. For example, considering a free jet in water, with centerline velocity $U_j = 0.1 \text{ m/s}$ and exit diameter $d_j = 5 \cdot 10^{-2} \text{ m}$ (resulting in $Re = U_j d_j / \nu \approx 5000$), the characteristic timescale is $\tau_0 = d_j / U_j = 0.5 \text{ s}$, while the Kolmogorov timescale is $\tau_\eta \approx Re^{-1/2} \tau_0 = 7.1 \cdot 10^{-3} \text{ s}$, corresponding to a frequency of $f_\eta \approx 140 \text{ Hz}$. Assuming that a limit value for the time interval to sweep the volume is one half of the Kolmogorov timescale, if the repetition rate of the laser is 10 kHz only up to 17 planes can be imaged (considering that for each plane two images with the desired time separation have to be captured). However, this strict requirement can be relaxed if the focus of the investigation is on the large scale motion.

A slightly different implementation is documented by David et al (2012). A volume of $185 \times 140 \times 83 \text{ mm}^3$ (the first two-dimension being an average value, while the third one is limited by depth of field requirements) in the wake of a flapping wing is discretized in 100 parallel equidistant planes, illuminated with a time-interval of 25 ms , corresponding to 4 kHz . The particle images are captured by a high speed Photron camera with 1024×1024 pixels resolution. Successively, the images are back-projected into the volume to build a pseudo-3D object, interrogated by 3D cross-correlation.

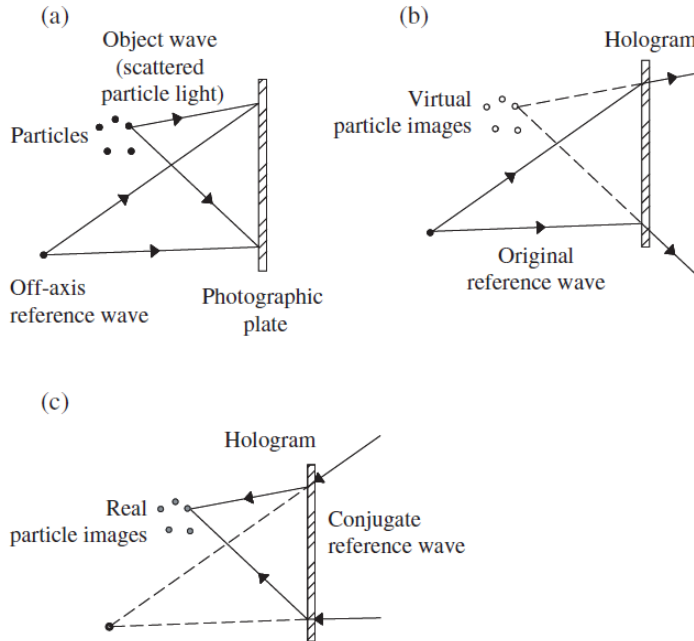


Fig. 1.7 Schematic of optical off-axis setups in particle holography. (a) Recording of a hologram of a particle field; (b) reconstruction of a virtual image; (c) reconstruction of a real image (Hinsch 2002).

The final result is a 3D3C velocity field with interestingly low uncertainty (the measured local divergence, which should be zero since the flow regime in incompressible, lies in the range $0.002 - 0.06 \text{ pixels/pixel}$). An example of vorticity magnitude contour representation is reported in Fig. 1.6.

1.3.2 Holographic PIV

Holographic image recording is performed without lenses, giving in practice infinite depth of field. This feature made the holographic principle (see Collier et al 1971 for reference) very appealing for the extension of PIV to the 3D field (Hinsch 1995, 2002). The schematic of the application of the holographic principle to PIV is reported in Fig. 1.7, and summarized in the following.

The particles are illuminated with a laser light wave. The light scattered by the particles is combined on a recording medium with a reference wave from the same laser source (while in planar PIV lasers are preferable because of the short pulse duration and high power, but still optional, in holography the requirement of coherent light really forces to use laser light). The produced interference pattern (*hologram*) is then used to reconstruct the image by illumination with a conjugate reference wave. The original image is generated by diffraction of the reference wave by the interference pattern on the recording medium.

The spatial resolution of the technique is related to the fringe spacing, given by $\lambda/(2s\sin\alpha)$, where α is half of the angle between the direction of propagation of the two waves. The fringe spacing sets the resolution requirement for the recording medium. This requirement can be relieved by reducing the angle α .

The main issue of using a photographic medium is that a very limited number of holograms can be captured (for each realization the film has to be changed). In order to allow statistics analysis the technique can be implemented on CCD cameras, where the reconstruction is performed numerically. This variation of the technique is also known as Digital Holographic PIV (DHPIV). DHPIV is mainly limited by the resolution of CCD cameras: generally the pixel pitch lies in the range $6 - 10\mu\text{m}$, allowing the detection of less than $200\text{lines}/\text{mm}$. In this case, using an in-line system (i.e. the object beam and the reference beam coincide) could be the only option. For in-line systems the main limitations are the small numerical aperture (providing very elongated particles along the optical axis of the holographic system) and the strong presence of *speckle noise*, arising from the interference of scattered light waves between themselves.

1.3.3 Defocusing PIV

Though traditionally included in the classification of photogrammetric systems (i.e. system using multiple camera views to determine the three-dimensional features of an object), the Defocusing PIV is an elegant solution which requires, as a matter of principle, a single camera to determine a virtual multiple-camera view. The idea, originally proposed by Willert & Gharib (1992), is based on defocusing imaging using a three-holes aperture. The working principle is illustrated in Fig. 1.8 in the simplified case of conventional imaging and with a two-holes aperture. For a conventional system a particle image outside of the focal plane is characterized by a larger diameter; for a defocusing system, instead, it will be imaged as a pair of particles, whose separation can be related to the distance from the reference focal plane. The introduction of a third hole eliminates the possible ambiguity of the particle position in front or behind the reference plane, as it would result in two opposite orientations of the corresponding imaged triangle.

In the implementation by Pereira et al (2000), the three-apertures system is replaced by three separate cameras arranged in the shape of an equilateral triangle, thus reducing the overcrowding of the images. The images are then processed to identify the triplet pattern and determine their spatial position. The displacement field can be determined either by 3D cross-correlation or by particle tracking. In the first case the object is discretized in voxels (the 3D equivalent of pixels, i.e. compact support functions with uniform value within a cubic domain and zero elsewhere) and the hyper-particles (i.e. the equivalent 3D particle images) are represented by an isotropic 3D Gaussian function.




Fig. 1.8 Schematic of the defocusing principle: **a)** conventional system **b)** defocusing system (Pereira et al 2000).

The limited solid angle subtended by the viewing rays determines a larger uncertainty on the depth position of the particles (10 times larger than the in-plane position, as documented by Pereira et al 2000). Furthermore, the image particles outside of the focusing plane will be larger, increasing the image density and limiting, as a matter of fact, the maximum allowed particle concentration. The need to minimize the number of overlapping particle images for the triangle pattern recognition also limits the maximum number of imaged particles for a successful reconstruction. Indeed the limited particle concentration and the 3D particle position identification step make the method more suitable for 3D particle tracking than for 3D cross-correlation.

1.3.4 3D Particle Tracking Velocimetry (PTV)

Particle Tracking Velocimetry relies on particles position identification and tracking of their motion within subsequent frames. The most appealing feature of 3D PTV is that tracking of the particles over long sequences allows to measure particle trajectories and Lagrangian properties (since the particles are illuminated within a 3D domain, one can expect that the effect of particles escaping from the measurement region is very limited), such as acceleration (Voth et al 1998), vorticity dynamics (Lüthi et al 2005), pair dispersion in turbulent flows (Bourgoin et al 2006), and many others.

The locations of the particles are identified by observing the pattern with multiple cameras from different directions; generally 3 cameras suffice to the purpose when the particle density is limited. Due to calibration errors and uncertainty in the particle image location in the images (pixelization,

inhomogeneous scattering conditions between the cameras, aberrations, etc.) the lines of sight of the same particle might not intersect within the volume. The best estimate of the particles location is obtained by inverting the generalized mapping function (1.7) for the several cameras, and solving the system of equations:

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = F_i^{-1} \begin{pmatrix} X_i \\ Y_i \end{pmatrix}, \quad i = 1, \dots, N_{cam} \quad (1.11)$$

The system (1.11) can be linearized by a Taylor series expansion truncated at the first order. The solution can be pursued with the standard numerical methods for over-determined systems of linear equations, for example a least square approach.

The spatial resolution and accuracy of 3D PTV are limited by the maximum allowed seeding density, both due to the finite probability of overlapping particles (affecting the reliability of the particles location algorithm) and due to the occurrence of spurious matchings, named *ghost particles* (Maas et al 1993; see Sec. 2.2 for further details).

The probability of overlapping particles can be easily estimated by knowing the source density N_s (see (1.4)) and modeling the process of imaging with a Poisson probability distribution. In this case, the probability for each pixel of being illuminated by a particle, or by more than one, is given respectively by:

$$Pr\{1 \text{ particle}\} = \exp(-N_s) \cdot N_s \quad (1.12)$$

$$\begin{aligned} Pr\{2 \text{ or more particles}\} &= 1 - Pr\{1 \text{ particle}\} - Pr\{0 \text{ particles}\} = \\ &= 1 - \exp(-N_s) \cdot [N_s + 1] \end{aligned} \quad (1.13)$$

The source density can be immediately associated to the image density expressed in number of particles per pixel (N_{ppp}) by observing that (1.4) is actually the product of the particle area $\pi d_\tau^2/4$ and the term $C\Delta z_0/M_0^2$. By assuming that M_0 is the average value of the magnification over the observed volume, and observing that the particle concentration is the total number of particles N_p divided by the size of the volume $\Delta x_0 \Delta y_0 \Delta z_0$, it follows immediately that the term $C\Delta z_0/M_0^2$, normalized with the pixel pitch d_p , is the number of particles divided by the area on which the particles are imaged, i.e. the number of particles per pixel. As a consequence, imposing d_τ^* as the particle image diameter in pixels (i.e. $d_\tau^* = d_\tau/d_p$):

$$N_s = N_{ppp} \cdot \frac{\pi d_\tau^2}{4 d_p^2} = N_{ppp} \cdot \frac{\pi d_\tau^{*2}}{4} \quad (1.14)$$

As for the ratio between the number of ghost and true particles for a three camera system, it can be estimated using the following relation, derived in Sec. 2.2.1:

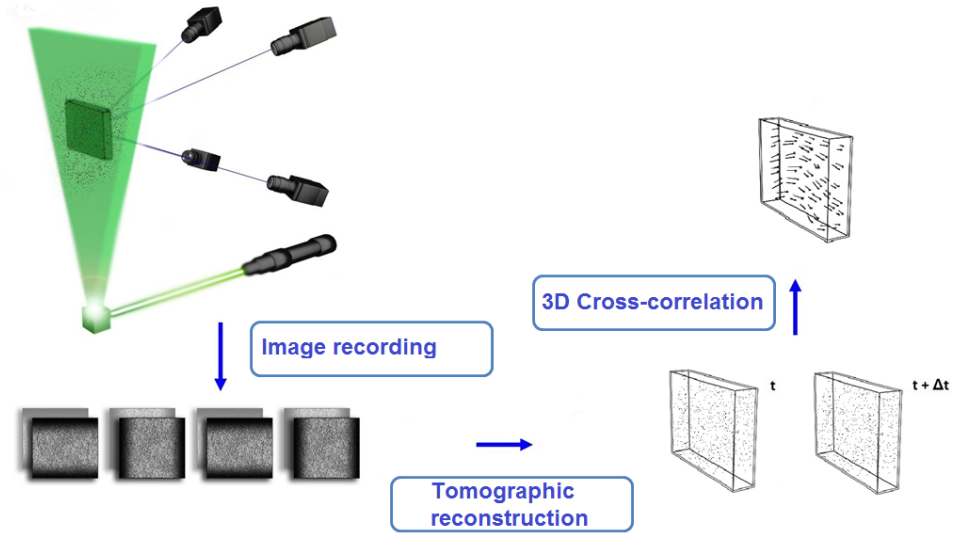


Fig. 1.9 Working principle of tomographic PIV (adapted from Elsinga et al 2006a).

$$\frac{N_g^{N_{cam}}}{N_p} = (N_{ppp})^{N_{cam}-1} d_\tau^* \left(\frac{\Delta z_0 M_0}{d_p} \right) \left(\frac{\pi d_\tau^{*2}}{4} \right)^{N_{cam}-2} = \frac{4}{\pi d_\tau} (N_s)^{N_{cam}-1} \Delta z_0 M_0 \quad (1.15)$$

where N_{cam} is the total number of cameras.

Maas et al. (1993) claimed $0.005ppp$ as an upper bound for the application of PTV with 3 cameras. Indeed, assuming a particle diameter of 3 pixels, the source density is equal to 0.035 and the probability of overlapping particles in that case would be 0.06%. Furthermore, assuming $\Delta z_0 = 20mm$, $M_0 = 0.1$ and $d_p = 10\mu m$, the ratio N_g/N_p is about to 0.1, i.e. of all the reconstructed particles about 10% will be spurious. Adding another camera, of course, reduces dramatically this percentage to less than 1%.

1.3.5 Tomographic Particle Image Velocimetry (Tomo-PIV)

The Tomo-PIV (Elsinga et al 2006a) relies on the reconstruction of the three dimensional distribution of scattering objects into a control volume. The reconstruction is achieved by analyzing the light intensity scattered by the particles illuminated by a pulsed light and imaged by multiple cameras from different viewing directions. The obtained light intensity field is discretized in cubic compact support functions named *voxels*, as they are, as a matter of fact, the 3D equivalent of the pixels. The full three components velocity field is then computed by 3D cross-correlation of the subsequent reconstructed volume. The working principle of the technique is sketched in Fig. 1.9.

The reconstruction problem consists in determining the best estimate of a 3D object using multiple projections. The analytical solution of the mathematic

problem dates back to the work by Radon (1917), who showed that a m -dimensional object can be reconstructed from an infinite set of $(m-1)$ -dimensional projections. In practice, a finite number of projections is available and the problem is ill-posed, leading to different techniques to determine approximate solutions. A more detailed description is provided in Chapter 2.

Since the technique does not rely on particles identification, Tomo-PIV can work with larger seeding density than 3D PTV. The parametric study by Elsinga et al (2006a) suggests that a satisfactorily accurate reconstruction can be achieved by a 4 cameras system with seeding density up to $0.05ppp$, with 3 pixels diameter particles (resulting in $N_s = 0.35$ according to (1.14)). The limitations in terms of depth of field (and, consequently, of $f_{\#}$) are the same of scanning PIV; the basic difference is the maximum amount of information that can be stored. Tomographic PIV has the advantage of allowing simultaneous illumination of the entire volume, while scanning PIV is limited by the scanning rate of the system. The number of imaged particles in both cases is limited by the relatively large particle diameter (depending on $f_{\#}$) and the finite probability of overlapping particle, i.e. of forming speckles. However, while Tomo-PIV records all the particle images at the same time, in scanning PIV the particle images are recorded in separate frames, enabling much better spatial resolution of the measurement. Indeed, supposing $F_0 = 1$, $F_I = 1$, the minimum size of the interrogation spot to satisfy the criterion $N_I F_0 F_I > 7$ in scanning PIV is given by:

$$N_I = N_{ppp} \cdot D_I^2 \quad (1.16)$$

resulting in $D_I \approx 12$ pixels. Instead, for a tomographic PIV system, the same relation has to be expressed in terms of ppv (particles per voxel):

$$N_{ppv} = \frac{N_{ppp}}{\Delta z_0 M_0} d_p \quad (1.17)$$

Similarly to (1.16), for a tomographic PIV system the minimum D_I is obtained by setting $N_I = 7$ in:

$$N_I = N_{ppv} \cdot D_I^3 \quad (1.18)$$

Supposing that $\Delta z_0 = 20mm$, $M_0 = 0.1$ and $d_p = 10\mu m$, (1.18) results in $D_I \approx 31$ voxels. Even accounting that the requirement $F_0 = 1$ can be fully satisfied in Tomo-PIV, while it is not accomplished in Scanning PIV in presence of out-of-plane motion, the difference is noticeable.

Straightforward substitution of (1.17) in the (1.18) leads to an estimate of the minimum physical size of the interrogation spot d_I (observing that $d_I = D_I d_p / M_0$):

$$d_I = \sqrt[3]{\frac{N_I \Delta z_0 d_p^2}{M_0^2 N_{ppp}}} \quad (1.19)$$




Fig. 1.10 Contour representation of d_l [mm] as a function of M_0 and Δz_0 for fixed $f_{\#}$, d_p , and image density equal to $0.05ppp$. The dashed-dotted curves represent the constraint due to finite depth of field ($\Delta z_0 = \delta z_0$), diffraction spot size ($d_{\tau}^* = 3$ pixels), and volume thickness larger than the interrogation spot size ($\Delta z_0 = d_l$). The area with colored contour lines is the valid area to place the design point for the Tomo-PIV experiment.

In Fig. 1.10 the contour representation of d_l as a function of M_0 and Δz_0 in the case of $0.05ppp$, $f_{\#} = 16$ and $d_p = 9\mu m$ is reported. The lower bound on the interrogation spot size is mainly determined by three constraints: diffraction spot size smaller than 3 pixels (in this case only diffraction imaging is considered); volume thickness smaller than the depth of field to have well-focused images (in practice, slightly out-of-focus particle images can be accepted, provided that the blur circle size is limited); volume thickness larger than the interrogation spot size to provide at least one vector. The area with contour lines indicates the region where the design point for the Tomo-PIV experiment can be placed. The first two constraints are $f_{\#}$ -dependent: a larger $f_{\#}$ relaxes the requirement on the depth of field while exacerbating the one on the particle diameter. In general, if high spatial resolution is required, decreasing $f_{\#}$ is advisable, though it limits the volume thickness (however, the minimum d_l is achieved for small Δz_0). Incidentally, a smaller particle diameter reduces N_s with the same N_{ppp} , thus increasing the quality of the reconstruction (provided that $d_{\tau}^* > 1$, see Fig. 2.3b).

A summary of applications of Tomo-PIV reported in the literature is presented in Tab. 1.1.

Chapter 1 – 3D Particle Image Velocimetry

| Experiment | Fluid, Re | Measurement domain | Illumination | Imaging system | d_I |
|---|----------------------------------|---|-----------------------------------|---|---|
| Wake of a cylinder (Elsinga et al 2006a) | Air, 2700 | 40x40x10mm ³ at 18.4vox/mm 0.05ppp | 400mJ/pulse at 1Hz | 4 cameras 1280x1024, $f_{\#} = 8$ | 41x41x21 vox 2.2x2.2x1.1mm 75% overlap |
| Finite cylinder obstacles (Hain et al 2008) | Water, 11000 | 80x50x12mm ³ at 50vox/mm | 150mJ/pulse at 10Hz | 4 cameras 4008x2672 $f_{\#} = 11$ | 48x48x48 vox 0.96x0.96x0.96mm 75% overlap |
| Time-resolved cylinder wake (Scarano & Poelma 2009) | Water, 1000 | 80x90x20mm ³ at 20vox/mm 0.025pp | 200mJ/pulse (double-pass) at 7Hz | 4 cameras 2048x2048 $f_{\#} = 11$ | 39x39x39 vox 1.95x1.95x1.95mm 75% overlap |
| Turbulent boundary layer/shock wave interaction (Humble et al 2009) | Air | 70x40x10mm ³ at 20vox/mm 0.05ppp | 400mJ/pulse at 10Hz | 4 cameras 2048x2048 $f_{\#} = 8 - 11$ | 42x42x42 vox 2.1x2.1x2.1mm 75% overlap |
| Drop coalescence (Ortiz-Dueñas et al 2010) | Water-glycerine | 80x90x20mm ³ at 12vox/mm 0.05pp | 30mJ/pulse at 1000Hz | 4 cameras 1024x1024 $f_{\#} = 8 - 11$ | 64x64x64 vox 5.3x5.3x5.3mm 75% overlap |
| Supersonic turbulent boundary layer (Elsinga et al 2010) | Water, $Re_{\theta} = 34000$ | 70x35x6.5mm ³ at 20vox/mm 0.05ppp | 250mJ/pulse | 4 cameras 2048x2048, $f_{\#} = 8 - 11$ | 40x40x40 vox 2x2x2mm 75% overlap |
| Rod-airfoil interaction (Violato et al 2011) | Air, 3500 | 50x50x3mm ³ at 16.6vox/mm 0.05ppp | 12mJ/pulse at 5000Hz | 4 cameras 1024x1024 $f_{\#} = 8$ | 47x47x19 vox 2.8x2.8x1.1mm 75% overlap |
| Circular and chevron jets (Violato & Scarano 2011) | Water, 5000 | 30x30x50mm ³ at 20vox/mm 0.04ppp | 25mJ/pulse at 1000Hz | 4 cameras 1024x1024 $f_{\#} = 32$ | 40x40x40 vox 2x2x2mm 75% overlap |
| Carotid artery bifurcation (Buchman et al 2011) | Water-glycerine, 360 | 22x24x6mm ³ at 40vox/mm 0.01ppp | 400mJ/pulse at 5Hz | 4 cameras 1280x1024, $f_{\#} = 5.6 - 8$ | 64x64x64 vox 1.6x1.6x1.6mm 75% overlap |
| Turbulent boundary layer (Atkinson et al 2011) | Air, 7800 | 60x60x9mm ³ at 20vox/mm 0.03ppp | 250mJ/pulse (double pass) at 1Hz | 4 cameras 2048x2048, $f_{\#} = 8$ | 64x64x64 vox 3.2x3.2x3.2mm 75% overlap |
| Turbulent channel flow (Schäfer et al 2011) | Air, 17000 | 18x18x7mm ³ at 96vox/mm 0.03pp | 30mJ/pulse | 4 cameras 2048x2048 $f_{\#} = 22$ | 96x96x96 vox 1x1x1mm 75% overlap |
| Time-resolved turbulent boundary layer (Schröder et al 2011) | Water, $Re_{\theta} = 2460$ | 63x68x15mm ³ at 11.7vox/mm 0.05ppp | 25mJ/pulse at 1000Hz | 6 cameras 1024x1024 $f_{\#} = 11$ | 32x32x32 vox 2.7x2.7x2.7mm 75% overlap |
| Wake of a trailing edge (Ghaemi & Scarano 2011) | Air, 38600 | 47x47x8mm ³ at 22vox/mm 0.065ppp | 13mJ/pulse (multi-pass) at 2700Hz | 4 cameras 1024x1024 $f_{\#} = 11-16$ | 32x32x32 vox 1.5x1.5x1.5mm 75% overlap |
| Zigzag boundary layer trip (Elsinga & Westerweel 2012) | Water, $Re_{\theta} = 116 - 184$ | 120x55x7mm ³ at 18.5vox/mm 0.03ppp | 250mJ/pulse at 2Hz | 4 cameras 2048x2048, $f_{\#} = 16$ | 28x28x28 vox 1.2x1.2x1.2mm 75% overlap |
| Pitching plate (Buchner et al 2012) | Water, 7500 | 85x120x12mm ³ at 25vox/mm 0.05ppp | 240mJ/pulse | 4 cameras 4008x2672 $f_{\#} = 11$ | 64x64x64 vox 2.6x2.6x2.6mm 50% overlap |
| Taylor-Couette flow (Tokgoz et al 2012) | Water, 1000-47000 | 40x20x10mm ³ at 27vox/mm 0.025ppp | 50mJ/pulse at 7.55Hz | 4 cameras 4800x3200 | 40x40x40 vox 1.5x1.5x1.5mm 50-75% overlap |

Table 1.1 Summary of Tomo-PIV applications reported in the literature.

Chapter 2 - Tomographic PIV

As outlined in the previous chapter, Tomographic Particle Image Velocimetry (Tomo-PIV) is perhaps the most promising among the possible candidates to play the role of the 3D equivalent of planar PIV. Though some particular experimental conditions might suggest using one technique more than the others, Tomo-PIV has shown in its early years of development great margins of improvement and remarkable flexibility (see Tab. 1.1 for an overview of applications reported in the literature). The technique relies on the reconstruction of the particles distribution within the measurement volume using multiple projections of it, and velocity measurement by 3D cross-correlation of the reconstructed volumes. Differently from the case of the standard Computed Axial Tomography (CAT), which is currently a well established technique in medical imaging and diagnostic, the object to be reconstructed is not steady. Consequently, the projections need to be simultaneously captured, unless one can scan the volume with ultra-high speed cameras and high precision mechanic shift systems. This option is currently unfeasible, thus simultaneous imaging of the particles pattern with several cameras is required. The solution of the ill-posed problem of the tomographic reconstruction for Tomo-PIV cannot be achieved with the standard Radon-Transform-based methods of the CAT. On the other hand, the objects to be reconstructed are much simpler, i.e. sparsely populated volumes with bright high spatial frequency spots on a dark background. In this scenario, iterative algebraic reconstruction techniques (Herman & Lent 1976) can be very effective. The features and the implementation of the algebraic reconstruction techniques are outlined in Sec. 2.1.

The performances of the algebraic methods in terms accuracy are strongly seeding-density-dependent. Indeed, the imaged particles are positioned into the measurement volume by taking into account the coherence of the reconstructed object with the ensemble of the projections; as a matter of fact, the framework is somehow similar (but more robust) to the PTV scenario, in which the 3D position of the particles is determined by triangulation. As a consequence, the problem of the ghost particles discussed in Sec. 1.3.4 still holds. Even though the cross-correlation approach is in general more robust than particle tracking with respect to random errors (for example loss of pairs, spurious matchings, etc.), the real issue is related to the bias error induced by the coherent motion of the ghost particles. This aspect is discussed in Sec. 2.2.

Though relatively “young”, Tomographic PIV with the aid of algebraic methods is quite well established among the scientific community. However, the scientific production on the development of solutions to push the limits of the technique is relevant. Many researchers are involved in a continuous *labor limae* to optimize the performances of the algebraic methods in terms of accuracy and hardware

requirements (see Sec 2.3 for an overview). The PIV analysis by cross-correlation is conceptually similar to the 2D case, though the implementation hides some harshness. The large amount of data and the relevant computational load require very efficient algorithms. The problem is outlined in Sec. 2.4, and some solutions are proposed in Chapter 5.

2.1 Algebraic reconstruction techniques for Tomographic PIV

2.1.1 The reconstruction problem

Tomography deals with the reconstruction of a m -dimensional object from a set of $(m - 1)$ -dimensional projections. A projection is defined as a line integral of the object to be reconstructed along a viewing direction. More specifically, in Tomographic PIV the object is the 3D distribution of the light scattered by the tracer particles $f(x, y, z)$ (where x, y, z are the spatial coordinates in a 3D reference system), while the recorded intensity values on the cameras constitute the set of projections $I(X, Y)$ (where X, Y are the coordinates in the 2D camera reference system). In the case of the iterative algebraic reconstruction techniques (Herman & Lent 1976), the object is discretized as a 3D array of voxels, as outlined in Sec. 1.3.5. The line integrals of the intensity along the lines of sight are discretized as a weighted sum of the voxel intensities:

$$I(X_i, Y_i) = \sum_{j=1}^{N_{vox}} w_{i,j} f(x_j, y_j, z_j) \quad i = 1, \dots, N_{pix} \quad (2.1)$$

The subscript i and j indicate the i -th pixel and the j -th voxel, respectively; N_{pix} and N_{vox} are the total number of pixels and voxels; $w_{i,j}$ is a weighting coefficient, determining the influence of the intensity of the j -th voxel on the intensity recorded on the i -th pixel. The (2.1) results in a system of N_{pix} linear equations with N_{vox} unknowns.

It can be readily shown that the problem is commonly underdetermined in the Tomo-PIV scenario. Consider for example a 4 cameras system with average magnification $M_0 = 0.1$ and pixel pitch $d_p = 10\mu m$, i.e. resolution of $10pix/mm$. Suppose that the illuminated region is $100 \times 100 \times 20mm^3$. the particles will be imaged on the projection of this region, i.e. approximately $1000 \times 1000pix$ (actually the projection is slightly larger, depending on the viewing angle; however, the order of magnitude of the number of illuminated pixels is very close to the one indicated above). On the other hand, since the spatial resolution of the volumetric discretization is set to be approximately the same of the imaging system, the number of voxels is $1000 \times 1000 \times 200$. As a consequence, the (2.1) is a system of $4 \cdot 10^6$ equations in $2 \cdot 10^8$ unknowns. The problem is largely underdetermined;

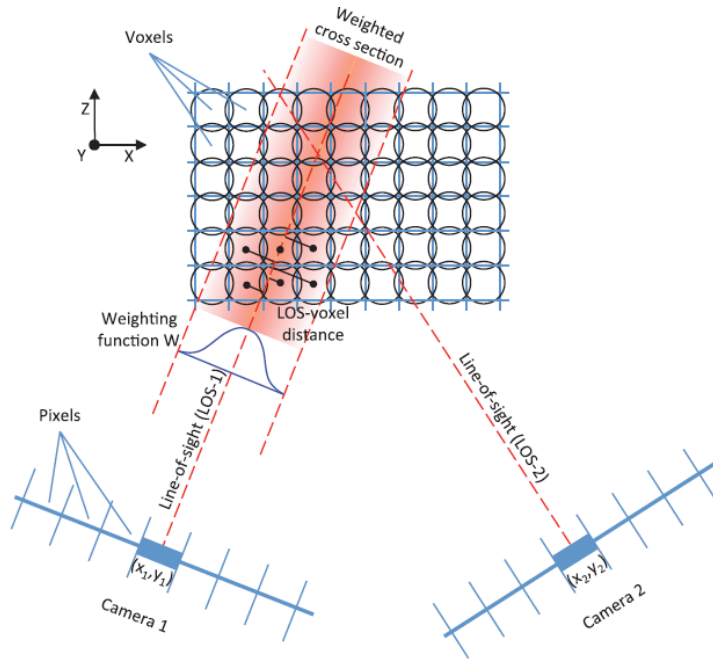


Fig. 2.1 Imaging model for the tomographic reconstruction. The weighting elements are non-zero in the shaded area; their values are a function of the distance between the line of sight and the center of the voxel. Spherical voxels are assumed to simplify the calculation of the weighting elements. (Scarano 2013).

however, some additional constraint can consistently reduce the dimension of the space of the solutions (for example in Tomo-PIV the solution is positive-definite).

Furthermore, the weighting matrix is extremely large (for the example indicated above, the number of weighting elements is $8 \cdot 10^{14}$, requiring $3200TB$ for the memory storage in single floating point precision), preventing the solution of the (2.1) with matrix approaches. Luckily, the matrix is also very sparse, since each pixel is influenced only by a small percentage of the voxels (i.e. those relatively close to the line of sight). This concept is illustrated in Fig. 2.1 in the simplified scenario of a 2D object imaged by two 1D cameras.

The weighting elements can be modeled and computed in a number of different ways. Most often, the weighting elements are computed as the intersecting volume between the voxels and the line of sight, expanded as a cylinder with cross-section equal to the pixel size. A further simplification consists in replacing the voxels with equivalent spheres, so that the intersected volume is a function of solely the distance between the axis of the line of sight and the center of the voxel, and the reciprocal orientation can be taken out of the process. An analytical solution is provided by Lamarche & Leroy (1990).

The most immediate approach is to set the weighting function linearly varying between 1 (line of sight passing through the center of the voxel) and 0 (no intersection between the line of sight and the voxel); however, one can more generally use more complicate weighting functions, such as Gaussian functions or B-Splines (see Hanson & Wecksung 1985 for more details). This procedure is commonly referred as *basis-function-method*. In all cases, the computational burden of calculating the triple integral to determine the weighting elements can be consistently reduced by using a pre-calculated look-up table (which entries are the distance between the line of sight and the center of the voxel, and the resolution ratio between voxels and pixels).

A slightly different approach consists in determining the weighting elements in terms of the distance between the image point corresponding to the center of the voxel and the center of the pixel. This method is easily implementable, as it requires only the straightforward application of the direct mapping function (Sec 1.2.1) and the calculation of the distance of two points in the image plane. For this reason, it will be referred as *direct-method* from this point on. As in the basis-function-method, one can choose generic functions of the distance and the resolution ratio. Furthermore, a more general approach consists in adapting locally and for each camera the functions for the weighting elements calculation, i.e. a variable Optical Transfer Function (OTF), as illustrated by Schanz et al (2013). The OTF can be locally set to compensate for optical distortions, aberrations, variable magnification, etc.

2.1.2 Algebraic reconstruction techniques

In the very first stages of development of Tomo-PIV, several approaches to solve the system (2.1) have been investigated. Elsinga et al (2006a) tested two iterative methods:

- Algebraic Reconstruction Technique (ART, Gordon et al 1970):

$$[f(x_j, y_j, z_j)]_{k+1} = [f(x_j, y_j, z_j)]_k + \mu \frac{I(X_i, Y_i) - \sum_{j=1}^{N_{vox}} w_{ij} [f(x_j, y_j, z_j)]_k}{\sum_{j=1}^{N_{vox}} w_{ij}^2} w_{i,j} \quad (2.2)$$

- Multiplicative Algebraic Reconstruction Technique (MART, Herman & Lent 1976):

$$[f(x_j, y_j, z_j)]_{k+1} = [f(x_j, y_j, z_j)]_k \left(\frac{I(X_i, Y_i)}{\sum_{j=1}^{N_{vox}} w_{i,j} [f(x_j, y_j, z_j)]_k} \right)^{\mu w_{i,j}} \quad (2.3)$$

The subscript k indicates the iteration number, and μ is a relaxation coefficient, that plays a significant role in the stability of the iterative process (in MART the stability criterion is $0 < \mu < 2$). In both cases one iteration is completed after that the update equation has been executed for all the projections (i.e. for all the pixels of the cameras set). Elsinga et al (2006a) documented a striking superiority




Fig. 2.2 Mechanism of formation of ghost particles. **a)** two particles imaged by two cameras lead potentially to two ghost particles (this is not exactly true in 3D, as the lines of sight can be non-coplanar. **b-c-d)** three possible solutions of the reconstruction problem (Elsinga et al 2011).

of MART in the framework of Tomo-PIV, as it operates as a logic AND operator (i.e. each voxel carries a non-zero intensity only and if only all its projections are non-zero); conversely, ART works as a logic OR operator (i.e. each voxel has a non-zero intensity if at least one projection is non-zero). Since the object to be reconstructed consists of small high intensity spots on a dark background, MART is certainly more suitable for this task. Actually, the constraint of positive-definite solution is automatically applied by the MART algorithm, as the projections in Tomo-PIV are strictly non-negative, and so are the voxels intensities. As a matter of fact, each zero-intensity projection set automatically to zero the intensity of all the voxels along the line of sight, thus consistently reducing the space of the solutions for the system (2.1). The MART procedure is characterized by a relatively fast convergence, as the reconstructed distributions do not change significantly after 5 iterations.

In the work by Atkinson & Soria (2007) a wider spectrum of algebraic techniques is considered, still resulting in MART being more suited for handling the problem of tomographic reconstruction in the Tomo-PIV scenario. To date, MART is considered the consolidated method in this framework.

The intensity fields reconstructed by MART are usually affected by artefacts mainly due to the ambiguity in the estimation of the particles correct position along the line of sight. Following Novara et al (2010), the main sources of error can be bundled in three categories:

- Ghost particles (Maas et al 1993), i.e. intensity blobs forming at the intersection of line of sight carrying non-zero values, but not corresponding to the position of actual particles. The mechanism of the ghost particles formation is sketched in Fig. 2.2 for the simplified case of a 2D slice reconstructed by two 1D cameras. Obviously the larger is the number of particles, the larger is the number of ambiguities in the reconstruction;

furthermore, the thicker is the volume, the higher is the number of positions that each true particle may occupy into the reconstructed volume, and accordingly the higher is the probability of ghost particles occurrence;

- Elongation of the reconstructed particles in the depth direction (especially when the total solid angle is too small), determining an increase of random errors in the measurement of the velocity component along the depth. On the other hand, a large viewing angle can lead to longer line of sights, enforcing the first source of error;
- Discretization artefacts due to under-sampling, arising when the diameter of the particles to be reconstructed is below 2-3 voxels.

The first source of error can be reduced by increasing the number of cameras, reducing the seeding density (at the expense of the spatial resolution), or using high accuracy reconstruction schemes (as the MTE-MART by Novara et al 2010, see Sec. 2.3.3), among the others. The effect of the solid angle is merely geometrical, and it can be compensated only by proper orientation of the cameras. As for the effect of the discretization, one should note that the particle images are very often diffraction limited spots, i.e. with nearly Gaussian shape. As a consequence, the reconstructed particles should have Gaussian shape as well. It should be reasonable to expect that applying some artificial diffusion with a 3D Gaussian filter in the iterative process of (2.3) should reduce the discretization error and increase the reconstruction accuracy. This aspect is briefly outlined in Sec. 2.3.2 and discussed in more detail in Chapter 4.

2.1.3 MART performances

Elsinga et al (2006a) assessed the performances of MART using 2D numerical simulations, in which a 2D object is reconstructed by 1D cameras. More recently 3D simulations have been conducted in several studies with the same frame of mind (see, for instance, Worth et al (2010) and de Silva et al (2012)), leading to more predictive results as they include the effect of non-coplanar views in the 3D scenario. The common conclusion is that the number of cameras and the tracers concentration are the leading parameters in determining the accuracy of the reconstruction (see Fig. 2.3, left, relative to 2D simulations taken from Scarano 2013). It comes with no surprise that these two parameters have a strong influence on the ghost particles number, as outlined later in Sec. 2.2.1. The accuracy is quantified in terms of the correlation factor between the synthetic generated distribution $f_{true}(x, y, z)$ and the reconstructed one $f_{rec}(x, y, z)$:

$$Q = \frac{\sum_{j=1}^{N_{vox}} f_{true}(x_j, y_j, z_j) f_{rec}(x_j, y_j, z_j)}{\sqrt{\sum_{j=1}^{N_{vox}} [f_{true}(x_j, y_j, z_j)]^2 \sum_{j=1}^{N_{vox}} [f_{rec}(x_j, y_j, z_j)]^2}} \quad (2.4)$$




Fig. 2.3 Quality of the reconstruction after 5 MART iterations as a function of the: particles concentration (left) for the case of $d_T^* = 3$; particles diameter (right) for the case of $0.05ppp$. (Scarano 2013).

Other experimental parameters might have a relevant impact on the reconstruction quality, such as the particles diameter (Fig. 2.3, right), the calibration uncertainty, the image quality and pre-processing, the optical aberrations, the solid viewing angle subtended by the cameras, etc. Most often, the quality drops down as the source density approaches values close to the speckle-mode (Adrian & Westerweel 2011), i.e. $N_s > 0.4$, regardless of the number of cameras, due to loss of the transmittivity of the medium and particles overlap along the lines of sight.

In general, Elsinga et al (2006a) proposed $Q > 0.75$ as a rule of thumb for a successful reconstruction. A commonly adopted upper bound for the image density is $0.05ppp$ (which does not affect much the spatial resolution, as it depends only weakly on the particle density; see for instance (1.19)).

2.1.4 Volume self-calibration

In the results from the synthetic experiments presented in Sec. 2.1.3 the weighting functions are obtained with ideally perfect correspondence between world coordinates and image coordinates. In the real experiments, this correspondence is achieved through the optical calibration procedure described in Sec. 1.2.1. The process requires the same software and hardware of the optical calibration for Stereo-PIV, but the effects of the calibration uncertainties are significantly different, thus affecting the procedure itself.

In Stereo-PIV the misalignment between the laser sheet and the calibration target is the most significant source of error, as discussed in Sec. 1.2.3. In Tomographic PIV this aspect is not relevant, provided that the volume swept by the

target in the calibration process covers reasonably well the illuminated volume. On the other hand, Tomo-PIV suffers of uncertainties due to inaccurate calibration plates, unstable calibration shift mechanism, loose connections in the cameras system, vibrations, optical distortions that are not accounted by the mapping function (this aspect is particularly critical when using the pinhole camera model), thermal deformations, etc. Indeed, Tomo-PIV requires that the mapping functions have uncertainties possibly below 0.4 pixels (Elsinga et al 2006a), which might be difficult to achieve in case of large volumes and significant aberrations (for example when imaging through optical windows). When the error is large, lines of sight relative to the same particle could not intersect within the volume, thus providing cancellation (or intensity subtraction) of true particles. In the very first applications the problem was tackled by smoothing with a Gaussian filter the original images, thus increasing the particles size and the probability of successful reconstruction in case of calibration errors. The immediate drawback is the reduction of the reconstruction quality due to the increased particle image diameter (see Fig. 2.3, right).

Subsequently, the development of a Volume Self-Calibration (VSC) technique by Wieneke (2008) determined a significant leap forward in this sense. The VSC is based on the correction of the mapping functions using the actual particles. The technique consists in locating the particles on the camera images and finding the 3D position of matching particles through triangulation, as in PTV (see Sec 1.3.4 for further details). The residual disparity obtained by computing the image distance between the projection of each particle and the correspondent positions on the camera images are used to correct the mapping functions. The application of this method works quite easily in case of very sparsely populated images (as in PTV), while the extension to Tomo-PIV images (with image density up to 0.1ppp) is not straightforward at all, and most often it is not possible to record a set of images with low density before the experiment. Wieneke (2008) proposed a clustering technique, consisting in an artificial reduction of the image density by taking only the brightest particles.

The procedure to identify the matching particles is outlined in Fig. 2.4. For each particle on the first camera image, the corresponding candidates for the matching on the second camera are those residing within a strip centered on the epipolar line (i.e. the projection of the line of sight of the first camera on the second camera) and half-width equal to the uncertainty ε_{max} (that has to be larger than the expected maximum calibration error). A first guess position is computed for each candidate, and then projected on the other cameras to find the corresponding particle, and, eventually, the 3D particle position by solving (1.11). Finally, the projection of the 3D particle position is compared with the respective image particle positions; the disparity vectors are displayed in a histogram map to separate true matchings from ghost particles. The disparity vectors are computed for a set of sub-volumes




Fig. 2.4 Particle triangulation procedure for the Volume Self-Calibration technique. (Wieneke 2008).

composing the measurement volume, and then used to correct the original mapping functions.

The VSC technique has shown to be able to reduce significant errors in the calibration. For example, in Chapter 8 an application is presented in which a jet issues from the bottom of a water tank into a plexiglass cylindrical chamber. Due to physical restriction the calibration is performed by moving the target in absence of the chamber; subsequently, the VSC is used to correct the misalignment of the lines of sight due to refraction effects induced by the plexiglass cylinder. A more detailed description of the implemented procedure and of the relevance of the corrections is provided in Sec. 8.2.3.

2.2 The problem of the ghost particles

The ghost particles constitute a severe source of error in the tomographic reconstruction. Even though a model to understand and quantify the ghost particles formation is of crucial importance, more interesting questions arise regarding their influence on the velocity measurement in terms of bias and random errors. For

example, Elsinga et al (2006a) documented an experiment on the wake behind a cylinder in which the ghost particles outnumbered the actual particles; the velocity measurement was still satisfactory nonetheless.

In Sec. 2.2.1 a model to estimate the number of ghost particles is presented. The model is slightly different from the one proposed by Elsinga et al (2006b) in its derivation, formulation and functional dependencies, and conceptually more similar to that by Wieneke (2008). In Sec. 2.2.2 the focus is on the effect of the ghost particles on the velocity measurement.

2.2.1 Estimate of the number of ghost particles

A ghost particle is formed in every occurrence of intersection of all the lines of sight in which a particle image is present. Consider, for example, an illuminated volume of size $\Delta x_0, \Delta y_0, \Delta z_0$ and a particle image of the first camera. The particle can be located anywhere along the line of sight, whose length is approximately Δz_0 (without affecting the generality of the problem the effect of the viewing angle is neglected). The possible candidates for the matching of a single particle of the first camera on the second camera are those included in a strip with length equal to the projection of the line of sight on the second camera (it can be estimated by multiplying Δz_0 for the average magnification M_0 and dividing by the pixel pitch d_p) and width equal to the particle image diameter d_τ^* . The number of candidates for the matching for each particle image of the first cameras in a 2 camera system can be statistically determined by multiplying this area for the particle image density. Normalizing with the number of true particles N_p one obtains:

$$\frac{N_g^{2cam}}{N_p} = N_{ppp} d_\tau^* \left(\frac{\Delta z_0 M_0}{d_p} \right) = \frac{4}{\pi d_\tau} N_s \Delta z_0 M_0 \quad (2.5)$$

where d_τ is the particle diameter obtained by (1.2), and (1.14) is applied.

Two projections are enough to define for each candidate a trial position in the 3D space. The trial position is projected onto the third camera to find the possible matchings. In this case the search area is a circle with diameter equal to the particle image diameter. The number of spurious matchings is statistically determined by multiplying the particle image density N_{ppp} for the search area, which is equal to the source density N_s . This leads to the general formula for a N_{cam} system:

$$\frac{N_g^{N_{cam}}}{N_p} = (N_{ppp})^{N_{cam}-1} d_\tau^* \left\{ \frac{\Delta z_0 M_0}{d_p} \right\} \left[\frac{\pi d_\tau^{*2}}{4} \right]^{N_{cam}-2} = \frac{4}{\pi d_\tau} (N_s)^{N_{cam}-1} \Delta z_0 M_0 \quad (2.6)$$

The term in curly brackets is the depth of the volume in voxels (provided that the resolution ratio between voxels and pixels is equal to 1), while the term in square brackets is the diameter of the particle image.




Fig. 2.5 Reoccurrence of ghost particles in the two exposures. Solid lines and dark blue particles indicate the first exposure; dashed lines and light blue particles relate to the second exposure. Dark and light red circles indicate ghost particles in the first and the second exposure, respectively. Left: ghost particle reoccurring in both the exposure. Right: ghost particle appearing only in the first exposure due to the presence of velocity gradients (Elsinga et al 2011).

2.2.2 The role of the ghost particles in the velocity measurement

The intuition might suggest that in Tomo-PIV the role of the ghost particles is marginal with respect to 3D-PTV, since MART iteratively damps their intensity; furthermore, cross-correlation is much more solid in terms of spurious matchings than particle tracking. These considerations are supported by the evidence that Tomo-PIV is able to work with particle image density more than 10 times larger than 3D-PTV. In many application reported in Tab. 1.1 the ratio (2.6) is close or even larger than 1. However, the contribution of the ghost particles to the cross-correlation maps is not exclusively dependent on their number *per se*. Elsinga et al (2011) observed that in some conditions ghost particles pairs between the exposures might form, and give an undesirable contribution to the cross-correlation map.

The process of ghost particles pairing is sketched in Fig. 2.5. Consider a set of 4 particles observed by 4 cameras, and suppose that the respective lines of sight intersect in a common point, determining the creation of a ghost particle. If the displacement along the direction normal to the viewing direction is nearly the same for the set of particles, a ghost particle will be formed in the second exposure by the same group of true particles. As a matter of fact, the ghost particle reoccurs in both the exposures in nearby locations, and its displacement is approximately the average displacement of the set of true particles. On the other hand, if the displacement is significantly different (more than one particle image diameter) for at least one particle of the set (i.e. if there is a significant velocity gradient) the ghost particles will not reoccur in the second exposure (see Fig. 2.5, right).

Elsinga et al (2011) draw two importance consequences:

- the coherent motion of the ghost particles might lead to velocity detection even outside of the illuminated volume. This problem can be avoided by identifying the illuminated region with the self-calibration technique (see Sec. 2.1.4) or by summing the particle intensities in the xy planes to determine the laser profile; the velocity vectors outside of this region are rejected;
- since the ghost particles displacement is an average over a set of particles whose reciprocal separation might be much larger than the interrogation spot size, the effect is a smoothing (say, a modulation) of the velocity field.

The second consequence is much more severe. Some of the most recent studies tackle the problem without focusing on the pairing, i.e. the bias and random errors due to the ghost particles are challenged by increasing the reconstruction accuracy (see for example Petra et al 2009, Novara et al 2010, Novara & Scarano 2012a, de Silva et al 2013) or by using hybrid PIV-PTV approaches (Novara & Scarano 2012b, Wieneke 2013). More recently, Discetti et al (2012) proposed a low cost Tomo-PIV setup, consisting of two independent tomographic systems with cheap single-shutter cameras. Since the ghost particles distribution is related to the camera orientations, this system actually provides reconstructed distributions in which the ghost particles do not reoccur coherently in the two exposures even in absence of velocity gradients. More details are provided in Chapter 6.

2.3 Improvements and advances on the reconstruction techniques

Since the birth of Tomo-PIV with the work by Elsinga et al (2006a), MART conquered the status of state-of-art reconstruction technique in the PIV community. Nevertheless, two points of weakness immediately captured the attention of several research groups, i.e. the large computational cost and the limits of MART in terms of allowed seeding density. The first concern drove some studies towards the development of more computationally efficient reconstruction techniques (see Sec. 2.3.1), mainly revolving around MART. The second issue stimulated the development of brand new approaches to tackle the mathematical solution of the problem of the tomographic reconstruction, moving away from the path of the traditional tomography. An overview of these advanced methods is provided in Sec. 2.3.2.

2.3.1 Fast reconstruction techniques

Several efforts have been made in order to reduce both the computational cost of the reconstruction and the memory storage requirements while retaining the same accuracy. Most of the suggested solutions are based on the sparsity of the

distributions to be reconstructed. Indeed, making reference to Fig. 1.2, 7-10 particles in each interrogation spot are sufficient to guarantee a valid detection probability above 95% (though the absence of *out-of-plane* motion can induce to slightly stretch this requirement). If a cube of $32 \times 32 \times 32$ voxels is considered, and proper $f_{\#}$ and resolution ratio are set in order to obtain a diameter of the particles of about 3 voxels, less than 1% of the voxels is characterized by a non-zero intensity value. Anyhow, as observed by Atkinson & Soria (2009), ambiguities in the reconstruction process and limited viewing angles determine the presence of noise and spurious voxel intensities, increasing the percentage of voxels with non zero intensity to approximately 5-10%.

Furthermore, the iterative nature of MART implies a certain dependence of the first guess on the rate of convergence. A uniform first guess is commonly assumed (Elsinga et al 2006a), but several alternative solutions have been proposed, mainly enjoying the possibility of identifying locations with zero intensity in advance. Worth & Nickels (2008) proposed a Multiplicative First Guess (MFG) to estimate the initial intensity field. MFG is mainly based on back-projection of the recorded images throughout the volume; the obtained fields are then multiplied in order to identify the locations with zero-intensity (or below a certain threshold), whose values will be no more updated during the process. Anyway, an efficient application of MFG coupled with MART requires the storage of at least one weighting element per camera and for each voxel; if sufficient computational resources are not available, the weighting elements have to be calculated on-the-fly, increasing considerably the processing time. A faster alternative solution, proposed by Atkinson & Soria (2009), is based on straightforward projection of the voxels on the image planes to determine whether their intensity is zero or not; the approach is named MLOS (Multiplicative Line Of Sight). Non-zero projections on each camera are multiplied to obtain a suitable first guess, successively refined by means of a MART variant with simultaneous update (SMART, Mishra et al 1999). The combination of MLOS with 40 SMART iterations enables the acceleration of the reconstruction process of 3.8 times for a seeding density of $0.05ppp$ in a 2D simulation for a 3 cameras setup, achieving the same accuracy of the standard method based on 5 MART iterations. However, this advantage quickly collapses by increasing the seeding density. Furthermore, both MFG and MLOS do not provide any attenuation of the intensity of the ghost particles.

A different solution is the Multi-Resolution MART proposed in this work (Chapter 3). The technique is still based on the exploitation of sparsity and quick first guess estimation, but the latter is slightly more elaborated. The first guess distributions are estimated by performing the tomographic reconstruction by MART with smaller spatial resolution (for example reduced by a factor of 2) while retaining the same voxel/pixel size ratio by binning the camera images. The obtained reconstruction are then interpolated on the final grid and subsequently refined by 2-3 further MART iterations. The real advantage of MR-MART with

respect to MFG and MLOS is that the first guess estimate already includes a damping process of the ghost particles, as it is MART-based. Furthermore, in some conditions the weighting matrix for the compressed configuration can be small enough to allow its storage in memory, thus consistently accelerating the first guess estimation. More details are reported in Chapter 3.

2.3.2 Developments on the accuracy of MART and advanced reconstruction methods

The mathematical recipe of MART dates back to the '70s for the general problem of tomographic reconstruction from few (and sometimes incomplete) projections. Interestingly enough, one can try to improve the accuracy of the technique by adding, somehow, information about the objects to be reconstructed, i.e. tracer particles (more exactly, in most of the applications the diffraction is dominant with respect to the geometric magnification in determining the particle image diameter; as a consequence, the reconstructed particles are *hyper-particles*, i.e. 3D blobs with diameter equivalent to the particle image diameter divided by the magnification).

In this thesis a method is proposed to tackle with the effects of erosion of MART on the shape of these 3D blobs and of larger uncertainty in determining the position of the particles along the depth direction. The method consists in applying artificial diffusion in between the MART iterations, i.e. the reconstructed distributions are slightly low-pass filtered after each MART iteration. The low-pass filter is expected to smear out some irregularities of the true particles, and to be quite effective in damping ghost particles, as their size is usually smaller, as observed by de Silva et al (2013). The filtering kernel should have size comparable to the particle image diameter. In Chapter 4 the effect of a Gaussian filtering window is investigated. This choice is extremely natural, as it is directly connected to the particle image shape. Furthermore, the effect of applying anisotropic filtering to reduce the uncertainty along the depth direction is investigated. This technique is named Spatial Filtering Improved Tomography (SFIT), and it is the object of Chapter 4.

More recently, de Silva et al (2013) proposed an *a posteriori* method following similar observations. The method is named Simulacrum Matching-based Reconstruction Enhancement (SMRE), as it is based on a template-matching scheme of true and ghost particles. The sum of the squared difference between the reconstructed object and a kernel with variable size and uniform intensity (the mean background intensity) is locally computed. The peaks above a threshold are considered true particles, while the intensity in the other regions is set to zero. It is not a surprise that, similarly to SFIT, the optimal size of the kernel is similar to (actually slightly larger than) the particle diameter.

Another path to increase the accuracy of the reconstructions stands in the computation of the weighting elements. In general the OTF to compute the weighting elements is chosen for simplicity to be uniform within the volume and for all the cameras. This simplification introduces a source of error in presence of magnification gradients, optical aberrations, distortions, blurring due to the limited depth of field, etc. Schanz et al (2013) proposed a locally adaptive OTF to compensate for these effects. The results show a promising improvement of the quality of the reconstruction for particle image density up to 0.07ppp.

An alternative approach to tackle the mathematical problem of tomographic reconstruction from few projections is explored by Petra et al (2007, 2009). The technique is named Successive Linearization Algorithm (SLA), based on the criterion of sparsity maximization of the distributions to be reconstructed. The results are promising, but an investigation of the experimental performances of the method is still lacking.

2.3.3 Multi-exposure methods

The great advantage of the CT scans with respect to the case of Tomo-PIV consists in the fact that the subject is steady and several views can be taken over time by moving the imager. In PIV the requirement of imaging the particles in the short pulse time limits the number of views to the total number of cameras. However, this handicap can be partially reduced by considering that in Tomo-PIV at least two exposures with (nearly) the same cloud of particles are required. Since the particles will move, the relative orientation between the tracers and the camera set is slightly different, thus enabling in some cases the extraction of more information on the particles distribution.

Novara et al (2010) introduced a Motion Tracking Enhanced (MTE) algorithm to build successive first guess distributions for the MART corrective procedure. The method is based on an iterative procedure in which first guess distributions for the two (or more) exposures, obtained by straightforward application of MART, are subsequently refined by the computation of pseudo-simultaneous objects from the recording views at different times. Indeed, the first guess objects are cross-correlated in order to have a rough estimate of the velocity field; then, considering for simplicity the case of two exposures, the intensity (say, the particles) of the first exposure is moved forward along the velocity field to build a pseudo-second exposure, and the same operation is performed on the second exposure. The pseudo simultaneous objects are then summed to the relative first guess exposures. In this scenario, the coherent component of the signal is strengthened, while the incoherent component (hopefully, the ghost particles) is weakened.

The principle is illustrated in Fig. 2.6 in the optimal case of velocity shear in the direction orthogonal to the viewing one. The results by Novara et al (2010) and by




Fig. 2.6 Left: two particles imaged by two cameras (red=first exposure; blue=second exposure) and relative ghost particles (light colors). The particles move according to the velocity field indicated with gray lines. Right: the true particles superimpose when summing the first exposure with the pseudo-first exposure, while the ghost particles do not (Novara et al 2010).

Novara & Scarano (2012a) show a significant improvement of the correlation signal and of the quality of the velocity measurement at particle image density up to $0.2ppp$, well beyond the values adopted in common practice. On the other hand, the potential quality improvement is strongly dependent on the velocity field itself, and in particular on the velocity gradients. Indeed, in case of small velocity gradients, the ghost particles will tend to add to the coherent component of the signal, as outlined in Sec. 2.2.2. In the particular case of no displacement, MTE-MART does not give any benefit at all.

2.4 Three-dimensional motion analysis

As mentioned in Sec. 1.1.4, the image deformation methods are considered the standard algorithms in planar PIV, as they largely compensate for the in-plane motion and velocity gradients. The potential of the deformation methods is even larger in the 3D scenario, since the ability to cope with the full velocity gradients allows the elimination of the loss of pairs due to the out-of-plane motion that affects the planar PIV. Since the introduction of the multi-grid multi-step window deformation algorithms by Scarano & Riethmuller (2000) many authors proposed improvements and updates. The skeleton of the algorithm is summarized in the following:

1. The predictor displacement field is calculated on a rather coarse grid with interrogation spot size chosen in agreement with the one-quarter rule (Keane & Adrian 1992). The operation consists in computing the normalized 3D cross-correlation (obtained by trivial extension to 3D of the (1.5));

2. The predictor displacement field is interpolated on each voxel in order to deform the volumes accordingly;
3. A grid refinement may be executed in the evaluation of the displacement field. The interrogation spot size can be progressively reduced, since the volume deformation compensates the loss of pairs effect due to the displacement and the velocity gradients within the interrogation spot;
4. A corrector displacement field is evaluated on the deformed volumes;
5. The true displacement is obtained by summing the corrector and a weighted average of the dense predictor over a prescribed region.

The steps 2-5 are repeated till convergence. For the current status of the hardware, the 2D motion analysis is not problematic, as large amount of samples captured with cameras with large format sensor can be easily processed with off-the-shelf computers in relatively short time. Adding the third dimension complicates the implementation, as the amount of data increases by two orders of magnitude (roughly the depth of the volume in voxels) and the computational cost increases both because the algorithm is based on 3D cross-correlation instead of 2D and the larger number of interrogation spots (one to two orders of magnitude more than the 2D case). Recently, different approaches for the computational cost reduction of 3D PIV have been proposed: the volume segmentation technique by Ziskin & Adrian (2011) is based on computing a fast reconstruction by a MLOS-based technique, and adjacent planes are summed to obtain a sufficiently high source density to perform the standard 2D correlation-based PIV algorithms; Bilsky et al (2011) propose multiple 2D projections of 2D cross-correlation maps of the projected interrogation volumes. However, a systematic study on the accuracy of these methods has not been conducted yet.

In this thesis the proposed approach is quite different, as it promises a computational cost reduction of up to two orders of magnitude without affecting the accuracy. The method is a combination of several solutions:

- Multi-resolution interrogation in the initial step of the process to get a quick estimate of the predictor;
- In the corrector estimation (step 4), since the predictor is already known, one could use a narrower search area of the peak (as already observed by Rohàly et al 2002), e.g. limiting the calculation of the (1.5) only to the neighbourhood of the peak itself;
- Sparse direct cross-correlation can consistently reduce the computational cost since the percentage of non-zero voxels is usually rather low;
- The PIV interrogation is usually conducted with overlapping interrogation windows (at least 50% to satisfy the Nyquist criterion). In 3D the margin of improvement by reduction of the number of redundant calculation in these region is very appealing. In Chapter 5 some solutions to optimize the algorithm are tested.

Chapter 3 - Multi-Resolution MART

In Tomographic PIV the construction of the first guess distributions for the iterative process of the reconstruction can considerably influence the rate of convergence of MART. As outlined in Sec. 2.3.1, to date the most effective techniques to speed-up the MART-based techniques rely on the quick identification of voxels with zero intensity. The Multiplicative Line Of Sight (MLOS) technique by Atkinson & Soria (2009) is widely used among the community for its intrinsic simplicity and effectiveness. However, one should not forget that, since the energy captured by the ghost particles is subtracted from the actual particles, a reliable first guess estimation process should not involve an increase of the number of the former ones, or pumping their intensity. Atkinson & Soria (2009) showed that the number of ghost particles is approximately the same for MART and straightforward application of MLOS, as it is determined by the experimental parameters affecting the ambiguities of the reconstruction problem (see Sec. 2.2.1), i.e. mainly the seeding density, the diameter of the particles and the depth of the illuminated volume. Differently from MLOS, MART is capable to provide an attenuation of the intensity of the ghost particles. For this reason, a first guess estimation procedure mainly based on the MART correction process seems to be more suitable to achieve the goal of reducing the computational effort of the reconstruction, without affecting the accuracy of the results. In this chapter an algorithm for a quick MART-based estimation of the first guess distributions is proposed. The algorithm is based on a multi-resolution approach, i.e. the distributions are initially reconstructed at a lower resolution (but still maintaining a pixel/voxel ratio approximately equal to 1 by pixels binning). Similarly to MLOS and MFG, the algorithm is still founded on the possibility of identifying in advance locations with zero intensity, but at the same time it provides attenuated ghost particles as it is based on the MART corrective procedure.

In Sec. 3.1 the Multi-Resolution algorithm is described and discussed in detail. A parametric study of the performances of the algorithm is proposed in Sec 3.2, mainly addressing the quality of the reconstruction and the reduction of both the computational cost and the memory storage requirements. Eventually, in Sec 3.3 a Tomo-PIV simulated experiment is presented.

The description of the technique and the results of the validations have been published in Discetti & Astarita (2012a).

3.1 Multi-Resolution approach to Tomo-PIV

The algorithm of the Multi-Resolution (MR) MART is outlined in the following. In the first step of the process, a coarser grid in the volume to be reconstructed is employed, lowering, as a matter of fact, the number of unknowns of the system (2.1). A factor of compression equal to 2 has been chosen as an upper limit in the present thesis, because of the unavoidable increasing noise due to discretization artefacts if a higher factor is adopted. As a matter of fact, the application of MR is particularly critical in case of smaller particles; however, the reconstruction quality can be improved by slightly smoothing the original images, e.g. with a simple Top Hat moving average filter with a 3x3 kernel. At the same time, the images are compressed with a binning-like procedure, retaining a resolution ratio between voxels and pixels possibly close to unity for two main reasons: voxels and pixels of equal size guarantee uniformity in the capabilities of sampling the shape of the particles both in the volume and in the camera planes; the number of pixels interested by the projection of the intensity of each voxel depends on the resolution ratio. Consequently, a unity value of the latter is the best compromise between memory storage reduction, decrease of the computational cost due to the projection process and the calculation of the weighting elements, and minimum decrease of the available number of equations (i.e. projections) of the system (2.1).

As shown in the flow chart in Fig. 3.1, the multi-resolution reconstruction scheme is structured in the following steps:

Step a: compression of the camera images by binning pixels in 2 x 2 kernels. The corresponding intensities values of each pixel on the compressed projections are determined through arithmetical average.

Step b: one or more MART iterations are performed onto the compressed configuration. The number of operations for each of the MART iterations is reduced by the cube of the factor of compression with respect to the original configuration.

Step c: the first guess distribution is interpolated onto the final grid, recovering the desired spatial resolution. A nearest neighbour approach is not advisable as it can cause deletion of the particle tails; linear interpolation is adopted, since it is the best compromise between computational cost and minimum expansion of the located non-zero positions.

Step d: supplementary MART iterations starting from the provided first guess are performed on the grid with the final resolution. The number of operations is proportional to $n_p \times N_{nz-voz} \times N_{cam}$ for each iteration, where n_p is the number of pixels influenced by the projection of a voxel, N_{nz-voz} is the total number of voxels with non zero intensity and N_{cam} is the number of cameras. The larger is N_{nz-voz} , the higher is the speed-up of the process.

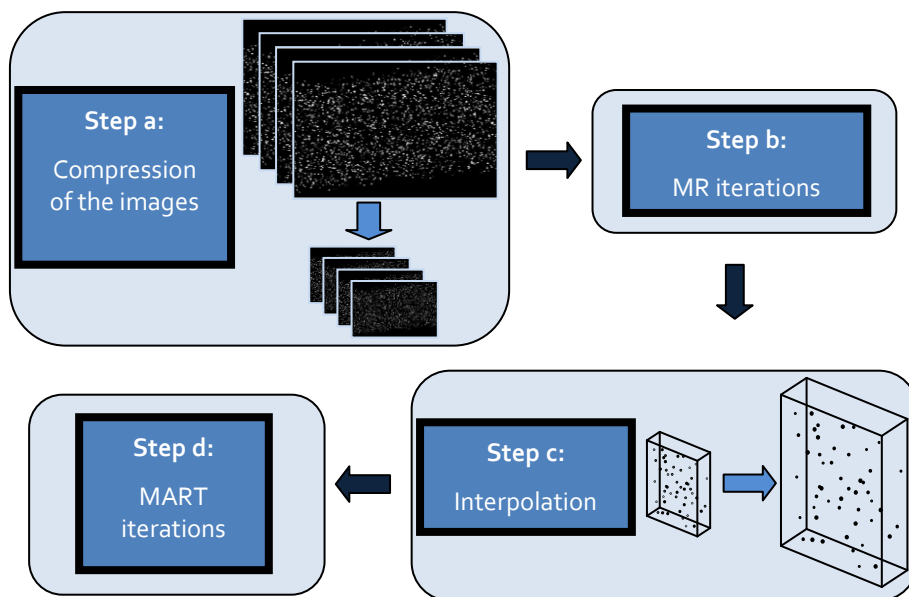


Fig. 3.1 Flow chart of the Multi-Resolution reconstruction process.

Some additional considerations regarding the weighting elements have to be drawn. Assumed that the complete storage of the weighting matrix, also taking into account its sparseness, is nearly impossible for the commonly adopted experimental configurations, one should take into consideration that it could be achievable for the compressed configuration. In this case, the MART iterations of the step *b* are performed considerably faster, not only for the reduction of dimension of the reconstruction problem, but mainly because weighting elements are computed only *una tantum* for an entire set of images.

Furthermore, in this last case it could be convenient to further accelerate the step *b* by initializing the distribution on the compressed configuration with MLOS. The subsequent MART iterations will provide the desired attenuation of the ghost particles intensity in the first guess for the step *d*. It is worth to underline that the identification of voxels with zero intensity on the compressed configuration can further reduce the weighting elements to be calculated in the final MART iterations, in most cases enabling the possibility to store them. Obviously, the weighting matrix has to be computed for every distribution of the full set, as long as the non zero voxels distribution is time-dependent; however, if the number of iterations of the step *b* is more than one, it could be convenient as well, since the number of weighting elements have to be computed only once for each analysis conducted on the compressed configuration. The same observation can be applied to the step *d*.

3.2 Parametric study of MR-MART

The performance assessment of the proposed technique consists of the reconstruction of a series of 3D computer generated particle distributions with a custom-made software. Spherical Gaussian particles with 3 voxels diameter are distributed into a $5 \times 5 \times 10 \text{ mm}^3$ volume, with a resolution of 20 vox/mm , resulting in a $100 \times 100 \times 200$ voxels control volume. The distributions are imaged by cameras with 200×200 pixels resolution, disposed in a linear arrangement in the horizontal mid-plane of the volume to be reconstructed, with different viewing directions ϑ . Three different cameras setups are tested, respectively 3 cameras ($\vartheta = -30^\circ, 0^\circ, 30^\circ$), 4 cameras ($\vartheta = -30^\circ, -10^\circ, 10^\circ, 30^\circ$), and 5 cameras ($\vartheta = -40^\circ, -20^\circ, 0^\circ, 20^\circ, 40^\circ$). The resolution ratio between pixels and voxels is set approximately equal to one throughout the volume; the projected volume is completely enclosed into all the image planes.

Such a small volume to be reconstructed is chosen in order to enable the storage of the weighting matrix even in the case of the adoption of the standard MART approach. In this way, for each type of process, only the time needed to perform the iterative procedure is accounted. Of course this approach strongly underestimates the capabilities of the MR approach when the amount of memory is sufficient to store the weighting matrix in the compressed configuration and/or in the original configuration after the estimation of the first guess.

The corresponding point of each voxel (x_j, y_j, z_j) on the camera planes (X_i, Y_i) is determined by means of the polynomial multi-dimensional mapping approach (Soloff et al. 1997) outlined in Sec. 1.2.1; no calibration error is introduced, since it affects in the same measure both the standard method and the multi-resolution approach. Thresholding is applied in order to eliminate the voxels with weaker intensity from the updating process and from the calculation of the corresponding weighting elements; since in the present simulation the maximum intensity is the same for all the particles, a unique threshold, equal to 0.1% of the peak, is applied to all the distributions during the updating iterative process, showing to be able to eliminate reconstruction artefacts whereas retaining a negligible cutting of the shape of true particles. The weighting elements of the final configuration are stored with a 3D extension of the well-assessed Yale format (Eisenstat et al 1982) for sparse 2D array storage. The 2D Yale format is founded on storing sparse matrices using three vectors: the first one contains the non-zero elements; the second one contains the indexes of column; the third one points to the first element of each row.

The source density is varied between 0.05 and 0.5; the equivalent image density N_{ppp} can be easily estimated by dividing the source density by the mean area of the particle images using the (1.14). The obtained results cannot be directly compared with those of the commonly adopted 2D simulations, unless a proper

evaluation of the effective source density is conducted. E.g., an image density of $0.05ppp$, which is widely accepted as a threshold value to achieve a reliable reconstruction, is equivalent to a source density of 0.15 in the 2D simulations of Elsinga et al (2006a) and Atkinson & Soria (2009).

The quality of the reconstruction is quantified in terms of the correlation between the reconstructed intensity fields, as in (2.4).

The relaxation parameter of the MART process (2.3) is set to 1; a uniform first guess intensity equal to 1 is imposed on the compressed configuration for the simple MR-MART, whereas, as stated in Sec. 3.1, an estimation of the initial distribution is deduced by MLOS for the modified version, referred as MLOS-MR-MART from now on.

The performances of MR-MART and MLOS-MR-MART are assessed in comparison to those of the commonly used method based on 5 MART iterations; the latter will be synthetically indicated as MART from this point on. The MR-MART approach is tested with variable number of iterations on the compressed configuration, while retaining the same total number of iterative steps equal to 5.

3.2.1 Quality of the reconstruction

The performances in terms of accuracy of the reconstruction of MR-MART in comparison to MART are quoted as a function of the seeding density in Fig. 3.2a for the 4 cameras setup. The effect of the number of iterations on the compressed configuration (from this point on referred as MR iterations) is also investigated; the execution of the algorithm with 4 MR iterations is not considered, since 1 single supplementary MART iteration on the final configuration is not able to recover a reliable estimation of the distributions.

A slight accuracy improvement can be observed throughout the whole range of tested seeding densities. One could reasonably expect a more relevant increase of the quality of the reconstruction, as long as the ambiguities in determining the depth position of the particles are reduced due to the lower number of voxels in the depth direction; it does not effectively happen, mainly because of the presence, in the first guess, of artefacts due to undersampling of the particles in the compressed configuration. The results of the different algorithms are evidently quite independent of the number of MR iterations, except for N_s exceeding 0.25; beyond this value of seeding density, the algorithm with 3 MR iterations performs slightly worse than those with 1 or 2 MR iterations, but still better than MART. The performances of MLOS-MR-MART in terms of accuracy are practically equal to that of the MR-MART algorithm, and for this reason they are not reported herein.

An overview of the trend of the quality of the reconstruction by varying the number of cameras is illustrated in Fig. 3.2b. The results for the processes with 3 MR iterations with uniform value or MLOS estimation of the first guess on the

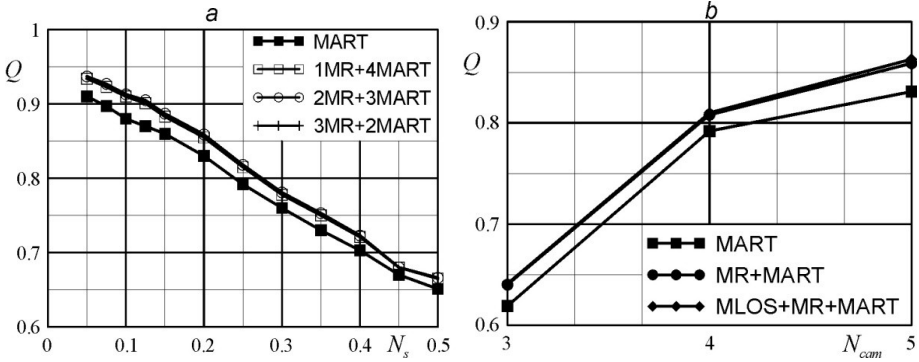


Fig. 3.2 Reconstruction quality as a function of: (a) the source density for a 4 cameras setup; (b) the number of cameras for $N_s = 0.25$ and 3 MR iterations.

compressed configuration are shown, with a fixed source density of 0.25, which could be considered a reasonable value for the adopted optical setup, since it results in a concentration of approximately 1.4 particles/ mm^3 , i.e. about 11 particles in a cubic interrogation volume of $40 \times 40 \times 40$ voxels, corresponding to $2 \times 2 \times 2 \text{ mm}^3$. MR-MART and MLOS-MR-MART provide higher quality of the reconstruction with respect to MART for all the tested cameras setups.

3.2.2 Acceleration of the process

As already stated before, only the time needed to perform the MART iterations will be considered. This implies a noticeable underestimation of the decrease of processing time (synthetically referred in the following as speed-up) of the algorithm in several conditions, as extensively explained in Sec. 3.1. Nevertheless, the speed of the reconstruction is remarkably increased, as illustrated in Fig. 3.3 for a 4 cameras setup; furthermore, the speed-up is guaranteed in the whole range of tested seeding densities, whose upper limit is well-beyond the capabilities of the ordinary reconstruction techniques. The speed-up with respect to MART is sensitive to the number of MR iterations and to the seeding density, both for MR-MART and MLOS-MR-MART.

For MR-MART (Fig 3.3a) at very low seeding densities, a lower number of MR iterations is preferred, since the relatively high percentage of voxels with zero intensity in the final configuration allows the fast performance of the subsequent final MART iterations. Conversely, increasing the seeding density (and, contemporaneously, the number of ambiguities generated in the reconstruction of the compressed volume), a higher number of MR iterations is preferable; a little decrease of the reconstruction quality is bearable, since MART performs still worse,




Fig. 3.3 Effect of seeding density and number of MR iterations on the speed-up for a 4 cameras setup: (a) MR-MART; (b) MLOS-MR-MART.

as illustrated in Sec. 3.2.1. In the range of source density covering 0.075 and 0.125 the processing time for MR-MART is rather independent of the number of MR iterations.

At the source density of 0.25, MR-MART provides a speed-up of about 2.4, 2.9 and 3.6 times with respect to MART for 1, 2 and 3 MR iterations respectively. In the same conditions MLOS-MR-MART (Fig 3.3b) guarantees slightly better performance, with a gain of a factor of 2.6, 3.3 and 4.4 for 1, 2 and 3 MR iterations respectively. MLOS-MR-MART consistently outperforms MR-MART at the lowest tested seeding density, while beyond the source density of 0.25 the differences between the two algorithms are relatively small; hence, if the weighting matrix in the compressed configuration can be stored, MR-MART is certainly preferable.

The speed-up could seem quite small if compared with that of 40 MLOS-SMART iterations, quoted in 8 times faster with respect to MART, without taking into account the computation of the weighting elements, in a 3D experimental configuration (Atkinson & Soria 2009). Anyway, those results are referred to a relatively low level of source density, since the experiments have been performed with a source density of about 0.13. For this level of image density, MLOS-MR-MART can provide an increase of the speed of reconstruction of 6.7, 8.1 and 10 times for 1, 2 and 3 MR iterations respectively. In the same configuration, both MR-MART and MLOS-MR-MART guarantee a factor of quality of the reconstruction of about 0.92 against 0.89 of the standard MART algorithm.

One should note that these results are not particle-size independent, since in case of smaller particles a smoothing of the images is required to recover the same accuracy performances of MART, as already stated in Sec. 3.1, sacrificing part of the sparseness of the original images.

3.2.3 Memory storage reduction

The reduction of the total number of weighting elements is a key point in the decrease of memory storage, since the size of the camera images and of the volume to be reconstructed is consistently smaller. As a matter of fact, for a volume discretization of $1000 \times 1000 \times 200$ voxels, with a four 1Mpixels cameras setup, the memory needed to store the weighting matrix, taking into account its sparseness, is about 58GB, while each of the volumes would need 800MB, and each of the camera images 4MB if all data are stored in single floating point precision. Incidentally, the number of elements of the weighting matrix can be easily estimated by considering that each voxel influences n_p pixels (typically 3×3 pixels) on each camera. It follows immediately that the number of weighting elements is $n_p \times N_{vox} \times N_{cam}$; for the aforementioned situation it is equal to $7.2 \cdot 10^9$. Considering that in a sparse format one has to store arrays with the indexes of rows, columns and z planes of the voxels, this leads to $28.8 \cdot 10^9$ elements, requiring about 115GB for the storage. On the other hand, using a 3D sparse Yale format, one can order the voxels (for example lexicographic order, and planes sorted along the z direction) and store only the columns indexes and a vector pointing to the first element of each row. This approach is much more efficient for memory access than simply storing the index of the voxel in the chosen order. Furthermore the memory required to store this second vector is usually very small, as it is composed by a number of elements equal to the product of the length and the depth of the volume. With this approach, the memory required is approximately 58GB, as stated before.

Consequently, the weighting elements represent more than 98% of the total data. Anyhow, when the number of weighting elements is reduced by a factor of 10 (for example using MR-MART or MLOS initialization), they still constitute about 90% of the total data amount, and if the memory storage reduction is equal to a factor of 20, their relative size is still approximately 80%. In the following only the reduction of the number of weighting elements will be considered as a parameter. Nevertheless, for the aforementioned reasons, the reported results are not so far from representing the real total memory storage reduction.

In this section only the algorithms with 3 MR iterations are considered, being the results quite independent of the number of iterations on the compressed configuration. As a matter of fact, a single iteration is sufficient to identify all the zero locations, while 1 or 2 additive iterations are able to reduce the intensity of the ghost particles, carrying the weakest of them below the threshold value. Anyhow, the further reduction in memory storage is not significant.

The results relative to a 4 cameras setup and simple MR-MART (Fig. 3.4) show that for the lowest levels of seeding density the memory storage reduction tends to the cube of the factor of compression, i.e. 8, as it was foreseeable. The storage of the weighting elements on the compressed configuration is really advantageous




Fig. 3.4 Effect of seeding density on the memory storage reduction for a 4 cameras setup for MR-MART and MLOS-MR-MART, compared with MLOS.

because it can be used for a full set of images; however, it is evident that it could be possible only if a relatively consistent amount of memory is available, and at the same time a low level of seeding is guaranteed on the whole sequence. In practice, this condition is not trivial to be achieved, and most often is not even desirable due to spatial resolution requirements.

In Fig. 3.4 the performances of MLOS-MR-MART are also quoted in comparison to the straightforward application of MLOS. For this setup, MLOS-MR-MART appears to outperform not only MR-MART but also the direct application of MLOS, for the higher tested seeding densities, thanks to the smoothing of the ghost particles due to MART corrective procedure in the MR iterations, and their consequent removal due to dropping below the threshold. The same threshold has been applied for both methods, meaning a stronger cutting for the MLOS-MR-MART method. Luckily, a more stringent threshold does not affect too much the result, because the effect of cutting the particle tails is minimum, thanks to the interpolation in the step c (see Sec. 3.1) of the process.

3.3 Validation via synthetic experiment

In this section a 3D synthetic experiment is presented. The flow motion induced by a vortex ring is investigated, following exactly Elsinga et al (2006a). A volume of $30 \times 30 \times 7 \text{ mm}^3$ is reconstructed with a resolution equal to 20 vox/mm , resulting in $600 \times 600 \times 140$ voxels (a resolution ratio between pixels and voxels of about 0.9 has been adopted). Four cameras are placed at a finite distance, with

magnification approximately equal to 0.12 throughout the volume; the mean particle diameter is set to 2.8 pixels, with a pixel pitch of $6.7\mu\text{m}$. The cameras are displaced in linear arrangement on a common horizontal plane, as in the performance assessment of Sec. 3.2, i.e. with $\vartheta = -30^\circ, -10^\circ, 10^\circ, 30^\circ$.

Both the reconstruction and the PIV volume processing have been carried out on a laptop machine using a single core of a 2.2 GHz T6660 processor; the available amount of RAM memory is 1.5GB. Only the algorithm with 3 MR iterations is considered, since the performance assessment of Sec. 3.2 has shown that it guarantees the best compromise between accuracy and speed of the reconstruction. Two different levels of seeding concentration are tested, i.e. 0.55 and 2.2 particles/ mm^3 (resulting in a source density of 0.075 and 0.3, respectively). The displacement field is calculated with a custom-made software, based on the 3D extension of the algorithm illustrated in Sec. 2.4. Iterative Discrete Volumes Offset (IDVO) or Volume Deformation Method (VDM) are adopted to enhance the correlation peak. In this section only linear interpolation is used for the deformation of the volumes, since the processing time is a crucial point of the analysis. Furthermore, for the sake of clarity, the solutions to accelerate the PIV analysis presented in Chapter 5 are not completely exploited here, except for the voxel binning in the predictor estimation. In this Multi-Resolution PIV interrogation the distributions are initially binned by a factor of four, in order to obtain a fast estimation of the predictor displacement field. In the subsequent iterations the factor of compression is relaxed (for example in this section it is halved in each step of the process, in order to better fit the performances of the FFTW algorithm (Frigo & Johnson 2005)). However, the reader should be aware that the processing time of the motion analysis can be easily reduced by a further order of magnitude by coupling the Multi-Resolution interrogation with efficient algorithms enjoying sparse cross-correlations and reduction of the number of redundant operations performed in case of overlapping interrogation spots. The topic will be discussed in more detail in Chapter 5.

3.3.1 Low seeding density simulations

First of all, one should inquire whether the memory storage of the weighting matrix is possible or not for the employed configuration. The memory demand can be estimated as in Sec. 3.2.3, resulting in about 15GB for the original configuration in single floating point precision. Since for the chosen seeding density and cameras setup the memory storage reduction can be quoted as about 6 times for MR-MART (see Fig. 3.4) the storage of the weighting elements is not possible for the adopted computer. Even in the compressed configuration, in which the memory storage is reduced by a factor equal to 8, their storage could not be achieved. On the other hand, the memory storage reduction of MLOS-MR-MART is more than 30 times.

Consequently, three approaches to the reconstruction are exploited: the standard method based on 5 MART iterations; MR-MART, with 3 MR iterations, and 2 final MART iterations, as presented in Sec. 3.2; MLOS-MR-MART, with initial estimation of the distribution by MLOS, and the same structure of MR-MART. For the first two methods, the weighting elements are calculated every time they are needed; MLOS-MR-MART, conversely, enables the possibility to calculate the weighting elements only once both for the compressed and the final configuration.

Confirming the results of Sec. 3.2, the quality factor of the reconstruction is 0.91 and 0.94 for MART and the multi-resolution approaches, respectively. The reconstruction of each volume with MART and MR-MART is performed in about 17 and 4 minutes, respectively, resulting in a speed-up of more than 4 times. MLOS-MR-MART, including the time to calculate the weighting elements on both the configurations, is able to perform the reconstruction in approximately 1 minute, i.e. 17 times faster than MART. The speed up is slightly lower than the one evaluated for the same source density in the performance assessment of Sec. 3.2 since having the same number of particles in a less deep volume implies the presence of a higher percentage of non-zero voxels to be updated. Anyhow, MLOS-MR-MART can stand approximately the foreseen speed-up thanks to the pre-calculation of the weighting elements.

The size of the interrogation volume (IV) for PIV processing is 64^3 voxels, i.e. cubes of $3.2 \times 3.2 \times 3.2 \text{ mm}^3$, containing on average 18 particles each one, with 75% overlap. The traditional process used in this section consists of the estimation of the predictor, followed by a single iteration both in the case of IDVO and VDM (cross-correlation is always performed with the aid of FFT). The multi-resolution process requires an additive iteration to perform the analysis on the original distributions. The results of the processing of the reconstructed distributions by MART and MR-MART are not reported herein, since there is no appreciable difference with those of MLOS-MR-MART.

The processing time of the traditional PIV algorithm and of the multi-resolution version are 5.3 and 2.7 minutes respectively, if IDVO is used; on the other hand the processing with VDM requires 7.5 minutes for the former, and 4.5 minutes for the latter. The results of the PIV interrogation in terms of accuracy for both the algorithms are also quoted in Tab. 3.1, in terms of standard deviation of the measurement errors (obviously only the u and the w component are considered, being the flow field axisymmetric). The mean error in the measurement of the w component μ_w of the velocity field is also reported, since a non negligible bias effect is present due to both the limited spatial resolution and the asymmetric distribution of w , as already observed by Elsinga et al (2006a).

In Fig. 3.5 the iso-vorticity surface, corresponding to a vorticity of $0.13 \text{ voxels/voxel}$, obtained by PIV multi-resolution interrogation of the distribution reconstructed by MLOS-MR-MART, is presented; 5 planes of velocity vectors (the

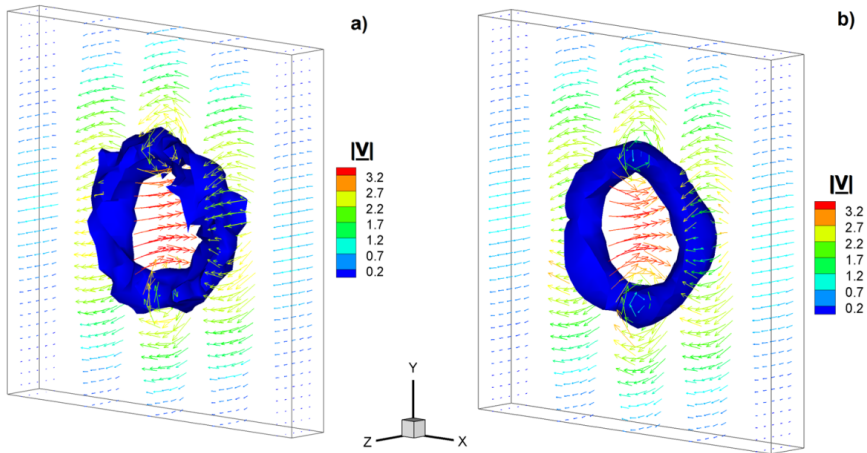


Fig. 3.5 Iso-vorticity surface (0.13voxels/voxel) for PIV multi-resolution processing with IDVO (a) and VDM (b) on the volume reconstructions obtained by MLOS-MR-MART in the case of low source density; 5 planes of velocity vectors (colour-coding indicates the magnitude in voxels of the velocity vector).

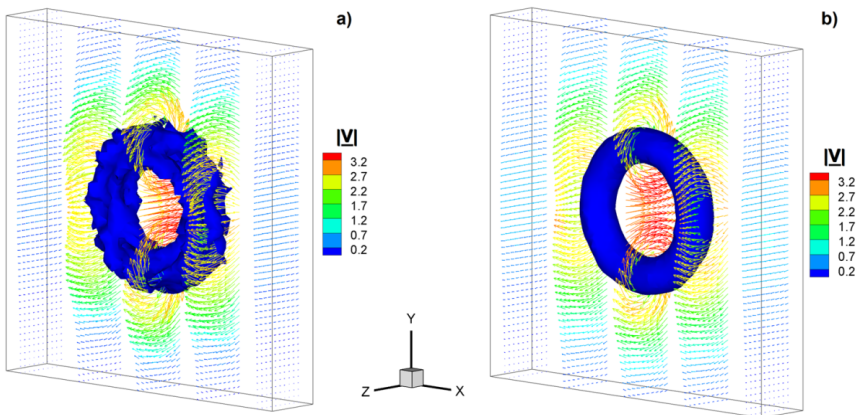


Fig. 3.6 Iso-vorticity surface (0.13voxels/voxel) for PIV multi-resolution processing with IDVO (a) and VDM (b) on the volume reconstructions obtained by MLOS-MR-MART in the case of high source density; 5 planes of velocity vectors (color-coding indicates the magnitude in voxels of the velocity vector).

colour-coding indicates the intensity of the velocity vector) are illustrated. Even if the spatial resolution is limited by the low seeding density, the main features of the field are captured, as highlighted by the relatively small standard deviations of the measurement error on the u and the w components of the displacement field; on the other hand, the estimation of the vorticity field is strongly affected by the amplification of the random errors, coupled with the low spatial resolution of the measurement.

| Reconstruction process | PIV process | Processing time [min] | $\sigma_u[\text{vox}]$ | $\mu_w[\text{vox}]$ | $\sigma_w[\text{vox}]$ |
|------------------------|-------------|-----------------------|------------------------|---------------------|------------------------|
| MART | IDVO | 39.3 | 0.22 | -0.07 | 0.30 |
| MART | MR-IDVO | 36.8 | 0.23 | -0.08 | 0.30 |
| MLOS-MR-MART | IDVO | 7.3 | 0.22 | -0.05 | 0.28 |
| MLOS-MR-MART | MR-IDVO | 4.7 | 0.23 | -0.06 | 0.28 |
| MART | VDM | 41.5 | 0.17 | -0.05 | 0.27 |
| MART | MR-VDM | 38.7 | 0.16 | -0.03 | 0.27 |
| MLOS-MR-MART | VDM | 9.4 | 0.17 | -0.03 | 0.27 |
| MLOS-MR-MART | MR-VDM | 6.6 | 0.16 | -0.03 | 0.27 |

Table 3.1 Performances of different reconstruction and interrogation algorithms, $N_s = 0.075$.

| Reconstruction process | PIV process | Processing time [min] | $\sigma_u[\text{vox}]$ | $\mu_w[\text{vox}]$ | $\sigma_w[\text{vox}]$ |
|------------------------|-------------|-----------------------|------------------------|---------------------|------------------------|
| MART | IDVO | 57.8 | 0.13 | -0.08 | 0.21 |
| MART | MR-IDVO | 51.7 | 0.13 | -0.11 | 0.21 |
| MLOS-MR-MART | IDVO | 23.8 | 0.12 | -0.08 | 0.20 |
| MLOS-MR-MART | MR-IDVO | 17.7 | 0.13 | -0.10 | 0.20 |
| MART | VDM | 59.7 | 0.09 | -0.04 | 0.14 |
| MART | MR-VDM | 52.5 | 0.08 | -0.01 | 0.13 |
| MLOS-MR-MART | VDM | 25.7 | 0.09 | -0.03 | 0.14 |
| MLOS-MR-MART | MR-VDM | 18.5 | 0.08 | -0.01 | 0.13 |

Table 3.2 Performances of different reconstruction and interrogation algorithms, $N_s = 0.3$.

3.3.2 High seeding density simulations

Increasing the seeding density, an improvement of spatial resolution is obtained at the expense of the computational cost, since the storage of the weighting elements is not achievable for all the tested reconstruction approaches on the employed computer because of the higher percentage of non zero voxels, due to both true particles and ambiguities in the reconstruction. In this section the weighting elements are calculated every time they are needed in the reconstruction process. The standard MART algorithm achieves a reconstruction quality of approximately 0.80, while for both MR-MART and MLOS-MR-MART the obtained quality is 0.82, again confirming the slight accuracy improvement. The reconstruction of each volume is performed in about 22, 8 and 5 minutes for MART, MR-MART and MLOS-MR-MART respectively. The increase of processing time with




Fig. 3.7 Scatter plot of the measurement error relative to u and w component of the velocity field, for IDVO (a) and VDM (b), PIV multi-resolution processing performed on the volume reconstructions obtained by MLOS-MR-MART.

respect to the simulations of Sec. 3.2.2 for the first two methods is due to a higher number of non-zero voxels; on top of this, the third method is also affected by the impossibility to store the weighting elements, unlike the simulation with low seeding density. Both for IDVO and VDM the process consists of 3 iterative steps: initially an IV of 64^3 voxels is employed, with 75% overlap; in the second step the IV is reduced to 42^3 voxels, while retaining the previous grid distance of 16 pixels between the vectors; finally, an iteration with interrogation volumes of 42^3 voxels (i.e. cubes of $2.1 \times 2.1 \times 2.1 \text{ mm}^3$, containing on average 18 particles, as in the simulations of Sec. 3.3.1) with 75% overlap is performed. In this paragraph, the multi-resolution PIV process is carried out with the same total number of steps of the traditional algorithm.

The processing time for traditional PIV is about 13 and 15 minutes with IDVO and VDM respectively; the multi-resolution approach, on the other hand, is performed approximately 1.9 times faster for both the types of process. Even in the case of high seeding density, the full Tomo-PIV process is executed more than 3 times faster than the standard algorithm if the reconstruction is performed with MLOS-MR-MART and PIV interrogation is carried out with the multi-resolution algorithm, both in the case of employment of IDVO and VDM. An overview of the results in terms of processing time and accuracy performances is reported in Tab. 3.2.

In Fig. 3.6 the results obtained by the PIV multi-resolution interrogation of the distribution reconstructed by MLOS-MR-MART are presented. The processing with VDM better captures the expected torus, and the velocity field appears much less noisy, as confirmed also by the data of Tab. 3.2 and by the scatter plot representation of the measurement error on the u and w components of the displacement field, shown in Fig. 3.7.

Chapter 4 - Spatial Filtering Improved Tomographic PIV

Tomographic reconstruction accuracy is of fundamental importance to obtain reliable three-dimensional three-components velocity field measurements. The pursuit of accurate reconstructions for larger particle concentrations is central in many recent investigations. To date, the most promising approaches are those based on multiple exposures, outlined in Sec 2.3.3. Furthermore, the advanced methods by Petra et al (2007, 2009), based on sparsity maximization, could potentially outperform the standard algebraic techniques. In both cases, the accuracy improvement is not priceless in terms of computational cost.

In this chapter an easily implementable modified version of MART is proposed, allowing a remarkable improvement of the accuracy of the tomographic reconstruction without significantly increasing the computational cost. The idea is based on the observation that ghost particles are often irregularly shaped, smaller and weaker than the true particles in the reconstructed distributions. Furthermore, the discretization artefacts may lead to an erosive effect of the shape of the true particles. In this sense, an artificial diffusion applied on the reconstructed distributions can be beneficial, as it damps the ghost particles and regularizes the shape of the true ones. The reconstruction algorithm proposed in Sec. 4.1 is based on spatial filtering of the distributions in between the MART iterations. For this reason, it is named Spatial Filtering Improved Tomographic MART (SFIT-MART).

In addition to this, the uncertainty in the reconstructed distribution could be not equally distributed in all the directions; in fact, physical constraints often lead to an elongation of the reconstructed particles along the depth direction. The effects of a purposely introduced anisotropy in the spatial filtering are investigated in this chapter, since one can tamper with the different shape and size of the reconstruction artefacts, like true particles elongation and ghost particles. Since most often the imaged particles are diffraction-limited spots with a Gaussian shape, the most natural choice is to apply diffusion in the reconstructed distribution using filtering windows with a Gaussian distribution of the weights, and size comparable to that of the particles.

A parametric assessment of the SFIT-MART performances is carried out in Sec. 4.2.1. The spatial resolution improvement is quantified with simulated Tomo-PIV experiments on a sinusoidal displacement field with a set of wavelengths in Sec. 4.2.2. Eventually, the technique is validated using experimental data on decaying nearly isotropic fractal generated turbulence (Sec. 4.2.3). The description of the technique and the main results have been published in Discetti et al (2013a).

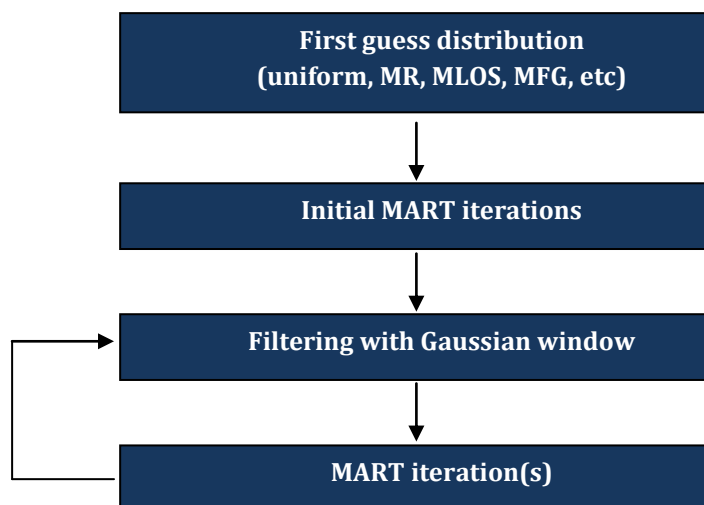


Fig. 4.1 Flow chart of the Spatial Filtering Improved Tomography (SFIT-MART).

4.1 Reconstruction quality enhancement by spatial filtering

The steps of the proposed method are summarized in the flow chart of Fig. 4.1.

1. A first guess for the distribution is built (uniform distribution equal to 1, MR, MFG, MLOS, etc.; the adopted method is of secondary importance in this framework);
2. One or two MART iterations are performed to refine the first guess distributions;
3. The reconstructed distributions are filtered by means of a Gaussian window to reduce the error due to voxel-based discretization and, as it will be discussed later, to apply a smoothing that, under proper choice of the shape of the filtering window, can be more effective on the reconstruction artefacts than on the true particles;
4. The procedure is repeated after each further MART iteration (the smoothing is not applied after the last iteration).

While the choice of the Gaussian filter is quite natural, as it is dictated by the feature of the objects to be reconstructed, the size and shape of the filtering kernel is a topic of discussion. In the present chapter, two options are tested: isotropic and anisotropic filtering.

Indeed, one can observe that the uncertainty in the reconstructed distributions is not equal along the spatial directions. The limited viewing angle that is usually




Fig. 4.2 Probability distribution function of the particles stretch ratio for the test layout illustrated in Sec. 4.2.1 with 4 cameras, source density $N_s = 0.5$ and a total viewing angle between the cameras $\phi = 60^\circ$.

employed in practice results in a larger uncertainty when determining the position of the particles in the depth direction. Consequently, both true and ghost particles are usually elongated in the depth direction.

A possible measure of the stretch of the particle in the z direction is given by the ratio of the standard deviations σ in the x and z direction, easily computed by local best fitting of the particles shape with a normal intensity distribution. Fig. 4.2 shows how this parameter is distributed in the case of the test layout illustrated in Sec. 4.2.1 for all particles, including ghost ones.

These observations suggest the possibility to tamper with the MART iterative algorithm by enjoying both the stretch ratio and the peculiar intensity distribution of true and ghost particles intensity (see Elsinga et al 2006a). In particular, one should transfer the energy of the intensity profile from the depth direction to the plane orthogonal to it. For this purpose the option of an anisotropic filtering of the reconstruction performed by MART with Gaussian windows orthogonal to the depth direction (for example $3 \times 3 \times 1$ voxels kernel, with the depth being the smallest dimension), having a kernel with size comparable to the particles diameter, is explored herein, in addition to the more intuitive isotropic filtering.

This approach is expected to provide a more intense smoothing of the intensity profiles of true and ghost particles in the planes orthogonal to the depth direction. Of course this reduces the intensity peak of both actual and ghost particles; on the other hand, the intensity of the ghost particles is reduced, while that of the actual

particles in the successive iterations is boosted up, recovering the energy subtracted by the ghost particles in the previous ones (provided that the smoothing is more effective on the ghost particles than on the true ones; otherwise, the linear update of MART will not provide any improvement on the quality of the reconstruction after the smoothing). The isotropic filter is expected to provide an improvement on the shape of the true particles, and to damp irregularly shaped ghost particles. The anisotropic filtering might provide an additional improvement when the ghost particles merge together with actual particles along the depth direction, or when they occur in form of streaks (but, on the other hand, it is less effective than the isotropic one in regularizing the shape of the true particles).

4.2 Performances of SFIT-MART

4.2.1 Effect of the experimental parameters

In this section, the performances are assessed by means of virtually simulated two-dimensional distributions of particles, reconstructed starting from 1D projections along different viewing directions, similarly to the approach used in other studies (Elsinga et al 2006a, Worth & Nickels 2008, Atkinson & Soria 2009). Gaussian particle images, with 3 voxels diameter (unless otherwise stated) and a maximum intensity of 200 counts, are generated at random locations by means of Monte Carlo technique in a $50 \times 10mm^2$ slice, discretized with a resolution of $20vox/mm$. Four 1020×1 pixel cameras, with a pixel pitch of $1/20mm$, are used. The cameras are placed at infinity, equally angularly spaced with symmetric arrangement; a uniform magnification is set equal to 1 over all the field of view.

The recorded intensities on the cameras are discretized at 12-bit levels. The tested source density N_s ranges between 0.05 and 0.85, corresponding to a concentration expressed in N_{ppp} varying between $0.017ppp$ through $0.282ppp$. However, one has to be careful in extrapolating the results to 3D since the 2D simulations do not account for cross-talk between planes. On the other hand, it is reasonable to assume that this issue affects in the same measure both MART and the proposed SFIT-MART. In order to get a proper scaling to the 3D application, the source density should be scaled roughly by a factor proportional to the particles diameter in pixel. The quality Q of the reconstruction is quantified in terms of the normalized factor of correlation between the reconstructed intensity field and the virtually generated distribution of particles, as in (2.4).

The SFIT method is applied after each MART iteration from the second one on; Gaussian filtering windows with 3×3 , 3×1 and 5×1 kernels and several values of the standard deviation (namely 0.5, 1 and 1.5 voxels; from this moment on the unit is not reported for simplicity) are applied.

In order to quantify the effect of the experimental parameters on the quality of the reconstruction, a test configuration characterized by 4 cameras with a total

viewing angle of $\phi = 60^\circ$ and $N_s = 0.45$ is considered as a reference. As for the filter, the considered standard is a kernel of 3×1 voxels with $\sigma = 1$. The relaxation parameter μ of the MART process (2.3) is set equal to 1. The weighting elements are computed as the area intersected by the lines of sight (a rectangle with width equal to the pixel size) and a circle with area equal to that of the voxels. The parameters used in the tests that have been performed are equal to the ones of the defined reference configuration unless otherwise stated.

Rate of convergence. The results in terms of accuracy of the reconstruction are reported as a function of the number of iterations in Fig. 4.3 for three levels of source density. In case of $N_s = 0.25$ the quality factor is increased by the anisotropic SFIT-MART of about 2% from iteration 5 on for both the tested values of standard deviation of the Gaussian filtering windows, $\sigma = 0.5$ and $\sigma = 1$; on the other hand, a too intense filtering is not advisable for a large number of iterations, since it seems to lead to divergence of the reconstruction from the exact solution. This problem can be prevented either returning to the more stable isotropic filter or using smaller σ , obtaining in this last case also a slightly larger quality factor. The divergence issue, however, is relevant just for small source densities which are not of interest for the standard Tomo-PIV applications.

A larger source density, of course, determines a reduction of the quality factor; on the other hand the gap between the standard MART based method and the proposed SFIT-MART technique increases. After 5 iterations, in case of $N_s = 0.45$, Q is approximately 0.83 for the reconstruction performed by MART, against 0.89 and 0.91 for the SFIT-MART algorithm, with anisotropic filtering and σ equal to 0.5 and 1, showing a relative improvement of 7.2% and 9.6%. In the same condition the isotropic filter provides a quality factor Q of approximately 0.88 and 0.89, respectively for σ equal to 0.5 and 1, slightly below the one provided by the anisotropic filter in both cases. The effect is even more consistent at very high source density. For $N_s = 0.65$, after the 5th iteration, the quality factor of the reconstruction by MART is 0.65, against 0.72 and 0.75 for the anisotropic SFIT-MART algorithm, with σ equal to 0.5 and 1, respectively. Again, the quality factor relative to the isotropic filter is slightly below the one provided by the anisotropic filter, in particular $Q = 0.71$, for $\sigma = 0.5$, and $Q = 0.73$, for $\sigma = 1$.

The SFIT-MART combination retains a good rate of convergence also after 5 iterations. After 10 iterations, in the case of $N_s = 0.45$, the best performances are achieved by the anisotropic filter with $\sigma = 0.5$ or $\sigma = 1$, for which Q is approximately 0.93. Under the same conditions, the isotropic filter provides a quality factor of about 0.92 while Q is about 0.86 when no filtering is applied. Finally, the anisotropic filter achieves higher Q after 10 iterations also for $N_s = 0.65$. In this case, applying SFIT-MART determines a quality factor of approximately 0.79 and 0.81 for anisotropic filtering and $\sigma = 0.5$ and $\sigma = 1$ respectively. The isotropic filter gives instead a slightly lower Q of about 0.78 and 0.80, while for MART $Q = 0.68$.

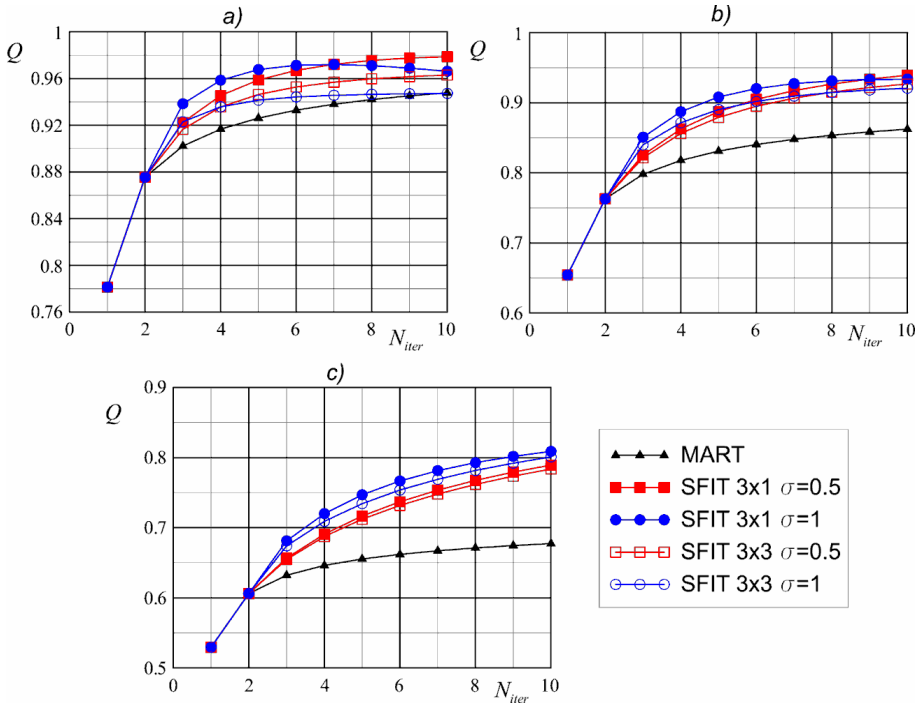


Fig. 4.3 Comparison between MART and the combination SFIT-MART (with both isotropic and anisotropic filters) in terms of Q as a function of the number of iterations N_{iter} . The results refer to a test layout with $\phi = 60^\circ$, 4 cameras and source density $N_s = 0.25$ (a), $N_s = 0.45$ (b), $N_s = 0.65$ (c).

The quality improvement might be due both to the regularization of the shape of the actual particles and to the suppression/damping of the ghost particles. In order to assess the effects of reduction of the ghost particles intensities the probability distribution functions (*pdf*) of the intensity of true and ghost particles are reported in Fig. 4.4. The results are reported for the case of 4 cameras, $N_s = 0.5$ and a viewing angle of 60° . The comparison of the distributions after 5 MART iterations and after SFIT-MART with anisotropic filtering 3×1 ($\sigma = 1$) shows that the ghost particles are weaker in the second case and the energy is re-distributed on the actual particles. The case of isotropic filtering provides similar results.

Effect of the source density. As outlined in the previous sub-section the anisotropic filter performs slightly better than the isotropic one. In Fig. 4.5 a comparison of three different filters in terms of quality factor Q is reported. It is clear that an anisotropic filter (in particular filtering windows of 3×1 or 5×1 pixels are considered) proves to be slightly more effective than an isotropic one at either low or high source density N_s . In particular, the isotropic filter appears to

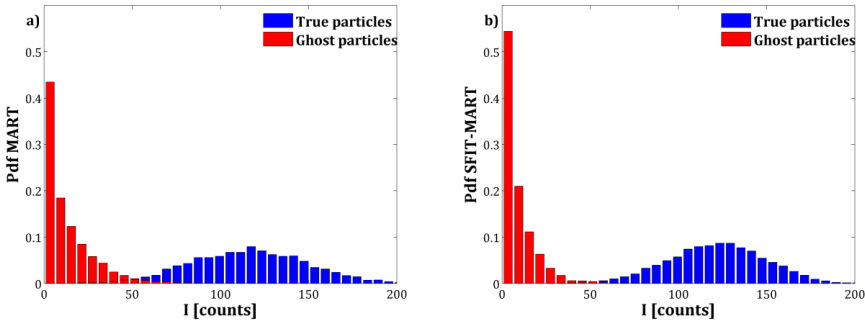


Fig. 4.4 Pdf of the particles peak intensity for the case of 4 cameras, source density $N_s = 0.5$ and a total viewing angle $\phi = 60^\circ$ after 5 MART iterations (a) or SFIT-MART with 3×1 Gaussian filtering (b).




Fig. 4.5 Comparison of Gaussian filters ($\sigma = 1$) characterized by different filtering windows size in terms of quality factor Q with varying source density N_s . The data are obtained after 5 iterations, with 4 cameras, total viewing angle $\phi = 60^\circ$.

be too intense especially at low source densities. Increasing N_s , the quality factor obtained with the isotropic filter approaches the one obtained with the anisotropic filter, but the former always provides slightly worse results.

Particle image and filter size. The influence of the particles size on the effectiveness of SFIT-MART is illustrated in Fig. 4.6. The performances of two Gaussian filters with kernel size of 3×1 and 5×1 are compared in terms of Q for variable particles sizes. It is possible to see that the method is effective provided that the particles diameter is at least of 2 pixels. Furthermore the effect of the standard deviation σ of the filters is shown. A kernel size of 5×1 appears to determine a lower or equal quality factor, when compared with the smaller one, for




Fig. 4.6 Quality factor after 5 iterations for two Gaussian filters with kernel dimensions of 3×1 (a) and 5×1 (b) voxels with varying standard deviation σ and the diameter of the particles. The data refers to an experimental setup with 4 cameras, $N_s = 0.45$, $\phi = 60^\circ$.




Fig. 4.7 Quality factor obtained using SFIT-MART (dashed and dotted lines for the anisotropic and the isotropic filtering, respectively) and MART (continuous lines) as a function of the viewing angle ϕ . These results are obtained using a Gaussian filter with $\sigma = 1$, 4 cameras and after 5 iterations.

all the tested configurations. This is in accordance with the intuition that the size of the optimal filtering should be comparable to that of the particles. However, the method seems to be only slightly sensitive to this parameter in the investigated range.

Solid viewing angle. In Fig. 4.7 the quality factor obtained using SFIT-MART (with isotropic and anisotropic filtering) is compared with the one obtained using the standard MART for different value of the total solid viewing angles ϕ . It can be seen that by increasing ϕ the absolute difference in the quality factor between the

three methods generally decreases, and a consistent accuracy improvement with respect to MART is retained in the entire range of investigation. This is understandable since the number of ghost particles and their actual distribution is strongly dependent on ϕ : the smaller the solid angle, the more elongated along the depth direction the true particles will be. The higher effectiveness of the method in case of small ϕ will actually strongly suggest the application of SFIT in cases in which the total angle subtended by the camera system is limited by physical constraints. As the solid angle approaches 90° , the performances of the isotropic and the anisotropic filter are practically equivalent, since the reconstructed particles are less elongated. In this scenario, the quality of the improvement is almost exclusively dictated by the filtering, independently of the shape of the filter.

4.2.2 Spatial resolution of SFIT-MART

A delicate aspect of tomographic PIV is connected to the coherent motion of the ghost particles between subsequent exposures, as outlined in Sec. 2.2.2. Since SFIT-MART has shown to increase the reconstruction quality, above all at high source density, where the degrading effect of the presence of ghost particles is overwhelming, one should expect a reduction of this modulation effect. This aspect is investigated by simulating a simplified tomographic PIV experiment, with the same layout of Sec. 4.2.1, i.e. 2D particles distribution observed by 1D cameras. The reconstruction is performed using 4 cameras with $\phi = 60^\circ$. A one-dimensional sinusoidal displacement field is imposed between two subsequent exposures:

$$u(x, z) = U \cdot \sin\left(\frac{2\pi}{\lambda} z\right) \quad (4.1)$$

where U is the amplitude, λ is the wavelength, z is the depth-volume coordinate, and x is the width-volume coordinate. The displacement along the depth direction is always set to zero; the amplitude U is equal to 2 voxels.

A 2D correlation analysis with iterative multi-grid and window deformation algorithm (see Secs. 1.1.4 and 2.4) is performed. The final interrogation window is 400×40 pixel (the width of the elongated interrogation window is set to ensure a sufficient number of particles; as a matter of fact, this does not change the generality of the analysis, since the displacement field is one-dimensional); the results are averaged on 50 independent realizations to reduce the effect of noise.

The spatial resolution performances are evaluated in terms of Modulation Transfer Function (MTF); the MTF of a classical top-hat moving average cross-correlation approach is practically coincident with that of a top hat moving filter. For the typical values of interrogation spot size W and λ in PIV applications it practically coincides with $\text{sinc}(W/\lambda)$. Since the displacement field is one-dimensional, W is the size of the interrogation window along the depth direction.

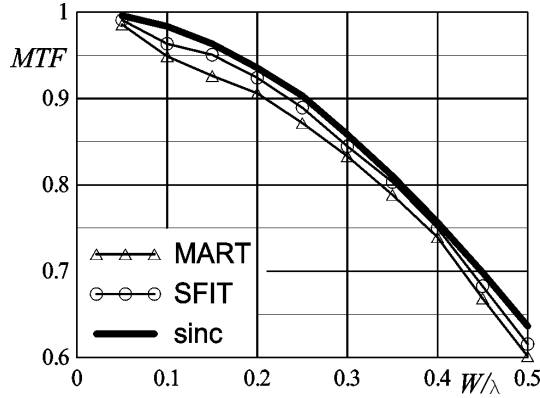


Fig. 4.8 Modulation Transfer Function at variable wavelength for $N_s = 0.45$, in the case of MART and anisotropic SFIT with standard deviation of the filtering window equal to 1.

Supposing that the interrogation process only modulates the exact displacement, as already shown by Astarita (2006) the MTF can be estimated as:

$$MTF(\lambda) = 1 - \sqrt{\frac{\sum_{i=1}^N (u_i - u)^2}{\sum_{i=1}^N (U \sin(2\pi z/\lambda))^2}} \quad (4.2)$$

In Fig. 4.8 the MTF is reported as a function of the spatial frequency W/λ in the case of a source density of 0.45; the performances of 5 MART iterations, and of SFIT-MART (with the filtering applied after each iteration from the second one on) are compared; in the case of SFIT-MART a Gaussian filtering window with a 3×1 kernel and a standard deviation equal to 1 is used. In this section only the anisotropic filtering is investigated, since it provides slightly better performances. However, the results with the isotropic filter are not expected to be significantly different, and for this reason they are not included herein. The results show a clear enhancement of the modulation in case of adoption of SFIT-MART; the effect is particularly pronounced at low frequencies, since at high frequencies the local relatively strong velocity gradient contributes to decouple the ghost particles in the two exposures by itself, as outlined in Sec. 2.2.2.

The MTF depends, of course, on the source density, since the number of actual particles is one of the main parameters in determining the percentage of ghost particles into the volume. This aspect is illustrated in Fig. 4.9, in which the MTF is plotted as a function of the source density for three different values of the wavelength of the sinusoidal displacement. The modulation is approximately constant at low source density, and the difference between the spatial resolution of MART and SFIT-MART is almost negligible, since the percentage of ghost particles is relatively small. An increase of the source density implies a stronger modulation effect on the reconstructions performed by MART with respect to those obtained by SFIT-MART.




Fig. 4.9 MTF at variable source density, for three different wavelengths in the case of MART (triangles) and anisotropic SFIT (circles) with standard deviation of the filtering window equal to 1.




Fig. 4.10 Standard deviation of the error on the x component of the displacement for $N_s = 0.30$ (a) and $N_s = 0.60$ (b).

The effect of the SFIT enhancement on the standard deviation of the error on the x component of the displacement field as a function of the frequency is illustrated for $N_s = 0.30$ and $N_s = 0.60$ (Fig. 4.10). The results clearly show that at low source densities the difference between MART and SFIT-MART is negligible. On the other hand, increasing the source density, a definite reduction of the standard deviation of the error is observed in the case of adoption of SFIT-MART for the reconstruction. Regarding the behaviour of the standard deviation of the error on the z component, no net difference between the two approaches for all the tested seeding densities is detected.




Fig. 4.11 Schematic drawing of the fractal square grid geometry (see Sec. 7.1.1 for further details).

4.2.3 Experimental application to fractal generated turbulence

The applicability of SFIT-MART is tested with an experiment performed in the open circuit wind tunnel facility of Arizona State University. The flow field under investigation is the decaying nearly isotropic turbulence generated by a square fractal grid (i.e. a grid with a square pattern repeated at increasingly smaller scales, see Sec. 7.1 for further details). The fundamental elements of the experimental apparatus are introduced in this section; a more insightful description is provided in Sec. 7.2. The square fractal grid is placed at the inlet of the test section of the wind tunnel, after the contraction.

The flow is seeded with olive oil particles with $1\mu m$ mean diameter, generated by a Laskin nozzle. The particles are illuminated by a double pulse Nd:Yag laser with maximum energy of $100mJ/pulse$; the illuminated region is a slab with the larger dimensions along the streamwise and the vertical crosswise direction, and a thickness of approximately $3mm$.

The projections of the distribution of particles are recorded by 4 TSI POWERVIEW™ Plus 11MP camera with 4008×2672 pixel resolution, with a pixel pitch of $9\mu m$. The cameras are placed approximately $550mm$ downstream of the grid in a linear arrangement, angularly equally spaced, spanning an angle of 80° . The cameras are equipped with Nikon objectives with a focal length of $60mm$ when focused at infinity, and $f_{\#} = 16$. The imaging system is set to have an average magnification of approximately 0.31, resulting in an image resolution of about $35pix/mm$. The diffraction-limited particle diameter is of approximately 3 pixels.

The calibration target for the optical calibration is made of a glass substrate with a $100 \times 100mm^2$ grid of black markers, with diameter of $250\mu m$ and spacing of $1mm$, with 0.2% tolerance. The target is then translated within the range $\pm 6mm$ with respect to the reference plane. A nonlinear regression algorithm is used to

obtain pinhole-based mapping functions (Tsai 1987) from the object space to the image plane (see Sec. 1.2.1). The volumetric self-calibration described in Sec. 2.1.4 is applied to reduce any residual calibration error and disparity between the cameras.

The images are pre-processed to eliminate the background and reduce the image noise. A four steps pre-processing is implemented: historical background removal by taking the minimum in each pixel over an ensemble of 500 realizations; instantaneous residual background removal by subtracting from the images a low-pass filtered image (a Gaussian filter on a kernel 30×30 and standard deviation equal to 5 is used); thresholding of the images to remove the residual noise due to high frequency background fluctuations (1% of the maximum particle intensity); 3×3 Gaussian smoothing with $\sigma = 0.5$ to restore the particle tails deleted by the thresholding.

The reconstructed volume is $50 \times 60 \times 4 \text{ mm}^3$ (the depth dimension is slightly larger than the laser slab thickness to ensure that all the imaged particles will be reconstructed), discretized with 35 vox/mm (i.e. an average resolution ratio of approximately one between pixels and voxels is maintained).

The reconstruction is performed using a camera system composed only of 3 out of 4 available cameras (the two external cameras and one internal camera) by a standard 5 MART iterations algorithm and the proposed SFIT-MART. The filtering step is applied from the second iteration on; a Gaussian filtering window with a kernel of $3 \times 3 \times 1$ voxels (the third one is the depth dimension) and with standard deviation equal to 1 is used (the difference of the results obtained with the isotropic filtering are not significant, and for this reason are not included herein; this is actually in line with the results of Fig. 4.7, since the angle subtended by the cameras is 80°).

From this moment on the distributions reconstructed by the 4 cameras system with the standard MART algorithm will be considered as “ground truth”, i.e. the quality factors can be calculated by assuming that the reconstruction performed by 4 cameras will lead to nearly exact result. Since the tomographic PIV system is operated at seeding density of around 0.045 ppp (particles per pixel), the numerical simulations by Elsinga et al (2006a) suggest that this hypothesis is quite satisfactorily verified. In order to avoid confusion, this relative quality factor will be indicated with the symbol Q^* . The seeding density is systematically underestimated due to the occurring of overlapping particles, the actual N_{ppp} is expected to be slightly larger (Novara & Scarano 2012a detect an underestimation of about 10%; however, in this application the depth of the volume is quite small, suggesting that the probability of overlapping particles along the lines of sight is consistently smaller).

The relative quality factor Q^* is reported in Fig. 4.12. The results after 5 iterations indicate that MART has approximately reached an asymptotic value of




Fig. 4.12 Relative quality factor of the tomographic reconstruction for MART and SFIT-MART with a 3 camera system.




Fig. 4.13 Mean intensity profile of the reconstructed distributions along the z direction.

0.64, while SFIT-MART achieves a quality factor of approximately 0.73 with a residual rate of increase. The simulations of Sec. 4.2.1 have shown that further iterations would determine an improvement of the quality factor for SFIT-MART; on the other hand, the computational cost and the limited amount of improvement suggest considering 5 iterations as a good standard for the tomographic reconstruction both for MART and SFIT-MART.

The profile of the mean intensity of the reconstructed distribution along the z direction can be considered an optional indicator of a successful reconstruction (Elsinga et al 2006b). The profile provides information about the mean intensity of




Fig. 4.14 Normalized factor of correlation and fluctuating velocity field for MART (left) and SFIT-MART (right) for a portion of the slice $z = 0\text{mm}$.




Fig. 4.15 Pdf of the normalized factor of correlation and fluctuating velocity field for MART (3cam) SFIT-MART (3cam) and the reference reconstruction with MART (4cam).

solely the ghost particles (outside the illuminated volume) and of the sum of ghost and true particles (within the illuminated volume). The ratio of the peak of the profile and the average value of the plateau outside the volume can be considered as a signal to noise ratio for the reconstruction. The results in Fig. 4.13 show a slight improvement of the quality of the reconstruction (the ratio peak/plateau increases from about 7 for the standard MART to approximately 7.4 for SFIT-MART). The result for the 4 camera system, providing a peak/plateau ratio of about 18, is included for comparison. The reconstructed distributions are then processed with the efficient multi-pass volume deformation algorithm outlined in the Chapter 5. The final interrogation volume size is $48 \times 48 \times 48$ voxels, with a vector spacing of 12 voxels (i.e. 0.34mm). The final steps of the process are performed with the aid of




Fig. 4.16 Pdf of the divergence (normalized with the square root of the characteristic turbulent rate of strain) of the velocity field for MART, SFIT-MART and the reference velocity field (MART with 4 cameras).

| Method | $\sigma_{(u-u_{ref})}$ [vox] | $\sigma_{(v-v_{ref})}$ [vox] | $\sigma_{(w-w_{ref})}$ [vox] |
|-----------|------------------------------|------------------------------|------------------------------|
| MART | 0.175 | 0.160 | 0.285 |
| SFIT-MART | 0.150 | 0.148 | 0.257 |

Table 4.1 Standard deviation of the velocity measurement difference with respect to the reference velocity field.

direct correlation with a narrow search area on the cross-correlation map. In this scenario, the standard definition of the signal to noise ratio as the ratio between the first and the second peak on the cross-correlation map cannot be applied. On the other hand, a reliable instrument is the normalized factor of correlation, ranging between 0 and 1 (Astarita 2009). A higher normalized factor of correlation gives direct information on the better quality of the measurement.

The factor of correlation distribution is reported in Fig. 4.14 as a contour plot for a portion of the middle slice of the interrogation volume; the fluctuating $u - v$ velocity field is plotted for comparison. In general, SFIT-MART yields higher normalized factor of correlation, and slightly lower outliers percentage (for the case presented in Fig. 4.14, the interrogation performed by SFIT-MART determines 88% of correct vectors, while MART achieves 85%; the validation is based on the universal median criterion, Westerweel & Scarano 2005).

The pdf of the factor of correlation obtained over an ensemble of 100 images, and computed only on valid vectors, is reported in Fig. 4.15 to further support this

assertion. The application of SFIT-MART skews the pdf towards larger values of the factor of correlation. The average factor of correlation is 0.54 and 0.59 for MART and SFIT-MART, respectively (for the case of the interrogation performed on the reference objects reconstructed by MART using 4 cameras, the averaged normalized factor of correlation is equal to 0.62).

The standard deviations of the difference between the reference velocity field and the velocity fields measured by interrogating the distribution reconstructed by MART and SFIT-MART are reported in Tab. 4.1. As expected, in both cases the error is larger along the depth direction; the accuracy improvement obtained by using SFIT-MART (about 10% reduction of the standard deviation of the error) is clearly detectable.

A more general indicator of the quality of the velocity measurement in case of incompressible flow regime is the standard deviation of the divergence. The probability distribution of the divergence for a single realization is reported in Fig. 4.16. The divergence is presented in non-dimensional form; the reference value is the square root of the characteristic value of the turbulent rate-of-strain tensor (expressed in a synthetic way in the figure as the square root of the ratio of the turbulent dissipation ε and the kinematic viscosity ν). This characteristic value is obtained by taking the median value of the highest 50 local peaks of the local turbulent dissipation within a single realization for the case of the reference velocity field. The standard deviation of the non-dimensional divergence is equal to 0.048 and 0.037 for the case of the standard MART with 3 and 4 cameras, respectively. With the SFIT-MART technique the standard deviation is reduced to 0.042.

Chapter 5 - Efficient 3D PIV interrogation algorithms

In the early stages of development of Tomographic PIV, most of the attention has been devoted to the tomographic reconstruction, as it was recognized to be the bottleneck in the processing algorithm in terms of computational cost. The 3D PIV interrogation process has received limited attention, as it is a conceptually trivial extension of the standard 2D interrogation algorithms. However, the recent achievements outlined in the previous chapters have changed the scenario. In addition to this, the development of more robust (but more computationally intensive) PIV interrogation algorithms (for example the adaptive PIV interrogation by Novara et al 2013, or the multi-frame pyramid correlation by Sciacchitano et al 2012) require new efficient algorithms to reduce the computational burden. Furthermore, one should note that 3D cross-correlation can also be employed to improve the quality of the tomographic reconstruction itself (see for example the MTE-MART discussed in Sec. 2.3.3). As a consequence, the PIV algorithm might repeatedly appear in the Tomo-PIV process. It is clear that a reduction of the computational load of 3D PIV can contribute not only to the diffusion of Tomo-PIV as a standard tool for 3D3C velocity measurements, but also to increase the accuracy of the tomographic reconstruction itself with a limited impact on the total processing time.

The optimization of 3D PIV algorithm efficiency is the object of the present chapter. A number of solutions are illustrated in Sec. 5.1, enabling a more efficient calculation of the velocity field without any significant loss of accuracy. The proposed algorithm combines the multi-resolution interrogation concept (see Chapter 3 for the application to the tomographic reconstruction) with sparse data storage, cross-correlation and interpolation. Furthermore, three different approaches to reduce the number of redundant calculations for overlapping windows are presented, based on pre-calculations of the contributions to the cross-correlations coefficients along segments, planes or blocks. The performances of the different processing solutions are tested in Sec. 5.2 in different scenarios with progressively tuned complexity. Finally, in the case of block cross-correlation, the degrading effects due to poor discretization of the weighting windows are discussed in Sec. 5.3. The results discussed in this chapter are published in Discetti & Astarita (2012b).

Again, volume deformation methods will be considered the reference processing algorithm.

5.1 Fast correlation computation

The structure of the interrogation algorithm is the same proposed in Sec. 2.4. For clarity a more detailed description is reported in this section.

1. The predictor displacement field is calculated on a rather coarse grid for a quick estimation of the velocity field. Since the displacement is not known *a priori*, the interrogation spot size (from this moment on indicated with the acronym IV, which stands for Interrogation Volume) has to be large enough to compensate for the motion of the particles within the volume. In 2D PIV it is commonly accepted that the interrogation spot size should be about 4 times the maximum expected displacement (the so called *one-quarter* rule by Keane & Adrian 1992), so that the loss of pairs due to the particles displacement is less than 25%. In general, as outlined in Fig. 1.2, it should be $N_I F_0 F_I > 7$. However, since there is no effect of out-of-plane motion in 3D PIV, one should simply set $N_I F_I > 7$. The most likely displacement of the particles is determined by the analysis of the cross-correlation map between subsequent exposures. The normalized 3D cross-correlation operator is defined as:

$$\Phi(\Delta X, \Delta Y, \Delta Z) = \frac{\sum_{i,j,k}^{D_I} w_{i,j,k}^2(A_{i,j,k} - \mu_A)(B_{i+\Delta X, j+\Delta Y, k+\Delta Z} - \mu_B)}{\sqrt{\sum_{i,j,k}^{D_I} w_{i,j,k}^2(A_{i,j,k} - \mu_A)^2 \sum_{i,j,k}^{D_I} w_{i,j,k}^2(B_{i+\Delta X, j+\Delta Y, k+\Delta Z} - \mu_B)^2}} \quad (5.1)$$

where A and B are the intensities of the IVs, μ_A and μ_B are their relative mean intensity values, and w is a weighting function. The (5.1) is a generalized 3D version of (1.5), as it includes also the possibility of using weighting windows in the cross-correlation step to tamper with the impulsive response of the algorithm. The operation can be conducted with the aid of FFT (Fast Fourier Transforms), requiring $3 \times 3D_I^3 \log_2 D_I^3$ operations (two direct and one inverse Fourier transforms) instead of D_I^6 by directly computing (5.1), i.e. using direct cross-correlation. This estimation includes only the computation of the numerator of (5.1). In the case of adoption of weighting windows, the overhead for the normalization can be more significant, as the term $\sum_{i,j,k}^{D_I} w_{i,j,k}^2(B_{i+\Delta X, j+\Delta Y, k+\Delta Z} - \mu_B)^2$ requires the computation of a cross-correlation. In practice, the term can be replaced by $\sum_{i,j,k}^{D_I} w_{i,j,k}^2(B_{i,j,k} - \mu_B)^2$ with acceptable approximation;

2. A *dense predictor* is calculated by interpolating the velocity field on each voxel. The dense predictor is then used to deform the volumes (in the present algorithm the volumes are displaced symmetrically to achieve second-order accuracy, Wereley & Meinhart 2001). In this step the performances of the velocity interpolation scheme may have an influence on the spatial resolution of the algorithm (Astarita 2008);

3. Once the predictor displacement is known, volume deformation can compensate the particles displacement, determining a higher valid detection probability. Since F_I is nearly 1, and the velocity gradients within the IV are partly compensated by the volume deformation, one can obtain a high detection probability with fewer particles, i.e. smaller interrogation spots. A grid refinement may be executed in the evaluation of the displacement field;
4. A corrector displacement field is evaluated on the deformed volumes by applying the (5.1). In the corrector estimation one could use a narrower search area of the peak (as already observed by Rohàly et al 2002), e.g. limiting the calculation only to the cross-correlation coefficients in the neighbourhood of the peak itself. As a limiting case, only the coefficients that are necessary to compute the Gaussian interpolation of the peak can be evaluated. In this case, direct cross-correlation is much faster and more accurate than FFT, since no periodicity is imposed. Of course, this approach works only if the predictor is accurate within ± 0.5 pixels on the finer grid, otherwise the full cross-correlation map should be calculated (or the search radius should be enlarged);
5. The true displacement is obtained by summing the corrector and a weighted average of the dense predictor over a prescribed window. This operation is necessary as the so-called *local approach* (i.e. summing the local predictor with the corrector) might lead to instability of the algorithm if large spatial frequencies are present in the flow field (unless an appropriate weighting window is used in steps 1 and 4, as proposed by Nogueira et al 1999). One should note that this operation requires that the dense predictor computed in step 2 has to be stored in memory (otherwise, it could be computed on-the-fly). The required memory can be consistently large (if single floating point precision is used, 4 bytes time the number of voxels are required for each velocity component of the dense predictor, i.e. for a $1000 \times 1000 \times 200$ voxels discretization, 2.4Gb are needed).

The part of the process from step 2 to step 5 is iteratively repeated until a satisfactory convergence is reached. Evidently, the predictor and corrector estimations have to be accelerated with different solutions, since in the first case FFT is to be preferred, while in the second case direct correlation can be more appealing.

5.1.1 Multi-resolution predictor estimation

Similarly to the solution proposed in Chapter 3, voxel binning can consistently accelerate the predictor computation. The binned 3D distributions are built by clustering voxels in $C \times C \times C$ kernels, being C the binning factor. The number of multiplications to be performed is reduced to $o((D_I/C)^3 \log_2(D_I/C)^3)$; e.g. if the IV is a cube of 64^3 voxels, a $2x$ and a $4x$ binning enable a processing time reduction (in

the following referred as speed-up) of the order of 10 and 100 times, respectively. Of course a certain overhead due to the computation of the binned volumes has to be considered, i.e. a number of additions equal to the number of voxels. However, the ratio between the number of operations to compute the FFT and that to bin the volumes is $o(\log_2(D_l/C)^3)$ if no overlapping windows are used, and it increases with the cube of the reciprocal of the overlap; furthermore additions may be computed faster than multiplications, so that the overhead is almost negligible.

5.1.2 Fast corrector computation with sparse direct correlations

Direct cross-correlation can be used to compute the coefficients only in the neighbourhood of the peak, e.g. the peak itself and the 6 coefficients to find the peak with sub-pixel accuracy using a Gaussian fit along the three directions. This reduces the number of multiplications to be performed to $7D_l^3$, versus $o(9D_l^3 \log_2(D_l))$ in case of adoption of the FFT, and with the same overhead for the normalization (i.e. compute the mean and the covariance of the IV, as in (5.1)); e.g. an acceleration of about 7.7 times is obtained in case of IV with linear dimension of 64 voxels.

Direct cross-correlation can be further accelerated observing that the distributions to be interrogated are very sparse. Considering the situation outlined in the example of Sec. 2.3.1, i.e. 10 particles per IV with the linear dimension of 32 voxels and particle occupying a kernel of $5 \times 5 \times 5$ voxels, less than 1% of the voxels carries a non-zero intensity; considering ambiguities in the tomographic reconstruction (depending mainly on the imaging system configuration, the depth of the volume to be reconstructed, the number of cameras and the seeding density), this percentage is seldom higher than 5%, since thresholding is usually applied to accelerate the reconstruction step. For planar PIV and analogous conditions, 24% of the pixels have non-zero intensity; in the 2D scenario the overhead to handle data in sparse matrix is not negligible and it complicates the algorithm with a worthless acceleration. In 3D, the very low percentage of non-zero voxels makes it worth the effort.

One aspect making it difficult to implement sparse cross-correlation with both the methods is connected to the memory access to the corresponding element in the second volume to each non-null element of the first one. For this reason the 3D extension of the Yale format (see Sec. 3.2) is of fundamental importance, as it allows for efficient data access.

5.1.3 Efficient corrector computation with overlapping windows

In common practice overlapping windows are used (overlap typically ranges between 50% and 75%). In this case, a relevant percentage of operations are repeated, and a strategy to reduce the number of redundant calculations has to be assessed. In the following, 3 approaches to reach this goal are proposed, each one relying on the observation that the cross-correlation coefficients for each IV can be

obtained by summing contributions of sub-volumes (Rohàly et al 2002). It is convenient to split the contributions to the correlation coefficients, decomposing (5.1) as in the following (for the sake of simplicity the weighting window is supposed to be constant and equal to 1, i.e. a top hat):

$$\Phi(\Delta X, \Delta Y, \Delta Z) = \frac{\sum_{i,j,k}^{D_I} A_{i,j,k} B_{i+\Delta X, j+\Delta Y, k+\Delta Z} - \mu_A \mu_B D_I^3}{\sqrt{(\sum_{i,j,k}^{D_I} A_{i,j,k}^2 - \mu_A^2 D_I^3) \cdot (\sum_{i,j,k}^{D_I} B_{i+\Delta X, j+\Delta Y, k+\Delta Z}^2 - \mu_B^2 D_I^3)}} \quad (5.2)$$

For each IV of linear dimension D_I , the following terms have to be calculated:

$$\mu_A = \frac{\sum_{i,j,k}^{D_I} A_{i,j,k}}{D_I^3} \quad (5.3)$$

$$\mu_B = \frac{\sum_{i,j,k}^{D_I} B_{i,j,k}}{D_I^3} \cong \frac{\sum_{i,j,k}^{D_I} B_{i+\Delta X, j+\Delta Y, k+\Delta Z}}{D_I^3} \quad (5.4)$$

$$\sum_{i,j,k}^{D_I} A_{i,j,k}^2 \quad (5.5)$$

$$\sum_{i,j,k}^{D_I} B_{i,j,k}^2 \quad (5.6)$$

$$\sum_{i,j,k}^{D_I} A_{i,j,k} B_{i+\Delta X, j+\Delta Y, k+\Delta Z} \quad (5.7)$$

The contributions of the formulae (5.3)-(5.7), as proposed by Rohàly et al (2002), can be computed on sub-volumes composing the IV (the geometry of the sub-volumes can be arbitrarily chosen, i.e. the sum on the three indexes are extended to the relative dimensions of the sub-volume in the three spatial directions). The algorithm is, thus, composed of three steps:

- The sums (5.3)-(5.7) are pre-calculated for each sub-volume (in this case the sums have to be generalized to a region with generic shape);
- The various contributions (5.3)-(5.7) relative to all the sub-volumes constituting the IV are summed up;
- The cross correlation coefficient is evaluated by using (5.2).

Imposing the symbols $W_Q = \sum_{i,j,k}^{D_I} w_{i,j,k}^2$, $W_{QA} = \sum_{i,j,k}^{D_I} w_{i,j,k}^2 A_{i,j,k}$ and $W_{QB} = \sum_{i,j,k}^{D_I} w_{i,j,k}^2 B_{i+\Delta X, j+\Delta Y, k+\Delta Z}$, in case of adoption of weighting windows, (5.1) reduces to:

$$\Phi(\Delta X, \Delta Y, \Delta Z) = \frac{\sum_{i,j,k}^{D_I} w_{i,j,k}^2 A_{i,j,k} B_{i+\Delta X, j+\Delta Y, k+\Delta Z} - \mu_A W_{QB} - \mu_B W_{QA} - \mu_A \mu_B W_Q}{\sqrt{(\sum_{i,j,k}^{D_I} w_{i,j,k}^2 A_{i,j,k}^2 + \mu_A^2 W_Q - 2\mu_A W_{QA}) \cdot (\sum_{i,j,k}^{D_I} w_{i,j,k}^2 B_{i+\Delta X, j+\Delta Y, k+\Delta Z}^2 + \mu_B^2 W_Q - 2\mu_B W_{QB})}} \quad (5.8)$$

The number of elements to be pre-calculated for each sub-volume and then summed up increases accordingly.




Fig. 5.1 Comparison between the standard Blackman weighting window and its piecewise version for several overlap values (nbl stands for n blocks).

5.1.3.1 Block cross-correlations

The most intuitive solution implies the pre-calculation of the terms (5.3)-(5.7) on cubic (or, in general, parallelepipedal) sub-volumes (in the following called *blocks*), whose dimension could be set as the greatest common divisor between the IV linear dimension and the grid distance (this second parameter is replaced by the overlapping part of the interrogation volumes when the overlap is smaller than 50%; e.g. if the linear dimension of the IV is 64 voxels, and the overlap is 25%, the grid distance is equal to 48 voxels, and the overlapping part is 16 voxels, so that the maximum possible linear dimension of the pre-calculated blocks is 16 voxels). The idea is equal to that of Roth & Katz (2001), with the difference that this method is employed only to calculate a very limited number of coefficients and no truncation is performed, i.e. the multiplications are performed in single precision floating point format (instead of the single-bit parallel multiplication algorithm proposed by Roth & Katz 2001) to avoid degrading effects on the accuracy of the results. Performing the calculation of the displacement map in case of overlap ranging between 25% to 75% has practically the same computational cost of the case of non-overlapping windows, i.e. overlap is introduced without any significant change of the processing time.

The main drawback regards the implementation of weighted cross-correlation: the weighting window can be replaced by a piecewise weights distribution for each IV (i.e. each block contribution is weighted with a constant value, for example the average of the weighting window on that block). Of course, this approach is reliable only in case of small ratio between the linear dimension of the blocks and that of the IV (i.e. highly or barely overlapped IV); a brief discussion is provided in Sec. 5.3. An example of block-version of the Blackman window is provided in Fig. 5.1 for different overlap percentages; in this case the choice is to weight each block with

the value of the original window discretized with a number of points equal to the number of blocks.

5.1.3.2 Segment or rectangular based cross-correlations

The choice of the sub-volume shape on which performing the pre-calculation is clearly arbitrary and, while the choice of a cubic block enables to reduce the time needed to perform the final sums, it does not allow the precise use of weighting windows. A way to perform all the calculations required by (5.8) is to pre-calculate the sums (5.3)-(5.7) (and the other sums required for the introduction of the weighting window, obtained by splitting (5.8) in its basic contributions) along two (e.g. index i and j) of the three indexes, i.e. to use rectangular (or better a parallelepiped with a dimension equal to one voxel) shaped sub-volumes. In this case, separable weighting windows (built as the product of three weighting windows for the three directions, i.e. $w_{i,j,k} = w_i \cdot w_j \cdot w_k$) can be correctly used.

The algorithm can be better understood by first considering the two dimensional case. As shown in Fig. 5.2a, where the actual IV is shown with a shaded square, the pre-calculation of the sums is performed along columns (indicated with blue rectangles in the figure, while green squares identify the pixels) and stored in a temporary array (schematized with the top rectangle in the figure); the substantial difference with respect to the block cross-correlation case is that the sums are only evaluated on a single row of interrogation volumes and successively the elements of the array are summed up to complete the process (shaded box on the top rectangle of Fig. 5.2a and, for the second IV of the row, in Fig. 5.2b). The inclusion of a separable weighting window is possible since the sums are split for the two indexes.

The calculation of the cross-correlation coefficient in the following rows of interrogation volumes (Fig. 5.2c-d) is performed with the same principle. Since the pre-calculated sum of the previous row of interrogation volumes is not used in the following one a significant overhead is introduced with respect to block cross-correlation. Actually in the evaluation of a complete two dimensional map of cross correlation coefficients the computational burden scales linearly with the overlap.

The extension to the three dimensional case can be made with two approaches that differ in the way in which the pre-calculation is performed. In the first one, called in the following 2D DC, the pre-calculation of the sums is performed along rectangular sub-volumes (i.e. by varying two indexes in the pre-calculation step) while in the second one (1D DC) segments (i.e. by varying only one index in the pre-calculation). In the 2D DC approach the pre-calculated sums are stored in an array and the algorithm is a very simple extension of the two-dimensional one, the only difference being that the overhead scales quadratically with the overlap. On the other hand with the 1D approach the pre-calculated sums have to be stored in a

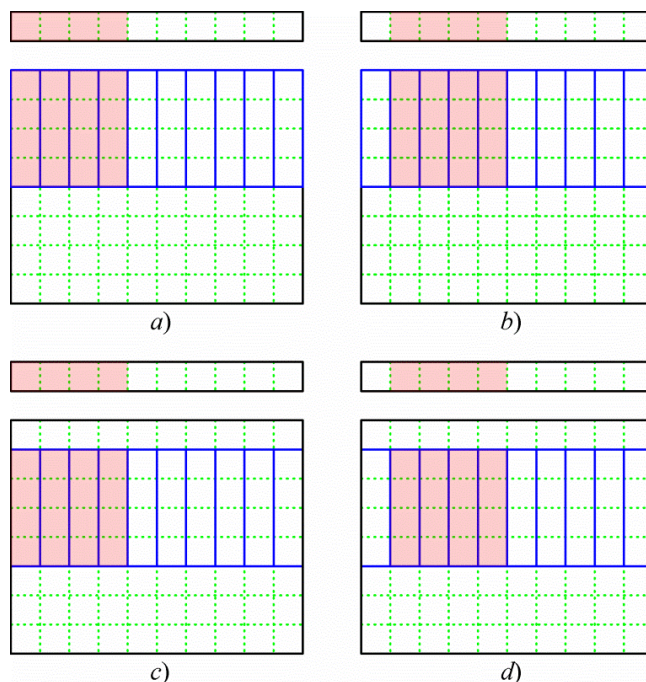


Fig. 5.2 Segment-based direct cross-correlations. The big black rectangles represent the images and the top rectangle is a temporary array. Blue rectangles indicate the columns on which sums are pre-calculated; shaded squares refer to actual IV and green squares identify the pixels. First (a) and second (b) IV on the first row of interrogation volumes; first (c) and second (d) IV on the second row of interrogation volumes.

two-dimensional array but the overhead still scales linearly with the overlap. For both approaches the size of the temporary arrays can be reduced by using cycling indexes.

5.1.3.3 Wider search area: block FFT

Due to modulation effects or inaccuracy of the estimated predictor, it might happen that the correlation peaks of the corrector displacement field fall above ± 0.5 pixels, and a great number of full correlation maps have to be computed by using FFT. This occurs especially in regions of strong velocity gradients, or low signal to noise ratio. One possible solution is to enlarge the search area of the peak, i.e. to compute a wider zone of the correlation map. Unfortunately, if k is the search radius, the number of coefficients to be computed is roughly proportional to k^3 (the peak can be detected if it is included in the zone $\pm(k-0.5)$). A solution to this problem is to compute the blocks of the correlation maps using FFT (as in Rohály et al 2002); the blocks are stored in memory and then summed to obtain the search area of the correlation map for each IV, as in 5.1.3.1.

This approach suffers of bias effects due to the imposed periodicity; the aspect is particularly critical because of the small size of the IV (e.g. for IV of 64^3 voxels, and 75% overlap, the blocks are only 16^3 voxels). This prevents the application to small highly overlapped IV and, in general, a bias correction is of fundamental importance. In the present work the bias correction is performed by multiplying the coefficients of the correlation maps with the inverse of a triangular window (Raffel et al 2007).

Often during the grid refinement part of the interrogation algorithm the modulation associated to both the broad dimension of the IV in the predictor estimation (Astarita 2007) and to the interpolation of the velocity field (Astarita 2008) on the refined grid makes it difficult to have small residual displacements. In these cases the block FFT approach is particularly effective since it avoids the need to recalculate the full cross-correlation map when the corrector residual displacements are larger than ± 0.5 voxels.

5.2 Performance assessment

The performances are assessed in four different layouts: synthetic distributions and small uniform sub-pixel displacement; shear displacement on distributions reconstructed by 5 MART iterations; synthetic distributions of a simulated jet profile; real images of a swirling jet. Each layout provides information on different features of the algorithm: the small uniform displacement highlights the performance limit of the algorithm; the shear displacement layout on reconstructed distribution focuses on the effects of the ghost particles in presence of spatial velocity gradients; the synthetic generated jet profile tests the algorithm in presence of tuned strong local gradients (in this case the effect of the ghost particles in modulating the velocity gradient along the depth direction provides undesirable anisotropy of the performances of the algorithm; consequently the analysis is performed on the original synthetic distributions); the real images of a swirling jet bundle all these aspects together.

The tests are carried out on a computer using a single core (so that the parallelization coding effects are reduced) of a 3.07 GHz i7 processor (with the exception of the last case, in which all 4 cores have been employed). In many cases the results will be presented in terms of speed-up with respect to a standard reference method (i.e. the full FFT analysis performed on the same grid and with the same IV size); this speed-up is only relative to the predictor and corrector evaluation steps (i.e. the processing time to compute the dense predictor field, to interpolate the volumes and to execute the validation is neglected).

From this moment on, the search radius for the direct correlation and block FFT based algorithms will be set to 1 and 2 voxels, respectively.

5.2.1 Test case 1 - Uniform displacement

Synthetic distributions of spherical particles with Gaussian profile and 3 voxels diameter are interrogated. A volume of $40 \times 40 \times 10 \text{ mm}^3$ is discretized with 20 vox/mm resolution, resulting in $800 \times 800 \times 200$ voxels; the IV is 40^3 voxels. A uniform displacement of 0.2 voxels is imposed, in order to assure that all the displacements are within 1 pixel search radius and, therefore, direct correlations always succeed in finding the correlation peak. Seeding concentrations of 1.25 and $3.75 \text{ particles/mm}^3$ (resulting in approximately 0.5% and 1.5% of voxel with non-zero intensity respectively when a threshold at 0.5% of the peak intensity is applied) are tested. 20 iterations are executed to obtain well converged time statistics.

The speed-up with respect to the equivalent process with computation of the full cross-correlation map by FFT is reported in Fig. 5.3. Standard direct correlations, without the aid of sparse matrices and redundancy avoidance, enable a processing time reduction of approximately 4 times, irrespective of the IV overlap. Among the others schemes reducing the redundant calculations and enjoying the sparsity of the distributions, block direct correlations are very effective when the overlap ranges between 25% and 75%. Indeed, in this case the number of operations in the cross-correlation computation step, being almost independent of the overlap, is much smaller than that relative to both 1D DC and 2D DC. In case of 75% overlap, the processing time is reduced of about 800 and 400 times for the lowest and the highest tested density, respectively, with respect to the standard processing by FFT. For the highest tested overlap the most performing approach is 1D DC that enables to have, for the smaller source density, a speed-up of over 1000 times.

The Block FFT algorithm, with a search radius of 2 pixels, is slower than the standard method based on FFT when the IV are only slightly overlapped, mainly because of the overhead due to summations of the maps. Increasing the overlap boosts up the efficiency of the Block-FFT approach; however, the speed-up is lower than that of the methods based on sparse direct cross-correlation and redundancy avoidance. Nevertheless, as already said, having a broader search area is still an advantage when the predictor is modulated, since a smaller number of cross-correlation maps have to be re-computed.

As shown in Fig. 5.4, where the normalized processing time with respect to the time to compute the standard FFT interrogation without overlapping windows is plotted as a function of the number of vectors per linear dimension of the IV, the methods with pre-calculations along segments (1D DC) or planes (2D DC) perform better than the one with block cross-correlations only when the overlap is limited or very high. As expected, 1D DC is normally faster than 2D DC, since the redundancy of operations is minimized. The number of operations varies almost

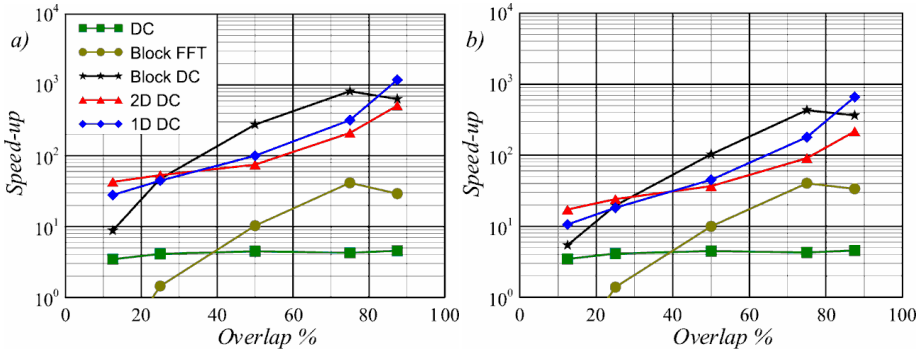


Fig. 5.3 Speed-up vs. overlap percentage for 1.25 (a) and 3.75 (b) particles/mm³.




Fig. 5.4 Normalized time (with respect to time to compute the standard FFT interrogation without overlapping windows) as a function of the number of vectors for linear dimension of the IV for the lowest tested seeding density.

linearly with the number of vectors to be computed for each linear dimension of the IV in the case of 1D DC, while for 2D DC the dependence is approximately quadratic.

5.2.2 Test case 2 - One-dimensional shear displacement

In this second case a volume with the same geometric features of the Test case 1 is built starting from four independent views (4 cameras are placed on a horizontal xz plane, with a uniform angular displacement of 15°). The magnification is approximately 0.135, and it is almost uniform throughout the volume. The imaging is simulated with a pixel pitch of 6.67 μm , and the $f\#$ is set to 13, so that the diameter of the particles is about 2.9 pixels. A custom-made software reconstructs the intensity distributions by 5 MART iterations, with an initial uniform first guest. The volume is discretized with 20 vox/mm , so that the resolution ratio is close to

| Step | IV size [vox] | Grid | Overlap | CC evaluation method |
|------|---------------|--------------------------|---------|----------------------|
| 1 | 64^3 | $24 \times 24 \times 5$ | 50% | FFT-Based |
| 2 | 40^3 | $39 \times 39 \times 9$ | 50% | FFT-Based |
| 3 | 40^3 | $77 \times 77 \times 17$ | 75% | FFT-Based |

Table 5.1 Example of standard process based on FFT.

| Step | Binning factor | IV size [vox] | Effective IV size [vox] | Grid | Overlap | CC evaluation method |
|------|----------------|---------------|-------------------------|--------------------------|---------|----------------------|
| 1 | 4x | 16^3 | 64^3 | $24 \times 24 \times 5$ | 50% | FFT-Based |
| 2 | 2x | 20^3 | 40^3 | $39 \times 39 \times 9$ | 50% | DC |
| 3 | 1x | 40^3 | 40^3 | $77 \times 77 \times 17$ | 75% | DC |

Table 5.2 Example of fast process with multi-resolution interrogation and direct cross-correlation on the corrector computation.

unity. Five different values of the source density (1.4), ranging between 0.1 and 0.5 (resulting in a percentage of non-zero voxels in the reconstructed distributions varying between 0.3% and 3.8% after the application of a threshold equal to 0.1% of the peak intensity of the generated particles), are tested. A one-dimensional displacement along the x direction is imposed, varying linearly between 0 and 20 voxels along the y direction.

The interrogation is executed in three steps, as reported in Tab. 5.1: initially IVs of 64^3 voxels with 50% overlap are employed; the grid is then refined to 40^3 voxels IV and 50% overlap; in the final iteration the overlap is increased to 75%. From now on, the process involving FFT in the evaluation of the cross-correlation map in all the three steps will be referred as "standard process". Direct sparse cross-correlations are executed for the other algorithms from the second step on. The structure of the algorithm in case of initial 4x binning is reported in Tab. 5.2 for example. A distinction is made between the size of the IV in the cross-correlation computation, and the effective one (say the equivalent IV on the original distributions). In case of failure of the peak identification in the corrector computation (i.e. the peak is not collocated within the search radius) the whole correlation map is computed again with the aid of FFT.

Firstly, the speed-up is quantified in the case of process without binning for the different approaches for the fast corrector computation. As illustrated in Fig. 5.5, the speed-up is consistently lower than the one obtained in the Test case 1, due to




Fig. 5.5 Speed-up as a function of the source density with direct cross-correlation from the second interrogation step on.




Fig. 5.6 Speed-up as a function of the source density with direct cross-correlation from the second interrogation step on, and multi-resolution predictor estimation: initial 2x binning (a) and 4x binning (b).

the necessary initial step performed by FFT. The slightly higher efficiency of 1D DC and 2D DC for the lowest seeding densities is due to a different exploitation of the sparsity of the distributions; e.g. segments of voxels with zero intensity are easily identified in the 1D DC method, and a certain number of useless multiplications are skipped, while an entire block with all zero intensity voxels is more difficult to be encountered; of course, this occurs more often at low seeding density. Increasing the seeding density, block direct cross-correlation performs better than the other methods, enabling a speed-up of about 18 times.

Adopting a multi-resolution predictor estimation can further accelerate the process: e.g. 2x and 4x binning are tested. The speed-up for the two different cases




Fig. 5.7 Standard deviation (in voxel) of the error as a function of the source density with and without 4x binning in the predictor estimation (right and left, respectively).

is quoted in Fig. 5.6 as a function of the source density. In case of 2x binning, only the first step is executed on binned distributions, while if 4x binning is employed, the binning factor is halved on the second iterations, and the original distributions are interrogated only in the final step (see Tab. 5.2 for details). For the lowest level of source density the small number of particles determines a higher probability of outliers occurrence in the predictor calculation, and, as a consequence, a higher percentage of correlation maps to be computed with the aid of FFT in the final iterations, reducing, as a matter of fact, the obtainable speed-up. In the case of $N_s = 0.1$, this percentage is 3.7% in step 2 and 2.7% in step 3 when a 4x binning is applied (in the case of 2x binning the percentage is 1.1% for step 2, and 0.5% in step 3). In all the other cases, the percentage is below 0.5% in the step 2, and negligible in the last step.

One should question whether the speed-up is costless in terms of accuracy or not. The obtained standard deviation of the error by using the standard process, or a process based on full (i.e. without thresholding) or sparse direct cross-correlations on the corrector estimation, is reported in Fig. 5.7. At low seeding density, direct correlations works better than FFT because the displacement is unbiased; increasing the mean number of particles per IV, this effect becomes less relevant. Sparse cross-correlation introduce a different source of error, mainly due the applied thresholding to increase the percentage of voxels with zero intensity; in case of predictor estimation with 4x binning, this aspect is more critical, i.e. too intense binning can compromise the accuracy or slow down the convergence of the interrogation process.




Fig. 5.8 Jet profiles for $\alpha = 2$ (a), $\alpha = 10$ (b). Exact profiles (black curve) are compared with profiles obtained by the fast process with Top Hat (TH) and Blackman (BL) weighting window in cross-correlation and dense predictor averaging.

5.2.3 Test case 3 - Circular jet

For this test case 26500 spherical Gaussian particles with 3 voxels diameter are distributed in a volume of $19.2 \times 19.2 \times 57.6 \text{ mm}^3$, again with a resolution of 20 vox/mm (i.e. $384 \times 384 \times 1152$ voxels, with approximately 1.5% of voxels with non-zero intensity, and about 17 particles per IV, being the final IV a cube of 48 voxels for each linear dimension). The displacement field of a circular jet directed along the z direction with a peak displacement W_{max} of 16 pixels is simulated. The used equation of the jet-like profile is:

$$w(x, y, z) = \frac{W_{max}}{2} \left[1 + \operatorname{atanh} \left(\alpha \left(\frac{1 - \sqrt{x^2 + y^2}}{r} \right) \right) \right] \quad (5.9)$$

where α is a parameter to modulate how steep is the descent from the peak velocity to zero and r is the station where the jet reaches a velocity that is equal to $W_{max}/2$ (set to 96 voxels).

The structure of the processing algorithm is reported in Tab. 5.3. Since the flow field is characterized by strong gradients (whose intensity is controlled by α), the modulation effects influence the predictor estimation, determining the corrector peak to be above the search radius of ± 0.5 pixels for a wide number of IVs. For this reason an hybrid method has been tested, i.e. the cross-correlation is evaluated using block FFT (BFFT from this point on) in the refinement steps, and one of the proposed solutions for the fast final corrector computation (for this test-case the 1D DC method is adopted; the lower efficiency of the method with respect to that of the block direct cross-correlations is of relative importance, since most of the processing time is related to the peaks falling above the search radius, i.e. the full correlation maps to be computed again with FFT; on the other hand, 1D DC enables to use rigorously weighting windows).

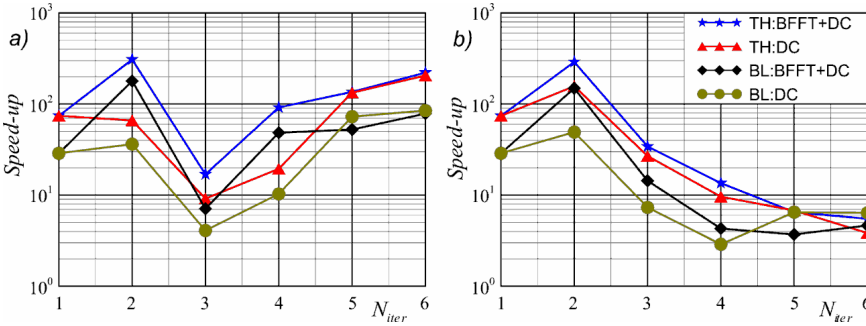


Fig. 5.9 Speed-up for each step as a function of the processing step for $\alpha = 2$ (a) and $\alpha = 10$ (b). TH and BL indicate Top Hat and Blackman weighting window for the cross-correlation map, respectively.

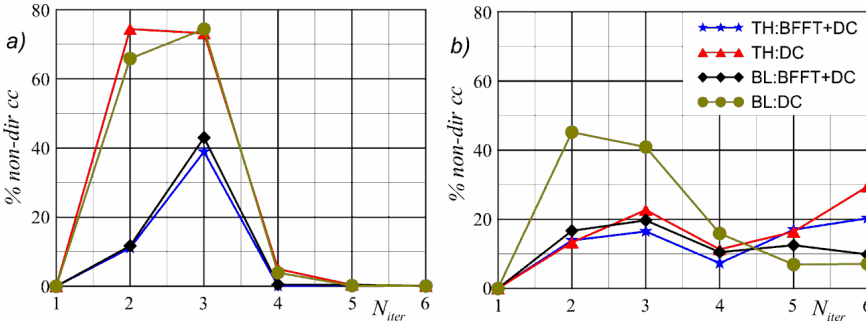


Fig. 5.10 Percentage of re-computed full cross-correlation maps with FFT for $\alpha = 2$ (a) and $\alpha = 10$ (b). TH and BL indicate Top Hat and Blackman weighting window for the cross-correlation map, respectively.

| Step [N_{iter}] | Binning factor | IV size [vox] | Effective IV size [vox] | Grid | Overlap | CC evaluation method |
|---------------------|----------------|-----------------|-------------------------|----------|---------|----------------------|
| 1 | 4x | 24 ³ | 96 ³ | 4x4x12 | 0% | FFT-Based |
| 2 | 4x | 24 ³ | 96 ³ | 7x7x23 | 50% | BFFT/DC |
| 3 | 2x | 24 ³ | 48 ³ | 15x15x47 | 50% | BFFT/DC |
| 4 | 1x | 48 ³ | 48 ³ | 29x29x93 | 75% | BFFT/DC |
| 5 | 1x | 48 ³ | 48 ³ | 29x29x93 | 75% | DC |
| 6 | 1x | 48 ³ | 48 ³ | 29x29x93 | 75% | DC |

Table 5.3 Processing algorithm: in the refinement section the CC evaluation method can be different from that of the iterations on the final grid. BFFT and DC stand for block FFT and 1D direct cross-correlation, respectively.

The velocity profiles of the jet-like velocity field are reported in Fig. 5.8. In the case of $\alpha = 2$ the results show that with a standard approach without using weighting windows (say, a top hat moving average, indicated with TH) the obtained profile is only slightly modulated with respect to the original one, the peak displacement being 2% smaller; the application of a Blackman (BL) weighting window in both cross-correlation and dense predictor averaging (over a window of the same size of the IV) reduces this modulation to less than 0.5%. When $\alpha = 10$ the modulation by the TH approach is much more evident; the profile obtained by using BL in the process is closer to the exact one.

The speed-up for each step (i.e. the reduction of the processing time relative to the time needed to execute the iteration with the full FFT approach) of the process is reported in Fig. 5.9. The first step is performed about 75 times faster thanks to 4x binning; when a Blackman weighting window is applied, the speed-up is of about 27 times with respect to the standard FFT approach. For the case of $\alpha = 2$, and no windowing applied, the approach based on block FFT during the grid refinement is consistently faster, since a broader search area increases the number of "fast-detected" peaks; this occurs due to modulation effects, placing the corrector displacement peak beyond the limit of the search radius of 0.5 pixels, as shown in Fig. 5.10, where the percentage of re-computed full cross-correlation map due to direct correlation failure is reported as a function of the iteration number (note that in the first iteration all the correlations are executed using the FFT, i.e. 100% of the correlation is non-direct; on the other hand, since no repetition of the calculation is performed, all the curves start fictitiously from 0 at the first iteration). Altogether, when no weighting windows are applied, the processing time is reduced by 18 and 11 times for the method employing BFFT in the first steps, and the other based only on direct cross-correlation with pre-calculation of sums along segments, respectively. When weighting windows are applied, the speed-up reduces to a factor of 12 and 9 for the two cases. However, one should consider that if the standard process with application of weighting windows in the computation of FFT is considered as a reference, the process with weighting window is accelerated by 27 and 21 times, respectively.

The case of $\alpha = 10$ is quite different, since the region affected by velocity gradients is smaller, but the slope of the velocity profile is much higher. In this case for all the 4 presented methods a not negligible percentage of non-direct cross correlations is still present even in the last iterations because of the modulation of the velocity profile; this reduces the maximum allowed speed-up of the cross-correlation algorithm to 6.4 and 5.9 times for the methods without the aid of weighting windows, with or without BFFT, respectively (the corresponding speed-up in case of filtering of the cross-correlation map with a Blackman weighting window is 3.9 for both the methods). One should note that the percentage of non-direct cross correlations for the method without using weighting windows increases with the number of iterations; this occurs because of aliasing. This aspect

is made clear in Fig. 5.10b. The approach with block FFT implies a higher percentage of correlation maps to be recomputed in the 5th step (when direct cross-correlation is employed) due to the effects of the imposed periodicity, reducing the accuracy of the estimated displacement field.

These issues are less critical if a Blackman weighting window is employed in computing the cross-correlation coefficients for all the proposed approaches. In this case, the frequency response is more similar to that of an ideal low pass filter (Astarita 2007), and no aliasing is present, so that the percentage of non direct cross-correlations decreases in the final iterations.

5.2.4 Experimental test case – swirling jet

The performances on a real test case are assessed by using the database of a tomographic PIV experiment performed in the Jet Tomography Facility (JTF) of TU Delft by Ianiro et al (2011). The flow field of a swirling jet at $Re = 1000$ and swirl number equal to 0.4 (see Ianiro & Cardone 2012 for definition) is investigated by using three Imager Pro HS 4M cameras observing the light scattered by $56\mu\text{m}$ polyamide particles dispersed in an octagonal water tank (80cm height, 60cm diameter). The light source is a Quantronix *Darwin-Duo* Nd:YLF laser, allowing $2x25\text{mJ}/\text{pulse}$ at 1kHz. Nikon objectives with $f = 105\text{mm}$ and $f_{\#}$ set to 32 are used; the optical magnification is approximately $M = 0.4$, with an average resolution of about $20\text{pix}/\text{mm}$. The particle image density is about 0.045ppp , resulting in a source density of 0.34. More details on the experimental apparatus can be found in Violato & Scarano (2011).

A volume self-calibration (Wieneke 2008) is executed to reduce the calibration error; 5 MART iterations reconstruct the light intensity distribution in the volume, discretized in $768 \times 768 \times 768$ voxels (i.e. a cube of 38.4mm for each side). Both the operations are performed by the *Davis 7.4* software.

The structure of the interrogation algorithm is summarized in Tab. 5.4-5.5 for the cross-correlation with top hat moving average, and with the aid of a Blackman weighting window (in both cases the dense predictor is averaged over a 48^3 voxels window). Since the flow field to be measured is challenging, different combinations of solutions are tested: block FFT in the refinement part of the process, and direct correlations in the final iterations (BFFT+DC); only direct correlation (DC) or block FFT (BFFT) during all the process, except for the predictor estimation. Generally speaking, one can choose between BFFT and DC for the different stages of the process. As shown in Tab. 5.4-5.5, it is proposed to use different solutions for the refinement steps and for the final iterations or the same solution for both the stages of the process. The initial part of the process is further accelerated by binning the distributions to be interrogated by a factor of 4 in the first two steps, and of 2 in the third one.

| Step | Binning factor | IV size [vox] | Effective IV size [vox] | N_{vect} | Overlap | CC evaluation method | Processing time for the standard FFT approach [s] |
|------|----------------|---------------|-------------------------|------------|---------|----------------------|---|
| 1 | 4x | 16^3 | 64^3 | 1.3k | 0% | FFT-Based | 2.6 |
| 2 | 4x | 16^3 | 64^3 | 12.2k | 50% | Method 1 | 20.9 |
| 3 | 2x | 32^3 | 64^3 | 103.8k | 75% | Method 1 | 178.8 |
| 4 | 1x | 48^3 | 48^3 | 250k | 75% | Method 1 | 238 |
| 5 | 1x | 48^3 | 48^3 | 250k | 75% | Method 2 | 239 |
| 6 | 1x | 48^3 | 48^3 | 250k | 75% | Method 2 | 239 |

Table 5.4 Processing algorithm: in the refinement section the CC evaluation method can be different from that of the iterations on the final grid.

| Step | Binning factor | IV size [vox] | Effective IV size [vox] | N_{vect} | Overlap | CC evaluation method | Processing time for the standard FFT approach with weighting windows [s] |
|------|----------------|---------------|-------------------------|------------|---------|----------------------|--|
| 1 | 4x | 24^3 | 96^3 | 1.3k | 33% | FFT-Based | 44 |
| 2 | 4x | 24^3 | 96^3 | 12.2k | 67% | Method 1 | 317 |
| 3 | 2x | 48^3 | 96^3 | 103.8k | 83% | Method 1 | 2662 |
| 4 | 1x | 72^3 | 72^3 | 250k | 83% | Method 1 | 2348 |
| 5 | 1x | 72^3 | 72^3 | 250k | 83% | Method 2 | 2349 |
| 6 | 1x | 72^3 | 72^3 | 250k | 83% | Method 2 | 2349 |

Table 5.5 Processing algorithm in the case of adoption of a Blackman weighting window in the cross-correlation step: in the refinement section the CC evaluation method can be different from that of the iterations on the final grid.

The results in terms of reduction of processing time are reported step by step in Fig. 5.11a. Binning the distributions by a factor of 4 provides an acceleration of about 63 times of the predictor estimation; the second step is even faster when BFFT is employed, while direct correlation is slowed down by a relevant percentage of non-direct correlations to be performed (see Fig. 5.11b). The obtained speed-up

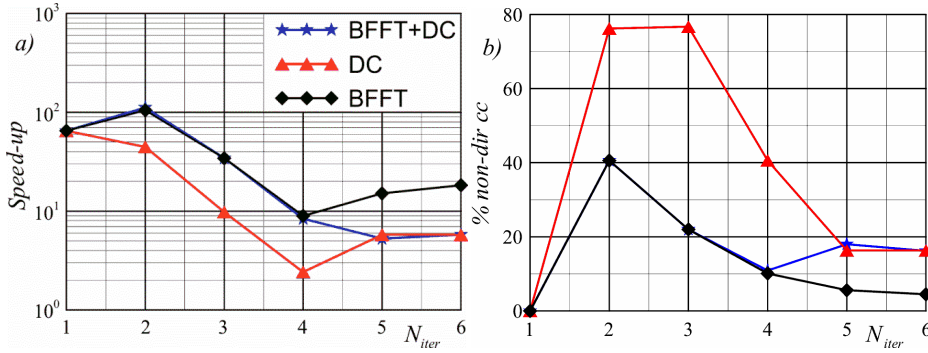


Fig. 5.11 Top hat moving average approach: **a)** Speed-up for each step as a function of the processing steps; **b)** percentage of non-direct correlations.




Fig. 5.12 Instantaneous vorticity magnitude isosurfaces ($0.15 \text{ voxels/voxel}$, in red) and of Q-criterion (in blue); planes of velocity vectors (color-coding indicates y-component). **a)** BFFT + DC **b)** BFFT **c)** BFFT + DC and **d)** standard process with filtering with Blackman window in the correlation step.




Fig. 5.13 Examples of profiles of the streamwise velocity component along the crosswise direction.

decreases in the refinement part as the binning factor is reduced up to unity. In the final iterations acceleration is recovered thanks to the results convergence on the final grid, reducing the percentage of non-direct correlations to be executed with the aid of FFT. In this sense, BFFT performs better than direct correlation thanks to a wider search area for the correlation peak. The overall speed-up is about 4.9, 6.6 and 11.8 times for DC, BFFT+DC and BFFT, respectively.

On the other hand, one should consider the accuracy and reliability of the results. The instantaneous vorticity fields reported in Fig. 5.12 testify that BFFT provides noisier results because of artefacts derived from the imposed periodicity on small IV, while DC on the final correction steps provides better convergence to slightly smoother results. However, the standard deviation of the difference between the two velocity field is rather small (0.07 pixels and 0.09 pixels for the u and v component, respectively).

An improvement in spatial resolution can be obtained adopting a Blackman weighting window in the cross-correlation step (see Fig. 5.12c). Of course, this increases the computational cost since more terms have to be computed to obtain the cross-correlation coefficients. However, if compared to the standard process based on FFT, the interrogation is executed 1.3 times faster with the approach BFFT+DC, with much better results. Considering that an "equivalent standard process", based on FFT weighted with a Blackman weighting window (Fig. 5.12d) is executed 10 times slower than the standard process with top hat filtering, this involves an acceleration of 13 times of the process BFFT+DC with respect to the respective standard one, with roughly the same results. The comparison between Fig. 5.12c and 5.12d points out that the proposed algorithm does not lead to a reduction of the accuracy; on the contrary, when DC is used in the last iterations of the process, the results are not affected by the artefacts due to the imposed periodicity when using the FFT. In these two last cases, the standard deviation of

the difference between the two velocity field is slightly larger than the previous comparison (0.09 pixel and 0.12 pixel for the u and v component, respectively), but still within the expected measurement accuracy.

In Fig. 5.13 the profiles of the streamwise velocity component along a straight line parallel to the x axis are reported for the same velocity fields illustrated in Fig. 5.12. The chosen profile placed is not exactly corresponding to the diameter of the nozzle, so that the modulation effects are clearly highlighted using peaks with different intensity and width. As expected, the process with windowing enables to have an improvement in the spatial resolution, being the profiles less affected by modulation issues.

5.2.5 Conclusions and guidelines

The results show that the processing time can be reduced by roughly two orders of magnitude; on the other hand, the interrogation acceleration is very sensitive to the complexity of the flow field and to the quality of the distributions to be interrogated, since most of the processing time is connected to the number of cross-correlation maps to be recomputed with the aid of FFT because of the failure in the peak identification within a narrow search area.

When dealing with the presented solutions to accelerate the process, the experimenter has to be aware that, in order to obtain the maximum speed-up, the process must be tuned with respect to the flow field under investigation. The results highlight the following points:

- When the overlap ranges between 25% and 75%, block cross-correlations provides the highest speed-up with respect to the standard process (even 800 times in the case of 100% of successful peak search within the area of 1 pixel); on the other hand, weighting windows cannot be rigorously applied in the cross-correlation step;
- The obtainable speed-up is flow field – dependent, in the sense that the peak search within a narrow area during the iterations can fail in regions with intense curvature of the velocity profile (e.g. shear layers). This aspect is particularly critical in the first steps of the process, when larger IV and smaller overlap values are used, causing a possible modulation of the velocity field. Generally speaking, no universal guideline can be given on the search radius choice for the peak detection; the application to the extremely complicate flow field of a swirling jet suggests that a search radius of 2 pixel suffices in determining an acceptably low rate of failure in the peak detection throughout all the process, retaining at the same time a relevant reduction of the computational cost;
- Direct correlation has always to be preferred to FFT in terms of accuracy. However, block FFT has shown to be very appealing in terms of computational cost reduction when the search radius has to be larger than 1

pixel. On the other hand, the experimental results suggest that performing the final iterations using direct correlation can recover the loss of accuracy due to the imposed periodicity artefacts.

5.3 Effects of poor discretization of the weighted cross-correlation

It has been proved that weighting windows can be very useful for a wide range of purposes: Nogueira et al (1999) introduced a weighting window to stabilize the interrogation process based on image deformation methods, and in subsequent papers (Nogueira et al 2001, Lecuona et al 2002) it is highlighted that scales smaller than the interrogation windows can be followed if correctly sampled by the particle motion; Astarita (2007) reviewed the performances of a wide range of weighting windows and developed a theoretical model to determine the MTF of the interrogation process when weighted dense predictor averaging is employed to stabilize the process; Astarita (2009) used the weighting windows feature to assess an adaptive PIV interrogation method; Novara et al (2013) made large use of spatially oriented Gaussian weighting kernels in the implementation of a 3D PIV adaptive interrogation algorithm.

However, the adoption of weighting highly-overlapping windows in cross-correlation computation introduces an unacceptable increase of processing time (see Sec. 5.1.3), urgently imposing to use the efficient “redundancy-free” algorithms presented in Sec. 5.1. In most of the applications the overlap ranges between 50% and 75%, so that the temptation of using the very efficient block cross-correlation can be compelling. On the other hand, in this case the weighting windows cannot be rigorously applied since they have to be replaced by a piecewise version, since each block can be weighted only by a single value.

The aim of this section is to outline that poor discretization of the weighting windows can lead to undesirable effects on the MTF of the process, even unexpected instability. The stability and the accuracy of the process are discussed using the theoretical model by Astarita (2007); eventually, the algorithm is tested on 3D simulated distributions of particles.

5.3.1 The theoretical model for the MTF estimation

The algorithm presented in Sec. 5.1 is considered as a reference in this section. By indicating with a the modulation associated with the step 1, with b the modulation due to predictor averaging (step 5) and with c the modulation due to interpolation (step 2), Astarita (2007) predicts the modulation at the generic k -th iteration as:

$$m^k = m^{k-1}bc + a(1 - m^{k-1}c) \quad (5.10)$$

Recalling that $m^0 = a$, Astarita (2007) introduces a stability criterion:

$$-1 < c(b - a) < 1 \quad (5.11)$$

Using (5.10)-(5.11) one can predict the MTF of the interrogation algorithm at the generic iteration k , and the range of wavelengths causing instability (if any). In the following large use of these instruments will be made.

5.3.2 Spatial resolution and stability of block weighting windows

In this section the MTF of block weighting windows is assessed. It will be assumed in the following that no modulation occurs due to interpolation (i.e. $c = 1$); this condition is closely met in case of small grid distance and high accuracy interpolation scheme, e.g. B-Splines of high order, see Astarita (2008). In section 5.3.2.1 the modulation of solely the cross-correlation step will be considered, i.e. the modulation in case of the so called local approach, in which the corrector is directly summed to the predictor, without any spatial averaging of the dense predictor. In section 5.3.2.2 the stability of the algorithm in case of weighted predictor averaging will be considered.

5.3.2.1 Stability of the process with local approach

When no dense predictor averaging is performed (i.e. $b = 1$, as proposed by Nogueira et al 1999), and no modulation occurs due to interpolation ($c = 1$), the relation (5.11) reduces to the requirement of positive a for a stable process.

The MTF of a triangular, a Blackman, a Gaussian window (with parameter $\alpha = 2$) and the weighting window proposed by Nogueira et al (1999) are compared for different values of the overlap percentage, i.e. number of blocks (obviously, the case of 50% overlap cannot be considered, since in this case every weighting window is reduced to a top hat). The MTFs are plotted in Fig. 5.14 as a function of the normalized frequency, i.e. the ratio between the IV size (indicated in this section with the letter W , and equal to 60 pixels for this investigation) and the wavelength λ .

The triangular function, and the weighting window proposed by Nogueira et al (1999), in their original versions, ensure the stability even with the local approach, since their MTF is definite semi-positive; the Blackman and the Gaussian window, on the other hand, are to be coupled with some dense predictor averaging, due to slightly negative values for high frequency, mainly due to discretization effects. However, when the discretization is very poor, as in block cross-correlation, even the triangular function and the function proposed by Nogueira et al (1999) present unexpected negative lobes in its MTF, causing instability in local approach application. The high frequency fluctuations are due to the abrupt variations of the weighting function between neighbouring blocks.

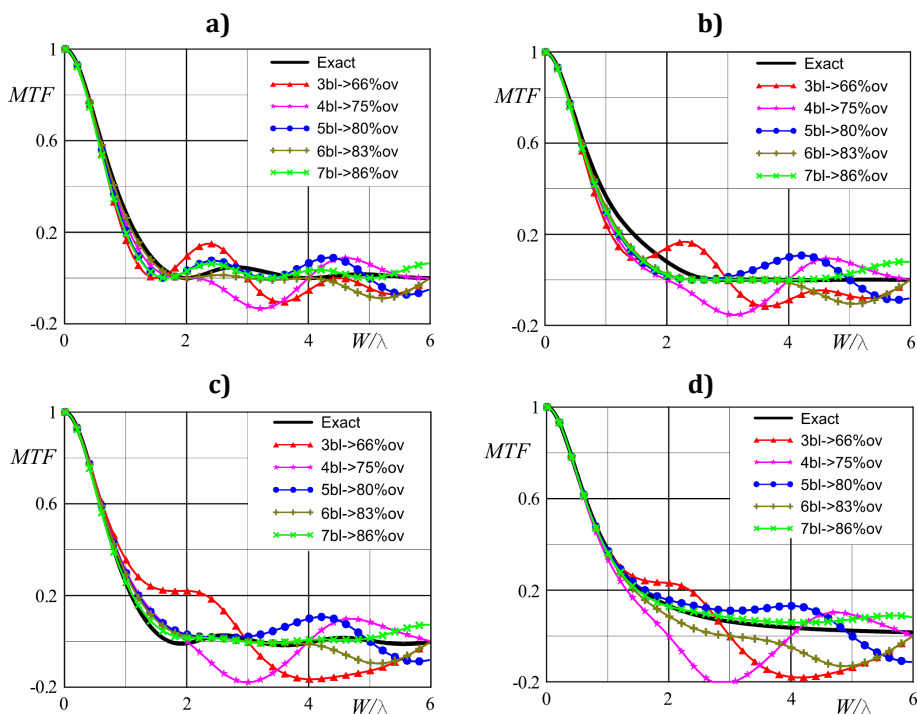


Fig. 5.14 Comparison of the MTF of weighting windows and their piecewise version for several overlap values (nbl stands for n blocks): **a)** triangular; **b)** Blackman; **c)** Gaussian, $\alpha = 2$; **d)** Nogueira et al (1999).

Otherwise stated, the poor discretized weighting functions can be considered as sum of properly weighted top hat windows, contributing with their own frequency response in building the MTF of the window. E.g., in case of 3 blocks discretization, the weighting window can be considered as a sum of a pedestal, with the height equal to that of the two external blocks, and a central top hat, one third of the total window wide (see Fig. 5.1). The MTF of the 3-blocks window is the sum of two aliased *sinc* functions: the first one, due to the pedestal, disappearing very quickly as the normalized frequency increases, the second one due to the small central kernel, introducing high frequency fluctuations.

In case of 3 and 4 blocks the performances are very similar, with wide oscillations and negative lobes: this is due to the fact that, when using an even number of blocks, the two central blocks are weighted with the same value (i.e. as a matter of fact the window is composed by only $n - 1$ blocks, the central one being twice as large as the other blocks). In all the reported cases using 80% overlap (i.e. 5 blocks) reduces the intensity of the fluctuations of the MTF, and determines negative lobes for wavelengths smaller than 5 times the IV size. Acceptably small fluctuations, however, are reached only when 7 blocks or more are used (i.e. about

86% overlap). Even in that case, the local approach is unstable; the stability of the process can be obtained by using weighted dense predictor averaging, i.e. $b \neq 1$.

5.3.2.2 Stability of the process with weighted predictor averaging

Differently from the cross-correlation step, in step 5 the weighting windows can be applied rigorously (the computational cost of dense predictor summing is very small if compared to the other steps of the process, so that there is no reasonable incentive to use block summing). The MTF, in this case, can be obtained by using (5.10), where a is the modulation due to piecewise weighting window, b is the modulation of the original window with the proper discretization, and c is again set to 1.

Recalling (5.11), in case of negative modulation in the cross-correlation step (some negative lobes have been observed in all the cases due to poor discretization) the process can be stabilized by a small positive or negative b . Different results can be obtained by using different windows in the dense predictor averaging; however, Astarita (2007) demonstrated that the main characteristics of the MTF are lead by the window chosen for the cross-correlation step. For this reason, from this moment on only the top hat window will be used for the dense predictor averaging, without any loss of generality. In Fig. 5.15 the MTFs for the case of the Gaussian window with $\alpha = 2$ after 3 and 100 iterations are reported. Three main discretization levels are investigated, i.e. 3, 4 and 5 blocks. The MTF of the original window is reported for comparison in Fig. 5.16.

Astarita (2007) observed that all the investigated weighting windows are stable when a predictor average with a top hat window of size $W_b \geq 3$ is applied. This is not the case of poorly discretized weighting window. A summary of the minimum size of the predictor average window in order to have a stable process is reported in Tab. 5.6; however it is important to point out that these results may be consistently higher if a different weighting window is used to average the dense predictor. The process is always stable when the size of the dense predictor averaging window is more than 7 pixels wide when a top hat moving average is applied.

The MTFs of the block cross-correlation approach show already undesired features after only 3 iterations: whichever is the chosen size for the predictor averaging window, all the curves present a wide negative lobe for $W/\lambda \geq 3$ and $2 \leq W/\lambda \leq 4$ in the case of 3 and 4 blocks discretization, respectively. In the case of 5 blocks discretization the negative lobe affects only very small scales, and it is of moderate intensity. Increasing the number of iterations, the curves relative to $W_b = 3$ and $W_b = 5$ rapidly diverge for the 3 and 4 blocks discretizations, being the process unstable. In the case of 5 blocks discretization, even if the process is stable

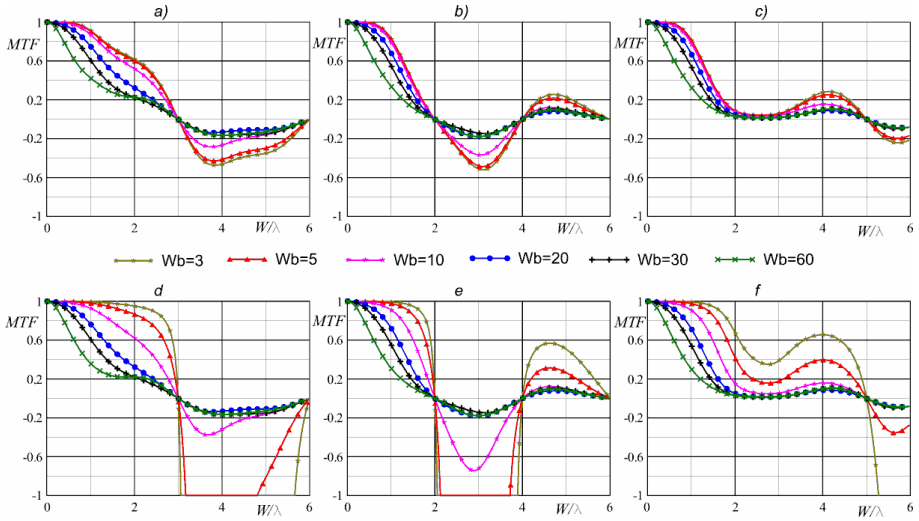


Fig. 5.15 MTF of a Gaussian window ($\alpha = 2$) as a function of the normalized spatial frequency for different size W_b of the window over which the predictor is averaged. Results reported after 3 iterations (**a,b,c**) and 100 iterations (**d,e,f**) in the case of 3, 4 and 5 blocks from the left to the right.

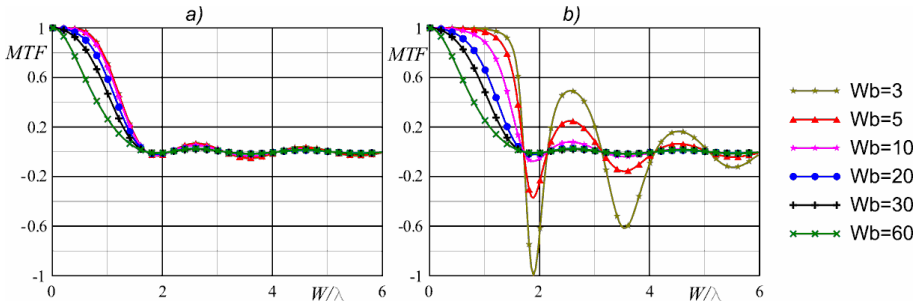


Fig. 5.16 MTF of the original Gaussian window ($\alpha = 2$) as a function of the normalized spatial frequency for different size W_b of the window over which the predictor is averaged. Results reported after 3 iterations (**a**) and 100 iterations (**b**).

| | 3 bl | 4 bl | 5 bl | 6 bl |
|---------------------------|------|------|------|------|
| Gaussian ($\alpha = 2$) | 5 | 7 | 3 | 3 |
| Blackman | 5 | 6 | 3 | 4 |
| Triangular | 5 | 6 | 3 | 3 |

Table 5.6 Minimum size (pixel) of the dense predictor averaging window W_b in order to have a stable process.

for $W_b = 3$, the MTF has a negative peak of about -2.7 at $W/\lambda \approx 5.6$. However, one should note that $W_b = 3$ is not advisable also for the original window, since the MTF presents wide oscillations, with a negative peak of approximately -1 at $W/\lambda \approx 1.9$.

When $W_b > 20$, the MTF is substantially unchanged in all the cases, i.e. a good convergence is reached after few iterations. The case of moderately small W_b (i.e. $W_b = 10$) is slightly different. While for the original window the MTF is monotonically decreasing for small normalized frequency, and then slightly oscillating around zero with decreasing amplitude, in the case of 3 and 4 blocks discretization wide negative lobes are present, with non negligible intensity. The situation is much worse in the case of 4 blocks discretization, being the negative peak equal to about -0.8 at $W/\lambda \approx 1.8$. For the 3 blocks discretization the negative peak is much less intense (approximately -0.4). On the other hand, when 5 or more blocks are used and $W_b = 10$, the region of negative MTF is practically the same of the cases with $W_b \geq 20$. This behavior is quite independent of the chosen original window.

5.3.3 Performance assessment by 3D simulations

The performances of the process with the block weighted cross-correlation are assessed by using virtually generated distributions of particles. Two different layouts are investigated: zero-displacement on a $240 \times 240 \times 240$ voxels volume, and one-dimensional sinusoidal displacement field on a $480 \times 480 \times 120$ voxels (the displacement is along the x direction, while the gradient is imposed on the y direction, with wavelength variable between 600 and 40 voxels). Particles with Gaussian shape, 200 counts of peak intensity and e^{-2} diameter of approximately 3 voxels are randomly distributed within the volumes; a Gaussian noise with a mean of 5 counts and a standard deviation of 2 counts is superimposed to the intensity distributions. In both cases the particle concentration is $6.5 \cdot 10^{-4}$ particles per voxel, i.e. approximately 41.6 particles in an interrogation region of $40 \times 40 \times 40$ voxels. The large number of particles ensures that the investigated wavelengths in the second test layout are correctly sampled. The size of the interrogation window is 60 voxels (accordingly, in the case of the overlap of 66%, 75% and 80% the size of each block is 20, 15 and 12 voxels, respectively).

The accuracy is reported in terms of total error δ :

$$\delta = \sqrt{\frac{1}{N} \sum_{i=1}^N (u_i - u)^2} \quad (5.12)$$

where N is the number of computed vectors, u_i is the i -th measured vector, and u is the correct displacement value. The first layout is particularly suitable to measure the accuracy of the process when a high accuracy interpolation scheme is used in step 2 of the algorithm (in the present work the velocity field is interpolated with a




Fig. 5.17 Total error as a function for the number of iterations. Left: 4 blocks discretization (75% overlap) and different size of the dense predictor averaging window. Right: $W_b = 6$ and variable overlap.

B-Spline of 10^{th} degree), since the bias error is very small and it is not influenced by the noise level (Astarita 2006); in this case the average total error is practically coincident with the error at zero displacement.

As in Sec. 5.3.2, a Gaussian filtering window with parameter $\alpha = 2$ is adopted in the cross-correlation, and a top hat moving window with variable size W_b is used (2, 6, 10, 20, 30 and 60 voxels).

The spatial resolution is assessed in terms of the MTF (4.2).

5.3.3.1 Precision

In Fig. 5.17 the total error is reported as a function of the number of iterations to assess the convergence of the process (for clarity the symbols are reported only each 5 iterations). In Fig. 5.17, left, the results in the case of 75% overlap (i.e. 4 blocks discretization) show that in the case of $W_b \geq 20$ the process converges quite rapidly (the average total error is substantially unchanged after 5 iterations). The cases with $W_b = 2$ and 6 are within the range of instability of the method; as expected the error diverges, even if the divergence is quite slow in the second case (the modulation in the step 2 of the process due to the relatively large grid distance of 15 voxels can significantly restrict the instability range). In the case of $W_b = 10$ the error is still increasing after 20 iterations, even if the rate of increase of the error is relatively slow.

In Fig. 5.17, right, the total error is reported for $W_b = 6$ for the three different discretizations of (3, 4, 5 blocks) and for the exact discretization over the size of the window (in this last case, the process is theoretically stable, Astarita 2007). The rate of increase of the error is consistently higher in the case of the process with block-discretized weighting windows in the cross-correlation step; the error is

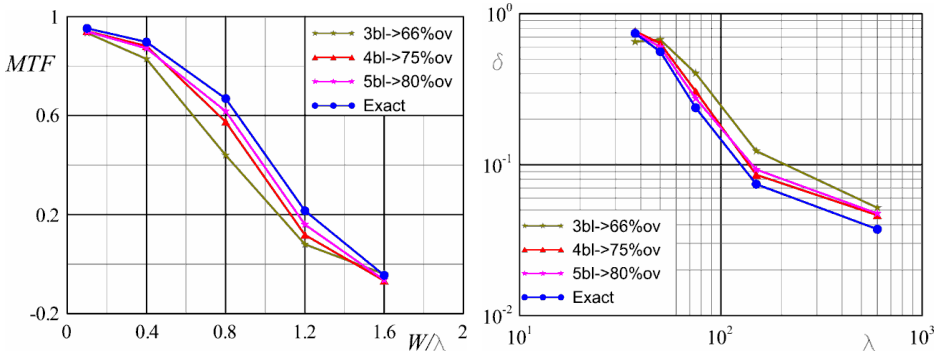


Fig. 5.18 Modulation Transfer Function (left) and total error (right) as a function of the normalized frequency for $W_b = 6$.

about 2 times higher with respect to the process with correct application of the weighting window after only 4 iterations, while it is more than 2.5 times higher after 15 iterations.

5.3.3.2 Spatial resolution

The MTF for the case of $W_b = 6$ is reported in Fig. 5.18, left, for different levels of discretization of the weighting windows. One should note that the measured MTF is much lower than the theoretical estimate reported in Sec. 5.3.2. The main reason for this discrepancy is the sensitivity to the noise effects of (4.2); the higher values of the theoretical MTF in the case of the 3-blocks discretization at higher frequency determine a stronger contamination of the MTF due to the noise (i.e. the measured MTF is consistently lower than the one relative to different levels of discretization). In addition, one should note that the theoretical estimated MTF does not take into account that the particles are randomly distributed, and they may not be able to properly sample the signal in the entire volume. However, this aspect is supposed to determine similar degrading effects to the MTFs of all methods, regardless of the discretization level of the weighting window.

The total error, reported in Fig. 5.18, right, as a function of the wavelength, is determined both by the presence of noise and by the modulation of the velocity profile. This latter effect is dominating, and determines a higher error for the process with 3 blocks discretization, except for the smaller tested wavelength (in this case the MTF for the process with 3 blocks discretization is the highest one). In the case of 4 and 5 blocks discretization, the performances are very similar; however, since the results are evaluated after only 20 iterations, and the process with 4 blocks discretization in this configuration is unstable (for $c = 1$), the difference between the two methods can be asymptotically larger.

5.3.4 Conclusions and guidelines

The results clearly underline that a poor discretization of the weighting window can determine a consistently different behaviour in terms of impulsive response of the algorithm, determining instability even when it is not theoretically expected. The local approach is never recommended, unless the number of blocks used for the discretization is consistently high (at least 10 blocks, i.e. 90% overlap); on the other hand, in this case, the block cross-correlation has been proven to provide a lower speed-up with respect to other approaches (pre-calculation of the contribution to the cross-correlation coefficients along segments or planes, instead of blocks), in which the weighting windows can be applied correctly.

On top of this, for a proper stabilization the process with block-discretization requires larger windows for the dense predictor average with respect to the case of the exact discretization of the weighting window over the whole size of the interrogation volume.

The virtual simulations highlight that the poor discretization of the weighting windows determines a higher sensitivity to random noise, i.e. a larger total error. This is due to the strong fluctuation of the MTF at high frequency, generated by the abrupt variations within the weighting window. In summary, in the application of block weighted cross-correlation two rules of thumbs can be proposed:

- When it is possible, *use a odd number of blocks*: this is because, when using an even number of blocks, the central blocks are equally weighted, determining in the frequency response a larger negative MTF for frequencies ranging between $0.5N_{blocks} \leq W/\lambda \leq N_{blocks}$ (the two central blocks of the weighting function can be considered a top hat filter with size equal to twice the size of each block);
- Use wide dense predictor averaging windows (for the top hat approach $W_b \geq 10$ is suggested when the discretization is performed on 5 or more blocks).

Chapter 6 - Low-cost Tomo-PIV systems

Tomo-PIV requires a significant effort in terms of cost of the experimental setup, since at least 3-4 cameras are required for a successful tomographic reconstruction. In particular double-shutter cameras for PIV applications are much more expensive than common cameras. However, the requirement of double-shutter cameras, which is very strict in planar PIV due to the need to acquire two separate images from the same point of view for the image cross-correlation, ceases to exist in a Tomo-PIV setup (in principle it can be eliminated also for planar PIV at the expense of the simplicity of the experimental setup and procedure, using two cameras and the plane identification method proposed by Discetti & Adrian 2012). In fact the 3D light intensity distributions can be reconstructed by observing the two exposures with independent imaging systems. Thus, while at least three double-shutter (or high speed) cameras are required in the standard (or high speed) Tomo-PIV implementation, the approach proposed herein requires the use of two imaging systems of three (or more) cheaper single-shutter cameras. The cost of the experimental setup can be strongly reduced since typically a single-shutter camera can be up to 10 times cheaper than a double-shutter camera with the same imaging quality (i.e. similar pixel size and quantum efficiency). Some guidelines on the possible arrangements of the proposed experimental setup are provided in Sec. 6.1.

Moreover in this scenario the ghost particles are not expected to coherently contribute (and hence, produce bias errors, see Sec. 2.2.2) to the cross-correlation map, since their distribution depends on the arrangement of the imaging system. In other words, the same particles pattern will produce approximately the same number of ghost particles (Sec 2.2.1), but their locations in the two exposures are uncorrelated, as explained in Sec. 6.1. In this framework it is reasonable to assume that the effectiveness of the Motion Tracking Enhanced MART (MTE-MART, Novara et al 2010) is maximized, independently of the flow field features.

A numerical analysis of the spatial resolution (in terms of Modulation Transfer Function MTF) is presented in Sec. 6.2 using 2D numerical simulations. 3D simulations of synthetic experiments with a jet-like flow field with and without velocity gradients in the depth direction are described in Sec 6.3 in order to assess the effect of MTE-MART.




Fig 6.1 Schematic of imaging systems for a jet-like flow field: **a)** Traditional Tomo PIV setup with 3 double-shutter cameras (in red and blue); **b)** Low Cost Tomo-PIV setup with 3 single-frame cameras for the first exposure (in red) and 3 single-frame cameras for the second exposure (in blue); **c)** Low Cost Tomo-PIV setup with 2 double-shutter cameras (in red and blue) and 2 single-frame cameras (for the first and second exposure in red and blue respectively).

6.1 Proposed experimental setup

In Tomo-PIV, the set of cameras to be used for the imaging of each exposure of the illuminated volume (i.e. the 2D projections for the volume reconstruction) is not required to be the same. If two different imaging systems are used for the reconstruction of the first and the second exposure, the required number of cameras is doubled. This introduces a complication on the experimental setup (more optical accesses, Scheimpflug camera mounts, optics, etc. are needed); on the other hand, one can purchase much cheaper single-shutter cameras, since the requirement of double-shutter cameras does not hold anymore.

In particular the simplest idea for the experimental setup is the use of 2 systems of three single-shutter cameras for the reconstruction of the first and the second exposure (see e.g. Fig. 6.1 where a schematic of possible imaging systems for a jet-like flow field is presented). Another possible implementation consists of two double-shutter cameras (available in all the labs in which stereoscopic PIV experiments are already performed) together with two single-shutter cameras, obtaining two imaging systems composed by the two double-shutter cameras and by the first and the second single-shutter camera for the first and the second exposure, respectively.

6.1.1 Suppression of the bias effect of the coherent ghost particles motion

The phenomenon of the ghost particles production and their effect on the measured flow field for the traditional Tomo-PIV setup and for the proposed low-cost system are sketched in the schematic of Fig. 6.2 (for the simplified case of 1D projections with two cameras). The actual and ghost particles are represented as filled and empty circles, respectively. Red and blue colours are used for the first and

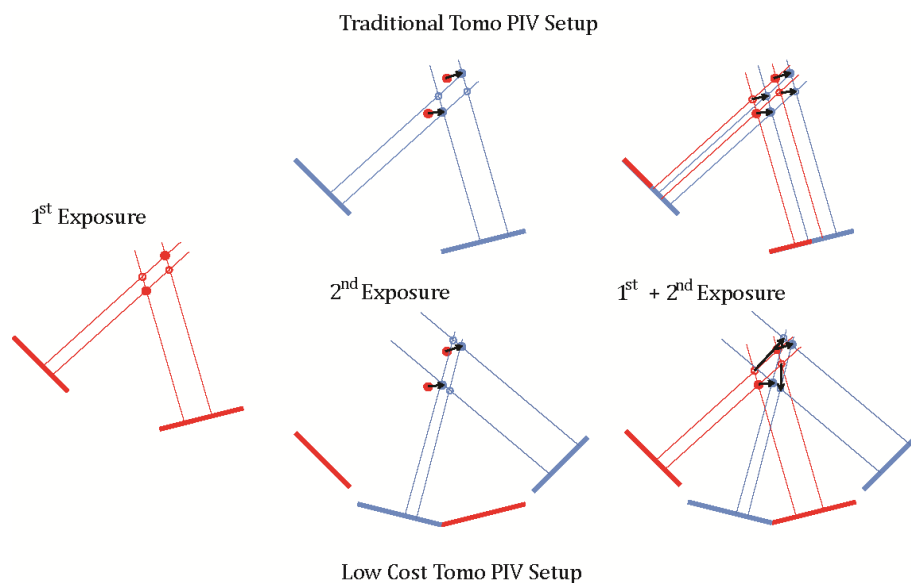


Fig. 6.2 Schematic of the ghost particles formation and motion in the traditional and in the Low Cost Tomo PIV setup.

the second exposure. The schematic highlights that in the traditional Tomo PIV setup, as described by Elsinga et al (2011), ghost particles formed from the same set of actual particles in both the reconstructed volumes are characterized by a resulting displacement that is approximately the average displacement of the set of the associated actual particles (see Sec. 2.2.2 for further details), causing a modulation of the velocity gradients. It has to be pointed out that the coherent motion of the ghost particles determines a non negligible contribution in building higher (but, unfortunately, biased) cross-correlation peaks. This aspect is particularly dangerous, as it might result in misleading interpretation on the quality of the measurement. In the proposed Low Cost Tomo-PIV setup (from now on referred as with the acronym LC, i.e. Low Cost) the ghost particles formed in the two subsequent exposures by the same set of actual particles are expected to randomly contribute to the cross-correlation maps. In fact, as shown in Fig. 6.2 (bottom), the number of ghost particles is the same of the standard Tomographic PIV system, as it depends mainly on the number of cameras and actual particles; however, their locations in the two exposures are completely uncorrelated with the actual particles displacement field. Accordingly, the bias effect is completely removed, at the expense of a reduction of the signal-to-noise ratio.

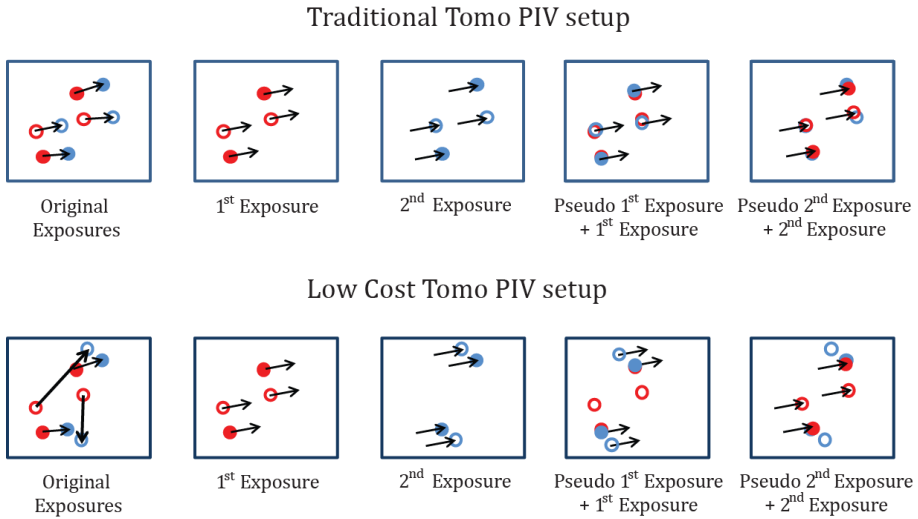


Fig. 6.3 Schematic of the working principle of MTE in case of weak velocity gradients along the depth direction for the traditional TomoPIV setup and the LC system.

As a matter of fact, in this framework the ghost particles contribution to the cross-correlation maps can be imagined as an out-of-plane motion (without any implication on perspective effects), determining a loss of pairs between the two exposures. A more complete model for the prediction of the number and intensity of the ghost particles, in this scenario, would potentially lead to an analytical predictability of the performance of the 3D cross-correlation interrogation via introduction of a novel loss of pairs factor in the analysis conducted by Keane & Adrian (1992).

6.1.2 Accuracy improvement via MTE-MART

The non-coherent ghost particles motion can be exploited for the suppression of the reconstruction artefacts. The MTE is particularly well suited to perform this task, since it is designed to eliminate inconsistent features of subsequent exposures. The working principle of MTE and its high effectiveness for the low cost Tomo-PIV setup are highlighted in Fig. 6.3, where the same particle pattern and camera configuration of Fig. 6.2 are used. In principle, a first guess for the distributions is obtained using, for example, the MART procedure. A rough displacement field is obtained by cross-correlating the first guess distributions. Subsequently, a pseudo-1st exposure is built by deforming the 2nd exposure according to the displacement field interpolated on each pixel (the same procedure is applied on the other exposure). The pseudo 1st exposure is summed to the original one; in this step, the actual particles are supposed to be strengthened,

while the ghost particles are not if, and only if, they do not move coherently in the two exposures.

In the standard implementation of Tomo PIV the coherent motion of ghost particles strongly limits the effectiveness of MTE. In fact Novara & Scarano (2012a) show that, when the velocity gradients are limited, only a moderate reduction in ghost intensity is expected by MTE. This aspect is clearly evident in the top part of Fig. 6.3. In the case of the traditional Tomo PIV setup, in fact, by summing the pseudo 1st exposure to the 1st exposure the ghost particles have nearly the same displacement of the actual particles and the resulting guess for the next MART iterations still contains the ghost particles unabated. On the other hand, in the case of the LC setup, the MTE is evidently effective independently of the flow features (i.e. the velocity gradients along the cameras viewing direction). Of course, the execution of the MTE algorithm implies an increase of the computational cost; however, one can easily reduce the processing time of the 3D PIV processing algorithm of one or two orders of magnitude with the solutions presented in Chapter 5.

6.2 Numerical simulations

In this section the performances of the LC setup are assessed in terms of spatial resolution and potential improvement on the accuracy with the MTE-MART. The spatial resolution is assessed similarly to the approach in Sec. 4.2.2 imposing a sinusoidal displacement on 2D distribution of particles reconstructed with 1D cameras. The accuracy of the method and the test on the effectiveness of the MTE-MART for the LC system is tested in the 3D scenario with synthetic particles distributions and the jet-like displacement field modelled by (5.9).

6.2.1 Spatial resolution

The volume considered for the present analysis is a slice of $100 \times 12 \text{ mm}^2$ discretized with $20 \text{ vox}/\text{mm}$. The synthetic imaging system for the traditional Tomo-PIV arrangement is composed by three cameras oriented at -45° , 0° and 45° . The proposed LC system is composed by 3+3 cameras oriented at -45° , -10° and 30° for the first exposure and -30° , 10° and 45° for the second exposure. The results are compared also with the case of a setup composed by six double-shutter cameras in the same overall arrangement of the LC system. Gaussian particles with a diameter of 3 voxels and maximum particle intensity of 200 counts are randomly distributed in the slice to be reconstructed. A sinusoidal displacement field, with $U = \sin(2\pi z/\lambda)$ and $V = 0$ is imposed; the wavelength λ is varied between 24 and 240 pixels. The present simulation is performed at three levels of source density N_s (namely 0.15, 0.25 and 0.35), defined as (1.4); the only difference is that the area is

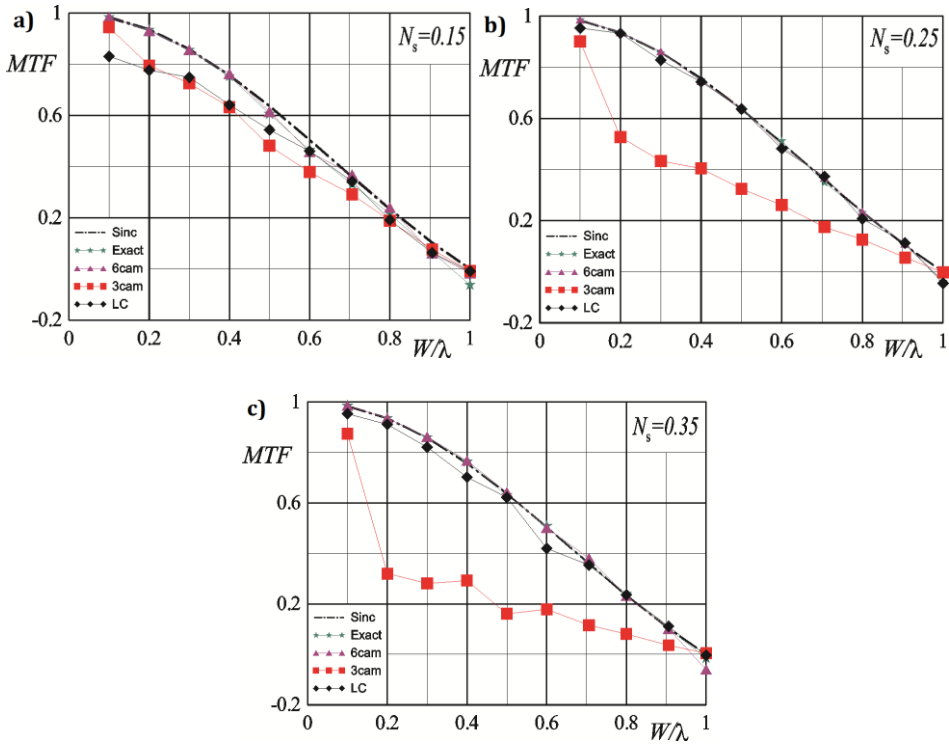


Fig 6.4 Modulation Transfer Function plotted against the normalized frequency for the 6cam, 3cam and LC systems: **a)** $N_s = 0.15$ **b)** $N_s = 0.25$ **c)** $N_s = 0.35$. The MTF of the process on the exact distributions and the sinc function are included for reference.

not circular, but it is the area occupied on a linear sensor (the term $\pi d_t^2/4$ has to be replaced by simply d_t). The Modulation Transfer Function is evaluated with the (4.2) on the detected flow field averaged over 500 images. The resolution ratio between pixels and voxels is set equal to 1; the reconstruction is performed by 5 MART iterations with a relaxation coefficient of 1.

Since the flow field is one-dimensional, the PIV process can be performed with interrogation windows elongated along the x -direction (in the present case 500×24 voxels with 75% overlap in the gradient direction) in order to guarantee a sufficient number of particles in every interrogation window.

The theoretical MTF of a standard PIV approach without applying weighting windows (and for the commonly adopted values of the normalized frequency defined as the ratio of the interrogation window size W along the depth direction and the wavelength λ) is that of a top hat filter, i.e. a *sinc* function. Thus the *sinc* profile is used as a reference for the MTF reported in Fig. 6.4.

For $N_s = 0.15$ (Fig. 6.4a) no significant differences are found between the proposed methodology and the 3 cameras (*3cam*) system, while the MTF relative to

the 6 cameras (*6cam*) system is practically the same of the one obtained by interrogating the original images (indicated as *exact* in the legend). One should note that for large wavelengths the 3cam system performs better than the LC one, since it benefits of the ghost particles contribution to the cross-correlation (the bias effect is less important, due to the relatively low gradient along the depth direction), reducing the measurement noise. It is important to point out that the equation (4.2) is sensitive to noise, i.e. a MTF lower than 1 is estimated in presence of noise even if the mean profile is not modulated; when the number of actual particles per interrogation window is relatively small, the noise induced by the uncorrelated ghost particles leads to a much larger random error, determining a lower MTF for the LC system. By increasing the sine wavelength, the bias effect due to the ghost particles motion modulates the gradient in the case of the 3cam system, while the LC system, as expected, appears to be less affected. At larger source densities (see for example $N_s = 0.25$, Fig. 6.4b) the MTF of the LC system strongly benefits of the higher number of true particles, and it is comparable with that of the 6cam system, while the MTF of the traditional 3cam system is consistently lower (the higher is the source density, the higher is the number of ghost particles; accordingly the bias effects are more intense); at source density $N_s = 0.35$ (Fig. 6.4c) this effect is significantly more evident.

In Fig. 6.5a-c the averaged flow fields and the correlation maps relative to a 500×100 pixels region in the centre of the slice are presented for the condition with $\lambda = 248$ voxels and $N_s = 0.35$. The LC system is able to obtain the same mean flow field obtained by the 6cam system, even if the slightly lower MTF testifies an increase of the measurement noise on the single realization. The presence of the noise is exemplified in the correlation maps. The LC system produces a normalized correlation peak of 0.22 corresponding to a signal to noise ratio (defined as the ratio of the first and the second tallest peaks in the cross-correlation map) of 2.8, which is much lower with respect to the signal to noise ratio achieved by the 6cam system (approximately 8.2) and by the 3cam system (about 5.6). It has to be remarked that the 3cam system, despite of the high signal to noise ratio (which might lead to a misleading interpretation on the quality of the results), is affected by bias effects due to the correlation of the ghost particles (see Fig. 6.5, middle-left). On the other hand, the noise level for the LC system can be easily reduced by applying the MTE MART, as discussed in the Sec. 6.1.2 and illustrated in Sec. 6.2.2.

Interestingly enough, the results of Figs. 6.4-6.5 reveal that *the quality of the reconstruction does not tell everything about the accuracy and the spatial resolution of the measured velocity field*. In fact, the quality of the reconstructed slices for the 3cam and the LC system is exactly the same for the two exposures (about 0.81, 0.59 and 0.47 for the three tested levels of source density), but the final results are dramatically different. In other words, the statement on the quality factor " $Q > 0.75$ is a good indicator for the quality of the results" is flawed.

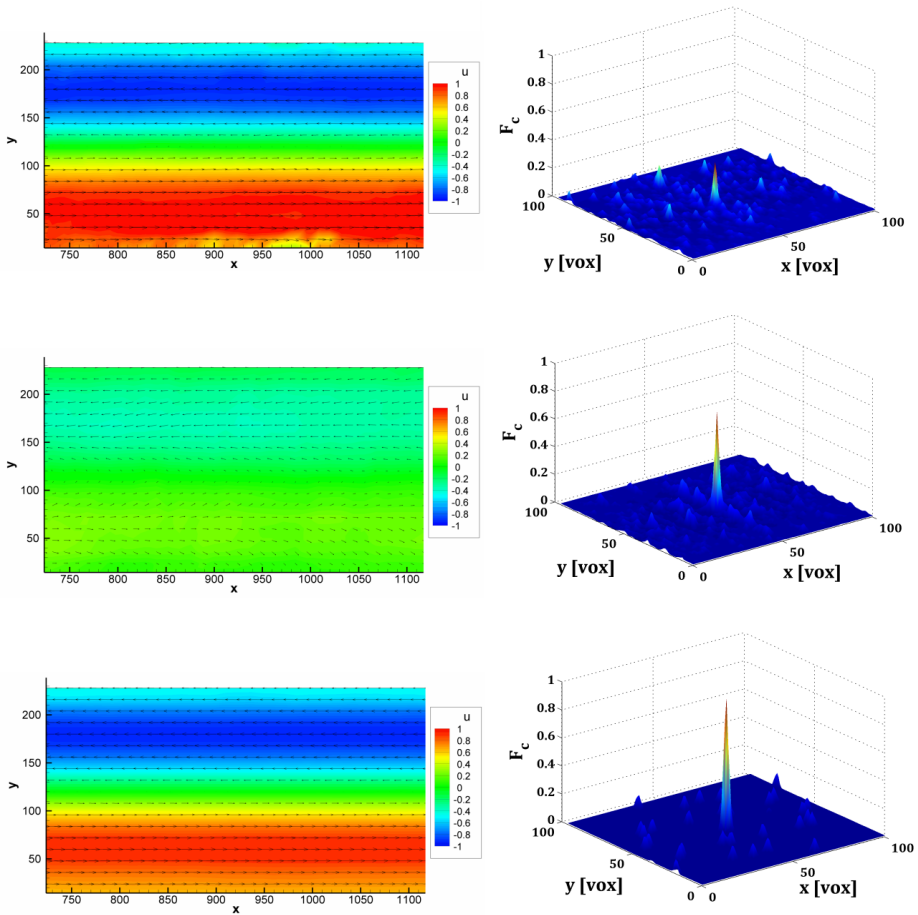


Fig 6.5 Evaluated flow field and example of correlation map at $\lambda=248$ pixels and $N_s=0.35$: LC (top), 3cam (middle), 6cam (bottom).

6.2.2 Measurement noise and MTE

In this sub-section numerical simulations on 3D synthetic particle distributions reconstructed by 2D cameras are presented in order to analyze the effect of the proposed imaging setup on the measurement noise and the possibility to improve the quality of the results with MTE-MART. The reconstructed region for the present analysis is a volume of $40 \times 40 \times 15 \text{ mm}^3$ with a spatial discretization of 20 vox/mm ($800 \times 800 \times 300$ voxels).

The synthetic imaging system for the reference Tomo-PIV system is composed by 6 cameras, disposed on two horizontal groups (as in Fig. 6.1b); on each horizontal group the cameras are angularly equally spaced by 30° and the two horizontal systems describe an angle of 60° , so that, in total, an angle of 60° is

enclosed by the outer cameras in both directions. From now on, the two horizontal groups will be referred as rails on which the cameras are placed (this terminology actually resembles the possible practical implementation of the system).

The performances of three systems are compared: 6 double-shutter cameras system; three double-shutter cameras (2 outer cameras of one rail and one on the centre of the other rail as in Fig. 6.1a) and a LC system with 2 imaging systems similar to that used for the 3 double-shutter cameras (i.e. the imaging systems for the first and the second exposure describe two triangles, with the vertices on the opposite sides as in Fig. 6.1b).

Gaussian particles with 3 voxels diameter and a maximum particle intensity of 200 counts are randomly distributed in the volume to be reconstructed. A jet-like displacement field is imposed, according to (5.9). Two cases are discussed, i.e. a jet-like profile directed along the z and y directions. In the first case there is no velocity gradient along the depth direction z , thus the MTE is expected to be nearly ineffective in the traditional 3cam layout, conversely to the case of the proposed LC system; in the second case the gradient along the depth direction determines the MTE to be highly effective for both systems.

The present simulation is performed at a source density of 0.24 corresponding to 0.035 $\rho\rho\rho$. The resolution ratio between pixels and voxels is set equal to 1. The first MTE step is applied after 5 iterations; subsequently, a MTE step is applied after each 3 MART iterations.

The PIV process is executed using 40 x 40 x 40 voxels interrogation windows with 75% overlap. Since the aim of the present simulations is to estimate the overall accuracy of the proposed method no ensemble averaging is performed.

The flow fields relative to a jet aligned with the z direction (without velocity gradients in the main camera viewing direction) are presented in Fig. 6.6. Both the flow fields obtained with the LC and with the 3cam systems (Fig 6.6a and Fig. 6.6c, respectively) are strongly affected by the measurement noise; however, the 3cam system exhibits also a modulation of the velocity field, which is not present in the case of the LC system. This aspect is testified by the scatter plot of the error along the x and the z direction reported in Fig. 6.7b. The elongated asymmetric shape of the scatter plot for the 3cam system is due to the modulation of the velocity profile.

With the application of the MTE, the quality factor of both the LC and 3cam systems increases significantly (Fig. 6.7a); in this case the MTE is less effective for the 3cam system, as expected. The qualitative inspection of Fig. 6.6b and 6.6e reveals that the measurement quality of the LC system is comparable to that of the 6cam system. Despite of a very high quality factor (around 0.86 after 3 MTE steps, against 0.91 achieved by 6cam), the LC system still exhibits some marginal noise effects, probably due to the fact that the particles are reconstructed with different shapes in the two reconstructed objects. This aspect is more evident in Figs. 6.7c and 6.7d, where the root mean square of the error for the u and w velocity

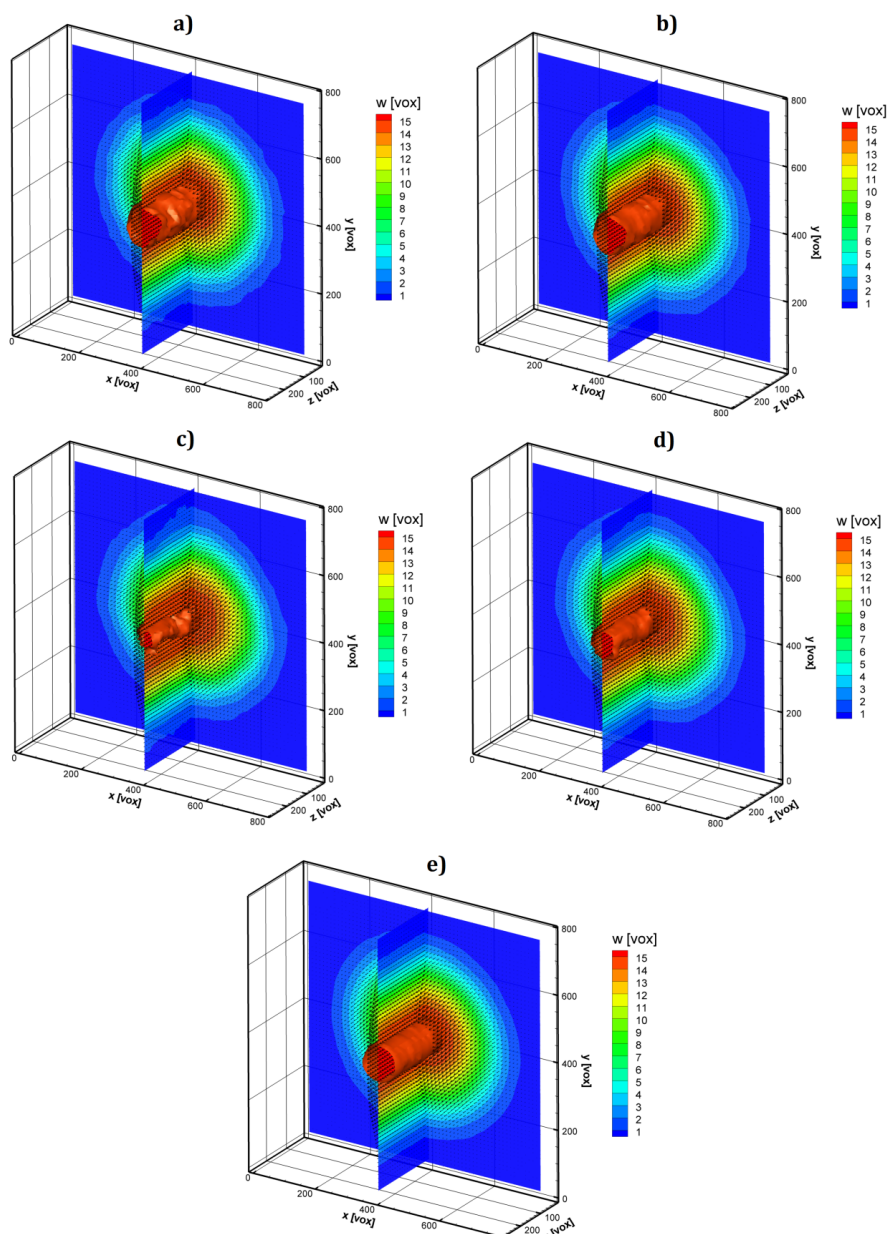


Fig 6.6 Numerical simulations on a jet aligned with the z direction – iso-surface ($w=15$ voxels) and contour with vector arrows of the displacement along the z direction: **a)** LC **b)** LC with 3 MTE steps **c)** 3cam **d)** 3cam with 3 MTE steps **e)** 6cam.

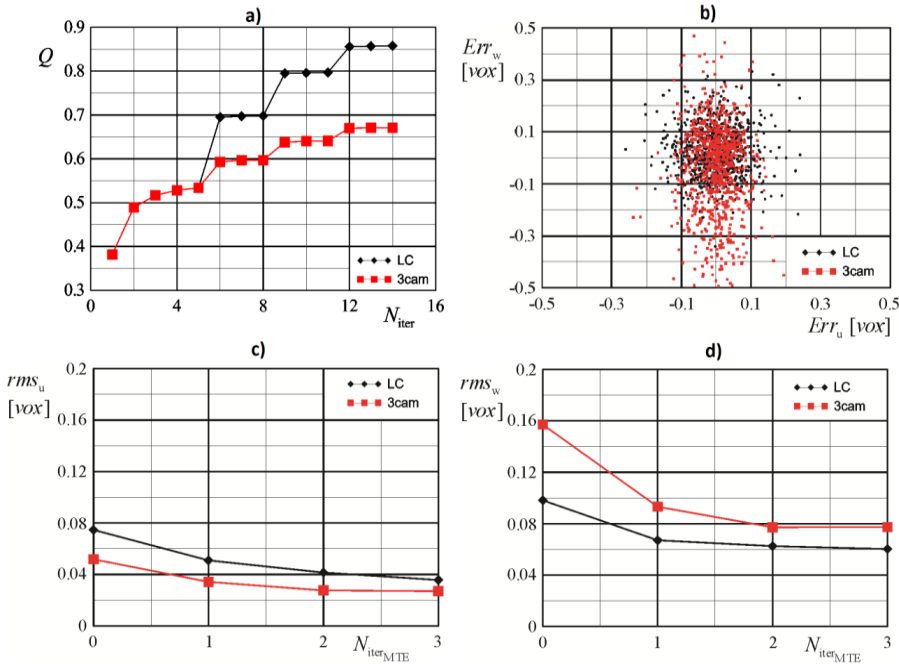


Fig 6.7 Numerical simulations on a jet aligned with the z direction – Comparison between the 3cam and the LC system: **a)** Quality factor Q versus number of MART iterations N_{iter} ; **b)** Scatter plot of the velocity error in the x and z direction; **c)** root mean square of the error on the u velocity component versus the number of MTE steps; **d)** root mean square of the error on the w velocity component versus the number of MTE steps.

components are plotted versus the number of MTE iterations for the LC and for the 3cam system (u, v, w are the velocity components along x, y, z , respectively). In the case of the u component the LC system exhibits a higher level of noise with respect to the 3cam system; on the other hand, if the w component is considered, the root mean square of the error rms_w in the velocity estimation for the 3cam system is much larger, as it is characterized by the bias effect of the coherent ghost particles motion.

The flow fields relative to a jet aligned with the y direction (with velocity gradients in the camera viewing direction) are presented in Fig. 6.8. Also in this case both the flow fields obtained with the LC and with the 3cam systems (Fig. 6.8a and Fig. 6.8c, respectively) are characterized by a relatively strong measurement noise. Furthermore, in agreement with Elsinga et al (2011), the 3cam system exhibits a much larger modulation. In this case, however, the application of the MTE determines a significant improvement of the quality factor of both the LC and 3cam systems; the presence of a velocity gradient along the depth direction increases the

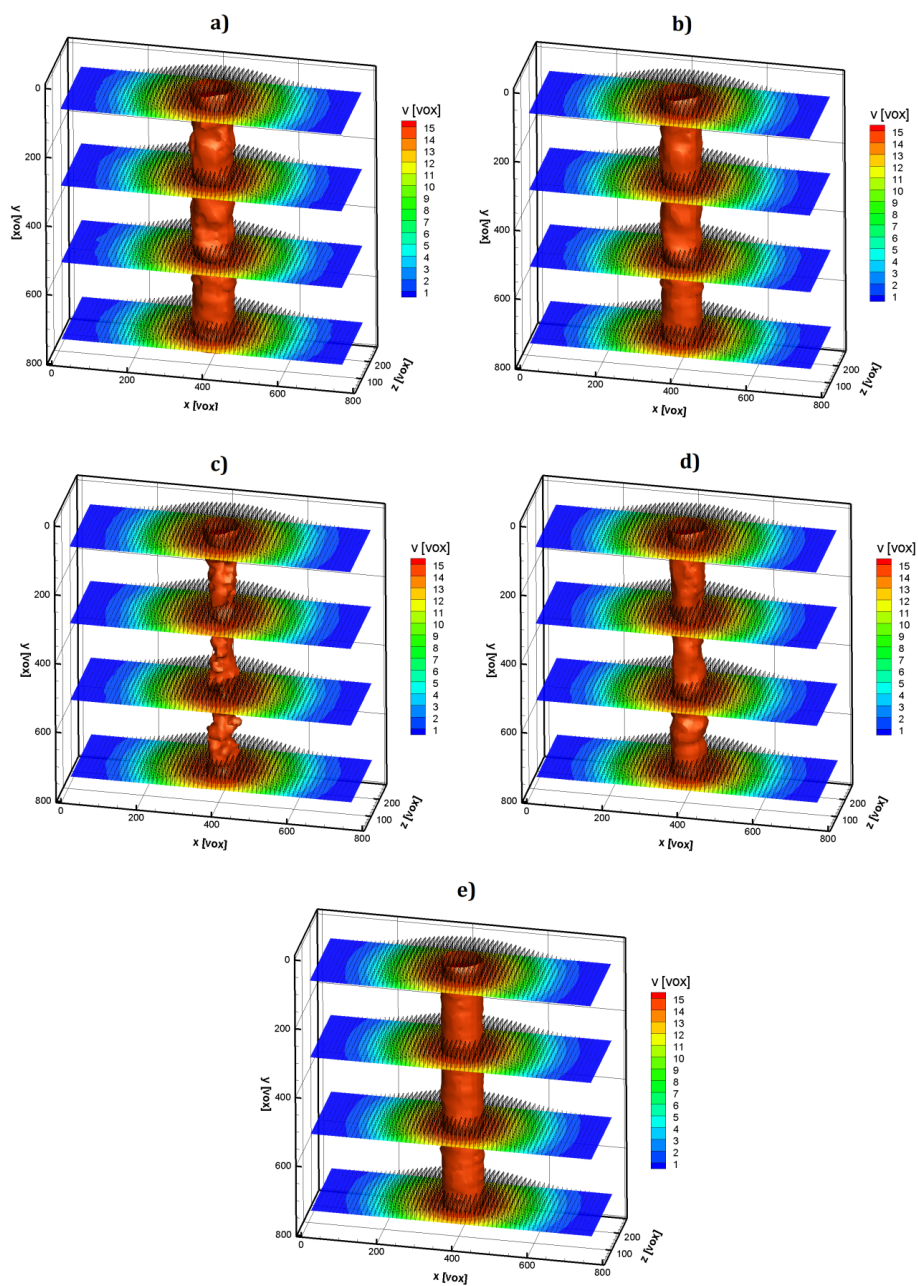


Fig. 6.8 Numerical simulations on a jet aligned with the y direction – iso-surface ($v=15$ voxels) and contour with vector arrows of the displacement along the y direction: **a)** LC **b)** LC with 3 MTE steps **c)** 3cam **d)** 3cam with 3 MTE steps **e)** 6cam.

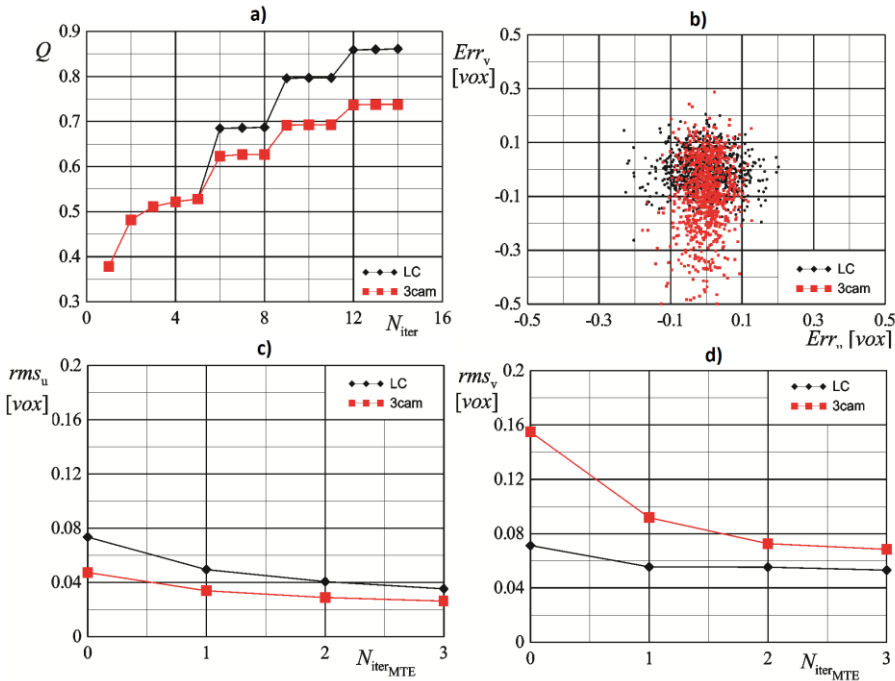


Fig 6.9 Numerical simulations on a jet aligned with the y direction – Comparison between the 3cam and the LC system: **a)** Quality factor Q versus number of MART iterations N_{iter} ; **b)** Scatter plot of the velocity error in the x and y direction; **c)** root mean square of the error on the u velocity component versus the number of MTE steps; **d)** root mean square of the error on the v velocity component versus the number of MTE steps.

effectiveness of MTE also in the traditional Tomo PIV configuration (compare, for instance, the quality achieved by the 3cam system in Fig. 6.7a and 6.9a). However, the MTE proves to be still more effective for the LC with respect to the 3cam system: due to the strong modulation of the velocity field evaluated with the 3cam system, the MTE MART produces a particles distribution that appears filtered in the depth direction. The MTE MART is able to reduce the initial modulation of the flow field but the final velocity field evaluated is still affected by a significant bias. Also in these conditions, despite of a very high quality factor, the LC system still exhibits some noise artefacts.

The measurement noise is quantified in Figs. 6.9c and 6.9d, where the root mean square of the error in the u and in the v velocity component is plotted versus the number of MTE iterations for the LC and for the 3cam system. For the u component the LC system exhibits again a higher level of noise with respect to the 3cam system; indeed, the exact u component is identically zero in the entire measurement domain, so the error is only due to random noise, and there are no bias effects. Again, the root mean square of the error on the v component for the

3cam system is larger, as it is characterized by the bias due to the coherent motion of the ghost particles (see, for instance, the asymmetric shape of the scatter plot in Fig. 6.9b).

6.3 Summary and conclusions

The proposed low-cost experimental setup does not compromise the accuracy of the measurements as confirmed by the parametric study on synthetic images. In fact, according to the 2D simulations, the effect of modulation due to ghost particles is strongly reduced. On the contrary, the measurement noise is found to be slightly larger, due the different reconstructed particles shape between the first and the second exposure and due to the random effect of the ghost particles on the correlation maps.

The spatial resolution of the LC system is comparable to that of a traditional tomographic PIV system with the same total number of cameras, even though the quality of the reconstruction for the single exposures is much lower. The inconsistent features of the subsequent exposures can be easily eliminated using the MTE-MART, which is highly effective regardless of the flow field features in this scenario. Furthermore, MTE-MART is expected to be more reliable, since the estimated velocity field for the volume deformation is less affected by modulation effects than the case of the traditional Tomographic PIV system.

Chapter 7 - The decay of fractal-generated turbulence

In turbulence theory very few are the properties believed to hold universally. The independence of the turbulent kinetic energy dissipation rate ε of the kinematic viscosity ν in the limit of infinite Reynolds number Re certainly belongs to this very exclusive class. This hypothesis is one of the most relevant consequences of Kolmogorov's *universal equilibrium theory* of small scales structure (Kolmogorov 1941). The assumption is based on the observation that as the Re increases (i.e. inertia effects become prevalent on viscous effects) the universal tendency of every flow is to develop smaller and smaller structures, i.e. sharp gradients, keeping the turbulent kinetic energy dissipation at a constant value. In the generic picture of turbulence, it is thought that inertia forces have the average tendency to spread the turbulent energy up the spectral pipeline from the low-wavenumber energy-containing scales to the high-wavenumber dissipative scales (Kolmogorov 1941). In this scenario, since the turnover time of the small scales is much smaller than that of the large scales (the two characteristic times scale as $Re^{-1/2}$), it is reasonable to assume that the motion of the small scales is statistically independent of the relatively slow motion of the large scale turbulence and of the mean flow. This leads to the anomalous idea that the small scales motion is determined only by the rate of energy transferred by the large scales, and by viscosity; furthermore, the dissipation rate can be assumed equal to the rate of energy transferred from the large scales, leading to the hypothesis that ε depends entirely on the large scales turbulence and is independent of viscosity. Taylor (1935) proposed the scaling $\varepsilon = C_\varepsilon k^{3/2}/L$, where k is the turbulent kinetic energy, L is an integral lengthscale, and C_ε is a constant of proportionality, which value is different for different classes of flows (Burattini et al 2005, Mazellier & Vassilicos 2008), provided that the Re number is sufficiently high. This scaling is often referred as *dissipation anomaly*.

Tennekes & Lumley (1972) define the dissipation anomaly as a cornerstone assumption of turbulence theory. It is central in many turbulence phenomenologies and modelling (Batchelor 1953, Townsend 1956, Pope 2000), such as one-point closures (as the widely used $k - \varepsilon$ model) or two-point closures (as in large eddy simulation). Strong proofs of its validity have been provided over the last decades: Sreenivasan (1984) observed in wind tunnel turbulence generated by bi-plane square grids that C_ε is a constant for $Re_\lambda = \sqrt{k}\lambda/\nu > 50$ (where λ is the Taylor microscale and Re_λ is a Taylor lengthscale-based Reynolds number); direct numerical simulation (DNS) data confirm the constancy of C_ε , perhaps with a slightly larger minimum Re_λ (Sreenivasan 1998, Burattini et al 2005). Evidences

against the dissipation anomaly have been reported in the past decades (Bevilaqua & Lykoudis 1978, Tong & Warhaft 1994, Lavoie et al 2005), mainly addressed by the persistence of coherent structures at all flow scales, determining long-remembered effects of inlet and boundary conditions. However, the first dramatic departures from the constancy of C_ε in experimental measurements are those reported by Seoud & Vassilicos (2007) and Mazellier & Vassilicos (2010). They observed a significant variation of C_ε in the turbulence generated in wind tunnels by fractal square grids (i.e. grids with a square pattern repeated at different scales), following approximately $1/Re_\lambda$. Furthermore a dependence of C_ε on a global Re determined by the inlet conditions has been detected. Subsequently the analysis by Valente & Vassilicos (2012) on the far-wake of regular grids highlighted that $C_\varepsilon \approx Re_M^{1/2}/Re_\lambda$ (where Re_M is the mesh-length M based Re number, $Re_M = U_\infty M/\nu$, with U_∞ being the inlet velocity), thus implying a more general behaviour, not limited to the case of the square fractal grids.

The importance of the empirical law $C_\varepsilon \approx const$ in most of the models and the theories of homogeneous and inhomogeneous turbulence, and such a dramatic violation, raise new questions to be addressed and new accurate and reliable models to be developed. So far, the evidences of violation of the dissipation anomaly have been detected by pointwise hot-wire anemometry measurements, characterized by very high spatial resolution (though relying on Taylor's frozen turbulence hypothesis) but based on small-scale isotropy assumptions to estimate the dissipation. However, if Kolmogorov's universal equilibrium theory is questioned, a true complete dissipation measurement (i.e. estimation of the full velocity gradient tensor) seems more suitable to perform the task. From this point of view Tomographic Particle Image Velocimetry appears to be the most promising technique, providing full 3D3C velocity field measurements. On the other hand, Tomo-PIV suffers of limited spatial resolution, mainly related to the maximum allowed particles concentration within the volume, thus raising questions about its possible applications in the measurement of small-scale turbulence.

In this chapter an experimental investigation of the freely decaying turbulence generated by fractal grids with square pattern is described. The objective is to assess the capability of Tomo-PIV in a scenario in which hot-wire anemometry is certainly the workhorse. The 3D3C measurement does provide an instrument to widely investigate homogeneity and isotropy; furthermore, the availability of all the 9 components of the velocity gradient tensor enables the possibility to perform a real measurement of the turbulent dissipation without invoking the small scale isotropy assumption.

The general features and the state of the art in the research of the wake of fractal objects are firstly reviewed in Sec. 7.1. The experimental setup and the Tomo-PIV system are described in detail in Sec. 7.2. Tomo-PIV is then applied in




Fig. 7.1 From left to right, cross, I and square fractal grids (Hurst & Vassilicos 2007).

Sec. 7.3 to extract information on the homogeneity and the isotropy, on the turbulent lengthscales, and on the turbulent dissipation. In this last case, the performances of a set of methods are addressed.

7.1 Wake of fractal objects

A number of turbulence theories relate the vortex cascade process with an underlying fractal/multifractal structure (see the works of Novikov 1971, Mandelbrot 1974 and 1982, Sreenivasan & Meneveau 1986, Gouldin 1987). However, the research on the turbulence generated by fractal elements has started only recently. Queiros-Conde & Vassilicos (2001) and Staicu et al (2003) measured the turbulence statistics in the wake of fractal tree-like generators. Though their investigations were not conclusive in distinguishing the effects of the finite size of the fractal tree and their self-similar structure, interestingly enough they pointed out how turbulence with much higher intensity can be generated by using a fractal stirrer. This stimulated on one side the development of fractal forcing techniques for numerical simulations (Mazzi et al 2002, Mazzi & Vassilicos 2004), while on the other side it led to the first systematic investigation of the wind tunnel turbulence generated by fractal grids (Hurst & Vassilicos 2007). Hurst & Vassilicos tested a total of 21 planar fractal grids from 3 different families: fractal cross grids, fractal I grids and fractal square grids (see Fig. 7.1). The interest in the first two families of fractal grids in the following years has been quite limited. Geipel et al (2010) used fractal cross grids to improve the turbulent mixing in opposed jet flows; Kinzel et al (2011) applied the same type of grid to increase the turbulence intensity and the local Re in shear-free turbulence under the influence of system rotation. Krogstad & Davidson (2012) and Krogstad (2012) investigated the near field decay of fractal cross grids by hot-wire anemometry and Laser Doppler Anemometry (LDA), and compared the decay rates and the degree of homogeneity with the turbulence generated by conventional grids.

The turbulence generated by planar fractal I grids has not been studied in deep after the survey by Hurst & Vassilicos (2007). On the other hand, the wake of similar fractal planar and non-planar fractal I trees has been investigated by Chester et al (2007) and Chester & Meneveau (2007), mainly aiming to a better understanding of the interaction of turbulent flows with boundaries characterized by multiple lengthscales using the Renormalized Numerical Simulation technique. Fractal trees with similar features have been investigated also by Bai et al (2012) with Particle Image Velocimetry (PIV). In their investigation the simplicity of the description of the multiscale nature of fractal trees well fits with the problem of modelling the multiple scales of sparse vegetation canopies.

Far more interest has been devoted on the family of square fractal grids (better described in Sec. 7.1.1). In this last case, Hurst & Vassilicos (2007) documented very unconventional and interesting properties, outlined in the Secs. 7.1.2-5.

7.1.1 Square space-filling fractal grids

The geometry of the square fractal grids (Fig. 7.2) is completely defined by the following parameters:

- number of iterations N_f ;
- the length L_0 and thickness t_0 of the bars of the first iteration;
- the ratio of the length and the thickness of subsequent iterations R_L and R_t (the length and the thickness at the j -th iteration are defined by the relations $L_j = R_L^j L_0$ and $t_j = R_t^j t_0$; most often the thickness ratio t_r of the first and the last iteration is used instead of R_t);
- The number B_j of patterns at each iteration.

In their pioneering work Hurst & Vassilicos (2007) used space-filling fractal grids, i.e. with fractal dimension $D_f = \log(B)/\log(1/R_L) = 2$ ($B = 4$ in the case of square fractal grids, R_L is set equal to $1/2$).

Hurst & Vassilicos (2007) proposed the following relations (holding in the case of square test section, with linear size T) to estimate the blockage ratio σ_B and effective meshlength M_{eff} :

$$\sigma_B = \frac{L_0 t_0 \sum_{j=0}^{N-1} 4^{j+1} R_L^j R_t^j - t_0^2 \sum_{j=0}^{N-1} 2^{2j+1} R_t^{2j-1}}{T^2} \quad (7.1)$$

$$M_{eff} = \frac{4T^2}{P} \sqrt{1 - \sigma_B^2} \quad (7.2)$$

where P is the fractal perimeter length of the grid. It is worth noting that, differently from the classical grids, fractal grids do not have a definite mesh size. Hurst & Vassilicos (2007) pointed out that, while it is intuitive to grasp the equivalent meshlength as proportional to the area of the test section T^2 divided by the fractal perimeter, such a definition would not take into account the effect of the blockage




Fig. 7.2 Geometry of a space-filling square fractal grid.

ratio. For this reason they introduced in (7.2) the corrective term $\sqrt{1 - \sigma_B^2}$ (in the case of classical rectangular grids, M_{eff} returns the meshlength of the equivalent regular grid).

7.1.2 The wake interaction lengthscale

From this point on, the letters x and y will indicate the longitudinal streamwise direction and the vertical crosswise direction; the z axis completes the left-handed Cartesian reference system. The relative velocity components are referred with the symbols U, V, W , respectively. The angled brackets $\langle \dots \rangle$ indicate the operation of ensemble averaging. The fluctuating velocity components, obtained with a Reynolds decomposition (i.e. subtraction of the respective time-averaged velocity components) are referred with the lower case letters u, v, w . The symbols u', v', w' indicate the root mean square of the turbulent fluctuations (i.e. $u' = \sqrt{\langle u^2 \rangle}$, etc).

The turbulence intensity increases in a protracted production region, peaks at a streamwise distance x_{peak} related to the geometrical features of the grid, and then decays with an unusually fast rate that could be fitted with an exponential curve instead of the well-documented power law decay for decaying homogeneous isotropic turbulence. Hurst & Vassilicos (2007) conjectured $x_{peak} = 75t_{min}T/L_{min}$, where t_{min} and L_{min} are the thickness and the length of the smallest square iteration, respectively. Most importantly, they pointed out that, since the position of the peak of the turbulence intensity can be tuned by changing the geometrical features of the grid, a complete understanding of the scaling of x_{peak} is of fundamental importance, especially in terms of potential application in combustion and industrial mixing.




Fig. 7.3 Scheme of the wakes interaction in the lee of fractal grids (Mazellier & Vassilicos 2010).

Subsequently, Mazellier & Vassilicos (2010) proposed a different scaling based on the interaction of the wakes of the different bars. Differently from the case of classical rectangular grids, the bars of the grid have different size and spacing. As a consequence, their wakes interact at different streamwise locations. Following Townsend (1956), the typical wake width l of the iteration j scales as $\approx \sqrt{t_j x}$, where x is the abscissa in the streamwise direction. Supposing that the scaling does hold with reasonable approximation in the case of non-isolated wakes, for each iteration one can calculate a characteristic first-interaction abscissa, such that $L_j \approx \sqrt{t_j x_j^*}$, i.e. $x_j^* \approx L_j^2/t_j$. A scheme of the wake interaction for a square fractal grid (SFG) is reported in Fig. 7.3.

In the case of the fractal grids, the last wakes to interact are those generated by the largest iteration of the grid. For this reason Mazellier & Vassilicos (2010) introduced the *wake-interaction lengthscale*:

$$x^* = \frac{L_0^2}{t_0} \quad (7.3)$$

The wake-interaction lengthscale can be defined also for regular grids (RGs), where L_0 is replaced by the meshlength. The data collected by Jayesh & Warhaft (1992) in the region $1 \leq x/M \leq 30$ (where M is the grid meshlength) highlight the existence of a highly inhomogeneous turbulence in the near-field ($x/M < 3$, corresponding to approximately $x_{peak}/x^* \approx 0.55$, only slightly larger than $x_{peak} \approx 0.45x^*$, typical of SFGs), where production is dominant, with a peak of turbulence intensity and the subsequent well-know power law decay. The ratio x_{peak}/x^* is slightly larger for RGs than for SFGs. Valente & Vassilicos (2011) argue that in the case of SFGs (and in general FGs) the presence of a set of wake-

interaction lengthscales is the responsible for this discrepancy: when the largest scales start to interact, the smallest scales have already merged upstream, determining a more consistent turbulent diffusion and grow rate of the largest wakes.

7.1.3 Quasi-isotropic and homogeneous “high” Re_λ turbulence

Grid turbulence has been widely used to generate nearly homogeneous and isotropic turbulence under well-controlled conditions (Batchelor 1953, Townsend 1956). However, the Re_λ of the generated turbulence is not high enough neither for industrial applications nor for turbulence theory conclusive studies. From this point of view SFG are very appealing, since they generate consistently higher Re_λ than the regular grids with equivalent blockage ratios under similar conditions (Comte-Bellot & Corrsin 1966), and comparable to the case of jet grids (Gad-El-Hak & Corrsin 1974) and active grids (Makita 1991, Mydlarski & Warhaft 1996), while retaining a good level of homogeneity, except for the persistence of pressure and transverse energy transport, as reported by Valente & Vassilicos (2011). On the other hand, acceptable homogeneity and isotropy are achieved further downstream. The quantification of the homogeneity and the isotropy of a flow with such interesting features is of central importance. Seoud & Vassilicos (2007) argue that *“a turbulence with kinetic energy that decays exponentially whilst its integral length scales remains constant is either nonisotropic and/or nonhomogeneous or, if it is homogeneous and isotropic, is such that the kinetic energy dissipation rate per unit mass, ε , is not equal to $C_\varepsilon (u')^3 / L$ [...] If this fractal-generated turbulence is found to be homogeneous and isotropic far downstream where it is decaying, then we will be forced to face the question of whether our fractal stirrer modifies the turbulence so deeply that it modifies the relation between kinetic energy dissipation rate and Reynolds number”*. Hurst & Vassilicos (2007) quantified the large scales isotropy by measuring the ratio of the root mean square of the streamwise and crosswise fluctuations u'/v' , showing that along the centreline the ratio is contained in the range $1.1 \div 1.25$ in the decay region. This value is slightly larger than the one obtained by RGs with a contraction as in Comte-Bellot & Corrsin (1966) (more recently Antonia et al 2009 have shown that a small secondary contraction can push the isotropy ratio very close to 1) but it is in line with the case of active grids (Mydlarski & Warhaft 1996), which produce turbulence with comparable Re_λ .

Seoud & Vassilicos (2007) concentrated their attention on the decay region, observing that the turbulence production (measured along y -profiles) is always below 30% of the dissipation beyond the turbulence peak. Mazellier & Vassilicos (2010) observed that the flow in the near-field is strongly inhomogeneous. The mean streamwise component is characterized by a jet-like behaviour along the centreline due to the non-uniform distribution of the blockage ratio of the grid; on the other hand, it could be considered “wake-like” in the wake of the bars of the

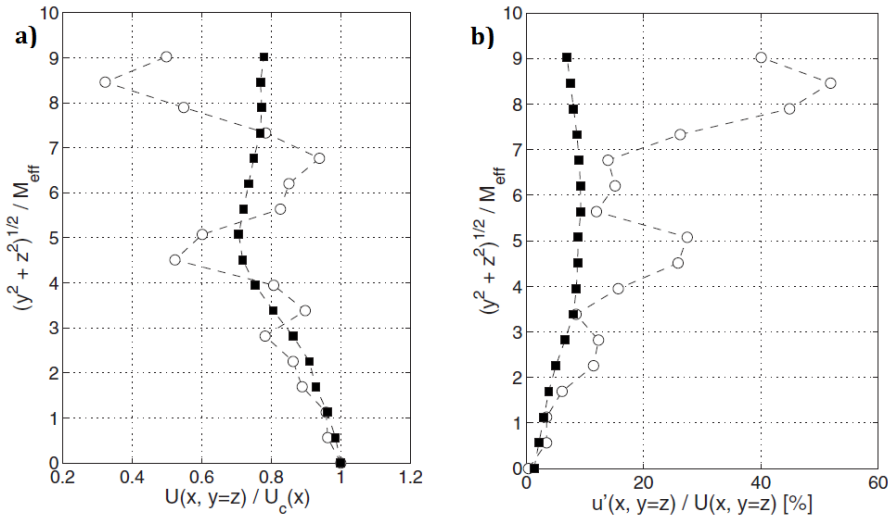


Fig. 7.4 Diagonal mean profiles of **a)** the velocity (normalized with the centerline velocity) and of **b)** the turbulence intensity for: $\circ x/x^* \approx 0.02$ $\blacksquare x/x^* \approx 0.15$. (Mazellier & Vassilicos 2010).

largest iterations. This trend is evident in Fig. 7.4, where the velocity and turbulence intensity profiles are plotted, and it tends to be smeared by the turbulent diffusion along the streamwise coordinate. Interestingly enough, for the same reason the centreline is almost turbulence free in the near-field. Mazellier & Vassilicos (2010) also documented that beyond $x/x^* \approx 0.6$ the ratio of the mean velocity and turbulence intensity measured on the centreline and in the wake of the corner of the second iteration of the grid is substantially equal to 1, and the flow can be considered homogeneous. Furthermore they argued that, since the time scales of the mean velocity gradients $(\partial U/\partial x)^{-1}$ and $(\partial U/\partial y)^{-1}$ are at least one order of magnitude larger than the time scale of the energy-containing eddies, these velocity gradients can be considered negligible.

7.1.4 Power-law or exponential decay?

A quite surprising feature of SFGs in the pioneering study by Hurst & Vassilicos (2007) is the unusual decay of the turbulence intensity along the streamwise direction, which appeared at odds with the classic relations of the power-law decay. In fact, their experimental data support the idea that the Taylor microscale beyond the turbulence intensity peak is independent of t_r , and, with good approximation, of the streamwise coordinate (though Valente & Vassilicos 2011 acknowledged a slow increase of λ in the far-wake of the grid, even if still incompatible with a power-law decaying turbulence). In a picture involving large and small scales isotropy and negligible turbulence production, according to Taylor’s hypothesis (1935):

$$-\frac{3}{2}U \frac{du'^2}{dx} = \frac{15u'^2\nu}{\lambda^2} \quad (7.4)$$

Assuming that the turbulence decay is modeled by a power-law $u'^2 \sim (x - x_0)^{-n}$ (with x_0 usually referred as “virtual origin”), it follows immediately from (7.4) that in the case of freely decaying homogeneous isotropic turbulence:

$$\lambda^2 \sim \nu(x - x_0)/U \quad (7.5)$$

This is at odds with the observed constancy of λ . By fitting the data with (7.5) Hurst & Vassilicos (2007) obtained a virtual origin very far upstream of the grid; by applying this virtual origin to find the exponent of the power-law decay of the turbulence intensity, they obtained $n = 6.75$, which gives a much steeper decay than any measured power-law decay. On the other hand, by setting the virtual origin equal to zero, and fitting the Taylor microscale trend with a law $\lambda \sim (x - x_0)^s$, they obtained $s \approx 0.1$, which is incompatible with the expected exponent $\frac{1}{2}$ in the power law decay.

For this reason Hurst & Vassilicos (2007) proposed an exponential decay, confirmed by Mazellier & Vassilicos (2010) in the form:

$$\frac{u'^2}{U^2} = A \exp \left[-B \left(\frac{x}{x^*} \right) \right] \quad (7.6)$$

A and B are dimensionless parameters. The virtual origin is arbitrarily set to zero, since it does not affect the value of B , but only the value of A , which plays only a role in shifting up or down the fitting curve.

Interestingly enough, George & Wang (2009) predicted the possibility of such exponential decay for homogeneous and isotropic turbulence in the case of time-independent characteristic lengthscale by using the only assumption of *equilibrium similarity* (George 1992). George & Wang (2009) demonstrated in this scenario the existence of two possible solutions of the spectral energy equation (Batchelor 1953): 1) an inviscid solution, leading to $u'^2 \sim x^{-2}$; 2) a viscous solution in which the turbulence decays exponentially. In this last case, the theory predicts that both the integral and the Taylor scales must be constant during the decay (and so does their ratio), and the energy spectra at different stages of the decay must collapse if normalized with respect to u'^2 and λ . This last assertion, descending directly from the assumption of single lengthscale self-preserving decay, is at odds with Kolmogorov’s idea of the presence of an inner and an outer scale.

However, Mazellier & Vassilicos (2010) recognized that turbulence generated by SFGs does not exactly follow the prediction of George & Wang (2009). As for example, the ratio of the integral and the Taylor scales is not constant throughout the decay, but is slightly increasing further downstream; furthermore, though the exponential decay fits reasonably well the data, the dependence of λ with the inlet Reynolds number does not. They proposed an alternative form to the kinetic energy decay, that is both compatible with the exponential decay proposed by George &

Wang (2009) and the power-law decay. Such a hybrid behaviour can be motivated by the fact that the range of scales that are simultaneously excited by the fractal stirrer is still of finite size, while the single lengthscale decay would perhaps involve simultaneous excitation of the entire range of turbulent scales down to the Kolmogorov scale.

7.1.5 The violation of the dissipation anomaly

As already stated in the *incipit* of this chapter, an immediate consequence of the scaling of the dissipation $\varepsilon = C_\varepsilon k^{3/2}/L$ is that $L/\lambda \propto Re_\lambda$. Valente & Vassilicos (2011) reviewed the rate of variation of the turbulent kinetic energy k and the integral lengthscale L along the streamwise downstream distance for a variety of high-Reynolds number self-preserving shear flows and in wind tunnel turbulence (wakes, jets, grid turbulence, and many others; see Tab. 1 of Valente & Vassilicos 2011 for reference). They observed, from the data reported in Comte-Bellot & Corrsin (1966) and Tennekes & Lumley (1972), that both the k and L streamline evolutions are modelled with power laws:

$$k \sim U_\infty^2 \left(\frac{x-x_0}{L_B} \right)^{-n} \quad (7.7)$$

$$L \sim L_B \left(\frac{x-x_0}{L_B} \right)^{-m} \quad (7.8)$$

L_B is a macroscopic scale depending on the flow, U_∞ is the inlet velocity and x_0 is a virtual origin. Assuming that the Taylor microscale can be inferred from the scaling by Taylor (1935) $\varepsilon \sim \nu k/\lambda^2$, and that the dissipation anomaly is valid, i.e. $\varepsilon \sim k^{3/2}/L$, quite surprisingly for this large variety of flows the following two relations do hold (defining $Re_0 = U_\infty L_B/\nu$ as an inlet-condition-based global Re number):

$$\frac{L}{\lambda} \sim \sqrt{Re_0} \left(\frac{x-x_0}{L_B} \right)^{(m/2)-(n/4)} \quad (7.9)$$

$$Re_\lambda = \frac{\sqrt{k}\lambda}{\nu} \sim \sqrt{Re_0} \left(\frac{x-x_0}{L_B} \right)^{(m/2)-(n/4)} \quad (7.10)$$

Astonishingly, regardless of the value of the exponents m and n , the ratio L/λ scales exactly as Re_λ . This relation actually reflects the Richardson-Kolmogorov phenomenology: the higher is the Re_λ , the larger is the ratio between the large scales of the energy-containing range and the small scales of the dissipative range.

The turbulence generated by SFGs is characterized by a dramatic departure from this scaling. Seoud & Vassilicos (2007) documented essentially that L/λ is constant along the streamwise direction while Re_λ significantly decays. This behaviour is incoherent with $C_\varepsilon \approx const$ and at odds with the dissipation anomaly. Indeed, Seoud & Vassilicos (2007) observed that a scaling $C_\varepsilon \sim 1/Re_\lambda$ actually well fits the data for SFG generated turbulence. Mazellier & Vassilicos (2010) confirmed

this finding, adding that L/λ and Re_λ are increasing function of the inlet Reynolds number Re_0 .

Valente & Vassilicos (2011) compared the behavior of SFG and RG by taking hot-wire anemometry measurements with different probes to validate the independence of the results on the spatial resolution. Although the Re_λ achieved by RG is close to the lower hand of the range of validity of the Richardson-Kolmogorov phenomenology, their data support to some extent the constancy of C_ε and the proportionality of L/λ with Re_λ . In the case of SFG, instead, they confirmed the observation by Seoud & Vassilicos (2007) and Mazellier & Vassilicos (2010) that roughly $C_\varepsilon \propto Re_\lambda^{-1}$.

Subsequently, Valente & Vassilicos (2012) investigated the decay of the turbulence generated by SFG and RG with large meshlength (in one case as large as the first iteration of the grid). Interestingly enough, their data show that the proportionality $C_\varepsilon \propto Re_\lambda^{-1}$ is not a unique feature of SFGs, but it also applies in a wide region in RGs generated turbulence provided that $0.5 \leq x/x^* \leq 2$. In this region $C_\varepsilon \propto Re_M^{-1/2}/Re_\lambda$ and $L/\lambda \propto Re_M^{-1/2}$, independently of Re_λ .

The importance of the cornerstone assumption $C_\varepsilon \approx const$ in turbulence theory and simulations, and the repeated violations of it in the more generalized scenario of regular grid turbulence, raise new questions on its universality and applicability in many other turbulent flows.

7.1.6 Questions raised and role of 3D data

The investigation of fractal generated turbulence, above all in the peculiar case of SGFs, is still at an early stage. Though the bulk of the work and results obtained using hot-wire anemometry are quite well supported by DNS simulations (see Nagata et al 2008a,b, Laizet & Vassilicos 2011), even if with limited Reynolds number, the scientific community is starving for 3D experimental data to achieve a better understanding of the underlying physics of fractal-generated turbulence and of its astonishing features. This investigation is actually to be contextualized in a wider framework, in which several turbulent flows can be thought, to some extent, as a combination of the Richardson-Kolmogorov phenomenology and a different scenario in which the dissipation anomaly does not hold. Actually, the turbulence generated by SFGs can be a mixture of these phenomenologies as well.

Three-dimensional three-components velocity measurements can provide a more complete picture of this peculiar class of turbulent flows. As for example, the validity of the assumption of small-scale isotropy in the turbulence generated by SFGs is still an open question; the most thorough analyses on this topic (Seoud & Vassilicos 2007, Valente & Vassilicos 2011) are performed with x-wire probes, with a separation between the wires of approximately 10 times the Kolmogorov scales, raising some doubts on the measurement of coherence spectra at small scales (see Fig 7 in Seoud & Vassilicos 2007 or Fig.4 in Valente & Vassilicos 2011). Similarly,

assuming small scales isotropy in the measurement of the turbulent dissipation rate, even if hardly questionable, may leave room to some perplexity. The assessment of a technique for the complete measurement of all the components of the stress-strain tensor contributing to the dissipation is strongly needed.

7.2 Experimental details

7.2.1 Experimental apparatus and design point

The experimental investigation is conducted in the open circuit wind tunnel of the Laboratory for Energetic Flow and Turbulence of the Arizona State University. The measurements are performed in a 1.5m long square test section of width $T = 0.154m$. The square fractal grid is placed at the inlet of the test section, immediately after a 5.2:1 contraction. The tested grid is composed of 4 iterations, with thickness ratio of the grid t_r equal to 13. The geometrical parameters of the grid are reported in Tab. 7.1.

A global Reynolds number (7.11) is defined by measuring the reference velocity U_r on the tunnel centreline approximately 400mm downstream of the outlet of the contraction; the reference length is the equivalent mesh-length, defined as in Hurst & Vassilicos (2007), and approximately equal to 15.3mm. In the present work, U_r is set equal to 11.5m/s, resulting in a global Reynolds number of approximately 11.5×10^3 :

$$Re = \frac{U_r M_{eff}}{\nu} \quad (7.11)$$

A Laskin nozzle generates seeding olive-oil particles with $1\mu m$ mean diameter and concentration of approximately $50 \text{ part}/\text{mm}^3$. The particles are illuminated by a double pulse Nd:Yag laser with maximum energy of $100 \text{ mJ}/\text{pulse}$, pulse duration of 8 ns and pulse frequency of 2 Hz . The laser is placed on a horizontal breadboard below the tunnel. A vertical laser slab is obtained by using a 45° mirror, a circular aperture and a diverging cylindrical lens. The illuminated region is a slab oriented along the streamwise and vertical crosswise directions, and a thickness of approximately 2.5 mm .

As in Sec. 4.2.3, the imaging system is composed of 4 TSI POWERVIEW™ Plus 11MP camera with 4008×2672 pixel resolution and pixel pitch of $9\mu m$. The cameras are disposed approximately in a linear arrangement, angularly equally spaced, spanning an angle of 80° , as shown in Fig. 7.5. Even if not optimized (the views provide more diversified information if a 3D configuration is used), this system is chosen due to its intrinsic simplicity and uniform scattering conditions. The imaging system is placed in 3 locations downstream of the grid, corresponding approximately to $x/x^* = 0.52, 0.61, 0.70$. The cameras are equipped with Nikon

| t_r | σ | $M_{eff}[mm]$ | $t_o[mm]$ | $t_{N-1}[mm]$ | $L_o[mm]$ | $L_{N-1}[mm]$ | $x^*[mm]$ |
|-------|----------|---------------|-----------|---------------|-----------|---------------|-----------|
| 13 | 0.34 | 15.3 | 8.12 | 0.62 | 80.8 | 10.1 | 801 |

Table 7.1 Geometric parameters of the tested grid.




Fig. 7.5 Experimental setup: top-view of the cameras arrangement and of the illuminated volume location into the test section of the wind tunnel.

objectives with a focal length of $60mm$ when focused at infinity, set at $f_{\#} = 8$, and Tamron 2x teleconverters to approximately double the image distance, thus resulting in a higher resolution with respect to the test case presented in Sec. 4.2.3, and effective $f_{\#}$ equal to 16. On the other hand, the object distance is slightly larger due to the limit in physical access with the 4 cameras. The average magnification is approximately 0.45, resulting in a resolution of about $50pixels/mm$.

The diffraction-limited particle diameter d_{τ}^* is 3.3 pixels, which is slightly larger than the commonly used diameter of 3 pixels used in Tomo-PIV experiments. This would lead the design point to be located outside the valid area of the diagram in Fig. 1.10. However, considering that the volume to be reconstructed is quite thin, the imaging system is composed of 4 cameras and the accuracy of the tomographic reconstruction can be improved using the SFIT-MART technique outlined in Chapter 4, a slightly larger source density can be considered acceptable.




Fig. 7.6 Contour representation of d_l [mm] as a function of M_0 and Δz_0 for fixed $f_{\#}$, d_p , and image density equal to 0.05ppp. The dashed black curves represent the constraint due to finite depth of field ($\Delta z_0 = \delta z_0$), diffraction spot size ($d_{\tau}^* = 3$ pixels and $d_{\tau}^* = 3.5$ pixels), and volume thickness larger than the interrogation spot size ($\Delta z_0 = d_l$). The area with colored contour lines is the valid area to place the design point for the Tomo-PIV experiment. The area with dashed contour lines is the area where the design point can be conditionally placed. The red dot is the design point for the experiment under analysis.

The design chart and the experiment design point are reported in Fig. 7.6. The limit of $d_{\tau}^* = 3.5$ pixels is introduced for reference, and the contour-lines in the region comprised in the range $3 < d_{\tau}^* < 3.5$ are dashed in order to identify that the design point can be conditionally located in the area (where the conditions to be met are relative to solutions devoted to provide a decent accuracy of the tomographic reconstruction even with a source density slightly larger than usual). As outlined in Fig. 7.6, under the imaging conditions of the experiment, d_l should be set to about 0.6mm in order to have in average 10 particles in each interrogation region (i.e. about $30 \times 30 \times 30$ voxels).

7.2.2 Calibration, tomographic reconstruction and PIV interrogation

The calibration target for the optical calibration is the same described in Sec. 4.2.3, i.e. glass substrate with a $100 \times 100 \text{mm}^2$ grid of black markers, with diameter of $250 \mu\text{m}$ and spacing of 1mm (0.2% tolerance in centring). The target is translated

with a motorized stage within the range $\pm 4\text{mm}$ with respect to the reference plane. The calibration images are analyzed with a template-matching technique to identify the calibration markers. A nonlinear regression algorithm is used to obtain pinhole-based mapping functions (Tsai 1987, see Sec. 1.2.1) from the object space to the image plane. The volumetric self-calibration described in Sec. 2.1.4 is applied to reduce any residual calibration error and disparity between the cameras. After the application of the volume self-calibration the standard deviation of the calibration error is less than 0.02 pixels for each camera of the set, with a maximum error of 0.1 pixels.

Since the thickness of the illuminated volume is smaller than the test case presented in Sec. 4.2.3, a higher signal/noise ratio is detected on the original images, thus allowing a less intense pre-processing. In this case the pre-processing is performed in three steps: historical background removal by taking the minimum in each pixel over the ensemble of 1000 realizations; sliding minimum subtraction with a kernel of 13×13 pixels to eliminate any fluctuating background; slight particles sharpening with a Laplacian-based filter to increase the contrast (this step is required due to the large particles diameter).

The reconstructed volume is $30 \times 30 \times 3\text{mm}^3$ (i.e. slightly larger than the illuminated region in order to allow the reconstruction of all the imaged particles); the average magnification of 0.45 leads to approximately $50\text{vox}/\text{mm}$ for a resolution ratio between voxels and pixels equal to 1. This results in a reconstructed volume of $1500 \times 1500 \times 150$ voxels. A set of 1000 realizations is reconstructed for each of the three streamwise locations.

The reconstructed distributions are then cross-correlated with the efficient multi-pass volume deformation algorithm outlined in the Chapter 5. The interrogation is performed using a Blackman weighting window in both the cross-correlation step and the dense predictor averaging (see Secs. 5.1-5.3 for more details on the algorithm) in order to achieve a MTF resembling that of a low pass filter. The final interrogation volume size is $48 \times 48 \times 48$ voxels (i.e. $0.96 \times 0.96 \times 0.96\text{mm}^3$), with 75% overlap (thus resulting in a vector spacing of 12 voxels, i.e. 0.24mm). Apparently, the nominal size of d_l is larger than the predicted value of 0.6mm for the design point; however, it has to be taken into account that the Blackman window is characterized by an equivalent noise bandwidth (ENBW, i.e. the “(band)width of a rectangle filter with the same peak power gain that would accumulate the same noise power”, Harris 1978) of 1.68. This lead to an “effective d_l ” equal to $0.96/\text{ENBW} = 0.57\text{mm}$, which is very close to the design value. However, this value does correspond to a significantly different impulse frequency response than the equivalent top hat moving average filter, as shown in Fig. 7.7, where the MTF of the Top Hat (TH) window with $D_l = 30$ voxels is compared with those relative to the adopted Blackman window with $D_l = 48$ voxels and variable size of the window for the dense predictor averaging W_b . The data are obtained




Fig. 7.7 MTF as a function of the normalized frequency D_1/λ (with D_1 set to 48 voxels) for: Top Hat moving average with window size equal to 30vox (dashed line); Blackman windowing with window size equal to D_1 and $W_b = 12(\blacksquare), 24(\blacktriangle), 48(\bullet)$ voxels.

with the same theoretical method outlined in Sec. 5.3 (the number of iterations is set equal to 3). Defining a cut-off wavelength Λ such that $MTF(D_1/\Lambda) = 0.5$, it results that for the TH windowing $\Lambda \approx 40$ voxels, while for the Blackman windowing Λ depends on W_b ($\Lambda \approx 42, 34, 30$ voxels for $W_b = 48, 24, 12$ voxels respectively). The TH approach is affected by a steep decay of the MTF beyond $\Lambda = 40$ voxels and by large negative lobes for wavelengths smaller than 30 voxels (i.e. beyond the Nyquist limit), thus determining significant aliasing effect. In this analysis the cross-correlation is performed using a Blackman windowing with dense predictor averaging on $W_b = 24$ voxels, that provides low values of the MTF beyond $D_1/\lambda = 2$ (thus significantly smoothing random noise fluctuations) and no negative lobes.

7.3 Results

7.3.1 Homogeneity assessment

The streamwise and cross-stream profiles of the mean streamwise velocity component are reported in Fig. 7.8. The most upstream measurement region is relatively close to the expected location of the centreline turbulence intensity peak (i.e. approximately $0.5x^*$). The relatively strong mean flow gradients along the streamwise and the cross-stream directions are evidence of non-negligible turbulence production, as discussed by Valente & Vassilicos (2011). As the flow develops downstream, the jet-like behaviour due to the non-uniform distribution of

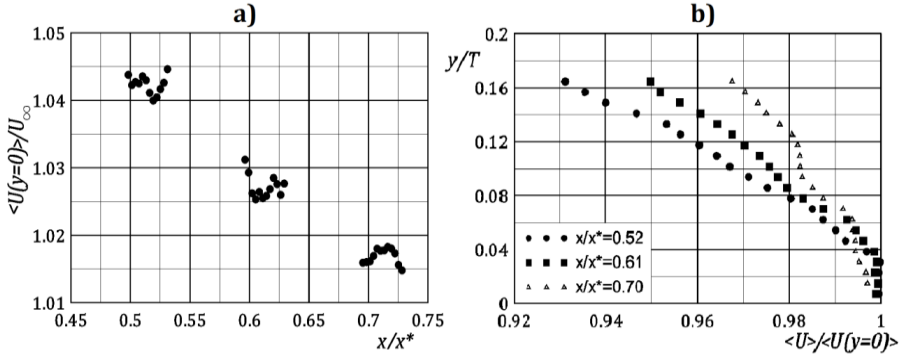


Fig. 7.8 Profiles of the mean streamwise velocity component: **a)** longitudinal profile of the centerline velocity; **b)** crosswise vertical profiles at three different longitudinal locations. (The profiles are sub-sampled for clarity; 9 points for each subset of 10 points are skipped for the longitudinal profile, 5 out of 6 for the crosswise profile).

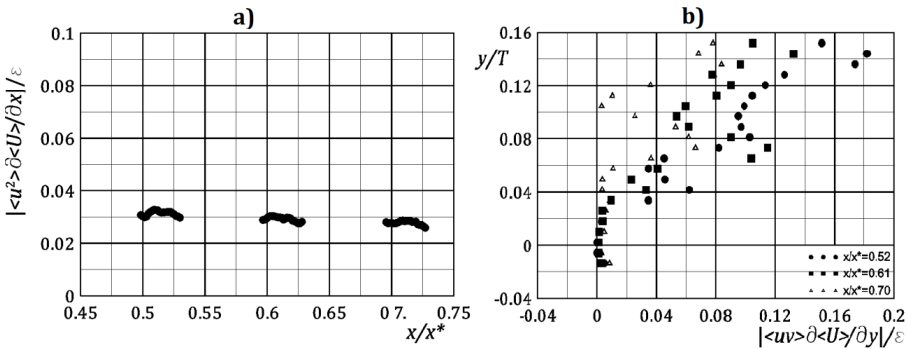


Fig. 7.9 Profiles of the normalized turbulent production terms: **a)** longitudinal profile of $|\langle u^2 \rangle \partial \langle U \rangle / \partial x|$; **b)** crosswise vertical profiles of $|\langle uv \rangle \partial \langle U \rangle / \partial y|$ at three different longitudinal locations. (The profiles are sub-sampled for clarity; 9 points for each subset of 10 points are skipped for the longitudinal profile, 5 out of 6 for the crosswise profile).

the blockage ratio of the square fractal grids is smoothed by the turbulent transport of momentum, and the cross-stream profile of the mean streamwise velocity component rapidly becomes flatter.

In order to assess the degree of inhomogeneity of the fractal generated turbulence, the production terms in the turbulent kinetic energy balance are evaluated and plotted in Fig. 7.9. In particular, the longitudinal evolution of the term $|\langle u^2 \rangle \partial \langle U \rangle / \partial x|$ along the centerline, and the crosswise distribution of $|\langle uv \rangle \partial \langle U \rangle / \partial y|$ at three streamwise locations, are analyzed after normalization with the local value of the measured dissipation (see Sec. 7.3.4). The results show that

the turbulent production along the centreline is not significant, as it is well below 4% of the dissipation over the entire range of investigation, in agreement with Seoud & Vassilicos (2007). On the other hand, the term $|\langle uv \rangle \partial \langle U \rangle / \partial y|$ indicates a significant lack of homogeneity by moving away from the centreline towards the wake of the largest iteration (that is placed outside of the observation region, at $y/T = 0.26$). However the inhomogeneity is significantly less pronounced at $x/x^* > 0.7$, where it never exceeds 10% of the dissipation.

It is worth to point out that these results are not enough to ensure that the decay of fractal generated turbulence is not conditioned by inhomogeneity. Considering the turbulent kinetic energy equation (here the Einstein notation is used, with the symbology $x_1 = x$, $x_2 = y$, $x_3 = z$, $U_1 = \langle U \rangle$, $U_2 = \langle V \rangle$, $U_3 = \langle W \rangle$, $u_1 = u$, $u_2 = v$, $u_3 = w$):

$$U_l \frac{\partial k}{\partial x_l} = -\langle u_i u_j \rangle \frac{\partial U_i}{\partial x_j} - \frac{\partial}{\partial x_l} \langle u_l k \rangle - \frac{\partial}{\partial x_l} \frac{\langle u_l p \rangle}{\rho} + \nu \frac{\partial^2 k}{\partial x_m \partial x_m} - \nu \left\langle \frac{\partial u_i}{\partial x_l} \frac{\partial u_i}{\partial x_l} \right\rangle \quad (7.12)$$

The terms on the right-hand side are referred as production, triple-correlation transport, pressure transport, viscous diffusion and dissipation. Tomo-PIV does not allow the measurement of the pressure transport, while the computation of the triple-correlation transport is extremely delicate, as it involves derivatives of third-order fluctuations products. However, the results of the hot-wire investigation by Valente & Vassilicos (2011) ensure that these terms are usually negligible, or approximately constant along the longitudinal direction (for example, this is the case of the transverse triple correlation transport), thus the decay is insensitive to their presence.

7.3.2 Large and small scale isotropy

In the investigation of regular grid turbulence the detection of large scale anisotropy has often been performed by measuring the so-called *isotropy factor*, i.e. the ratio of the mean square of the streamwise to cross-stream velocity fluctuations.

Hurst & Vassilicos (2007) applied the same criterion to the turbulence generated by square space-filling fractal grids, and concluded that along the centreline the isotropy factor is contained within the range 1.1-1.25. The mean isotropy factor is only slightly higher than the case of regular grids after a contraction (Comte-Bellot & Corrsin 1966), and it is of the same order of magnitude in the case of active grids (Mydlarski & Warhaft 1996). On the other hand, the inhomogeneous distribution of the blockage ratio for fractal generated turbulence causes a strong variation of the isotropy factor along the cross-stream direction, as shown in the u'/v' profiles in Fig. 7.10. The isotropy factor decreases along the longitudinal direction, approaching a value of approximately 1.1 at $x/x^* > 0.7$. The profiles along the crosswise vertical direction, on the other hand, show a significant

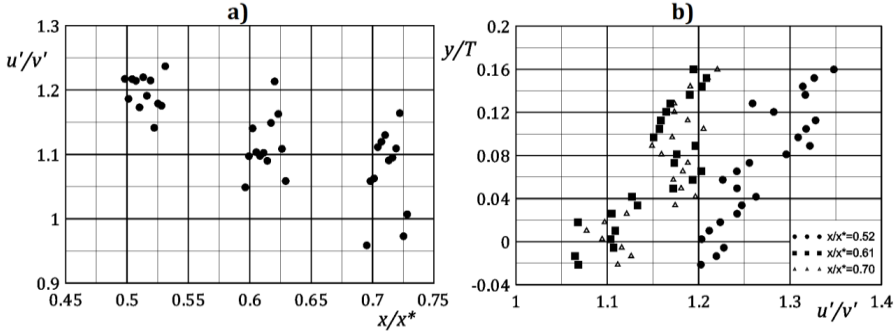


Fig. 7.10 Profiles of the isotropy factor along the longitudinal direction (a) and the crosswise vertical direction at three different longitudinal locations (b). (The profiles are sub-sampled for clarity; 9 points for each subset of 10 points are skipped for the longitudinal profile, 5 out of 6 for the crosswise profile).

persistence of the anisotropy when moving towards the wake of the first iteration of the grid. Again, the diffusion progressively smears the anisotropy moving downstream; for $x/x^* > 0.7$ the isotropy factor attains a maximum value of 1.2 in the measurement region. It is evident that the flow remembers the inhomogeneity and anisotropy effects induced by the non-uniform distribution of the grid blockage ratio for a long time during the decay.

Another possible path to assess the isotropy of fractal generated turbulence is followed by Valente & Vassilicos (2011), which measured the two-point longitudinal and transverse autocorrelation functions, and used the relations of isotropic turbulence to test for large and small scale isotropy. The autocorrelation functions are obtained from the two-point correlation tensor, defined as follows:

$$R_{ij}(\underline{x}, \underline{r}) = \langle u_i(\underline{x})u_j(\underline{x} + \underline{r}) \rangle \quad (7.13)$$

Let the separation be only in the longitudinal direction (i.e. $r_1 \equiv |r|$), $f(x_1, r) = R_{11}(x_1 + r, 0, 0)/u'$ denote the non-dimensional longitudinal two-point correlation, and $g(x_1, r) = R_{22}(x_1 + r, 0, 0)/v'$ be the transverse two-point correlation function measured from the experimental correlations. Then, by the assumption of isotropic turbulence:

$$g_{iso}(x_1, r) = \frac{1}{2r} \frac{\partial}{\partial r} [r^2 f(x_1, r)] \quad (7.14)$$

The longitudinal and the transverse correlation functions are plotted for $x/x^* = 0.52$ and 0.61 in Fig. 7.11. The transverse correlation function has been normalized with respect to v' instead of u' in order to compensate for the large-scale anisotropy, and a parabolic fit has been applied to the peak of both the longitudinal and the transverse correlation function to remove the effect of the bias due to random noise (this aspect is discussed in more detail in Sec. 7.3.3). The

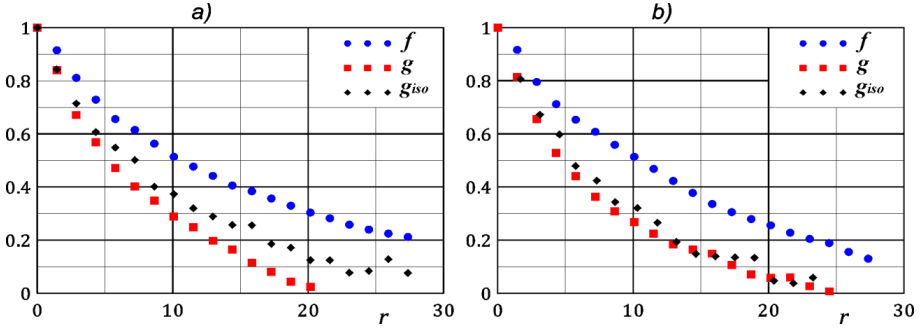


Fig. 7.11 Longitudinal and transverse autocorrelation functions for $x/x^* = 0.52$ and $x/x^* = 0.61$. (The functions are sub-sampled for clarity; 4 points for each subset of 5 points are skipped).

comparison in Fig. 7.11 between the measured transverse correlation function and the isotropic estimate (7.14) highlights a modest agreement, even if it improves moving downstream along the tunnel centreline. The results clearly indicate an acceptable degree of small-scale isotropy when r is smaller than 5mm (i.e. approximately the Taylor lengthscale, see Sec. 7.3.3). When moving downstream, as the Taylor lengthscale increases, the maximum lengthscale for which local isotropy appears valid increases.

7.3.3 Longitudinal integral and Taylor lengthscales

The integral lengthscale L_{ij} and the Taylor microscale λ_{ij} are calculated using the following expressions:

$$L_{ij}(\underline{x}) = \frac{1}{\langle u_i^2 \rangle} \int_0^\infty R_{ii}(\underline{x}, \underline{r}) dr_j \quad (7.15)$$

$$\lambda_{ij}(\underline{x}) = \sqrt{\frac{\langle u_i^2 \rangle}{\left. \frac{\partial^2 R_{ii}(\underline{x}, \underline{r})}{\partial r_j^2} \right|_{r=0}}} \quad (7.16)$$

In these expressions, and unless otherwise stated, the Einstein convention is not used. The subscript i indicates the velocity component, and the subscript j refers to the direction of separation.

The simultaneous measurement of both the integral scale and the Taylor microscale is a challenge for PIV, because one must observe a field of view with a width of at least 6-7 integral lengthscales to accurately estimate L , and at the same time have high enough resolution to resolve the Taylor microscale. This is inherently difficult because of the relatively low dynamic spatial range of the two-pulse digital PIV technique, and it is expected to be even more challenging in Tomo-PIV, with the resolution being limited by the maximum allowed seeding density.




Fig. 7.12 Longitudinal integral lengthscale along the centerline. The Tomo-PIV data are compared with the PIV data set by Discetti et al (2013b). The TPIV data are taken with a step of 10 vectors for clarity.

de Jong et al (2009) have shown that, with PIV data, a good estimate of the longitudinal integral lengthscale L_{11} , can be obtained even if the extension of the field of view is of the same order of magnitude of L_{11} by fitting a decaying exponential function to the tail of the longitudinal two-point correlation function. Considering that the integral lengthscale is of the same order of magnitude of the effective meshlength (i.e. about $15mm$), and that the observation region is $30 \times 30 \times 3mm^3$, one can reasonably expect to be able to measure L_{11} in the left half of the volume for the three streamwise locations. In order to calculate the integral lengthscale, the streamwise autocorrelation function is computed over all the points in the left half of the field of view, and then integrated as in the (7.15). An exponential function is fitted to the tail of the autocorrelation function by using the last 20 points (considering that 75% overlap is used, it corresponds to 5 statistically independent vectors) to elongate the region of integration. The results are compared with those obtained by Discetti et al (2013b) with planar PIV in the same conditions. The significant difference is that in the experiments by Discetti et al (2013b) the field of view is much larger ($75 \times 50mm^2$), thus reducing the contribution of the extended exponential tail of the autocorrelation function to less than 4% (for the current Tomo-PIV data the correction to the measured integral is up to 50% of the integral lengthscale).

The results, obtained by averaging L_{11} over 40 rows of vectors, i.e. a region of approximately $10mm$ of extension in the cross-stream direction, are plotted in Fig. 7.12. Tomo-PIV data show a more significant scatter, both due to non perfect convergence (the PIV statistics are obtained over an ensemble of 5000 independent realizations, i.e. 5 times the number of snapshots reconstructed and cross-correlated with Tomo-PIV) and to larger uncertainty introduced by the large contribution of the extended exponential tail to the integral (7.15). While PIV data

show a slow but sensitive growth of the integral lengthscale along the streamwise direction (in agreement with Valente & Vassilicos 2011), the extraction of the same information from the scatter of Tomo-PIV is not straightforward; however, the order of magnitude is correctly evaluated.

As previously mentioned, the Taylor microscale is estimated by computing the second derivative of the two-point correlation function at zero offset. The estimate of the two-point correlation in the origin is affected by the random error, which is in general uncorrelated for non-overlapping interrogation windows (Adrian & Westerweel 2011). This means that noise affects the estimate of the first 4 points of the two-point correlation when 75% overlap is employed. In particular, for $r_1 = r_2 = 0$ the expected value in the case of zero offset of the estimate of the two-point correlation $\tilde{R}_{11}(\underline{x}, \underline{0})$ is:

$$\tilde{R}_{11}(\underline{x}, \underline{0}) = u'^2 + \sigma_\varepsilon^2 \quad (7.17)$$

where σ_ε is the variance of the random error. For this reason, a parabolic fit is applied in proximity of the origin:

$$\tilde{R}_{11}(\underline{x}, [r_1, 0, 0]) = u'^2 \left(1 - \frac{r_1^2}{2\lambda_{11}^2}\right), \quad \frac{r_1}{\lambda_{11}} < 1 \quad (7.18)$$

A supplementary benefit of the application of the (7.18) is that an estimate of both u'^2 and λ is provided; in particular, the former can be used to define a signal-to-noise ratio:

$$SNR(\underline{x}) = \frac{u'^2}{\sigma_\varepsilon^2} = \frac{u'^2}{\tilde{R}_{11}(\underline{x}, \underline{0}) - u'^2} \quad (7.19)$$

The SNR computed with (7.19) is a good *a-posteriori* indicator of the quality of the results. The median of the SNR along the crosswise direction is considered at each streamwise location, and plotted in Fig. 7.13. As expected, the SNR decreases moving downstream, since the mean square of the velocity fluctuation decreases along the decay direction, while the measurement noise retains the same intensity.

The accuracy of the estimation of the Taylor lengthscale is strongly dependent on the upper limit used for the fitting: a higher number of points reduces the error in the fitting procedure, but determines a lower resolution (and, consequently, an overestimation of the Taylor microscale). The criterion for the definition of the upper limit is based on the minimum number of points to obtain a standard deviation smaller than 10% of the mean of the estimated value of the Taylor microscale in each column of vectors (40 rows are considered in this analysis). In all the experiments a maximum of 8 points (equivalent to about $2mm$) suffices to this task. Furthermore, by estimating the mean value over each column of vectors by using only 6 and 1 points, a range of uncertainty of about 15%-20% is estimated.




Fig. 7.13 Signal to noise ratio as a function of the streamwise location. The data are taken with a step of 10 vectors for clarity.




Fig. 7.14 Longitudinal Taylor microscale along the centerline. The Tomo-PIV data are compared with the PIV data set by Discetti et al (2013b). The TPIV data are taken with a step of 10 vectors for clarity.

The data are compared with the PIV results obtained by Discetti et al (2013b), and plotted as a function of the streamwise coordinate in Fig. 7.14. The Tomo-PIV results evidently suffer for a lack of spatial resolution if compared with the PIV data set. As a matter of fact, even though the average magnification is similar, in the PIV experiment by Discetti et al (2013b) a much higher resolution can be achieved, with interrogation spots of 32×32 pixels ($630 \times 630 \mu\text{m}$ in the object space; actually, considering that a Blackman window is used, the effective size is obtained by dividing the real size by the ENBW, resulting in $375 \times 375 \mu\text{m}$). This striking difference is due to the fact that the PIV image density is not limited by the




Fig. 7.15 Ratio of the longitudinal integral lengthscale L_{11} and Taylor microscale λ_{11} as a function of the local Re_λ . Dashed line from (7.20) are included for reference.

requirements of sufficiently accurate tomographic reconstruction. For this reason Tomo-PIV measurement of the Taylor microscale are biased towards larger values. Again, the PIV results identify a significant growth of the Taylor microscale along the streamwise direction, while the scatter of the Tomo-PIV data does not highlight this trend. This effect might be also due to the fact that along the streamwise direction the Kolmogorov scale does effectively increase (see Sec. 7.3.4), thus partly compensating the lack of resolution and consequently reducing the gap between the PIV and the Tomo-PIV results.

7.3.4 Dissipation measurements

As previously stated, one intriguing feature of the turbulence generated by square space-filling fractal grids highlighted by the pioneering experiments of Hurst & Vassilicos (2007) is the apparent violation of the commonly accepted turbulence dissipation scaling $\varepsilon = C_\varepsilon k^{3/2}/L$. An immediate implication of this scaling, when combined the relation $\varepsilon = 10\nu k/\lambda^2$ (coming from the isotropy assumption in (7.4) and $k = 3u'^2/2$) is that:

$$\frac{L}{\lambda} = \frac{C_\varepsilon}{10} Re_\lambda \quad (7.20)$$

The (7.20) shows that, according to the so-called dissipation anomaly, the ratio L/λ should increase as Re_λ increases. The assertion (7.20) is tested in Fig. 7.15, where the ratio L_{11}/λ_{11} is plotted as a function of Re_λ . The function L_{11}/λ_{11} is underestimated due to the effects of limited spatial resolution discussed in Sec. 7.3.3; on the other hand, Re_λ is proportional to the Taylor microscale as well, and for this reason it is overestimated in the same measure, thus retaining the proportionality scaling (7.20). Dashed lines obtained from (7.20) with different values of the constant C_ε are included in the figure for reference. The plot indicates

that the dissipation anomaly does not hold, since the ratio is approximately constant along the streamwise direction while Re_λ decreases rapidly due to the fast decay of the turbulent kinetic energy.

A different way of testing the validity of the dissipation scaling is to determine C_ε from measurement of the kinetic energy, dissipation and integral length scale. The turbulent dissipation is notoriously difficult to be measured accurately in anisotropic and inhomogeneous flows, and for this reason three methods are implemented. In the first method, the dissipation is found directly from the measurements of the mean squared velocity gradient:

$$\varepsilon = 2\nu \langle s_{ij}s_{ij} \rangle \quad (7.21)$$

The turbulent rate-of-strain tensor s_{ij} is defined as follows:

$$s_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \quad (7.22)$$

The great advantage of 3D3C measurements is the availability of all the components of the velocity gradient tensor in the same location. The scenario is rather different with respect to the case of the hot-wire measurement performed by Seoud & Vassilicos (2007), in which the small-scale isotropy assumption is invoked, or by Valente & Vassilicos (2011), who used a x-wire with separation of 10 times the Kolmogorov scale, thus leaving some questions open about the validity of the dissipation measurement. On the other hand, the limited spatial resolution is an obstacle to the successful measurement of the spatial derivatives of the velocity fluctuations at the smaller turbulent scales. On top of this, in order to reduce the degrading effect of noise in the derivative computation, the velocity fluctuations fields are low-pass filtered by a local 2nd order polynomial fitting function on a 5 x 5 kernel.

The second method extracts the turbulent dissipation from the rate of decay of the turbulent kinetic energy, as is commonly done for homogeneous grid turbulence. In the present case, the inhomogeneous nature of the flow recommends to use the full turbulent kinetic energy balance. The kinetic energy balance (7.12) can be expressed on the centreline as in the following (supposing $\langle V \rangle = \langle W \rangle = 0$):

$$\varepsilon = -\langle U \rangle \frac{\partial k}{\partial x} - \left[\langle u^2 \rangle \frac{\partial \langle U \rangle}{\partial x} + \langle uv \rangle \frac{\partial \langle U \rangle}{\partial y} + \langle uw \rangle \frac{\partial \langle U \rangle}{\partial z} \right] - \left\{ \frac{\partial \langle uk \rangle}{\partial x} + \frac{\partial \langle vk \rangle}{\partial x} + \frac{\partial \langle wk \rangle}{\partial x} + \right. \\ \left. 1\rho\partial u p \partial x + 1\rho\partial v p \partial y + 1\rho\partial w p \partial z + \nu \partial^2 k \partial x^2 + \partial^2 k \partial y^2 + \partial^2 k \partial z^2 \right\} \quad (7.23)$$

The first term in the right hand side dominates the balance, so the others can be treated as small corrections in the evaluation of the dissipation. The term in curly brackets (i.e. the sum of the triple correlation transport and the pressure transport) is evaluated by using the second moment closure model suggested by Daly & Harlow (1970).

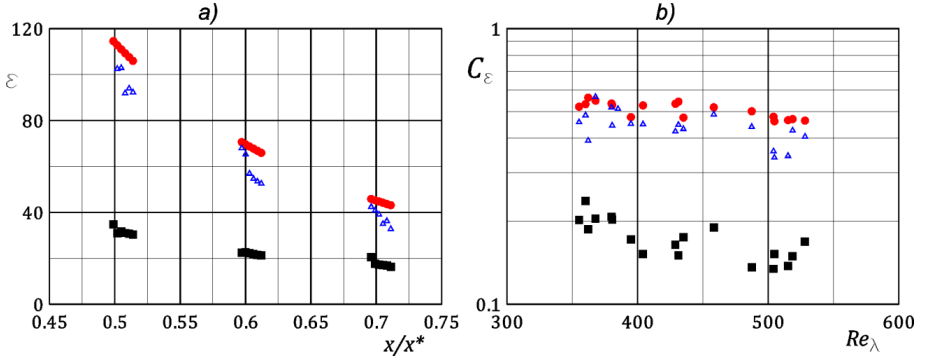


Fig. 7.16 Dissipation obtained with direct measurement \blacksquare , energy balance \bullet and Large Eddy PIV \triangle . **a)** Streamwise profile of the dissipation along the centerline; **b)** C_ε as a function of the local Re_λ .

The third method is often referred as *Large Eddy PIV* (Sheng et al 2000), as it assumes the equilibrium of the transfer of energy from the large lengthscales (production) to the small lengthscales (dissipation). The energy flux is evaluated in terms of the resolved rate of strain tensor \tilde{s}_{ij} and a subgrid stress (SGS) tensor τ_{ij} :

$$\varepsilon = -2\langle \tau_{ij}\tilde{s}_{ij} \rangle \quad (7.24)$$

The SGS must be modelled by a small-scale turbulence model, as in Large Eddy Simulations. In this case, a Smagorinsky eddy viscosity model (Smagorinsky 1963) is used:

$$\tau_{ij} = -C_S^2 d_l^2 \langle |\tilde{s}_{ij}| \tilde{s}_{ij} \rangle \quad (7.25)$$

where $C_S = 0.17$ is the Smagorinsky constant.

The dissipation obtained with the three methods is plotted as a function of the streamwise location x/x^* in Fig. 7.16. As expected, the direct measurement with (7.21) achieves the smallest values due to the lack of resolution at the smaller scales where the dissipation takes place. The Large Eddy PIV technique results in a significant improvement of the dissipation measurement. The gap between the direct measurement and the dissipation obtained with the energy balance decreases along the decay. Indeed, while the spatial resolution is the same, the Kolmogorov scale (and so the dissipative scales) increases along the decay, as it depends on $\varepsilon^{-0.25}$ ($\eta = (v^3/\varepsilon)^{0.25}$, Kolmogorov 1941). The Kolmogorov scale obtained using the dissipation measured with the (7.23) ranges between $74\mu m$ and $96\mu m$ (i.e. about 7.6 and 5.9 times the effective d_l , thus justifying the significant underestimation of the dissipation with the direct measurement).

The scaling $\varepsilon = C_\varepsilon k^{3/2}/L$ is applied to evaluate C_ε (Fig. 7.16b). The results from the direct measurement show a significant variation of C_ε along the decay, most likely to be addressed to the increasing Kolmogorov scale along the

longitudinal direction (consider that smaller Re_λ correspond to larger x/x^* , as the turbulent kinetic energy strongly decays while the Taylor microscale shows only a weak rate of growth. The variation of C_ε observed with the data obtained with the Large Eddy PIV method and the energy balance is less remarkable, but still relevant (about 20% over the investigated range). Such variation is substantially at odds with RANS modelling in which the semi-empirical quantity C_ε is assumed to be constant.

7.4 Final remarks

Tomographic PIV is extremely challenged in applications involving small-scale turbulence measurement. The standard planar PIV has progressively reduced the gap with hot-wire anemometry over the years in this kind of application, mainly thanks to the technological advancement (higher signal to noise ratio in the imaging, advanced interrogation algorithms, etc.); Tomographic PIV, instead, has to walk over a much steeper path before being well-suited in this experimental scenario.

Tomo-PIV has provided the unquestionable advantage of assessing the inhomogeneity of the flow in the decaying region by measuring all the three components of the velocity field, and all the nine components of the velocity gradient tensor (at least on the central plane of the volume, since the thickness is limited by the desired high resolution). The results show that the flow in the decaying region is inhomogeneous in all directions, but the strength of the inhomogeneity is weak and the turbulent production is limited. The large-scale anisotropy persists over the range under investigation, but the small scales, less than one Taylor micro scale in size, can be considered approximately isotropic for $x > 0.6x^*$. The comparison of the longitudinal and transverse correlation functions with the corresponding ones under the hypothesis of isotropic turbulence underlines a satisfactory agreement for scales smaller than the Taylor microscale, while for larger scales the agreement is quite poor.

Three different methods for the dissipation measurements have been compared: direct measurement, Large Eddy PIV and balance of the turbulent kinetic energy. The output of the first method is quite poor due to the limited spatial resolution. This aspect is reflected also in the overestimation of the Taylor microscale. According to Fig. 7.6, there is no much room for higher resolution. A larger magnification could be obtained by moving the limit on the particle size towards larger values of M_0 , for example by reducing $f_\#$. On the other hand, smaller $f_\#$ means stricter requirements on the depth of field, while larger M_0 requires higher Scheimpflug angles, thus complicating the focusing and determining wide gradients of illumination and magnification on the camera images. These factors

might affect the tomographic reconstruction quality, thus reducing the accuracy of the measurement.

Large Eddy PIV and measurement of the dissipation from the turbulent kinetic energy balance seems to confirm the remarkable finding that the scaling of the turbulent dissipation with the ratio of the turbulent kinetic energy and the large eddy turnover time is not universally valid for all boundary-free turbulent shear flows (Hurst & Vassilicos 2007). This peculiar feature could be attributed to the effect of the fractal generator on the cascade mechanism; perhaps this initial pumping of energy at a range of different length scales does not allow for the normal separation of the energy containing and dissipative scales so that the small scales do not have “room” (spectrally speaking) to develop the conventionally expected behaviour.

Chapter 8 - Jet flows past a sudden expansion

In turbulent diffusion flames the role of the large coherent structures in the mixing of the fuel and the oxidant has a leading importance in determining the combustion efficiency, the flame stability, the amount of soot and NO_x, and so on. The requirements depend strongly on the burner; for example, gas turbines require the minimization of the radiant heat transfer towards the walls, and consequently a very short residence time; on the other hand, in furnaces and kilns the radiance is the leading mechanism in heat transfer, thus making very desirable a long residence time. In this second case, considerable attention has been devoted to the development of devices with the aim of exciting and enhancing the large scale coherent structures embedded into the shear layer of jets. The organization of these structures plays a key role in the transport of mass and momentum in flames (Mungal et al 1991). These devices can be bundled in three categories: acoustic excitation (see Reynolds et al 2003 for a review), which can be obtained both by external excitation (for example using speakers) or by self-excitation, enjoying the coupling between flow instabilities and acoustic resonance; mechanically oscillating devices (see for example Simmons et al 1981), effective in cold environment but not well suited for the combustion environment, in which the high temperature prevents the use of fast moving parts; fluidically excited devices, like flapping jets (Mi et al 2001a) or precessing jets (Nathan et al 1998).

The precessing jet (PJ) is generated with an axisymmetric jet flowing through a circular nozzle and subject to an abrupt expansion in a chamber that is coaxial with the nozzle. The jet reattaches asymmetrically after the abrupt expansion, and on the opposite side fluid is entrained into the chamber from the ambient. Instantaneous asymmetries trigger a rotating pressure field, inducing a precession of the jet, i.e. a rotation of the jet axis around the nozzle axis. The PJ, in this sense, is different from swirl flows, in which the rotation of the jet occurs around its own axis; on the other hand, the two jet flows have some similarities, as precession has been observed in proximity of the jet axis for relatively strong swirl (the so called precessing vortex core; see Syred 2006 for a review). The PJ has shown extremely interesting features, determining a reduction of the global flame strain rates within the flames, thus leading to an increase of the volume of soot. The larger amount of soot enhances the radiative heat transfer, reducing the flame temperature and the production of NO_x (see for more details Newbold et al 2000 and Nathan et al 2006). PJs have found application in industrial processes in which the radiant heat transfer plays a leading role over convection, such as rotary kilns for the production of cement and lime. Extensive studies have addressed the main statistical features of PJs in the near field (within the chamber) and the external field; however, the flow

field is extremely challenging, being it strongly unsteady and three-dimensional. Furthermore, PJs belong to the class of bifurcating flow instabilities arising in symmetrical configuration in absence of initial bias, which are usually difficult to be modelled and numerically simulated.

Very limited attention has been provided on the topology of the field within the chamber, mainly with flow visualization (Nathan et al 1998), phase-averaged investigation with pointwise techniques (Wong et al 2003) or with numerical simulations (Guo et al 2001, Revuelta et al 2002, Revuelta et al 2004). Tomographic PIV, in this sense, can provide a leap forward in understanding the organization of the flow topology in the complicate scenario of an unsteady, intermittent and strongly three-dimensional turbulent flow field. The task of this chapter is to provide a description of the underlying phenomenology of the FPJ, and to provide a benchmark for numerical simulations, that are extremely challenged in this scenario.

An overview of the main features of the phenomenology of the PJ is provided in Sec. 8.1. The experimental setup for the Tomo-PIV investigation of the near field of a PJ is described in Sec. 8.2. The analysis is conducted on the instantaneous and average flow field features (Sec. 8.3) and using the Proper Orthogonal Decomposition to assess the role of the leading large coherent structures (Sec. 8.4). The experiments are carried out in two scenarios: no initial swirl; inlet swirling flow, with two levels of initial swirl.

8.1 Fluidic precessing jet nozzles

The jet precessing motion is defined as the rotation of the jet about an axis other than its own one in a gyroscopic-like motion. In this chapter the basic configuration with a round nozzle of diameter d issuing into a cylindrical chamber with diameter D and length L (the nozzle and the chamber are concentric) will be defined as Fluidic Precessing Jet (FPJ) nozzle.

The first systematic investigations of the precessing motion in axisymmetric chambers date back to the '80s, with the works by Hallett & Günther (1984) and Dellenback et al (1988). In both cases the precession is induced by an upstream swirl motion in the inlet flow. Hallett & Günther (1984) performed experiments with an expansion ratio $D/d = 2.22$, and observed a precessing motion in the case of low swirl number. On the other hand, for high swirl number the flow field is dominated by the vortex breakdown, inducing the presence of a central recirculation zone. Interestingly enough, this phenomenon occurs at lower swirl number than required in case of unconfined swirling jets. Dellenback et al (1988) added to this scenario, with a systematic analysis, the effect of upstream swirl (0 to 1.2) for $D/d = 1.94$ and Reynolds number (referred to the inlet diameter) varying in the range 30,000-100,000. They found out that the precession motion and the

mean swirl occur in opposite directions at low swirl, and in the same direction for high swirl.

Luxton & Nathan (1989) investigated the effect of an abrupt expansion on axisymmetric jets, showing that an upstream swirl is not required to generate a precessing motion, provided that the expansion ratio is sufficiently large. Since the generated flow field within the chamber leads to a precessing motion at the exit of the chamber, this large scale excitation mechanism immediately appeared interesting for combustion applications, thus leading to the commercialized device *Gyro-Therm*, patented by the University of Adelaide (Luxton et al 1987). Even though the patent is consolidated, still much has not been understood of the topology of the flow field of FPJ.

In this section an overview of the main results reported in the literature is provided (Sec. 8.1.1), including a detailed description of the main flow phenomenology, and of the effects of the geometrical and experimental parameters (Reynolds number, aspect ratio of the chamber, see Sec. 8.1.2). Furthermore, some insight is provided on the measurement of the precession frequency in Sec 8.1.3.

8.1.1 The phenomenology of precession in axisymmetric jets

The description of the phenomenology of precessing jet is extensively provided in the seminal work by Nathan et al (1998), and summarized in the following. In their work particular care is taken to ensure that the inlet flow is symmetric and free of bias to avoid any influence of residual or induced swirl within the chamber. The flow exhibits an intermittent behaviour, switching between a precessing jet mode and an axial mode. The visualization of the instantaneous pathlines of the two modes, obtained by collecting the light scattered by air bubbles seeding the flow, is reported in Fig. 8.1. It has to be noted that in the configuration tested by Nathan et al (1998) an exit lip is added to the chamber in order to enhance the deflection of the exit jet and increase the asymmetry, which is the driving force of the precessing motion.

In the precessing mode the flow field is characterized by a continuously unstable reattaching jet, with a strong recirculation region located on the opposite side of the chamber. A transverse pressure gradient in the outflow is established, thus determining a sharp deflection of the wall jet at the exit of the chamber. Since the flow is in a condition of neutral equilibrium, the effect of any asymmetry and/or the turbulence fluctuations induces the reattachment point to move; as the jet starts to rotate along one direction, the asymmetry of the flow entrained into the chamber induces the establishment of a rotating pressure field. Indeed, the recirculating fluid moves upstream within the chamber and swirls in the opposite direction to that of the precession, thus retaining the net angular momentum equal to zero. This aspect is sketched in Fig. 8.2.

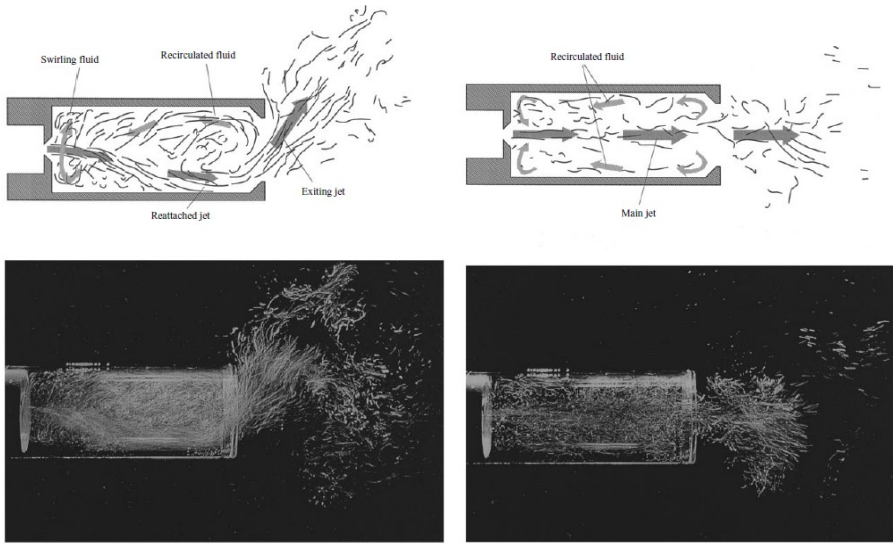


Fig. 8.1 Water flow visualization (bottom) and schematic interpretation (top) of the instantaneous pathlines in the precessing jet mode (left) and the axial mode (right). (Nathan et al 1998).




Fig. 8.2 Schematic interpretation of the interaction between swirling entrained fluid and precessing outflow (Nathan et al 2006).

The axial mode is characterized by a statistically axisymmetric spreading of the jet. Nathan et al (1998) noted that the flow does not reattach to the wall within the chamber, but it does interact with the lip, determining the occurrence of a wide recirculation region in proximity of the wall and the rising of large coherent structures (significantly larger than the case of the unconfined jet). Nevertheless, the flow visualizations by Nathan et al (1998) testify that the spreading rate of the jet outside the chamber is still much larger in the case of the precessing mode than

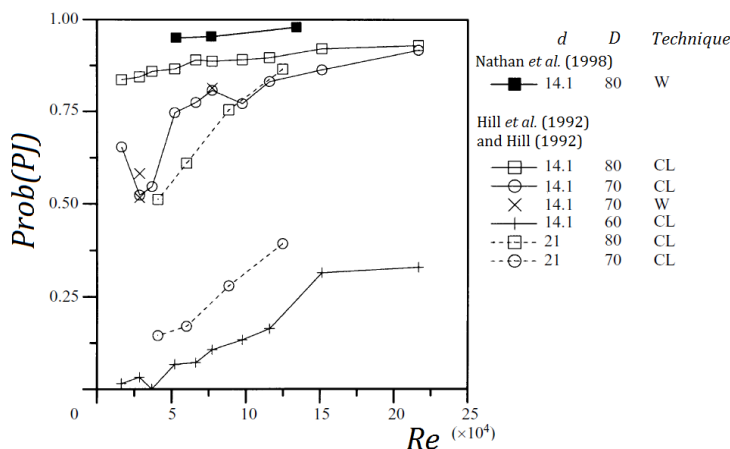


Fig. 8.3 Precession probability as a function of the inlet Reynolds number. The discrimination between the two modes is obtained by using the pdf of the pressure measurement at the wall (indicated with *W*), or the signal from a total pressure probe on the centreline (indicated with *CL*). (adapted from Nathan et al 1998).

in the axial one, thus inducing the wish to maximize the probability of occurrence of the precession with respect to the axial outlet.

8.1.2 The effects of geometry and *Re* on the precession probability

Several studies have concentrated their focus on developing solutions to favour a stable precessing motion. Nathan et al (1998), and many following studies of the same research group, fixed the expansion ratio $D/d = 5$, which has shown to be very favourable for the precession, and investigated the effects of the chamber aspect ratio, the Reynolds number and the chamber geometry on the probability of precession. Incidentally, a minor (but still significant) influence is due to the characteristics of the inlet flow; in fact, Wong et al (2004) verified that a smooth contraction nozzle results in a lower precession probability than inlet from an orifice plate or a long pipe. This effect is related to the different initial behaviour of the boundary layer, that for the nozzle with a smooth contraction inlet is characterized by a symmetrical shedding initial boundary layer, while the sharp-edged orifice generates an asymmetric shedding of roll-up structures, and the pipe jet produces randomly distributed large structures into the shear layer (Mi et al 2001b).

Nathan et al (1998) introduced an exit lip to increase the deflection of the exit wall jet, thus enhancing the transverse pressure gradient and the asymmetry in the entrained flow. They measured the precession probability as a function of the inlet-based Reynolds number for an aspect ratio of the chamber $L/D = 2.75$ by using the

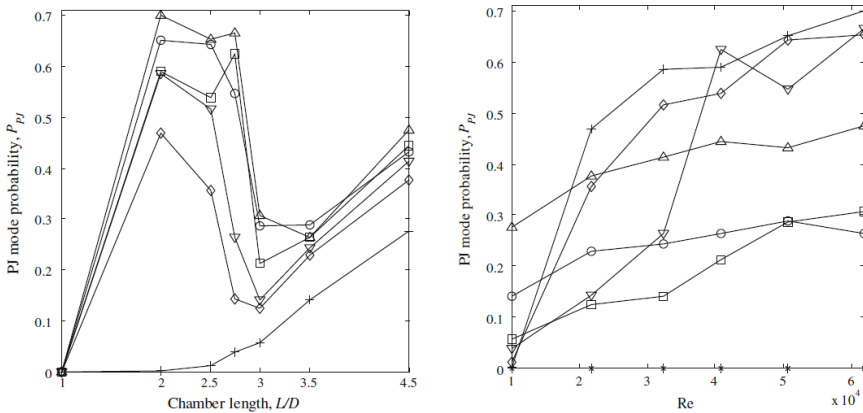


Fig. 8.4 Precession probability as a function of: chamber aspect ratio L/D (left); inlet Reynolds number (right). On the left: + $Re = 10000$, $\diamond Re = 21800$, $\nabla Re = 32400$, $\square Re = 40800$, $\circ Re = 50700$, $\triangle Re = 61900$. On the right: * $L/D = 1.0$, + $L/D = 2.0$, $\diamond L/D = 2.5$, $\nabla L/D = 2.75$, $\square L/D = 3.0$, $\circ L/D = 3.5$, $\triangle L/D = 4.5$. (Madej et al 2011).

probability distribution function of the wall pressure measurements. The results reported in Fig. 8.3 show that the precession probability is an increasing function of the inlet Reynolds number, and that FPJ nozzles with larger expansion ratio are more favourable to the precessing motion (within the range of the tested ratio, i.e. $3.3 < D/d < 5.7$).

More recently, Madej et al (2011) investigated the effects of variable aspect ratio and Reynolds number on the precession probability by using Stereoscopic PIV at the exit of a FPJ nozzle with $D/d = 5$ and no exit lip. The mode determination is performed using three criteria, based on the maximum outflow velocity position, centroid location of the volume flow rate and volume flow beyond a cut-off radius. The three methods lead with relatively low uncertainty to the results reported in Fig. 8.4. In the case of $L/D = 1$ the flow field does not exhibit the precessing mode in the range of tested Reynolds number. Chamber aspect ratios included in the range $2 \leq L/D \leq 2.75$ are very favourable for the precession, with probability of precessing mode increasing with the inlet Reynolds number. In all cases, the relative increase of the probability with the Reynolds number progressively decreases, as observed also by Nathan et al (1998). The probability significantly drops down for $L/D = 3$, with a partial recover for larger aspect ratios.

An improved version of the FPJ basic nozzle is obtained by adding a bluff body along the centreline of the chamber in proximity of the exit, as shown in Fig. 8.5. The centrebody ensures nearly 100% probability of precession, provided that the Reynolds number is sufficiently large (Wong et al 2004). The first extensive description of the flow field within the chamber in the chamber-centrebody-lip




Fig. 8.5 Sketch of the chamber-lip-centrebody configuration (Wong et al 2003).

configuration is provided by Wong et al (2003). They identified with phase-averaged Laser Doppler Anemometry (LDA) the asymmetry of the reattaching jet and measured a spreading rate significantly larger than turbulent free jets in the external field. The limits of the measurement technique, on the other hand, provide only phase-locked measurements, giving up any instantaneous field information. A more detailed outline of the external velocity field is obtained by Wong et al (2008) using phase-triggered PIV and surface visualization. The identification topology of the external velocity field goes beyond the scope of the investigation addressed in this chapter; the reader is referred to Wong et al (2008) for an exhaustive description.

8.1.3 Determination of the Strouhal number of the precession

Measuring the Strouhal number, i.e. a normalized frequency, associated with the precession, and its dependence on the geometric and flow parameters is of fundamental importance, as it influences the mixing outside the chamber, in the region where the flame will be located. Nathan & Luxton (1992a, 1992b) verified that, for the case of the FPJ, the following definition best collapses the data:

$$St_h = \frac{f_p h}{V_j} \quad (8.1)$$

where f_p is the precession frequency, h is the step height (that is equal to $[D - d]/2$) and V_j is the bulk velocity of the inlet jet. However, in the literature the definition St_d based on the inlet diameter d instead of the step height is most widely used, especially for the case of the structures embedded in shear layers. Nathan et al (1998) measured the precession frequency by identifying the peak in the power spectrum of the signal from a total pressure probe placed along the axis of the chamber, and collected data on the measured Strouhal number for earlier




Fig. 8.6 Schematic of the mechanical precessing nozzle (Mi & Nathan 2004).

investigations on FPJ and on a flapping jet into a rectangular chamber (Shakouchi 1981, 1989). Interestingly enough, in both cases the obtained Strouhal number ranges is 0.001-0.005, i.e. one order of magnitude below the typical values related to the shedding of structures embedded within the shear layer; furthermore, the results are obtained with different fluids (water, air) and Mach number varying in a significant range (0-0.7), thus excluding any acoustic coupling with resonant modes to be responsible of the precession.

More recently, Mi & Nathan (2004, 2006) investigated the influence of the chamber length, the inlet geometry and the Reynolds number on the precession frequency. They observed that St_d increases almost linearly with both the chamber length and the jet velocity. In addition to this, the effect of St_d is much more significant than that of Re_d (i.e. the Reynolds number based on the inlet diameter) in determining the mixing of the fluids outside the chamber.

Due to the intrinsic difficulty in isolating the precession features in a continuously intermittent flow field, in some explorative studies the effects of the precession frequency on the mixing characteristics have been quantified using a Mechanical Precessing Jet nozzle (MPJ, Schneider et al 1997). The MPJ is sketched in Fig. 8.6. While the MPJ has the unquestionable advantage of enabling the independent variation of the parameters influencing the initial conditions, it is not well suited for the applications in combustion due to the presence of moving parts;

furthermore, there are quite significant evidences of important differences between the FPJ and the MPJ. For example, the jet emerging from the FPJ is not circular, but “kidney-shaped” (Wong et al 2008, Madej et al 2011); furthermore, the MPJ generates a central recirculation zone in the near field, while the FPJ does not. For this reason a more robust method to obtain phase-locked measurements and isolate the precession frequency is provided in this chapter.

8.2 Experimental details

8.2.1 Water facility

The experiments are carried out in a water facility at the University of Naples Federico II. The jet is issued from a circular nozzle (with diameter $d = 20\text{mm}$) installed on the bottom of a nonagonal plexiglass tank (internal diameter 600mm , height 700mm); the plexiglass walls allow a complete access for both the illumination and the camera imaging. The jet expands into a cylindrical chamber, coaxial with the nozzle, with diameter $D = 100\text{mm}$ and length $L = 275\text{mm}$ (so that $D/d = 5$ and $L/D = 2.75$). While in many investigations an exit lip and/or a centre-body have been included into the chamber to favour the precession mode with respect to the axial mode and to condition the exit angle, in this fundamental study these solutions are not considered in order to assess the topology of the flow field without external conditioning.

Stabilized water supply of 2.3kg/s is provided upstream of the nozzle by a centrifugal pump and is laminarized by passing through flow-conditioning grids and honeycombs installed in the plenum chamber. The diameter of the plenum chamber is $5d$ and its length is $20d$, thus no significant effects of fluctuations, or bias due to swirl, are expected in the flow at the inlet of the chamber. In order to check that the effects of asymmetry and bias are negligible, the profiles of the velocity and the turbulent statistics at the exit of the nozzle have been analyzed with planar PIV for the case of outflow without the external chamber. The profiles of the longitudinal velocity V/V_0 (where V_0 is the centreline velocity) and of the root mean square of the longitudinal velocity fluctuations v'/V_0 (Fig. 8.7) outline that the flow can be reasonably assumed symmetrical and unbiased.

The experiments are carried out with three different outlet conditions: no swirl, low swirl (swirl number equal to 0.2) and high swirl (swirl number equal to 0.6). The swirl is generated with helical inserts as reported by Ianiro & Cardone (2012). In all cases the bulk velocity V_j of the jet entering the chamber is about 7.5 m/s , thus resulting in $Re_d = V_j d/\nu = 150 \cdot 10^3$. In this scenario the precessing mode is expected to considerably prevail on the axial mode; indeed, Madej et al (2011) indicated 84% of probability of precessing motion for $Re_d = 61,900$, $L/D = 2.75$ and $D/d = 5$; furthermore the probability increases with Re_d (even if only a weak dependence is reported in the literature).




Fig. 8.7 Profiles of the inlet velocity V/V_0 • and of the root mean square of the axial fluctuations v'/V_0 ■ (where V_0 is the centreline velocity) of the jet without confinement.

Unless otherwise stated, the letters U, V, W indicate the velocity components along the width, the height and the depth of the measurement volume (respectively X, Y, Z). The corresponding lower case letters u, v, w refer to the turbulent velocity fluctuations obtained by subtracting the mean velocity components from the instantaneous realizations. The angled brackets $\langle \dots \rangle$ indicate the operation of ensemble averaging over the set of realizations. Finally, the symbols u', v', w' are used to refer to the root mean square (*rms*) of the turbulent velocity fluctuations.

8.2.2 Tomographic PIV system

The flow is seeded with neutrally buoyant polyamide particles with average diameter of $56\mu\text{m}$, dispersed homogeneously within the facility with a concentration of approximately $0.15\text{ part}/\text{mm}^3$. Laser pulses are produced with a double-cavity Gemini PIV Nd:Yag system (532nm , $200\text{mJ}/\text{pulse}$, 5ns pulse duration). The exit beam of 5mm diameter is shaped into a parallelepiped volume; a knife-edged slit is placed along the laser path to set the thickness of the volume to about 34mm . The light scattered by the particles is collected by a tomographic system composed of four LaVision Imager sCMOS 5.5 megapixels cameras (2560×2160 pixels resolution, pixel pitch of $6.5\mu\text{m}$, 16bit resolution in intensity). The cameras are disposed in linear arrangement, as sketched in Fig. 8.8, covering an angle of 90° . The cameras are equipped with 100mm EX objectives, set at $f_\# = 16$ and equipped with Scheimpflug adaptors to obtain well focused particle images throughout the volume. The average magnification in the centre of the measurement volume is about 0.06, thus resulting in a depth of field of more than




Fig. 8.8 Sketch of the illumination and the camera arrangement of the Tomographic PIV system.

200mm according to (1.3), and a particle image diameter of 3.4 pixels according to (1.2).

Sequences of the tracers particles with time separation of $130\mu\text{s}$ are taken at a frequency equal to 10Hz . Considering the typical values of the Strouhal number (see Sec. 8.1.3), the acquisition frequency does not suffice for the temporal resolution of the vortices within the shear layer, but is certainly higher than the Nyquist frequency for the precession motion (that is expected to lie in the range $0.4 - 1.5\text{Hz}$).

8.2.3 Calibration correction

The three-dimensional mapping functions are obtained by recording images of a calibration target mechanically translated along the depth direction of the measurement volume, as described in Sec. 1.2.1, in the range $\pm 20\text{mm}$. The calibration is performed with a double/plane target, with white dots on dark background. The spacing of the markers on the same plane is 15mm ; the separation between the planes is 3mm . A three-dimensional polynomial function, 3rd order in x and y , and 2nd order in z , is fitted to the calibration points correspondence to build the mapping functions for each camera of the tomographic setup. A template-matching technique, with a cross-correlation based algorithm, is used to identify the location of the markers. The *rms* of the initial calibration error is quite high (about 0.8 pixels) due to the relatively large size of the dots (about 30 pixels on the camera planes).

The challenge in this application of the procedure resides in the impossibility to perform the calibration *in-situ* due to physical restrictions, i.e. the calibration is performed without the presence of the tube. For this reason the self-calibration is the keystone for a successful experiment, as it can correct the mapping functions to account for misalignment of the lines of sight due to the viewing angle through the Plexiglas chamber, determining refraction effects along the optical path.

An automatic iterative self-calibration procedure (see more details on the standard self-calibration in Sec. 2.1.4) has been implemented to gradually correct the misalignment and registration errors, which can be significant due to inserting the chamber after the calibration. The procedure is a multi-grid algorithm similar to the single-image self calibration implemented by Michaelis & Wolf (2011), in this case applied to 200 images. In each step the image density is artificially reduced by considering only the brightest particles on the original images. The iterative procedure strongly relies on the reliability of the measured disparity peaks at the previous step. For this reason a signal/noise (SN) ratio is defined as the ratio of the disparity peak and the second maximum within the map. A value of 1.5 is considered a reasonable threshold to validate the disparity peaks.

The procedure is outlined in the following:

- The misalignment due to net displacement between the cameras is corrected by computing disparity maps relative to the entire measurement volume. For this step the 5000 brightest particles of each image have been considered (i.e. about 4% of the total particles, equivalent to $0.002ppp$), with a search radius of 8 pixels. The maximum correction applied on the camera is of 0.6 pixels. The average SN is about 5;
- A refinement of the grid is performed: the measurement volume is divided in $3 \times 3 \times 3$ sub-volumes and the disparity maps are computed using the 10000 brightest particles. In this step, possible relative rotation between the cameras can be corrected. The corrections range between $-0.7 \div 0.7$ pixels (i.e. larger than the previous step), and this is due to the fact that the left and right sides of the measurement volume are much more affected by the optical distortions due to chamber. The disparity maps for one camera relative to the central plane are reported in Fig. 8.9 (top left). For all sub-volumes the SN is larger than 2, with peaks over 30.
- The grid is refined in 3 additional steps ($5 \times 5 \times 3$ volumes with 20000 particles, $7 \times 7 \times 5$ volumes with 40000 particles, $9 \times 9 \times 5$ volumes with all the detected particles). The relative disparity maps for one camera relative to the central plane are reported in Fig. 8.9. A significantly lower SN is observed on the left and right side of the volume, with some vector rejections. This is due to the effect of residual reflections, obscuring particles and consequently making it difficult the matchings on the different images. Furthermore particles are

elongated and distorted due to optical aberrations, thus increasing the uncertainty in position measurement.

At the end of the self-calibration procedure, the root mean square of the calibration error is brought down to 0.01 pixels, with a maximum error of 0.08 pixels, acceptable for a successful tomographic reconstruction.

8.2.4 Tomographic reconstruction and CC analysis

A measurement volume of $100 \times 250 \times 34\text{mm}^3$ (i.e. $1D \times 2.5D \times 0.34D$) is reconstructed using the Multi-Resolution algorithm outlined in Chapter 3, with MLOS initialization, 3 MART iterations on a binned $2\times$ configuration and 2 final MART iterations on the final resolution. A further accuracy improvement is obtained by applying the SFIT technique with anisotropic filtering on a $3\times 3\times 1$ kernel, with Gaussian distribution of weights and standard deviation equal to 1. The volume is discretized with $10\text{vox}/\text{mm}$, thus resulting in a reconstruction volume of $1000 \times 2500 \times 340$ voxels.

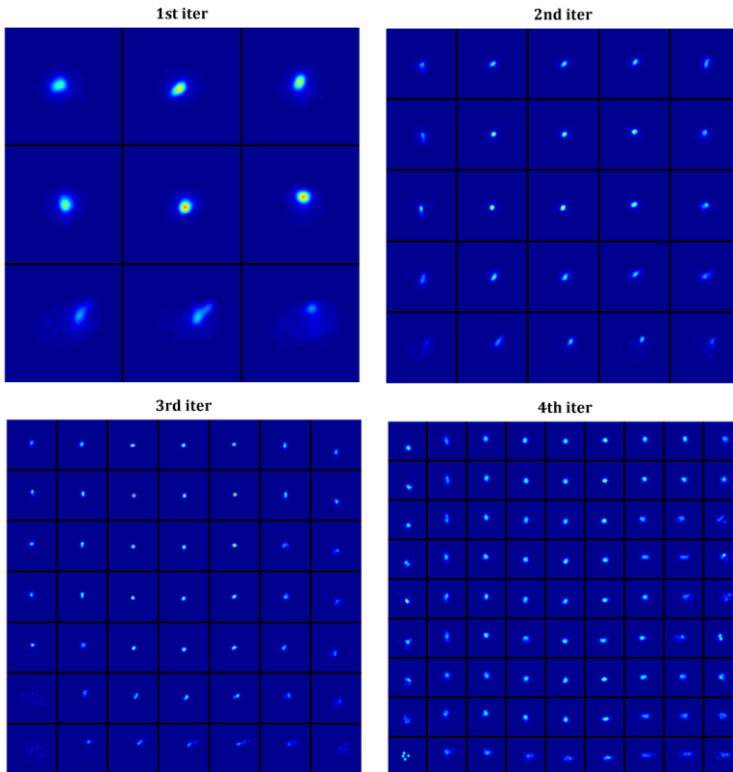


Fig. 8.9 Disparity maps on the middle plane of the measurement volume for one camera in the 4 steps of the self-calibration process.




Fig. 8.10 Intensity profile along the depth direction (left) and pdf of the divergence of the raw velocity field (right, continuous line) and of the filtered velocity field (right, dotted line).

The background on the raw images is eliminated by removing the historical minimum on each sequence of 500 images; the residual background intensity is deleted by performing a sliding minimum subtraction in both space (13×13 pixels) and time (5 samples). The quality of the reconstruction is evaluated *a-posteriori* by computing the sum of the intensity on the xy planes (i.e. the z -profile of the light intensity within the volume), as done in Sec. 4.2.3. The profile in Fig. 8.10 (left) shows that the ratio between the intensity in the illuminated region, which is due to both true and ghost particles, and the ghost intensity outside the volume, is quite limited, but still acceptable for the 3D velocity measurement.

The cross-correlation analysis is performed with the efficient algorithm based on direct sparse cross-correlations and redundancy avoidance outlined in Sec. 5.1. The final interrogation spot is 64^3 voxels (corresponding to $6.4 \times 6.4 \times 6.4 \text{ mm}^3$) with 75% overlap (thus resulting in a vector spacing of 1.6 mm). The uncertainty on the velocity measurement can be assessed by applying physical criteria, for example by computing the divergence of the velocity field, as in Sec. 4.3. The uncertainty in the divergence is both due to the measurement error on the velocity and the numerical truncation in the derivative calculation; however, the 75% overlap reduces this second source of error, thus making it possible to quantify with reasonable approximation the uncertainty on the velocity measurement using the standard deviation of the divergence. In this case, the data are presented in non-dimensional form by referring the data to the typical value of the vorticity within the shear layer in the case of the circular jet (0.2 voxels/voxel). For the raw velocity field the standard deviation is $0.028 \text{ voxels/voxel}$ (0.14 in the non-dimensional version); the uncertainty is reduced to $0.023 \text{ voxels/voxel}$ if a low-pass Gaussian filter on a kernel $3 \times 3 \times 3$ and standard deviation equal to 1 is applied.

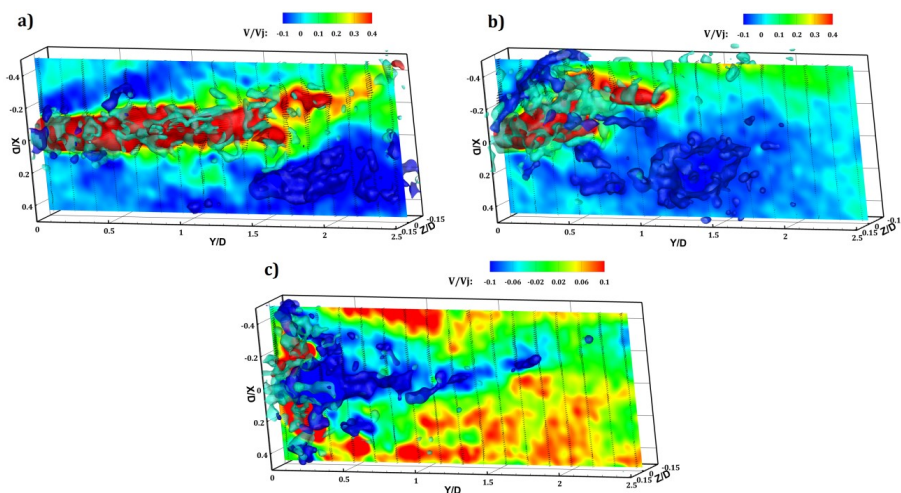


Fig. 8.11 Contour representation of the longitudinal velocity component V/V_j on the middle plane of the measurement volume, and iso-surfaces of $V/V_j = 0.4$ (red), $V/V_j = -0.1$ (blue) and $Q > 0$ (light blue) for the case of: **a)** $S = 0$; **b)** $S = 0.2$; **c)** $S = 0.6$.

8.3 Results

8.3.1 Instantaneous flow features

Snapshots of instantaneous realization of the velocity field for the three tested nozzles are presented in Fig. 8.11. The normalized longitudinal velocity component V/V_j is presented with a contour representation on the middle-plane of the measurement volume along the depth direction; iso-surface of positive and negative velocity, and of positive Q (where Q is the second invariant of the velocity gradient tensor; when $Q > 0$ rotation prevails on shear, so that it can be used as a vortex identification criterion; see Jeong & Hussain 1995 for further details) are also reported. The results are presented in non-dimensional form, using the bulk jet velocity V_j and the chamber diameter D as a reference.

In the case of the circular nozzle, i.e. $S = 0$, the jet asymmetrically attaches to the wall and a wide entrainment region is formed on the opposite side of that of impingement. Interestingly enough, the point of attachment is beyond $Y/D = 2$, in contrast with the surface visualization reported by Nathan et al (1998), reporting the impingement point to be located about at half-height of the chamber. Such a discrepancy is addressed to the presence of the exit lip, which exasperates the exit angle (and, consequently, the swirl number of the outlet of the chamber), thus inducing the attachment point to move upstream due to the stronger induced swirl. This hypothesis is confirmed by the analysis of the instantaneous flow field in case of $S = 0.2$, in which adding a swirl in the upstream flow determines both the

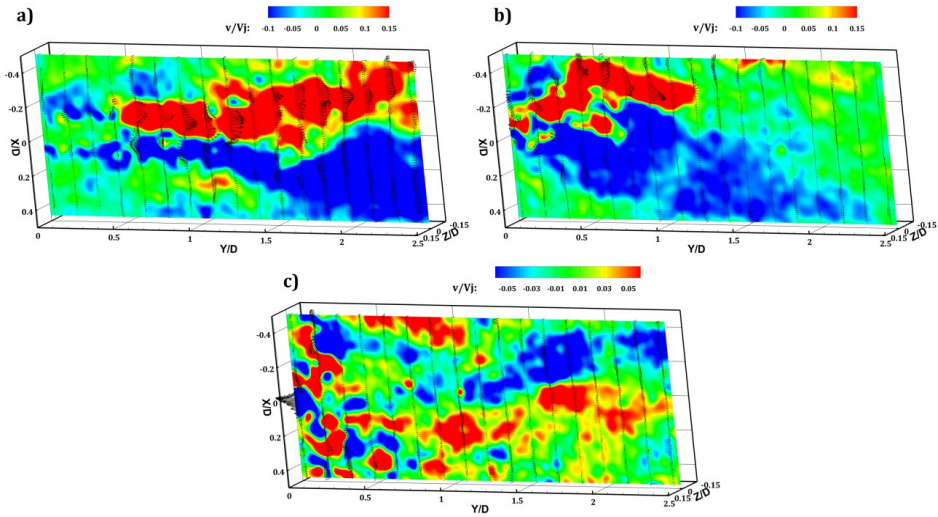


Fig. 8.12 Contour representation of the longitudinal fluctuating velocity component v/V_j (relative to the same instantaneous realization of Fig. 8.11) on the middle plane of the measurement volume: **a)** $S = 0$; **b)** $S = 0.2$; **c)** $S = 0.6$.

recirculation regions and the impingement point to be located closer to the inlet nozzle. In both cases the Q -criterion identifies the main vortical structures to be located within the shear layer of the inlet flow, appearing in the form of roll-up vortices generated by the Kelvin-Helmholtz instability. Furthermore, a significant recirculation region is identified in the upstream part of the chamber below the attachment point; the recirculation is remarkably more intense in the case of $S = 0.2$.

The instantaneous flow field for $S = 0.6$ is extremely different. In this case the flow field is dominated by the vortex breakdown, typical of flows with significant swirl. The central recirculation region is elongated along the chamber axis, leading to a substantial entrainment of fluid from the external ambient. These features were foreseeable, given the quite large swirl number. Indeed, in swirling flows the radial pressure gradient induces an intense expansion of the jet flowing out from the nozzle, thus determining an adverse axial pressure gradient; when the swirl number is sufficiently high, this phenomenon induces the formation of a central recirculation zone. The vortex breakdown has an intimate bond with another precessing motion, which shares similar features with the fluidic precessing jets, i.e. the precessing vortex core, in which the recirculation region precesses around the axis of symmetry of the jet as a helical structure. The reader is referred to more extensive reviews on this topic, provided by Lucca-Negro & O'Doherty (2001) and Syred (2006).

More interesting features can be highlighted by observing the fluctuating velocity fields obtained by a Reynolds decomposition, i.e. subtracting the ensemble average of the entire set of realizations from each velocity component. For the case of $S = 0$ and $S = 0.2$ the data of Fig. 8.12 highlight the presence of a pseudo-wall jet and a strong entrainment region on the opposite side of that of impingement. In absence of swirl the velocity fluctuations due to the precession are significant only for $Y/D > 0.5$, i.e. the jet is undisturbed and quasi-axisymmetric in its early region of development. It is reasonable to assume that this longitudinal abscissa can be considerably smaller in presence of an exit lip or an exit centrebody, thus influencing the spreading rate of the jet in the very near field. In presence of a weak swirl, strong longitudinal velocity fluctuations are evident up to the inlet nozzle, and this is due to the fact that the precession is induced (and, in a certain sense, powered) by the upstream swirl itself. Again, a wide recirculation region is observed, extending almost up to the nozzle inlet.

The case of $S = 0.6$ presents significant discrepancies with the two cases previously analyzed. The near field is characterized by extremely large velocity gradients and the dominance of small-scale turbulent structures due to the mixing effects induced by the vortex breakdown. Evidences of a large-scale precession of the vortex core are present in the upstream part of the chamber, where a fluctuating entrainment and a corresponding outlet are clearly distinguishable. Differently from the case of the circular nozzle and weak swirl, the entrainment and the fluctuating outlet occur much closer to the chamber axis.

8.3.2 Mean flow features

Even though the motion is extremely three-dimensional and intermittent, some interesting considerations can be drawn by observing the mean flow features. In Figs. 8.13-8.14 the profiles of the mean velocity components and of the root mean square of the relative velocity fluctuations are reported at three longitudinal location in the near field ($Y/D = 0, 0.25, 0.5$) for the cases of $S = 0$ and $S = 0.6$ (the results for $S = 0.2$ present qualitatively similar features with $S = 0$, and are not included for brevity).

The profile of $\langle W \rangle / V_j$ for $S = 0$ highlights the presence of a significant swirl induced within the lower region of the chamber. This effect is much stronger than the rate of entrainment of the jet in the near field, as testified by comparison with the profile of $\langle U \rangle / V_j$. It is important to note that the direction of swirl is connected to that of the precession. Physical arguments on the balance of angular momentum suggest that the swirl occurs in the direction opposite to that of precession (in agreement with the observations by Dellenback et al 1988). A more detailed discussion and evidences of this phenomenon are reported in Sec. 8.4.4.

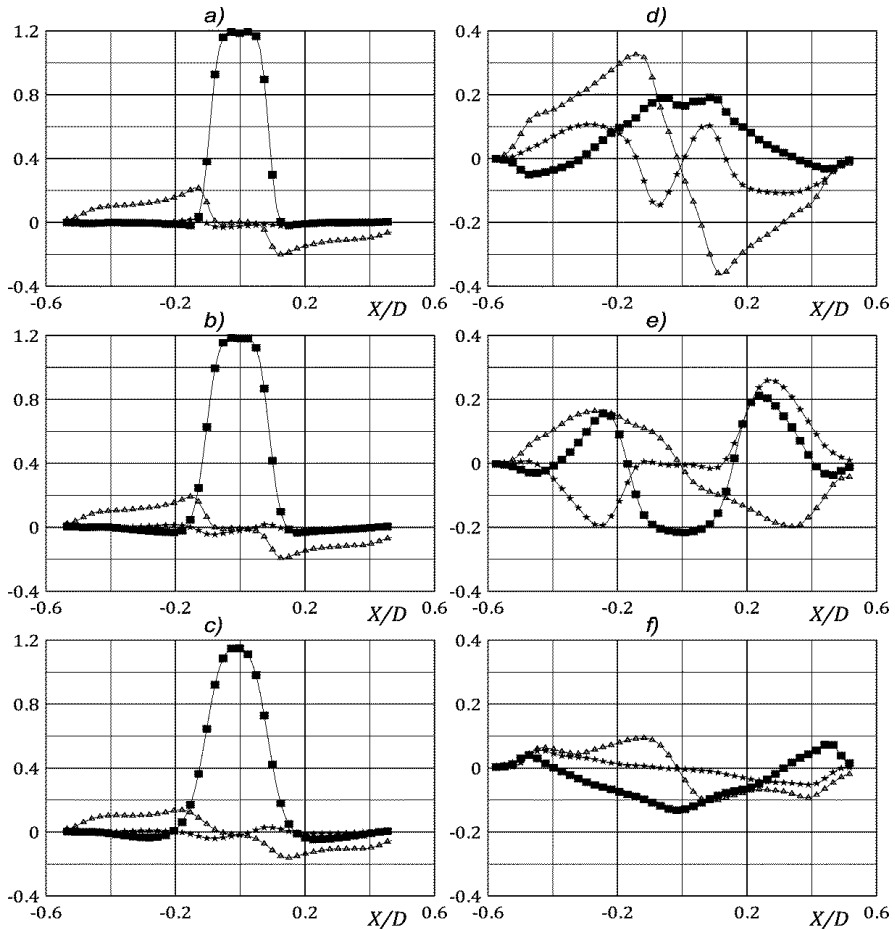


Fig. 8.13 Radial profiles of the mean velocity components for $S = 0$ (a,b,c) and $S = 0.6$ (d,e,f) for $Y/D = 0, 0.25, 0.5$ (top to bottom). $\star \langle U \rangle / V_j$ $\blacksquare \langle V \rangle / V_j$ $\triangle \langle W \rangle / V_j$. Symbols are placed each 4 computed vectors.

For the case of $S = 0.6$ the flow field is dominated by the vortex breakdown, as revealed by the profiles of $\langle V \rangle / V_j$. At the nozzle inlet ($Y/D = 0$) a positive longitudinal velocity is observed, with two weak regions of negative velocity due to the recirculation regions in proximity of the wall. Moving longitudinally along the axis of the chamber, a wide intense recirculation region is formed due to the vortex breakdown. The effects of the intense adverse axial and radial pressure gradients determine a rapid transfer of spectral energy from the large to the small scales of turbulence, thus enhancing the turbulent mixing at the expense of the mean flow.

The turbulent fluctuations profiles along the X direction are reported in Fig. 8.14. The profiles of the root mean square of the velocity fluctuations for the case

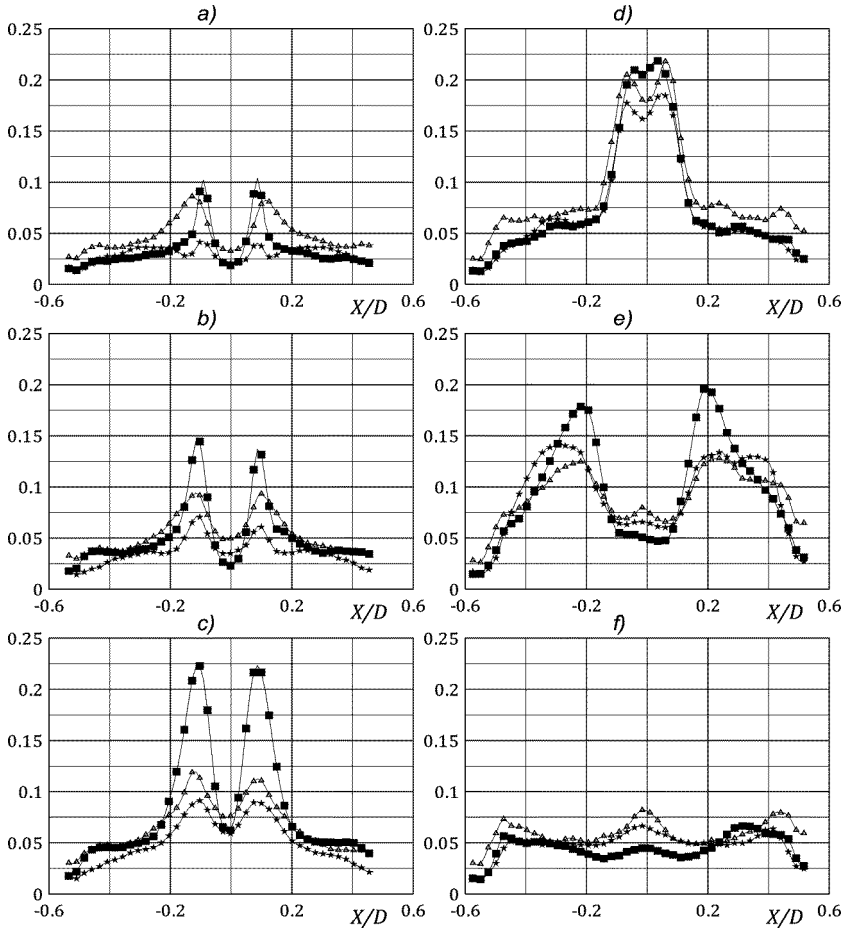


Fig. 8.14 Radial profiles of the rms of the normalized turbulent fluctuations for $S = 0$ (a,b,c) and $S = 0.6$ (d,e,f) for $Y/D = 0, 0.25, 0.5$ (top to bottom). $\star u'/V_j$ $\blacksquare v'/V_j$ $\triangle w'/V_j$. Symbols are placed each 4 computed vectors.

of $S = 0$ indicate, as expected, a predominance of the longitudinal fluctuations within the shear layer. In addition to this, the effect of the induced upstream swirl determines larger turbulent fluctuations w' than u' (the last one being mainly induced by the entrainment for the chosen profile). This gap persists over the entire near field region, while it disappears for $Y/D > 1$ (see Fig. 8.15), where the large scale mixing due to the precession motion is dominant. In case of high swirl the turbulence intensity is considerably larger at the inlet than in the case of $S = 0$. Moving downstream, the vortex breakdown determines a rapid decay of the intensity of the turbulent kinetic energy in the central recirculation zone.




Fig. 8.15 Longitudinal profiles of the normalized rms of the turbulent fluctuations along the centreline of the chamber for $S = 0$ (a) and $S = 0.6$ (b). $\star u'/V_j$ $\blacksquare v'/V_j$ $\triangle w'/V_j$. Data are plotted with a step of 8 vectors for clarity.

8.4 POD analysis and Low Order Reconstruction

As already outlined in the previous sections, the limited acquisition rate is not large enough to allow time-resolved measurements, but it does suffice to solve the large scale precessing motion, at least for $S = 0$ and $S = 0.2$ (for the case of high swirl this condition is not satisfied, since the flow field is dominated by the precessing vortex core and the vortex breakdown, whose typical Strouhal number is usually one-two orders of magnitude larger than that associated to the precession). On the other hand, the acquisition is not triggered by pressure signals (as done by Wong et al 2008, among the others), thus introducing the problem of the phase identification to obtained phase-averaged measurements. The task is quite challenging, since the flow field is intermittent and highly turbulent.

Proper Orthogonal Decomposition (POD) is a powerful instrument to extract information on the coherent structures in turbulent flows (Berkooz et al 1993). Furthermore, in strongly periodic flow fields, the flow phase information can be inferred using a relatively low number of POD modes (Ben Chiekh et al 2004, Van Oudheusden et al 2005, Meyer et al 2007). In this section the POD implementation is briefly outlined, and the most significant POD modes are illustrated for the three tested nozzles. A low order reconstruction method, using the most energetic modes, is implemented to extract phase-averaged information on the precessing motion for $S = 0$. Finally, the extracted phase data are used to estimate the precession frequency range.

8.4.1 POD implementation

The Proper Orthogonal Decomposition is a mathematical procedure aimed to identify an orthonormal basis using functions estimated as the solution of the

integral eigenvalue problem known as a Fredholm equation (see Berkooz 1993 or Sirovich 1987 for a more rigorous formulation). Suppose that a function $\underline{U}(\underline{x}, t)$ is approximated by:

$$\underline{U}(\underline{x}, t) = \langle \underline{U}(\underline{x}, t) \rangle + \underline{u}(\underline{x}, t) = \langle \underline{U}(\underline{x}, t) \rangle + \sum_{n=1}^{N_m} a_n(t) \underline{\varphi}_n(\underline{x}) \quad (8.2)$$

Without any loss of generality, let \underline{x} and t be the spatial and time coordinates, respectively. The symbols $\langle \underline{U}(\underline{x}, t) \rangle$ and $\underline{u}(\underline{x}, t)$ indicate the mean and the fluctuating part of the velocity field, respectively; the functions $\underline{\varphi}_n(\underline{x})$ constitute the decomposition basis of the fluctuating velocity field, and $a_n(t)$ are the time coefficients; the symbol N_m indicates the number of modes. Evidently, the solution for the decomposition (8.2) is not unique, as it depends on the chosen basis functions $\underline{\varphi}_n(\underline{x})$. Furthermore, another difficulty arises in determining the time coefficients given the set of basis functions.

One possible solution to determine the $a_n(t)$ consists in using orthonormal basis functions, i.e.:

$$\overline{\underline{\varphi}_i(\underline{x}) \cdot \underline{\varphi}_j(\underline{x})} = \delta_{ij} \quad (8.3)$$

where δ_{ij} is the Kronecker delta symbol, and the overbar indicates spatial integration over the measurement domain. This choice leads to:

$$a_n(t) = \overline{\underline{u}(\underline{x}, t) \cdot \underline{\varphi}_n(\underline{x})} \quad (8.4)$$

Therefore, in case of orthonormal basis functions, each coefficient $a_n(t)$ is dependent only of the pertaining $\underline{\varphi}_n(\underline{x})$.

The remaining problem is the criterion to find the set of orthonormal functions. Among the infinite possible solutions, POD aims to find the set of basis which is optimal in a least square sense. In the scenario of the formulation of the POD snapshot method by Sirovich (1987), it can be demonstrated that the POD modes are the eigenmodes of the two-point temporal correlation matrix \underline{R} :

$$\underline{R}\underline{\varphi} = \lambda\underline{\varphi} \quad (8.5)$$

with $R_{ij} = \langle \underline{u}(\underline{x}_i, t) \cdot \underline{u}(\underline{x}_j, t) \rangle$. Since R_{ij} is a non-negative Hermitian matrix, it has a complete set of non-negative eigenvalues, whose magnitude indicates the energy contribution of the respective eigenmodes. Indeed:

$$\overline{\langle \underline{u}(\underline{x}_j, t) \cdot \underline{u}(\underline{x}_i, t) \rangle} = \sum_{n=1}^{N_s} \lambda_n \quad (8.6)$$

The POD modes with the largest eigenvalues determine the most dominant flow field features in terms of energy. In the case under investigation the precessing motion is expected to give a dominant contribution, thus allowing the possibility of low-order reconstructions using only a small number of modes (as in Ben Chieck

2004), or build phase-averaged measurements by identifying the phase of each realization.

8.4.2 POD modes

The snapshots method provides, in principle, a number of modes equal to the number of snapshots (in this test case, 500 realizations). For clarity, in Figs. 8.16-8.18-8.20 the plot is limited only to the first 100 modes, containing the bulk of the energy. The eigenvalues are normalized with respect to their sum, representing the total turbulent energy of the fluctuations. Furthermore, the cumulative sum of the energy is reported, in order to identify the number of modes that significantly contribute to build up the decomposition of the velocity field.

In Fig. 8.16 the case of no upstream swirl is illustrated. Evidently, the first three modes constitute the most relevant contribution, as they contain respectively 20.8%, 10.9% and 7.5% of the energy, respectively (building in total 39.2% of the overall energy). The modes 4-9, contain in total about 10% of the energy; the modes 10-500 are associated each one with less than 1% of the total energy, thus they are most certainly associated with random fluctuations due to small scale turbulence and measurement noise.

The first three modes for the case of $S = 0$ are reported in Fig. 8.17, in which the contour representation of the v component, normalized with the bulk velocity, and the iso-surfaces of $v/V_j = -0.15$ and $v/V_j = 0.15$ are reported. The first and second modes are associated with the large scale precession, as they present an asymmetric outflow associated with an inflow on the opposite side of the chamber. This interpretation is confirmed by observing the instantaneous velocity fields reported in Fig. 8.12. The strong periodic component of the flow field enables the possibility to reconstruct the main topology of the precessing motion using only the first two modes (see Sec. 8.4.3). The difference in the energy pertaining to the first and the second mode is probably related to the geometry of the measurement volume, since it contains the entire inflow and outflow regions for the case of the first mode, while in the second mode the same two highly energetic regions are partly located outside of the observed volume in the far field. The third mode is quasi-axisymmetric, as it presents a strong on-axis outflow, associated to the axial mode. As a matter of fact, the ratio of the energy associated to the first two modes and the total energy of the first three modes can be interpreted as an “effective energy – probability” of precessing motion (about 81% for the case of $S = 0$).




Fig. 8.16 Energy distribution for the case of $S = 0$: **a)** normalized eigenmodes; **b)** cumulative energy.

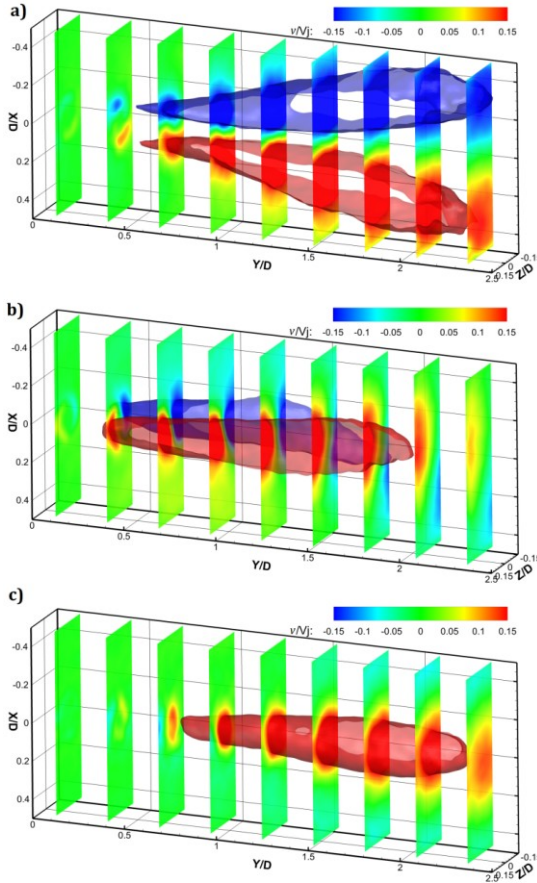


Fig. 8.17 Contour representation on XZ slices of the v/V_j , and iso-surfaces of $v/V_j = -0.15$ and $v/V_j = 0.15$ for the first (a), second (b) and third (c) POD mode for the case of $S = 0$.




Fig. 8.18 Energy distribution for the case of $S = 0.2$: **a)** normalized eigenmodes; **b)** cumulative energy.

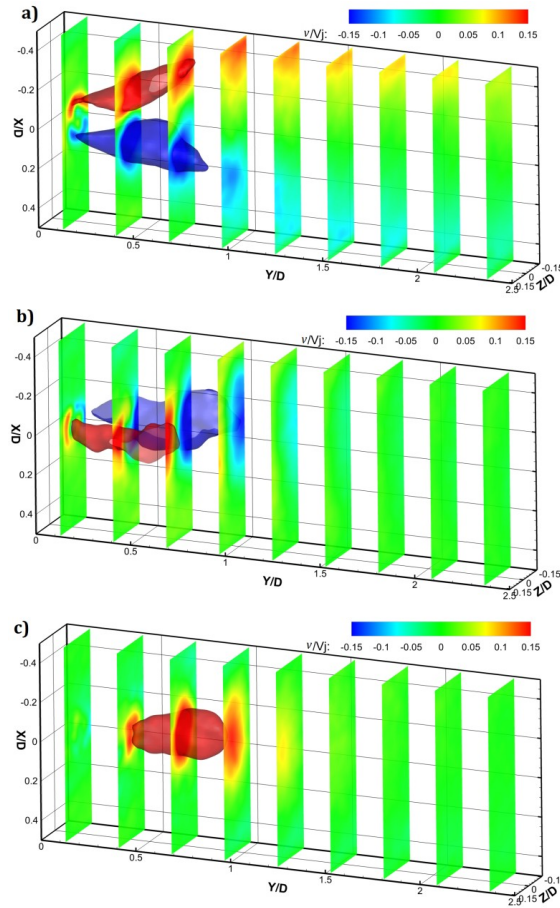


Fig. 8.19 Contour representation on XZ slices of the v/V_j , and iso-surfaces of $v/V_j = -0.15$ and $v/V_j = 0.15$ for the first **(a)**, second **(b)** and third **(c)** POD mode for the case of $S = 0.2$.




Fig. 8.20 Energy distribution for the case of $S = 0.6$: **a)** normalized eigenmodes; **b)** cumulative energy.

Also in the case of $S = 0.2$ the first three modes contain a significant part of the energy, i.e. 11.7%, 6.6% and 3.5% of the energy, respectively, as outlined in Fig. 8.18. In this case the energy of the first three modes constitutes only 21.8% of the total energy, mainly due to the effect of swirl, which determines a faster and stronger mixing, thus spreading part of the energy over a wide spectrum of modes. The first three modes are illustrated in Fig. 8.19 with the same layout of Fig. 8.17. The results confirm the presence of two dominant modes induced by the precession, and a third mode associated with the axial outflow. The analysis of the eigenvalues determines a precession probability of about 84%, which is only slightly larger than the case of the circular jet. This result indicates that the upstream swirl does not bias significantly the flow field towards a stable precessing motion with respect to a non swirling inlet.

As expected, the POD for the case of $S = 0.6$ presents significantly different features with respect to the two previously presented cases. A quite limited part of the energy is contained in the first four modes (3.2%, 2.9%, 1.9%, 1.5%, constituting about 9.5% of the overall energy, see Fig. 8.20). The spreading of the energy over the set of modes is due to the strong mixing effects of the vortex breakdown, determining an extremely fast decay of the turbulent energy along the spectral pipeline. Nevertheless the physical content of the first modes is still significant. Differently from the cases of $S = 0$ and $S = 0.2$, the first two modes (Fig. 8.21) present both an inflow and an outflow separated along the X direction, the main difference being the width of this separation (that is much smaller for the second mode). The second mode is related to a far-field precession effect induced by the precessing vortex core, while the first mode is similar to some extent to those of $S = 0$ and $S = 0.2$, but with inflow and outflow confined in the exit region. This suggests a more significant interaction with the external field, which goes beyond the topic of this investigation. The third and the fourth mode are relative to the recirculation region dominating the near field of the sudden expansion.

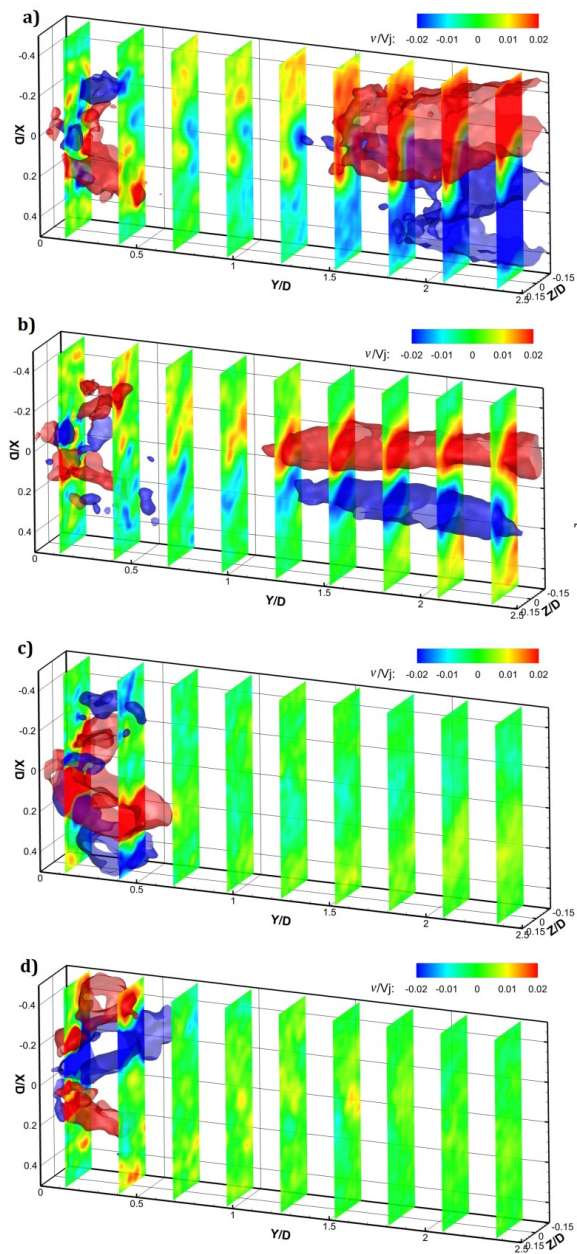


Fig. 8.21 Contour representation on XZ slices of the v/V_j , and iso-surfaces of $v/V_j = -0.02$ and $v/V_j = 0.02$ for the first (a), second (b), third (c) and fourth (d) POD mode for the case of $S = 0.6$.




Fig. 8.22 Scatter plot of the time coefficients for the first two modes in the normalized plane. The circumference with radius 1 is plotted for reference.

8.4.3 Low Order Reconstruction of the circular precessing jet

Following the work by Ben Chiekh et al (2004), in flow fields with relevant shedding phenomena one can use the POD modes to extract phase information by composing a subset of modes carrying the most significant part of the energy. In the attempt to reconstruct the features of the precession motion, a sort of phase-locked measurement is pursued using only the first two modes:

$$\underline{U}(\underline{x}, t) = \langle \underline{U}(\underline{x}, t) \rangle + \underline{u}(\underline{x}, t) \approx \langle \underline{U}(\underline{x}, t) \rangle + a_1(\vartheta) \underline{\varphi}_1(\underline{x}) + a_2(\vartheta) \underline{\varphi}_2(\underline{x}) \quad (8.7)$$

The coefficients $a_1(\vartheta)$ and $a_2(\vartheta)$ are both related to the precession phase angle (and, consequently, are not independent, even though they are statistically uncorrelated):

$$a_1(\vartheta) = \sqrt{2\lambda_1} \sin(\vartheta) \quad (8.8)$$

$$a_2(\vartheta) = \sqrt{2\lambda_2} \cos(\vartheta) \quad (8.9)$$

A method to confirm that the first two modes represent the coherent harmonic associated with the precession consists in observing the scatter plot of the time coefficients for the first two modes in Fig. 8.22. The scatter of points around the circle is due to the effect of turbulent fluctuations, which are randomly distributed over the cycles. Furthermore the samples with simultaneously large absolute value of a_3 (over 0.8 after normalization with the pertaining eigenvalue) and small a_1 and a_2 have been excluded from the scatter plot (the threshold for a_1 and a_2 is set to 0.3), as they are most likely representative of jet flowing in the axial mode. The samples isolated with this procedure are about 8% of the entire set.

The phase-averaged flow fields obtained with the Low Order Reconstruction (8.7) are reported in Fig. 8.23 for the case of $\vartheta = 0^\circ, 90^\circ, 180^\circ, 270^\circ$. The contour representation on the XZ central slice is blanked in case of $v/V_j < 0.01$, thus highlighting the outflow regions. The regions with stronger recirculation are identified with the isosurfaces $v/V_j = -0.05$. The flow field topology reveals the presence of a strongly asymmetric entrainment region located on the opposite side to that of the outflow in the far field of the chamber. This region is coupled with an extended recirculation region (characterized by weaker negative velocity with respect to the entrainment one) located below the region of attachment of the exiting jet. The two recirculation regions asymmetrically embrace the jet attached to the wall, and precess around it, thus determining a region of strong swirl in the near field, testified by the isosurfaces of the Y -component of the normalized vorticity $\omega_y D/V_j$. The positive ω_y and the streamline representation indicate a net rotation around the Y axis in the opposite sense to that of the precession (it will be shown in the next section that in the performed experiment the precessing motion, triggered with a sense of rotation determined by random asymmetries due to the turbulence or non perfect initial conditions, persists with that sense of rotation, even though occasionally switching to the axial mode, and this sense is opposite to that of swirl in the near field to balance the overall angular momentum). It is reasonable to conjecture that the asymmetry of the entrainment and the recirculation regions are the inertial driving force of the precession, thus making solutions like the exit lip or the exit centrebody much more effective than a weak upstream swirl in triggering and stabilizing the precessing mode.

8.4.4 Phase computation: precession frequency measurements for the FPJ

The effectiveness of the Low Order Reconstruction in identifying phase-average flow fields implies an immediate implication: the phase of a single realization can be extracted using the time coefficients of the first two modes for that snapshot. According to van Oudheusden et al (2005), the phase can be obtained by observing that:

$$\frac{a_1(i)}{\sqrt{2\lambda_1}} = r_i^2 \sin(\vartheta) \quad (8.10)$$

$$\frac{a_2(i)}{\sqrt{2\lambda_2}} = r_i^2 \cos(\vartheta) \quad (8.11)$$

An immediate consequence of (8.10-11) is:

$$\frac{a_1(i)^2}{2\lambda_1} + \frac{a_2(i)^2}{2\lambda_2} = r_i^2 \quad (8.12)$$

$$\vartheta = \arctan\left(\frac{\sqrt{2\lambda_2} a_1}{\sqrt{2\lambda_1} a_2}\right) \quad (8.13)$$

Chapter 8 – Jet flows past a sudden expansion

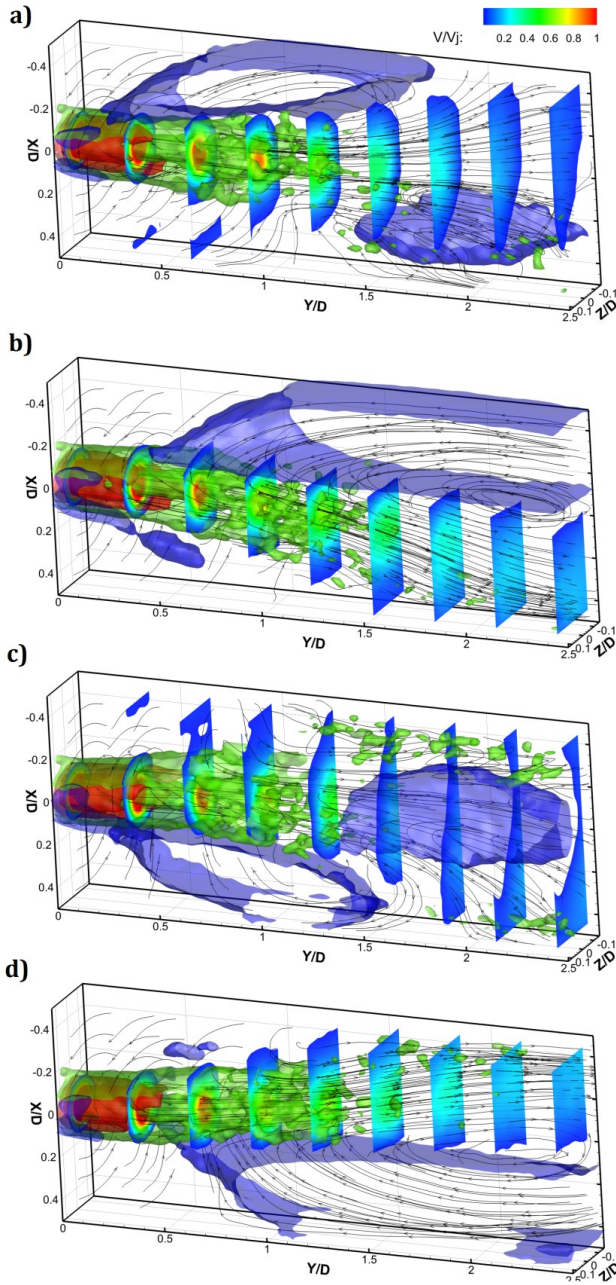


Fig. 8.23 Streamline representation of the phase-resolved flow fields by means of Low Order Reconstruction. **a)** $\vartheta = 0^\circ$, **b)** $\vartheta = 90^\circ$, **c)** $\vartheta = 180^\circ$, **d)** $\vartheta = 270^\circ$. Isosurfaces of $v/V_j = -0.05$ (blue), $\omega_y D/V_j = 3.8$ (red), $\omega_y D/V_j = 0.8$ (green). Contour representation of v/V_j on XZ planes (colouring blanked for $v/V_j < 0.01$).

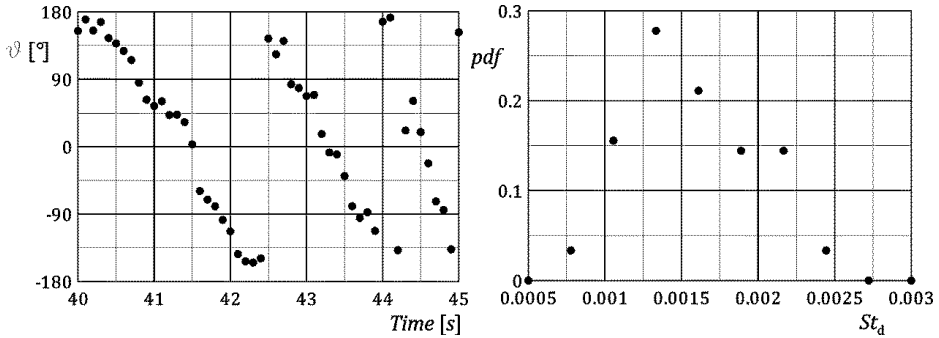


Fig. 8.24 Computed phase angles (left) for a 5 seconds time interval, and pdf of the Strouhal number based on the inlet diameter for $S = 0$ (right).

The (8.13) is applied on the samples to determine the phase; then the sequences are analyzed to identify patterns of regularly spaced phase angles. The results, reported in Fig. 8.24, highlight a regular tendency of the jet to precess with a negative angular velocity, i.e. in opposition with the Y axis. This result confirms the intuition that the upstream swirl generated in the near field (and in particular in the shear layer of the jet) is in the opposite direction to that of precession. The measured Strouhal number (based on the inlet diameter) presents a dispersion due to the intermittency of the phenomenon and the uncertainty in determining the exact phase shift in the sequence analysis. The median of the distribution is about 0.0015 (corresponding to a precession frequency of 0.6Hz), in agreement with the data reported in the literature.

8.5 Final remarks

Tomographic PIV has demonstrated to be a powerful instrument for the understanding of complex large-scale turbulence organization. The instantaneous and statistical analysis of the flow field with variable intensity of the inlet swirling motion leads to the following remarks:

- For a non-swirling inlet ($S = 0$) the driving force of the precession is the inertial effect induced by the asymmetric entrainment (on the opposite side of that of the outflow), triggered by the instantaneous turbulent fluctuations in the condition of neutral equilibrium of the jet attached to the chamber. The entrainment region is extended along the chamber and interacts with an asymmetric recirculation region, placed right below the exiting jet and extending down to the basis of the chamber. The asymmetry of these two regions determines an azimuthal pressure gradient, driving the precession. Furthermore, a swirling motion is imparted to the shear layer of the jet in the direction opposite to that of precession to balance the angular momentum;

- In case of weak swirl ($S = 0.2$) in the inlet flow, the phenomenology is similar to that of the circular jet, with the most significant difference being the attachment point, located more upstream with respect to the case of $S = 0$. The detected probability of precessing motion is only slightly larger than the case of the non-swirling inlet, thus testifying that the precession is not essentially driven by initial asymmetries; in this sense, the asymmetric entrainment plays a much more important role. Other solutions, like the exit lip or the centrebody obstruction, are more effective in maximizing the precession probability. However, a weak swirl might be useful in case of short chamber due to physical restriction (that can prevent the triggering of the precessing motion for non swirling inlet, as testified in Sec. 8.1.2), as it promotes the attachment of the jet in the upstream part of the chamber;
- The flow field for large swirl ($S = 0.6$) is extremely different, as it is dominated by the vortex breakdown and the precessing vortex core, localizing the entrainment region in proximity of the nozzle axis. The POD analysis has shown a significant interaction (and cross-talk between modes) between the external precession and the inner precessing vortex core. This mechanism complicates the predictability of the features of the flow field in the external part of the chamber, thus leaving this interpretation to future investigations.

CONCLUSIONS AND PERSPECTIVES

In this thesis the limits in the implementation and application of Tomographic Particle Image Velocimetry to turbulent flow investigations are addressed, and some solutions to improve the accuracy of the technique and to extend its range of applicability have been presented.

Advances in tomographic reconstruction

A multi-resolution method to accelerate the convergence of the reconstruction procedure has been proposed. The algorithm is based on a fast estimation of a first guess distribution for the voxels carrying non-zero intensity by solving the tomographic reconstruction problem discretized with a lower resolution. The multi-resolution technique results in a reduction of the computational cost by a factor ranging between 5 and 20 times, depending on the image density (and, consequently, on the sparseness of the distributions to be reconstructed). This speed-up is obtained without affecting the quality of the results; actually, the tomographic reconstruction accuracy is slightly improved, probably due to the smoothing induced within the reconstruction process by the binning procedure and the interpolation step when building the first guess on the final grid. Indeed, an induced diffusion within the reconstruction process has shown to be significantly beneficial in increasing the quality of the reconstructions. This observation has led to the development of the Spatial Filtering Improved Tomographic PIV proposed in this thesis. In particular, smoothing of the 3D intensity distributions within the iterative reconstruction process can reduce the error due to undersampling in the voxels discretization. Considering that the uncertainty in the reconstruction is usually anisotropic (most often the cameras of the tomographic system share the same depth direction, thus causing elongation of both true and ghost particles), the spatial filter can be properly shaped to contrast the anisotropy. The SFIT-MART technique has proved to provide a remarkable reconstruction accuracy improvement and a consistent reduction of the uncertainty of the velocity measurement (about 10%).

Optimization of 3D cross-correlation

The digital cross-correlation analysis of the tomographic reconstructed volumes is more consolidated, since it is conceptually a 3D extension of the well-assessed 2D algorithms for planar PIV analysis. On the other hand, the large amount of data and of operations required for the 3D analysis complicates the application of Tomo-PIV with off-the-shelf machines. In this thesis some solutions are proposed and combined to result in an efficient algorithm, based on sparse direct cross-

correlation and minimization of the number of redundant operations. The calculation of the predictor is significantly accelerated by binning the distributions; subsequently the measured velocity field is refined by the fast computation of the corrector velocity field. In this step, since the corrections are usually small (within 1 pixels, provided that the under-detection in the validation procedure and the modulation of the velocity field in both predictor computation and in the outliers replacement do not determine the need of a larger correction), one can use direct cross-correlation (that is well suited for the implementation of sparse algorithms) with a narrow search radius for the peak identification. Since usually overlapping interrogation regions are used, one can pre-calculate the contributions to the cross-correlation maps on sub-volumes, in form of blocks, planes or segments. The results have shown that this efficient algorithm can achieve a speed-up of more than one order of magnitude with exactly the same results. In other words, the data of a test campaign requiring one month for processing can be analyzed in three days or less. Said in a different way, large amount of data can be processed in a reasonable time on off-the-shelf computers instead of expensive (often time-shared) supercomputers.

Reduction of the hardware costs

Last, but not least, the application of Tomographic PIV is limited by the operative costs. Many laboratories are equipped with one or two cameras for planar PIV or Stereo-PIV applications; on the other hand, Tomo-PIV requires at least three cameras (preferably 4 or more) to achieve an acceptable quality of the reconstructions. The cameras for PIV applications are usually expensive not only for the high required sensitivity, but also due to the requirement of the double-shutter operating mode. To some extent Tomo-PIV provides the advantage of being based on cross-correlation of reconstructed objects, thus eliminating the requirement of double-shutter cameras. A "Low-Cost Tomo-PIV" has been proposed, based on using two or more independent tomographic systems, to reconstruct the single exposures; even though the number of cameras is larger than usual, one can use cameras with the same sensitivity of the common PIV cameras but much lower price since there is no need to work in double-shutter mode. The benefit is not limited to the cost reduction: indeed, the technique has shown to be completely unaffected by the bias effect due to the coherent motion of ghost particles, thus enabling an improvement of the quality of the results with multi-exposure methods much more remarkable (and displacement-field independent) than the case of the standard Tomo-PIV implementation.

Applications to turbulent flows: overview and perspectives

The application of the technique in two significantly different scenarios has provided an overview of the capabilities and limits of Tomo-PIV. The analysis of the

decaying fractal generated-turbulence has shown that the technique has a long path to walk before covering the gap with hot-wire anemometry and planar PIV in the measurement of the dissipation and the dynamics of small-scale turbulence. The main obstacle is the maximum allowed seeding density, which limits the achievable spatial resolution. Nevertheless Tomo-PIV has shown a great potential in quantifying the degree of isotropy and homogeneity even down to scales of the order of the Taylor microscale, leading to the conclusion that the turbulence generated by square space-filling fractal grids is with good approximation homogeneous and isotropic for $x/x^* > 0.6$ (with x^* being the wake-interaction lengthscale defined by applying the wake similarity for the largest bars of the grid). The possibility to measure all the components of the velocity gradient tensor without relying on any theoretical assumption opens fascinating horizons, but in the case under analysis the scatter of the data does not allow a conclusive statement on the violation of the dissipation anomaly.

On the other hand, Tomographic PIV has shown to be mature enough for the quantitative observation of the large scale organization of high Reynolds number turbulent in extremely complex 3D flows, such as the case of the precessing jet motion after an abrupt expansion in a cylindrical chamber and the interaction of this gyroscopic-like behaviour with an induced swirl in the inlet flow. The application of the Proper Orthogonal Decomposition on the 3D data provided by Tomographic PIV has demonstrated to be a very powerful instrument for the understanding of the topology of inhomogeneous turbulent flows and the underlying mechanism generating the precession. As a matter of fact, 3D data facilitate the interpretation of the velocity measurements in the complicate framework of an intermittent switching between two different behaviours, i.e. precessing motion or quasi axisymmetric expansion within the chamber. In case of weak (or even absent) upstream swirl, the driving force of the precession is the interaction between the asymmetric entrainment from the external ambient and the recirculation regions formed below the region of instantaneous reattachment of the jet flow. The upstream swirl does not significantly favour the precessing motion with respect to the axisymmetric expansion mode, but moves upstream the reattachment point, thus allowing the generation of the bi-stable behaviour in shorter chambers.

Future developments

The outlined achievements in terms of operative cost (hardware, computational resources) and reconstruction accuracy improvement are expected to have a significant impact, as they allow for a much wider and faster spreading of the technique among the scientific community. However, Tomo-PIV development is still at its infancy; there are great margins of improvement to annihilate the

Conclusions and perspectives

weaknesses of the technique and establish this new instrument as a solid tool for turbulent flows investigation.

The tomographic reconstruction is capturing an increasing interest. The single-exposure MART is quite well consolidated, and most probably a significant advancement will be obtained following a different path. Optimization methods based on sparsity maximization, hybrid Tomo-PTV methods, and other solutions, are still at an embryonic stage. Perhaps the most significant margins of improvement rely within multi-exposure methods, following the path of the MTE-MART by Novara et al (2010). Increasing the accuracy without affecting the computational burden (or possibly reducing it!) will most certainly stimulate the development of brand-new solutions.

Perhaps the most intriguing feature of Tomo-PIV is the possibility to exploit the complete 3D information (or 4D in case of time-resolved data) to increase the accuracy and robustness of the velocity measurement. Adaptive 3D interrogation enables a much larger margin of improvement with respect to the case of 2D data, as outlined by Novara et al (2013); furthermore, time-resolved measurements can exploit the temporal coherence to improve the solidity of the technique. In addition to this, multi-exposures system (as the triple pulse PIV described in Westerweel et al 2013) can be easily implemented in the case of the Low-Cost Tomo-PIV proposed in this thesis (for example, the first system can be composed by single-shutter cameras, and the second one by double-shutter cameras, recording the second and third frame), thus enabling acceleration measurements and/or significant reduction of the measurement uncertainty even in ultra-high speed applications.

REFERENCES

Adrian RJ (1984) Scattering particle characteristics and their effect on pulsed laser measurements of fluid flow: speckle velocimetry vs. particle image velocimetry. *Appl Opt* 23:1690-1691

Adrian RJ (1991) Particle-imaging techniques for experimental fluid mechanics. *Ann Rev Fluid Mech* 23:261-304

Adrian RJ, Yao CS (1984) Development of pulsed laser velocimetry (PLV) for measurement of fluid flow. In: Patterson G, Zakin J (ed) *Proceedings of the 8th biennial symposium on turbulence*, Rolla, Missouri, September 1984, 170–186

Adrian RJ, Westerweel J (2011) *Particle Image Velocimetry*. Cambridge University Press, UK

Antonia RA, Lavoie P, Djenidi L, Benaissa A (2009) Effect of a small axisymmetric contraction on grid turbulence. *Exp Fluids* 49:3-10

Astarita T (2006) Analysis of interpolation schemes for image deformation methods in PIV: effect of noise on the accuracy and spatial resolution. *Exp Fluids* 40:977-987

Astarita T (2007) Analysis of weighting windows for image deformation methods in PIV. *Exp Fluids* 43:859-872

Astarita T (2008) Analysis of velocity interpolation schemes for image deformation methods in PIV. *Exp Fluids* 45:257-266

Astarita T (2009) Adaptive space resolution for PIV. *Exp Fluids* 46:1115-1123

Astarita T, Cardone G (2005) Analysis of interpolation schemes for image deformation methods in PIV. *Exp Fluids* 38:233-243

Atkinson C, Soria J (2007) Algebraic reconstruction techniques for tomographic particle image velocimetry. *Proc. 16th Australasian Fluid Mechanics Conf. (Gold Coast, Australia)* 191–198

Atkinson C, Soria J (2009) An efficient simultaneous reconstruction technique for tomographic particle image velocimetry. *Exp Fluids* 47:553-568

Atkinson C, Coudert S, Foucaut JM, Stanislas M, Soria J (2011) The accuracy of tomographic particle image velocimetry for measurements of a turbulent boundary layer. *Exp Fluids* 50:1031-1056

Bai K, Meneveau C, Katz J (2012) Near-wake turbulent flow structure and mixing length downstream of a fractal tree. *Bound-Lay Meteorol* 143:285-308

Batchelor GK (1953) *The theory of homogeneous turbulence*. Cambridge University Press. Cambridge, UK

References

- Ben Chiekh M, Michard M, Grosjean N, Bera JC (2004) Reconstruction temporelle d'un champ aérodynamique instationnaire a partir de mesures PIV non résolues dans le temps. In: 9 Congrès Francophone de Velocimétrie Laser. Bruxelles, Belgium
- Berkooz G, Holmes P, Lumley JL (1993) The proper orthogonal decomposition in the analysis of turbulent flows. *Ann Rev Fluid Mech* 25:539-575
- Bevilaqua PM, Lykoudis PS (1978) Turbulence memory in selfpreserving wakes. *J Fluid Mech* 89:589-1978
- Bilsky AV, Dulin VM, Lozhkin VA, Markovich DM, Tokarev MP (2011) Two-dimensional correlation algorithms for tomographic PIV. Proc. of 9th International Symposium on Particle Image Velocimetry PIV '11, July 21st – 23rd, Kobe University, Kobe, Japan
- Bourgoin M, Ouellette NT, Xu H, Berg J, Bodenschatz E (2006) The role of pair dispersion in turbulent flows. *Science* 311:835-838
- Brücker C (1995) Digital-particle-image-velocimetry (DPIV) in a scanning light-sheet: 3-D starting flow around a short cylinder. *Exp Fluids* 19:255-263
- Buchmann NA, Atkinson C, Jeremy MC, Soria J (2011) Tomographic particle image velocimetry investigation of the flow in a modeled human carotid artery bifurcation. *Exp Fluids* 50:1131-1151
- Buchner AJ, Buchmann N, Kilany K, Atkinson C, Soria J (2012) Stereoscopic and tomographic PIV of a pitching plate. *Exp Fluids* 52:299-314
- Burattini P, Lavoie P, Antonia RA (2005) On the normalized turbulent energy dissipation rate. *Phys Fluids* 17, 098103
- Chester S, Meneveau C (2007) Renormalized numerical simulation of flow over planar and non-planar fractal trees. *Environ Fluid Mech* 7:280-301.
- Chester S, Meneveau C, Parlange MB (2007) Modeling turbulent flow over fractal trees with renormalized numerical simulation. *J Comp Phys* 225:427-448.
- Collier RJ, Burckhardt CB, Lin LH (1971) *Optical holography*. New York: Academic
- Comte-Bellot G, Corrsin S (1966) The use of a contraction to improve the isotropy of grid-generated turbulence. *J Fluid Mech* 25:657-682
- Daly BJ, Harlow FH (1970) Transport equations in turbulence *Phys. Fluids* 13 2634-2649
- David L, Jardin T, Braud P, Farcy A (2012) Time-resolved scanning tomography PIV measurements around a flapping wing. *Exp Fluids* 52:857-864
- de Jong J, Cao L, Woodward SH, Salazar JPLC, Collins LR and Meng H (2009) Dissipation rate estimation from PIV in zero-mean isotropic turbulence. *Exp Fluids* 46:499-515
- de Silva CM, Baidya R, Khashehchi M, Marusic I (2012) Assessment of tomographic PIV in wall-bounded turbulence using direct numerical simulation data. *Exp Fluids* 52:425-440
- de Silva CM, Baidya R, Marusic I (2013) Enhancing Tomo-PIV reconstruction quality by reducing ghost particles. *Meas Sci Technol* 24, 024010

- Dellenback PA, Metzger DE, Neitzel GP (1988) Measurements in a turbulent, swirling flow through an abrupt axisymmetric expansion. *AIAA J* 266:669-681
- Discetti S, Astarita T (2012a) A fast multi-resolution approach to tomographic PIV. *Exp Fluids*, 52:765-777
- Discetti S, Astarita T (2012b) Fast 3D PIV with direct sparse cross-correlations. *Exp Fluids*, 55(3):1437-1451
- Discetti S, Adrian RJ (2012) High accuracy measurement of magnification for monocular PIV. *Meas Sci Technol* 23, 117001
- Discetti S, Ianiro A, Astarita T, Cardone G (2012) On the development of a novel low cost high accuracy experimental setup for Tomographic Particle Image Velocimetry. 16th Int. Symp. on Applications of Laser Techniques to Fluid Mechanics, 9-12 July, Lisbon, Portugal
- Discetti S, Natale A, Astarita T (2013a) Spatial Filtering Improved Tomographic PIV. *Exp Fluids* DOI: 10.1007/s00348-013-1505-7
- Discetti S, Ziskin IB, Astarita T, Adrian RJ, Prestridge K (2013b) PIV measurements of anisotropy and inhomogeneity in decaying fractal generated turbulence. In press on *Fluid Dyn Res*
- Eisenstat SC, Gursky MC, Schultz MH, Sherman AH (1982) Yale sparse matrix package I: the symmetric codes. *Int J Numer Methods in Engin* 18:1145-1151
- Elsinga GE, Scarano F, Wieneke B, van Oudheusden BW (2006a) Tomographic particle image velocimetry. *Exp Fluids* 41:933-947
- Elsinga GE, van Oudheusden BW, Scarano F (2006b) Experimental assessment of tomographic PIV accuracy. 13th Int. Symp. on Applications of Laser Techniques to Fluid Mechanics, 26-29 June, Lisbon, Portugal
- Elsinga GE, Adrian RJ, van Oudheusden BW, Scarano F (2010) Three-dimensional vortex organization in a high-Reynolds-number supersonic turbulent boundary layer. *J Fluid Mech* 644:35-60
- Elsinga GE, Westerweel J, Scarano F, Novara M (2011) On the velocity of ghost particles and the bias errors in Tomographic PIV. *Exp Fluids* 50:825-838
- Elsinga GE, Poelma C, Schröder A, Geisler R, Scarano F, Westerweel J (2012) Tracking of vortices in a turbulent boundary layer. *J Fluid Mech* 697:273-295
- Elsinga GE, Westerweel J (2012) Tomographic-PIV measurements of the flow around a zigzag boundary layer trip. *Exp Fluids* 52:865-876
- Frigo M, Johnson SG (2005) The design and implementation of fftw3. *Proc IEEE* 93:216-231
- Gad-El-Hak M, Corrsin S (1974) Measurement of the nearly isotropic turbulence behind a uniform jet grid. *J Fluid Mech* 62:115-143
- Gao Q, Ortiz-Dueñas C, Longmire EK (2011) Analysis of vortex populations in turbulent wall-bounded flows. *J Fluid Mech* 678:87-123

References

Geipel P, Goh HKH, Lindstedt PR (2010) Fractal generated turbulence in opposed jet flows. *Flow Turbulence Combust* 85:397-419

George WK (1992) The decay of homogeneous turbulence. *Phys Fluids A* 4:1492-1509

George WK, Wang H (2009) The exponential decay of homogeneous turbulence. *Phys Fluids* 21, 025108

Ghaemi S, Scarano F (2011) Counter-hairpin vortices in the turbulent wake of a sharp trailing edge. *J Fluid Mech* 689:317-356

Giordano R, Astarita T (2009) Spatial resolution of the Stereo PIV technique. *Exp Fluids* 46:643-658

Gordon R, Bender R, Herman GT (1970) Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and x-ray photography. *J Theor Biol* 29:471-481

Gouldin FC (1987) An application of fractals to modeling premixed turbulent flames. *Combust Flame* 68:249-266

Guo B, Fletcher DF, Langrish TAG (2001) Numerical simulation of unsteady turbulent flow in axisymmetric sudden expansions. *J Fluids Eng* 123:574-587

Hain R, Kähler CJ, Michaelis D (2008) Tomographic and time resolved PIV measurements on a finite cylinder mounted on a flat plate. *Exp Fluids* 45:715-724

Hallett, WLH, Günther R (1984) Flow and mixing in swirling flow in a sudden expansion. *Can. J. Chem. Engng* 62:149-155

Hanson KM, Weckung GW (1985) Local basis-function approach to computed tomography. *Applied Optics* 24:4028-4039

Harris FJ (1978) On the use of windows for harmonic analysis with the discrete Fourier transform. *Proc IEEE* 66:51-83

Herman GT, Lent A (1976) Iterative reconstruction algorithms. *Comput Biol Med* 6:273-294

Hinsch KD (1995) Three-dimensional particle velocimetry. *Meas Sci Technol* 6:742-753

Hinsch KD (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

Huang HT, Fiedler HE, Wang JJ (1993) Limitation and improvement of PIV, part II. Particle image distortion, a novel technique. *Exp Fluids* 15:263-273

Humble RA, Elsinga GE, Scarano F, van Oudheusden B (2009) Three-dimensional instantaneous structure of a shock wave/turbulent boundary layer interaction. *J Fluid Mech* 622:33-62

Hurst D, Vassilicos JC (2007) Scalings and decay of fractal-generated turbulence. *Phys Fluids* 19, 035103

- Ianiro A, Violato D, Cardone G, Scarano F (2011) Time-resolved tomographic PIV measurements in swirling jets. Proc. of 64th Annual Meeting of the American Physical Society's Division of Fluid Dynamics. Baltimore (MD, USA), 20th-22th Nov
- Ianiro A, Cardone G (2012) Heat transfer rate and uniformity in multichannel swirling impinging jets. *Appl Therm Eng* 49:89–98
- Jambunathan K, Ju XY, Dobbins BN, Ashforth-Frost S (1995) An improved cross correlation technique for particle image velocimetry. *Meas Sci Technol* 6:507-514
- Jayesh, Warhaft Z (1992) Probability distribution, conditional dissipation, and transport of passive temperature fluctuations in grid-generated turbulence. *Phys Fluids* 4:2292-2307
- Jeong J, Hussain F (1995) On the identification of a vortex. *J Fluid Mech* 285:69-94
- Kähler CJ, Kompenhans J (2000) Fundamentals of multiple plane stereo particle image velocimetry. *Exp Fluids* 29:S70-S77
- Keane RD, Adrian RJ (1992) Theory of cross-correlation analysis of PIV images. *Appl Sci Res* 49:191-215
- Keane RD, Adrian RJ, Zhang Y (1995) Super-resolution particle image velocimetry. *Meas Sci Technol* 6:754-768
- Kinzel M, Wolf M, Holzner M, Lüthi B, Tropea C, Kinzelbach W (2011) Simultaneous two-scale 3D-PTV measurements in turbulence under the influence of system rotation. *Exp Fluids* 51:75-82
- Kolmogorov AN (1941) The local structure of turbulence in incompressible viscous fluid for very large Reynolds. *C. R. Acad. Sci. U.R.S.S.* 30:301-305
- Krogstad PA (2012) Turbulent decay in the near field of multi-scale and conventional grids. *Int J Heat Fluid Flow* 35:102-108
- Krogstad PA, Davidson PA (2012) Near-field investigation of turbulence produced by multi-scale grids. *Phys Fluids* 24, 035103
- Laizet S, Vassilicos JC (2011) DNS of fractal-generated turbulence. *Flow Turbul Combust* 87:673-705
- Lamarche F, Leroy C (1990) Evaluation of the volume of intersection of a sphere with a cylinder by elliptic integrals. *Comput Phys Commun* 59:359–369
- Lavoie P, Burattini P, Djenidi L, Antonia RA (2005) Effect of initial conditions on decaying grid turbulence at low Re . *Exp Fluids* 39:865-874
- Lawson NJ, Wu J (1997a) Three-dimensional particle image velocimetry: error analysis of stereoscopic techniques. *Meas Sci Technol* 8:894-900
- Lawson NJ, Wu J (1997b) Three dimensional particle image velocimetry: experimental error analysis of a digital angular stereoscopic system. *Meas Sci Technol* 8:1455-1464
- Lecordier B, Demare D, Vervisch LMJ, Réveillon J, Trinité M (2001) Estimation of the accuracy of PIV treatments for turbulent flow studies by direct numerical simulation of multi-phase flow. *Meas Sci Technol* 12:1382–1391

References

- Lecuona A, Nogueira J, Rodriguez PA, Santana D (2002) Accuracy and time performance of different schemes of the local field correction PIV technique. *Exp Fluids* 33:743-751
- Lucca-Negro O, O' Doherty T (2001) Vortex breakdown: a review. *Prog Ener Comb Sci* 27:431-481
- Luxton RE, Nathan GJ, Luminis Pty. Ltd (1987) Mixing using a fluid jet. Patent Application No. P14068/87, Australian Patent Office
- Luxton RE, Nathan GJ (1989) A precessing asymmetric flow field in an abruptly expanding axi-symmetric duct. Tenth Australasian Fluid Mechanics Conference, Melbourne, Australia, 2:11.29-11.33
- Lüthi B, Tsinober A, Kinzelbach W(2005). Lagrangian measurement of vorticity dynamics in turbulent flow. *J Fluid Mech* 528:87-118
- Maas HG, Gruen A, Papantoniou D (1993) Particle tracking velocimetry in three-dimensional flows. *Exp Fluids* 15:133-146
- Madej AM, Babazadeh H, Nobes DS (2011) The effect of chamber length and Reynolds number on jet precession. *Exp Fluids* 51:1623-1643
- Makita H (1991) Realization of a large-scale turbulence field in a small wind tunnel. *Fluid Dyn Res* 8:53-64
- Mandelbrot BB (1974) Intermittent turbulence in self-similar cascades: divergence of high moments and dimension of the carrier. *J Fluid Mech* 62: 331-358
- Mandelbrot BB (1982) *The Fractal Geometry of Nature*. Freeman, San Francisco
- Mazellier N, Vassilicos JC (2008) The turbulence dissipation constant is not universal because of its universal dependence on large-scale flow topology. *Phys Fluids* 20, 015101
- Mazellier N, Vassilicos JC (2010) Turbulence without Richardson-Kolmogorov cascade. *Phys Fluids* 22, 075101
- Mazzi B, Okkels F, Vassilicos JC (2002) A shell-model approach to fractal-induced turbulence. *Eur Phys J B/Fluids* 28:243-251
- Mazzi B, Vassilicos JC (2004) Fractal-generated turbulence. *J Fluid Mech* 502:65-87
- Meyer KE, Pedersen JM, Ozcan O (2007) A turbulent jet in crossflow analyzed with proper orthogonal decomposition. *J Fluid Mech* 583:199-227
- Mi J, Nathan GJ, Luxton RE (2001a) Mixing characteristics of a flapping jet from a self-exciting nozzle. *Flow Turb Combust* 67:1-23
- Mi J, Nathan GJ, Nobes DS (2001b) Mixing characteristics of axisymmetric free jets from a contoured nozzle, an orifice plate and a pipe. *J Fluids Eng* 123:878-883
- Mi J, Nathan GJ (2004) Self-excited jet-precession Strouhal number and its influence on downstream mixing field. *J Fluid Struct* 19:851-862
- Mi J, Nathan GJ (2006) The influence of inlet flow condition on the frequency of self-excited jet precession. *J Fluid Struct* 22:129-133

- Michaelis D, Wolf CC (2011) Vibration compensation for Tomographic PIV using single image volume self calibration. 9th Int Symp on Particle Image Velocimetry PIV'11, 21st-23rd July, Kobe (Japan)
- Mishra D, Muralidhar K, Munshi P (1999) A robust MART algorithm for tomographic applications. *Num Heat Transfer Part B* 35:485-506
- Mungal MG, Karasso PS, Lozano A (1991) The visible structure of turbulent jet diffusion: large-scale organization and flame tip oscillation. *Combust Sci Technol* 76:165-185
- Mydlarski L, Warhaft Z (1996) On the onset of high-Reynolds-number grid-generated wind tunnel turbulence. *J Fluid Mech* 320:331-368
- Nagata K, Suzuki H, Sakai H, Hayase Y, Kubo T (2008a) Direct numerical simulation of turbulence characteristics generated by fractal grids. *Int Rev Phys* 2:400-409
- Nagata K, Suzuki H, Sakai H, Hayase Y, Kubo T (2008b) Direct numerical simulation of turbulent mixing in grid-generated turbulence. *Phys Scr T* 132, 014054
- Nathan GJ, Luxton RE (1992a) Mixing enhancement by a self-exciting, asymmetric precessing flow-field. In *Transport Phenomena in Heat and Mass Transfer*. (ed. J. A. Reizes) 2:1297-1307. Elsevier.
- Nathan GJ, Luxton RE (1992b) The flow field within an axi-symmetric nozzle utilizing a large abrupt expansion. In *Recent Advances in Experimental Fluid Mechanics* (ed. F. G. Zhuang) 527-532. Academic.
- Nathan GJ, Hill SJ, Luxton RE (1998) An axisymmetric "fluidic" nozzle to generate jet precession. *J Fluid Mech* 370:347-380
- Nathan GJ, Mi J, Alwahabi ZT, Newbold GJR, Nobes DS (2006) Impacts of a jet's exit flow pattern on mixing and combustion performance. *Prog Energ Combust Sci* 32:496-538
- Newbold GJR, Nathan GJ, Nobes DS, Turns SR (2000) Measurement and prediction of NO_x emissions from unconfined propane flames from turbulent-jet, bluff-body, swirl and precessing jet burners. *P Combust Inst* 28:481-487 Part 1
- Nogueira J, Lecuona A, Rodriguez PA (1999) Local field correction PIV: on the increase of accuracy of digital PIV systems. *Exp Fluids* 27:107-116
- Nogueira J, Lecuona A, Rodriguez PA (2001) Local field correction PIV, implemented by means of simple algorithms, and multigrid version. *Meas Sci Technol* 12:1911-1921
- Novara M, Batenburg KJ, Scarano F (2010) Motion tracking-enhanced MART for tomographic PIV. *Meas Sci Technol* 21, 035401
- Novara M, Scarano F (2012a) Performances of motion tracking enhanced tomo-PIV on turbulent shear flows. *Exp Fluids* 52:1027-1041
- Novara M, Scarano F (2012b) Lagrangian acceleration evaluation for tomographic PIV: a particle-tracking based approach. 16th Int. Symp. on Applications of Laser Techniques to Fluid Mechanics, 9-12 July, Lisbon, Portugal

References

Novara M, Ianiro A, Scarano F (2013) Adaptive interrogation for 3D-PIV. *Meas Sci Technol* 24, 024012

Novikov EA (1971) Intermittency and scale similarity in the structure of a turbulent flow. *PMM J App Math* 35:231-241

Ohmi K, Li HY (2000) Particle-tracking velocimetry with new algorithms *Meas Sci Technol* 11:603-616

Ortiz-Dueñas C, Kim J, Longmire EK (2010) Investigation of liquid-liquid drop coalescence using tomographic PIV. *Exp Fluids* 49:111-129

Pereira F, Gharib M, Dabiri D, Modarress D (2000) Defocusing digital particle image velocimetry: a 3-component 3-dimensional DPIV measurement technique. Application to bubbly flows. *Exp Fluids* 29:S78-S84

Petra S, Schnörr C, Schröder A, Wieneke B (2007) Tomographic image reconstruction in experimental fluid dynamics: synopsis and problems. *Mathematical Modelling of Environmental and Life Sciences Problems* ed R A Bucuresti

Petra S, Schröder A, Schnörr C (2009) 3D tomography from few projections in experimental fluid dynamics. *Imaging Measurement Methods for Flow Analysis (Notes on Numerical Fluid Mechanics and Multidisciplinary Design)* 106:63-72

Pope SB (2000) *Turbulent flows*. Cambridge University Press. Cambridge, UK

Prasad AK (2000) Stereoscopic particle image velocimetry. *Exp Fluids* 29:103-116

Queiros-Conde D, Vassilicos JC (2001) Turbulent wakes of 3-D fractal grids. In *Intermittency in Turbulent Flows and Other Dynamical Systems*. Cambridge University Press, Cambridge, UK

Radon J (1917) Über die bestimmung von funktionen durch ihre integralwerte langs gewisser mannigfaltigkeiten. *Ber. Sachsische Akad. Wiss., Leipzig, Math Phys.* 69:262-267

Raffel M, Willert CE, Weerely ST, Kompenhans J (2007) *Particle Image Velocimetry: A practical guide*. Springer-Verlag. ISBN: 3-540-72307-2

Revuelta A, Sánchez AL, Liñán A (2002) Confined axisymmetric laminar jets with large expansion ratios. *J Fluid Mech* 456:319-352

Revuelta A, Sánchez AL, Liñán A (2004) Confined swirling jets with large expansion ratios. *J Fluid Mech* 508:89-98

Reynolds WC, Parekh DE, Juvet PJD, Lee MJD (2003) Bifurcating and blooming jets. *Ann Rev Fluid Mech* 35:295-315

Rohàly J, Frigerio F, Hart DP (2002) Reverse hierarchical PIV processing. *Meas Sci Tech* 13:984-996

Roth GI, Katz J (2001) Five techniques for increasing the speed and accuracy of PIV interrogation. *Meas Sci Technol* 12:238-245

Scarano F (2002) Iterative image deformation methods in PIV. *Meas Sci Technol* 13:R1-R19

- Scarano F (2013) Tomographic PIV: principles and practice. *Meas Sci Technol* 24,012001
- Scarano F, Riethmuller ML (2000) Advances in iterative multigrid PIV image processing. *Exp Fluids* 29:S51–S60
- Scarano F, David L, Bsibsi M and Calluaud D (2005). S-PIV comparative assessment: image dewarping+misalignment correction and pinhole+geometric projection. *Exp Fluids* 39:257-266
- Scarano F, Poelma C (2009) Three-dimensional vorticity patterns of cylinder wakes. *Exp Fluids* 47:69-83
- Schäfer L, Dierksheide U, Klaas M, Schröder W (2011) Investigation of dissipation elements in a fully developed turbulent channel flow by tomographic particle-image velocimetry. *Phys Fluids* 23, 035106
- Schanz D, Gesemann S, Schröder A, Wieneke B, Novara M (2013) Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. *Meas Sci Technol* 24, 024009
- Scheimpflug T (1904) Improved Method and Apparatus for the Systematic Alteration or Distortion of Plane Pictures and Images by Means of Lenses and Mirrors for Photography and for other purposes. GB Patent No. 1196
- Schneider GM, Hooper JD, Musgrove AR, Nathan GJ, Luxton RE (1997) Velocity and Reynolds stresses in a precessing jet flow. *Exp Fluids* 22:489–495
- Schröder A, Geisler R, Staack K, Elsinga GE, Scarano F, Wieneke B, Henning A, Poelma C, Westerweel J (2011) Eulerian and Lagrangian views of a turbulent boundary layer flow using time-resolved tomographic PIV. *Exp Fluids* 50:1071-1091
- Sciacchitano A, Scarano F, Wieneke B (2012) Multi-frame pyramid correlation for time-resolved PIV. *Exp Fluids* 53:1087-1105
- Seoud RE, Vassilicos JC (2007) Dissipation and decay of fractal-generated turbulence. *Phys Fluids* 19, 105108
- Shakouchi T (1981) An experimental study on the cavity type fluidic oscillator. *Res Rep Fac Engng Mie Univ* 6:1-13
- Shakouchi T (1989) A new fluidic oscillator, flowmeter, without control port and feedback loop. *J Dyn Sys, Meas, Control* 111:535-539
- Sheng J, Meng H, Fox RO (2000) A large eddy PIV method for turbulence dissipation rate estimation. *Chem Eng Sci* 55:4423-4434
- Smagorinsky J (1963) General circulation experiments with the primitive equation I the basic experiment. *Monthly Weather Review* 91:99-164
- Simmons JM, Platzer MF, Lai JCS (1981) Jet excitation by an oscillating vane. *AIAA J* 19:673–676

References

- Sirovich L (1987) Turbulence and the dynamics of coherent structures. *Quart Appl Math* 45:561-590
- Smith SW (1999) *The scientist and engineer's guide to digital signal processing*. California Technical Publishing, San Diego
- Soloff SM, Adrian RJ, Liu ZC (1997) Distortion compensation for generalized stereoscopic particle image velocimetry. *Meas Sci Technol* 8:1441-1454
- Sreenivasan KR (1984) On the scaling of the turbulence energy dissipation rate. *Phys Fluids* 27:1048-1051
- Sreenivasan KR (1998) An update on the energy dissipation rate in isotropic turbulence. *Phys Fluids* 10:528-529
- Sreenivasan KR, Meneveau C (1986) The fractal facets of turbulence. *J Fluid Mech* 173:357-386
- Staicu A, Mazzi B, Vassilicos JC, van de Water W (2003) Turbulent wakes of fractal objects. *Phys Rev Lett* 91:066306
- Syred N (2006) A review of oscillation mechanisms and the role of the precessing vortex core (PVC) in swirl combustion systems. *Prog Energy Combust Sci* 32:93-161
- Tanaka T, Eaton JK (2007) A correction method for measuring turbulence kinetic energy dissipation rate by PIV. *Exp Fluids* 42:893-902
- Taylor GI (1935) Statistical theory of turbulence. *Proc R Soc A* 151:421-444
- Tennekes H, Lumley JL (1972) *A first course in turbulence*. MIT Press
- Tokgoz S, Elsinga GE, Delfos R, Westerweel J (2012) Spatial resolution and dissipation rate estimation in Taylor-Couette flow for tomographic PIV. *Exp Fluids* 52:561-583
- Tong C, Warhaft Z (1994) Turbulence suppression in a jet by means of a fine ring. *Phys Fluids* 6:328-333
- Townsend AA (1956) *The structure of turbulent shear flows*. Cambridge Univ Press
- Tsai RY (1987) A versatile camera calibration technique for high accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses. *IEEE J Rob Autom* 4(RA-3):323-344
- Valente PC, Vassilicos JC (2011) The decay of turbulence generated by a class of multiscale grids. *J Fluid Mech* 687:300-340
- Valente PC, Vassilicos JC (2012) Universal dissipation scaling for non-equilibrium turbulence. *Phys Rev Lett* 108, 214503
- van Oudheusden BW, Scarano F, Van Hinsberg NP, Watt DW (2005) Phase-resolved characterization of vortex shedding in the near wake of a square-section cylinder at incidence. *Exp Fluids* 39:86-98
- Violato D, Moore P, Scarano F (2011) Lagrangian and Eulerian pressure field evaluation of rod-airfoil flow from time-resolved tomographic PIV. *Exp Fluids* 50:1057-1070

- Violato D, Scarano F (2011) Three-dimensional evolution of flow structures in transitional circular and chevron jets. *Phys Fluids* 23, 124104
- Voth GA, Satyanarayan K, Bodenschatz E (1998) Lagrangian acceleration measurements at large Reynolds numbers. *Phys Fluids* 10:2268-2280
- Wereley ST, Meinhart CD (2001) Second-order accurate particle image velocimetry. *Exp Fluids* 31:258-268
- Westerweel J (1993) Digital Particle Image Velocimetry – Theory and Application. PhD Thesis. Delft University of Technology
- Westerweel J (1997) Fundamentals of digital particle image velocimetry. *Meas Sci Technol* 8:1379-1392
- Westerweel J (2000) Theoretical analysis of the measurement precision in particle image velocimetry *Exp Fluids* 29:S3–12
- Westerweel J, Dabiri D, Gharib M (1997) The effect of a discrete window offset on the accuracy of cross-correlation analysis of digital PIV recordings. *Exp Fluids* 23:20–28
- Westerweel J, Scarano F (2005) Universal outlier detection for PIV data. *Exp Fluids* 39:1096-1100
- Westerweel J, Elsinga GE, Adrian RJ (2013) Particle Image Velocimetry for complex and turbulent flows. *Ann Rev Fluid Mech* 45:409-436
- Wieneke B (2005) Stereo-PIV using self-calibration on particle images. *Exp Fluids* 39:267-280
- Wieneke B (2008) Volume self-calibration for 3D particle image velocimetry. *Exp Fluids* 45:549-556
- Wieneke B (2013) Iterative reconstruction of volumetric particle distribution. *Meas Sci Technol* 24, 024008
- Willert CE (1997) Stereoscopic digital particle image velocimetry for application in wind tunnel flows. *Meas Sci Technol* 8:1465-1479
- Willert CE, Gharib M (1991) Digital particle image velocimetry. *Exp Fluids* 10:181-193
- Willert CE, Gharib M (1992) Three-dimensional particle imaging with a single camera. *Exp Fluids* 12:353-358
- Wong CY, Lanspeary PV, Nathan GJ, Kelso RM, O'Doherty T (2003) Phase-averaged velocity in a fluidic precessing jet nozzle and in its near external field. *Exp Therm Fluid Sci* 27:515–524
- Wong CY, Nathan GJ, O'Doherty T (2004) The effect of initial conditions on the exit flow from a fluidic precessing jet nozzle. *Exp Fluids* 36:70–81
- Wong CY, Nathan GJ, Kelso RM (2008) The naturally oscillating flow emerging from a fluidic precessing jet nozzle. *J Fluid Mech* 606:153–188

References

Worth NA, Nickels TB (2008) Acceleration of Tomo-PIV by estimating the initial volume intensity distribution. *Exp Fluids* 45:847-856

Worth NA, Nickels TB, Swaminathan N (2010) A tomographic PIV resolution study based on homogeneous isotropic turbulence DNS data. *Exp Fluids* 49:637–656

Ziskin IB, Adrian RJ (2011) Volume segmentation tomographic particle image velocimetry. Proc. of 9th International Symposium on Particle Image Velocimetry PIV '11, July 21st – 23rd, Kobe University, Kobe, Japan

LIST OF PUBLICATIONS

Journal papers:

Discetti S, Natale A, Astarita T (2013) Spatial Filtering Improved Tomographic PIV. *Exp Fluids*, doi: 10.1007/s00348-013-1505-7

Discetti S, Ziskin IB, Astarita T, Adrian RJ, Prestridge K (2013) PIV measurements of anisotropy and inhomogeneity in decaying fractal generated turbulence. *Fluid Dyn Res* (in press).

Discetti S, Adrian RJ (2012) High accuracy measurement of magnification for monocular PIV. *Meas Sci Technol*, 23, 117001, doi 10.1088/0957-0233/23/11/117001, ISSN 0957-0233.

Discetti S, Astarita T (2012) Fast 3D PIV with direct sparse cross-correlations. *Exp Fluids*, 55(3):1437-1451, doi 10.1007/s00348-012-1370-9, ISSN 0723-4864.

Discetti S, Astarita T (2012) A fast multi-resolution approach to tomographic PIV. *Exp Fluids*, 52(3):765-777, doi 10.1007/s00348-011-1119-x, ISSN 0723-4864.

Carlomagno GM, Discetti S, Astarita T (2011) Experimental assessment of a new heat flux sensor for measuring convective heat transfer coefficients. *QIRT Journal*, vol 8(1) pp. 37-49, ISSN 1768-6733.

Discetti S, Astarita T (2011) Multi-Resolution algorithms for fast PIV image processing. *Aerotecnica Missili & Spazio*, 90:3-11, ISSN 0365-7422.

Journal papers under review:

Discetti S, Ianiro A, Astarita T, Cardone G (2013) On a novel low-cost tomographic PIV system. Under consideration for publication in *Meas Sci Technol*.

Contributions in conference proceedings:

Discetti S, Astarita T, "Effect of poor discretization in weighted cross-correlation for 3D PIV". *Proceedings of 16th International Symposium on application of laser techniques to fluid mechanics*, July 9th-12th 2012, Lisbon (Portugal),

Discetti S, Ianiro A, Astarita T, Cardone G, "On the development of a novel low cost high accuracy experimental setup for Tomographic Particle Image Velocimetry". *Proceedings of 16th International Symposium on application of laser techniques to fluid mechanics*, July 9th-12th 2012, Lisbon (Portugal)

List of publications

Ceglia G, Discetti S, Astarita T, Carlomagno GM, "On the coherent structures motion in horizontal convection". Proceedings of 15th International Symposium on Flow Visualization, June 25th-28th 2012, Minsk (Belarus), 8pp, ISBN 9789856456759

Ceglia G, Discetti S, "An experimental analysis in horizontal convection with IR thermography". Proceedings of 11th International Conference on Quantitative InfraRed Thermography, June 11th-14th 2012, Naples (Italy), 8pp, ISBN 97888906484417

Discetti S, Astarita T, "Assessment of Spatial filtering Improved TomoPIV". Proceedings of Forum on recent developments in Volume reconstruction techniques applied to 3D fluid and solid mechanics, November 28th – December 1st 2011, Poitiers (France).

Discetti S, Astarita T, "Fast multi-resolution 3D PIV with direct correlations and sparse arrays". Proceedings of Forum on recent developments in Volume reconstruction techniques applied to 3D fluid and solid mechanics, November 28th – December 1st 2011, Poitiers (France)

Discetti S, Astarita T, "Horizontal natural convection PIV study". In Bulletin of the American Physical Society, vol 56(18), 1pp, Proceedings of 64th Annual Meeting of the APS Division of Fluid Dynamics, 20-22 November 2011, Baltimore (MD, USA).

Discetti S, Astarita T, Carlomagno GM, "Horizontal natural convection: a PIV study". Proceedings of 7th International Symposium on Stratified Flows, August 22-26th 2011, Rome (Italy), 8pp, ISBN 978-88-95814-49-0

Discetti S, Astarita T, "Spatial filtering Improved Tomographic PIV". Proceedings of 9th International Symposium on Particle Image Velocimetry PIV '11, July 21-23th 2011, Kobe (Japan).

Discetti S, Astarita T, "Fast Multi-resolution 3D PIV". Proceedings of 9th International Symposium on Particle Image Velocimetry PIV '11, July 21-23th 2011, Kobe (Japan).

Discetti S, Ziskin IB, Adrian RJ, Prestridge K, "PIV study of fractal grid turbulence". Proceedings of 9th International Symposium on Particle Image Velocimetry PIV '11, July 21-23th 2011, Kobe (Japan).

Discetti S, "PIV investigation of horizontal natural convection". Proceedings of 5th International Conference on Vortex Flows and Vortex Models, November 7-10th 2010, Caserta (Italy), ISBN 978-88-905218-6-7

Discetti S, Astarita T, "A fast reconstruction algorithm for Tomo-PIV". Proceedings of XVIII A.I.V.E.L.A. National Meeting, December 15-16th 2010, Roma (Italy), ISBN 9788896378526

Carlomagno GM, Discetti S, Astarita T, “Experimental assessment of a new technique for measuring heat transfer coefficients”. Proceedings of 10th International Conference on Quantitative InfraRed Thermography, July 27-30th 2010, Quebec City (Canada), ISBN 978-2-9809199-1-6.

Discetti S, Astarita T, “Acceleration of Tomo-PIV by multigrid reconstruction schemes”. Proceedings of 15th International Symposium on Applications of Laser Techniques to Fluid Mechanics, July 5-8th 2010, Lisbon (Portugal)

Carlomagno GM, Discetti S, Astarita T, “A novel technique for measuring convective heat transfer coefficients”. Proceedings of ASME-ATI-UIT 2010 Conference on Thermal and Environmental Issues in Energy Systems, May 16-19th 2010, Sorrento (Italy), 5pp. ISBN 978-884672659-9

Discetti S (2010) On a rapid algorithm for PIV image processing. VI AIAA-Pegasus Student Conference, April 28-30th 2010, Seville (Spain).

ACKNOWLEDGMENTS

"I am a part of all that I have met"

Alfred Tennyson

Achievements are the sweet melody of a noisy life. Since I was a kid, I was taught to fight to pursue whatever I wanted, and this has always been my religion. However, you can climb very high with your own arms, but sooner or later you end up feeling you cannot make it anymore by yourself. In those moments, the people around you determine how high you will bounce back after your falls.

Firstly, I wish to acknowledge Prof. Tommaso Astarita for his invaluable support. An impressively skilled person, a loyal friend, a patient guide, a model for what I want to be, and many many other things you meant to me. I would have hardly done anything without your continuous support and commitment in doing always the best thing for me. I really appreciate the trust you gave me since the very beginning. You blindly believed me in every situation. Very few are the students as lucky as I have been during my PhD.

I sincerely thank Prof. Carlomagno for all the experience, the knowledge and the support he transmitted me in every moment. You always had the right word to back up my growth as a man and as a researcher. Knowing that I could have always counted on you every time my ship was going off course really helped me to get through the many difficulties of the experimental research.

I would like to thank Prof. Adrian for hosting me twice in his lab at Arizona State University for a total period of one year. Thank you very much for the insightful discussions and the directions you continuously gave me to improve the quality of my research. The experience really enriched me from a cultural and human point of view. I enjoyed very much your companionship, hospitality, and prompt disposability in giving me everything I needed to make the best out of me.

I thank also all the faculty members I met during my PhD at the Department of Aerospace Engineering (DIAS) at the University of Naples. A special thank goes to Prof. Cardone, who was my tutor for my Bachelor's thesis. Actually, I don't think I will be here without the interest on fluid dynamics you excited in me in the very first stages of my career. I extend the thank to Prof. De Luca for the insightful discussions on instability in fluid dynamics, and his inspiring way of giving lectures, which I consider a model for my teaching commitments.

I would like to thank the colleagues/friends of the Laboratory of Gas Dynamics at the University of Naples. I wish to thank Giuseppe Ceglia, who I admire for the pragmatism and the astonishing way he can get through the variegated situations of

complicate experimental setups. I really acknowledge your sense of loyalty and the reciprocal respect and affection we give each other. I thank Andrea Ianiro, for the patience during the first year of my PhD in training me on the use of InfraRed Thermography, and for all the adventures we lived in these past years, sharing ideas, opinions, and a lot of reciprocal help during every activity we have been individually involved. I would like to thank also Giuseppe Sicardi, Annagrazia Orazzo, Michele Imbriale, Mario Panelli, Fortunato de Rosa, Carlo Greco, Francesco Avallone and Gioacchino Cafiero for the many enjoyable moments we shared together during these years.

Certainly I cannot forget the many undergraduate and graduate students I trained and led to the graduation during these years. Sometimes it has been tough to keep you straight on your route, but your work has always been valuable.

A special thank goes to group of students working in the Laboratory for Energetic Flows and Turbulence at Arizona State University. In particular, I would like to thank Liuyang Ding for the respect you always demonstrated me, the commitment while I was training you in your very first weeks in the lab, and your loyal friendship. Once you said that it is very hard to find a friend; well, I am disappointed I have found one that lives 10000km away from me. I also thank Rafeed Chaudhury, Jon Baltzer and Isaac Ziskin for your very nice companionship. You really made me feel home.

I also warmly thank LaVision GmbH for providing the experimental equipment for the precessing jet experiment.

Last, but most remarkably, I would like to thank my family, which supported me (with great patience and commitment) in reaching this goal. Thanks mom for all the sacrifices and the lovely way you dedicated your entire life to my (and my siblings') happiness. My achievements are your achievements. You have given me all the strength, courage, wilfulness, love, that I needed to overcome my limits and eventually find my way in this world. Many many thanks to my sister Carmen and my brother Gabriele. I have never dedicated to you all the time you deserved, but I am sure you know that I really love you. A thought goes necessarily to my father, who cannot be here to enjoy this success. No day has passed without thinking of you. I know you would be proud of me.

Dulcis in fundo, I wish to thank my lovely wife Mariella. You supported me along this difficult path, bearing the burden to cheer me up when I was down, sharing my emotions as they were yours, listening endless conversation on my infinite number of problems in the experiments, understanding the long periods of time in which I could not give you the attention you needed, and always being on my side even when I did not deserve it. Somebody wrote that we are just like angels with one wing, thus we cannot fly alone. This is certainly true, but I do believe that I would not be where I am now without you flapping my wing too when I was not strong enough to do it by myself.