

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/278392925>

A scanning PIV method for fine-scale turbulence measurements

Article in *Experiments in Fluids* · November 2014

DOI: 10.1007/s00348-014-1857-7

CITATIONS

33

READS

422

2 authors:



John M. Lawson

University of Southampton

24 PUBLICATIONS 218 CITATIONS

SEE PROFILE



James R. Dawson

Norwegian University of Science and Technology

97 PUBLICATIONS 2,086 CITATIONS

SEE PROFILE

A scanning PIV method for fine-scale turbulence measurements

John M. Lawson · James R. Dawson

Received: 2 September 2014 / Revised: 16 October 2014 / Accepted: 1 November 2014
© Springer-Verlag Berlin Heidelberg 2014

Abstract A hybrid technique is presented that combines scanning PIV with tomographic reconstruction to make spatially and temporally resolved measurements of the fine-scale motions in turbulent flows. The technique uses one or two high-speed cameras to record particle images as a laser sheet is rapidly traversed across a measurement volume. This is combined with a fast method for tomographic reconstruction of the particle field for use in conjunction with PIV cross-correlation. The method was tested numerically using DNS data and with experiments in a large mixing tank that produces axisymmetric homogeneous turbulence at $R_\lambda \simeq 219$. A parametric investigation identifies the important parameters for a scanning PIV set-up and provides guidance to the interested experimentalist in achieving the best accuracy. Optimal sheet spacings and thicknesses are reported, and it was found that accurate results could be obtained at quite low scanning speeds. The two-camera method is the most robust to noise, permitting accurate measurements of the velocity gradients and direct determination of the dissipation rate.

1 Introduction

Making measurements of the smallest scales of turbulent motion has been a long standing challenge in fluid

mechanics. It is a desirable aim, as the fine scales are responsible for (amongst other things) the dissipation of kinetic energy and chemical mixing (Tsinober 2009) and have an important role in phenomena such as cloud formation (Bodenschatz et al. 2010) and flame extinction (Sreenivasan 2004). The challenge stems from the fact that as the Reynolds number increases, the smallest length- and timescales become smaller and faster, requiring significant measurement or computational capability to capture even low Reynolds number flows.

For laboratory investigations of turbulence, it would be valuable to have an accurate, three-dimensional, three-component (3D–3C) velocimetry technique with good space and time resolution. Of particular importance is access to the full velocity gradient tensor (VGT), as it embodies the fine scales of turbulence and provides direct access to the turbulent kinetic energy dissipation rate (Wallace 2009). Because of the wide range of lengthscales involved, a large spatial dynamic range (SDR, the ratio of the smallest measurable lengthscale to the largest) is also valuable.

Hot wire probes with nine, twelve or even twenty wires (Wallace 2009) offer access to the full VGT with good temporal resolution, but are invasive, complicated to calibrate and only offer pointwise measurements. Dual plane PIV can offer a 3D–3C measurement with good accuracy, time and spatial resolution (Mullin and Dahm 2006), but the set-up is complex: it requires four cameras and some means of differentiating between planes, usually chromatically by using two different wavelength laser sources. Furthermore, the depth of the measurement volume is very limited (i.e. between two planes), and the full VGT is only available in one plane. Cinematographic stereo PIV can be used to yield snapshots of 3D–3C velocity fields (Ganapathisubramani et al. 2007), but is only applicable to convective flows as it relies upon Taylor's frozen turbulence hypothesis.

J. M. Lawson (✉)
Cambridge University Engineering Department, Trumpington
Street, Cambridge CB2 1PZ, UK
e-mail: jml70@cam.ac.uk

J. R. Dawson
Department of Energy and Process Engineering, Norwegian
University of Science and Technology, 7491 Trondheim, Norway
e-mail: james.r.dawson@ntnu.no

Scanning PIV and tomographic PIV/PTV are the only non-invasive measurement techniques which can offer practical, time-resolved, 3D–3C turbulence measurements. Tomographic PIV has been applied to numerous experimental investigations of turbulent flows (Worth and Nickels 2011; Elsinga et al. 2006; Schröder et al. 2008). Whilst the technique offers ready experimental access to the full 3D–3C velocity field, it suffers from some unfortunate disadvantages, particularly when accurate turbulence measurements are needed. For adequate tomographic reconstruction, constraints upon seeding density confer a restriction upon the spatial resolution which can be achieved (Scarano 2013) as the correlation volumes must contain a sufficient number of particles. Because the reconstruction problem is under-constrained, “ghost particles” are reconstructed, which introduce significant bias error and degrade spatial resolution (Elsinga et al. 2011).

Numerous variations on a technique broadly labelled scanning PIV, first introduced by Brücker (1995), have been developed over the past two decades. Whilst specific implementations differ, the overall concept is similar: one or more cameras capture particle images as a laser sheet is quickly moved across the measurement volume. The most common approach has been to perform stereoscopic PIV cross-correlations of images at different depths through the volume, which are then “stacked” together to recover the velocity field (Brücker 1996; Hori and Sakakibara 2004; Diez et al. 2011; Soodt et al. 2011). However, this suffers from the same disadvantages as stereo PIV: error in the out-of-plane component is greater than the in-plane components and resolution is constrained by the finite thickness of the laser sheet, with the added complexity of scanning.

To get around the limitations of conventional tomographic PIV, a number of authors have turned to laser scanning methods (Casey et al. 2013; Ponitz et al. 2012; Brücker et al. 2013). The approach is to perform tomographic reconstructions of the 3D particle field as the first step, then use 3D reconstructions (discretised into cubic voxels) to perform a velocimetry. Brücker et al. (2013) solve a tomographic reconstruction problem (although it is not expressed in such terms) for just one camera with many images taken at different laser sheet positions. Ponitz et al. (2012) and Casey et al. (2013) combine multiple-camera tomographic reconstructions independently calculated for each position of the (thin) laser sheet. Both techniques reduce the incidence of ghost particles and permit high seeding densities to be used. All techniques achieve some degree of super-resolution in the scanning direction: the spacing between laser sheets may be tens of voxels wide.

The main disadvantage of scanning techniques is that in practice, the flow speed and timescales must be quite slow in order to satisfy the “frozen flow” condition during

scanning. However, the advantages the technique provides in terms of seeding density may justify its use in cases where the flow can be quite slow moving and small-scale information is important. For example, scanning methods have already been applied to the far field of a turbulent jet (Hori and Sakakibara 2004; Diez et al. 2011; Casey et al. 2013), to study unsteady separation around a flapping wing (David et al. 2012) and the unsteady flow field in a model IC engine (Brücker 1997). The present work specifically addresses the issue of how fast one needs to scan.

This paper presents a variant of scanning PIV using a tomographic approach. Whilst the scanning tomographic technique has been shown to be viable, it is still in its infancy and there are a number of questions which existing literature cannot resolve. The interested experimentalist will wonder: does the technique provide measurements of sufficient quality to justify its use? Furthermore, there are questions about the details of applying scanning PIV. How many images must be taken per reconstruction? How far apart can the laser sheets be spaced? How thick must the laser sheet be? How quickly does one need to scan? How robust is the technique to image noise and misalignment? The first half of this paper aims to answer these questions. The second half demonstrates the high fidelity of the technique through its application to a laboratory study of homogeneous, axisymmetric turbulence at moderate Reynolds number.

The paper is structured as follows. Sections 2.1 and 2.2 outline the scanning tomographic reconstruction and self-calibration technique, which contain some novel aspects. Section 2.3 then explains the details of a parametric investigation designed to quantify the effect of different set-up parameters upon measurement accuracy. Section 2.4 gives details of a laboratory experiment in a von Kármán swirling water flow, where scanning PIV has been applied. Sections 3.1, 3.2, 3.3 and 3.4 present results of the parametric investigation, and Sect. 3.5 assesses the performance of the self-calibration scheme. Section 3.6 then presents a number of results from the laboratory experiment, which demonstrate the technique can achieve high fidelity, spatially resolved velocity and velocity gradient measurements.

2 Methods

2.1 Scanning PIV reconstruction

In this section, the principles of the scanning tomographic reconstruction are outlined in an approach which differs from previous techniques. The general set-up of the scanning PIV laboratory experiment is shown in Fig. 1. The numerical simulation uses a similar configuration. One or more cameras are positioned to observe the measurement

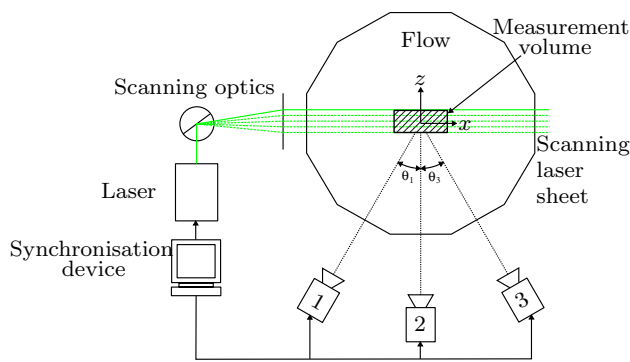


Fig. 1 Scanning PIV set-up, as used in laboratory experiment and in numerical simulation

volume as a laser sheet is traversed (“scanned”) rapidly across it. An image n is recorded for each of the $1, \dots, N_S$ positions of N_S near-parallel laser sheets, which overlap substantially.

The objective of the tomographic reconstruction is to find a scalar field $E(\mathbf{x})$, which represents the reflectance of some small lump of diffusely reflecting surfaces (a particle) at point \mathbf{x} . Thus, $E(\mathbf{x}) = 0$ when there is no particle at \mathbf{x} and is non-negative otherwise. Interference effects, absorption and specular reflection are neglected in this model. The reconstruction is made from a number of images $I_{n,j}(\mathbf{y}_j)$ recorded by cameras j . A given point in object space \mathbf{x} will influence image $I_{n,j}$ at a position \mathbf{y}_j , by an amount which is proportional to the strength of illumination and the reflectance at that point. The projection of point \mathbf{x} onto the image plane is modelled using a pinhole camera model (Hartley 2003):

$$\tilde{\mathbf{y}}_j \sim \mathbf{G}_j \mathbf{R}_j [\mathbf{I}_3 | -\mathbf{x}_{c,j}] \tilde{\mathbf{x}} \tag{1}$$

\mathbf{I}_3 is the 3×3 identity matrix and \mathbf{G}_j is a 3×3 upper triangular matrix which represents intrinsic camera parameters, such as focal length and magnification. The variables $\mathbf{x}_{c,j}$ and \mathbf{R}_j are extrinsic camera parameters which represent the camera’s position in object space and orientation with respect to the object space coordinate system, respectively. The subscript j corresponds to having different model parameters for each camera. $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{y}}_{c,j}$ correspond to \mathbf{x} and $\mathbf{y}_{c,j}$ in homopolar coordinates.

Using a Gaussian model of laser sheet intensity, one can model the illumination at point \mathbf{x} for the n th laser sheet as

$$f(\mathbf{x}, n) = \exp\left(-8(\mathbf{x} \cdot \mathbf{e}_n - d_n)^2 / w_n^2\right) \tag{2}$$

where w_n is the e^{-2} width of the sheet and \mathbf{e}_n and d_n give the equation of plane for the given laser sheet.

Using the pinhole camera model, any point \mathbf{y}_j viewed by camera j is the projection of a single ray, described by

$\mathbf{x}_{\text{ray}}(s)$ where s parameterises position along the ray. This yields a closed model for the signal recorded at a given point by a given camera:

$$I_{n,j}(\mathbf{y}_j) = \int_0^\infty E(\mathbf{x}_{\text{ray}}) f(\mathbf{x}_{\text{ray}}, n) ds. \tag{3}$$

Reconstruction of $E(\mathbf{x})$ is now a tomography problem, which can be solved efficiently using the MART algorithm (Herman and Lent 1976; Worth and Nickels 2008). For the single-camera case, rays are reconstructed individually upon a discrete grid defined by the intersection of planes of constant depth z with rays back-projected from each pixel (Fig. 2a). Let a given ray have intersections with planes of constant z which are indexed by m and write the reflectance at this location as E_m . This ray is always projected onto the same pixel on the camera, so there are N_S observations of this pixel’s brightness denoted ϕ_n . Using this discretisation, Eqs. (2) and (3) can be used to define a tensor of weights W_{nm} as in (4). The tomography problem is solved by finding the solution for E_m using the MART algorithm.




Fig. 2 Schematic representation of reconstruction grid for **a** single-camera reconstruction and **b** stereo reconstruction

$$\phi_n = \sum_m W_{nm} E_m \quad (4)$$

When all rays have been reconstructed, a resampling in planes of constant z (using a 10×10 Lanczos interpolation kernel) provides a reconstruction of the particle field on a regular cartesian grid. This extends the work of Brücker et al. (2013) to consider a more generalised pinhole projection model and permit the use of a MART solver.

For the stereo case, independent 2D reconstructions are performed on many different epipolar planes. Points on an epipolar plane will always be projected onto so-called epipolar lines in either camera, as shown in Fig. 2b. That is to say, both cameras view the epipolar plane edge-on. Equation (3) is cast as a 2D-tomography problem in this plane, with a grid defined by the intersection of rays from camera 1 with rays from camera 2. By indexing the rays from the first camera as $s = 1, \dots, N_1$ and the second as $t = 1, \dots, N_2$, a 2D tensor E_{st} can be built with $N_1 \times N_2$ elements which represents the planar reconstruction. The images can be interpolated, again using a 10×10 Lanczos kernel, to yield a set of observations $\phi_{1,ns}$ and $\phi_{2,nt}$ from the first and second cameras, respectively, corresponding to different sheet positions and rays. As before, the projection is represented as a linear operation on E_{st} , namely

$$\begin{aligned} \phi_{1,ns} &= \sum_t W_{nst} E_{st} \\ \phi_{2,ns} &= \sum_s W_{nst} E_{st} \end{aligned} \quad (5)$$

The 3D tensor W_{nst} represents the illumination of the point indexed by (s, t) for a sheet n , and the summation over s and t represents the integral of Eq. (3). The projection model assumes particles have the same reflectance when viewed from either camera. In practice, this is not the case, but images can be normalised as a preprocessing step to account for this. Equation (5) can then be cast into matrix form and solved using MART. The solutions are first resampled using linear interpolation on epipolar planes with a regularly spaced grid (x, z) for arbitrary y and then resampled again to provide reconstructions on a regularly spaced Cartesian grid (x, y, z) . This combines the multiple view problem discussed by Casey et al. (2013) and Ponitz et al. (2012) with the multiple sheet problem described by Brücker et al. (2013). In addition, this method is not restricted to the use of a simple camera pinhole model. Nonlinear distortion can be compensated for in the image plane through the use of a dewarping function applied before the images are interpolated.

2.2 Self-calibration

Whilst the single-camera reconstruction technique is relatively insensitive to calibration error, it is important for

the two-camera technique in a manner similar to regular tomographic PIV (Elsinga et al. 2006). To rectify this, a self-calibration technique similar to Wieneke (2005) was implemented. First, single-camera reconstructions of the same particle field are performed for both cameras and are cross-correlated with find a set of disparity vectors. Then, a correction is found to the camera pinhole model and laser sheet model that provides a best fit to the observed disparity vectors in the least-squares sense.

Consider the disparity $\delta(\mathbf{x}_1)$ at position \mathbf{x}_1 : particles that appear at \mathbf{x}_1 in the first camera reconstruction tend to appear at $\mathbf{x}_2 = \mathbf{x}_1 + \delta(\mathbf{x}_1)$ in the second reconstruction. Suppose one had a priori information that a given particle was recorded at position \mathbf{y}_1 on the first camera and at \mathbf{y}_2 on the second camera and furthermore that one also knew what time t_B during the acquisition sequence this particle was illuminated most brightly. This could be inferred using Eq. (2) from the variation in brightness of the particle over the course of a scan. The single-camera back-projection should reconstruct it at some position \mathbf{x}_1 in the first reconstruction and \mathbf{x}_2 in the second with any disparity δ indicating a miscalibration.

Having measured the disparity field, using the uncorrected projection model, one can find \mathbf{y}_1 , \mathbf{y}_2 and t_B for some imaginary particle at \mathbf{x}_1 in the first reconstruction and corresponding \mathbf{x}_2 in the second. The point \mathbf{x}_1 is arbitrary: it does not have to correspond to a real particle, but it does have a correspondence to position \mathbf{x}_2 . With this in mind, one could suppose that there exists a small correction to the model which can be written as

$$\begin{aligned} \tilde{\mathbf{y}}_j &\sim \mathbf{G}_j \mathbf{R}_j \boldsymbol{\Omega}_j [\mathbf{I}_3 | -\mathbf{x}_{c,j}] \tilde{\mathbf{x}}, \\ \mathbf{x}_j \cdot (\boldsymbol{\Omega}_s \mathbf{e}(t_B)) - (d(t_B) + \epsilon_n) &= 0 \end{aligned} \quad (6)$$

This correction represents a small rotation $\boldsymbol{\Omega}_j$ of either camera, a shift in sheet position ϵ_d and a small rotation of the laser sheet normal $\boldsymbol{\Omega}_s$. $\boldsymbol{\Omega}_j$ and $\boldsymbol{\Omega}_s$ are rotation matrices. Because t_B may not coincide exactly with when an image was taken, a linear interpolation provides the orientation of the sheet normal $\mathbf{e}(t_B)$ and sheet position $d(t_B)$ at this time.

For a given disparity model $(\boldsymbol{\Omega}_j, \boldsymbol{\Omega}_s, \epsilon_d)$, one finds the back-projection \mathbf{x}_j from either camera j of a particle 'observed' at \mathbf{y}_j and brightest at time t_B , which provides a model of the disparity vector δ . Thus, self-calibration is achieved by finding the best model parameters which match the observed disparity field $\delta(\mathbf{x})$ in the least-squares sense. As self-calibration requires only trivial computational expense and can rectify even large calibration errors well (as demonstrated in Sect. 3.4), its application is recommended.

2.3 Parametric investigation

A handful of parameters can be used to describe the scanning PIV set-up in Fig 1. These are the number of cameras

Table 1 Relevant scanning PIV parameters for the simulation and laboratory experiments

Parameter	Symbol	Numerical simulation	Laboratory experiment
Kolmogorov lengthscale (mm)	η_k	0.00287	1.02
Kolmogorov timescale (ms)	τ_k	0.0446	1.079
RMS turbulent velocity (mm/s)	u_{turb}	0.681	7.12
Kinematic viscosity (mm ² /s)	ν	0.000185	0.97
Taylor microscale Reynolds number	R_λ	433	219
Measured volume width	L_1/η_k	128	123
Measured volume height	L_2/η_k	128	121
Measured volume depth	L_3/η_k	32	23
Number of independent samples	N	10	1,003
Number of laser sheets	N_S	40	50
Image resolution (px)	–	1,024 × 1,024	1,024 × 1,024
Laser sheet spacing (voxels)	$M\Delta z$	Variable	4.75
Sheet e^{-2} width	$w/\Delta z$	Variable	4.2
Camera frame rate (kHz)	f_{scan}	Variable	2
Camera 1 viewing angle	θ_1	–30°	~ –30°
Camera 2 viewing angle	θ_2	0°	n/a
Camera 3 viewing angle	θ_3	30°	~ 30°
Laser sheet speed	u_s/u	Variable	157
Magnification (vx/mm)	M	6.35	7.65
PIV time separation	$\Delta t/\tau_k$	0.089	0.056
PIV window size (vx)	$W_1 \times W_2 \times W_3$	32 × 32 × 32	48 × 48 × 48
Seeding density (vx ⁻³)	–	23.7/32 ³	~ 25/32 ³
PIV window overlap	–	75 %	75 %
PIV vector spacing	$\Delta x/\eta_k$	1	1.53

C and their relative angles θ_j , the optical magnification M , the number of images per scan N_S , the frequency of image acquisition f_S , the seeding concentration ρ_S , the thickness of the laser sheet w and the spacing between consecutive sheets Δz . The characteristic velocity u of the flow also needs to be considered, which in this case is taken to be the RMS velocity in the z direction. One can construct a “sheet speed” $u_s = f_S \Delta z$ that defines the rate at which the laser sheet is scanned across the measurement volume. Six independent dimensionless groups arise out of these terms: $M\Delta z$, the sheet spacing (in voxels), $w/\Delta z$, the dimensionless sheet thickness, u_s/u , the dimensionless sheet speed, the dimensionless seeding concentration $\rho_S M^{-3}$ (particles/vx³) as well as C and θ_j .

To investigate this parameter space, numerical simulations of scanning PIV were performed under various conditions. Synthetic PIV images were generated for simulations of the scanning PIV configuration shown in Fig. 1. This was achieved by generating a random distribution of seeding particles, with a given seeding density, which were then advected without slip along Lagrangian trajectories in a simulation of isotropic turbulence ($R_\lambda = 433$) from the JHU turbulence database (Li et al. 2008). Scanning PIV was simulated by generating images of the particles

at different laser sheet positions using the EUROPIV Synthetic Image Generator (Lecordier and Westerweel 2004). Salient parameters of the reconstruction are defined in Table 1. The synthetic images were passed to the reconstruction routine; subsequent reconstructions were cross-correlated using an in-house code. The resultant velocity fields were then checked against their expected values. For each combination of parameters, ten pairs of scanned volumes were simulated. The same ten velocity fields were used to advect particles each time, so the only differences between runs are the set-up parameters. Issues related to statistical convergence are discussed in Sect. 3.1.

To test the self-calibration scheme, the effect of perturbations to a typical stereo scanning set-up where $M\Delta z = 4$ vx, $w/\Delta z = 4$, $N_S = 65$, and $\theta_j = \pm 30^\circ$ was tested using a Monte Carlo procedure. As in the PIV simulation, images of a synthetic particle field are generated from the perspective of the two stereo cameras. Then it is supposed that the orientation of the cameras and the position/orientation of the laser sheets has been measured inaccurately. This is simulated by handing a perturbed back-projection model to the reconstruction algorithm. Single-view reconstructions are made from either camera, which no longer overlap exactly because of the imposed

misalignment. The self-calibration scheme is then applied, and the disparity model tested to see how well disparity can be minimised. This was done using 1,000 Monte Carlo trials: 34 resulted in a disparity too large to measure reliably (over 32 voxels in any direction), so statistics of the remaining 966 trials are presented.

2.4 High Reynolds number mixing tank experiment

The present scanning PIV technique was specifically developed to study fine-scale turbulence experimentally at high Reynolds number in a purpose built large-scale mixing tank whose arrangement is colloquially referred to as the “French washing machine”. It consists of two counter-rotating impellers 1.6 m in diameter immersed in a dodecagonal Perspex tank filled with water as shown in Fig. 3. Baffles spanning the length of the tank are positioned at each vertex protruding a distance of 100 mm into the flow. The impellers generate a region of axisymmetric shearing near the middle of the tank, generating high levels of turbulence with a weak mean flow approaching zero near the centre of the tank. Since the largest lengthscale scales with the impeller radius (Douady et al. 1991), the Kolmogorov lengthscale is of the order of 1 mm at $R_\lambda \simeq 219$. This makes fully spatially resolved measurements possible with moderate levels of magnification (Worth 2010; Worth and Nickels 2011).

The experimental set-up is illustrated in Fig. 1 and is deliberately similar to the simulation experiment. A three-camera set-up was used, with the initial intention of comparing a stereoscopic arrangement (cameras 1 and 3) with a single-camera arrangement (camera 2). As such a comparison has been made redundant in this paper by the parametric investigation, only the results of stereo measurements made




Fig. 3 Schematic (a) and illustration (b) of the large mixing tank facility used in the laboratory experiment

by cameras 1 and 3 are presented. The cameras were Photron SA1.1s equipped with Sigma 105 mm macro-lenses and Scheimpflug adapters and are focused on a measurement volume at the centre of the mixing tank. The parameters of the laboratory scanning experiment are provided in Table 1. The experiment aimed to produce a large number of statistically independent samples of the turbulent velocity field, with accurate measurements of its spatial and temporal derivatives. To achieve this, each statistically independent sample consisted of a small time-series of ten vector fields, with time resolution determined by the PIV time separation.

At this point, it is emphasised that the laboratory experiment does not achieve a particularly high level of spatial resolution (the window size corresponds to around 6 mm). However, it does achieve a high seeding concentration (~ 25 particles/ 32^3 vx), which is necessary to capture fine-grained velocity field information smaller than the interrogation window size and contributes to the high fidelity of the velocity gradient measurements.

Throughout the experiment, the water temperature was maintained at 22 ± 1 °C. Polystyrene seeding particles (TS-20, Microbeads, AS) of density 1.05 kg/L and mean diameter $21.3 \mu\text{m}$ were used as flow tracers. Their size is two orders of magnitude smaller than the Kolmogorov scale, the Stokes number is very low (~ 0.015), and they occupy a volume fraction of ~ 2 ppm.

To generate approximately parallel laser sheets with control over laser sheet thickness, an optical set-up similar to Brücker et al. (2013) was used. This is illustrated in Fig. 4. The beam from a Darwin Duo pulsed Nd:YLF laser enters a set of three lenses which act as a beam expander, to control the laser sheet thickness. After passing through a spherical divergent lens (A), it reflects off a mirror (B) attached to a galvanometer (dynAXIS XS, Scanlab GmbH) and passes through a cylindrical convergent lens (C). Finally, a cylindrical divergent lens (D) helps expand the beam further into a laser sheet. The mirror surface (B) is placed at the focus of the convergent lens (C), such that as the laser beam is deflected by rotation of the mirror,

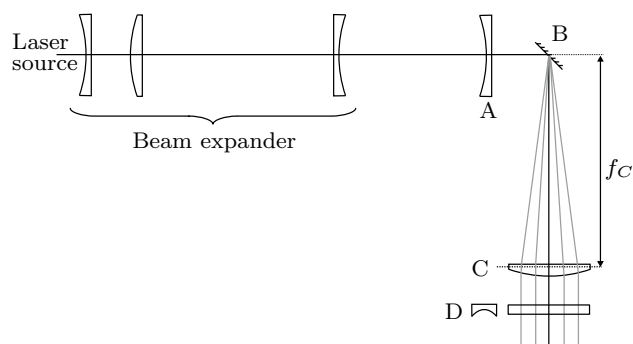


Fig. 4 Schematic of scanning optics used in the laboratory experiment

ensuring the laser sheets remain parallel. The divergent lens (A) was introduced to control the position of the beam waist induced at (C). By coordinating the steady rotation of the mirror with the laser triggering, parallel laser sheets can be generated which traverse the measurement volume at a constant rate.

The cameras were calibrated using a bundle calibration technique similar to Brücker et al. (2013). A calibration target at placed at an angle of 30° to the laser sheet was traversed across the measurement volume, permitting a pin-hole model to be fitted for both cameras and the position, thickness and orientation of the laser sheet during a scan to be measured. This information was then incorporated into the reconstruction model described in Sect. 2.1.

Using the self-calibration procedure described in Sect. 2.2, corrections to the projection model were generated for small batches of 80 samples at a time, which corresponds to around two hours worth of measurements. Each self-calibration used ten sets of images, from ten volumes in each batch, to construct an average disparity vector field between the views of cameras 1 and 3. The largest disparity tended to exist in the x direction, as large as $30 v_x$. The typical disparity observed in the y direction is much smaller (as large as $2.5 v_x$) and negligible in z . Most of the disparity can be accounted for by a drift in the z position of the laser sheets, which is probably due to thermal drift of the galvanometer's zero position. The disparity correction model was able to correct most of the disparity error: in the worst case, the RMS remaining disparity was 0.62 and 0.21 v_x in x and y , respectively.

2.5 PIV processing

For both the laboratory and simulation experiments, PIV processing was performed using an in-house code. Prior to cross-correlation, reconstructions are normalised with a nonlinear filter to minimise variations in particle image intensity, which can impair accuracy (Nobach 2011). This consists of a subtract sliding minimum filter over the lengthscale of a cross-correlation window, followed by normalisation by a local maximum value obtained from a sliding maximum filter whose size is the lengthscale of a particle. This code implements a 3D version of multiple pass cross-correlation with symmetric window deformation (Meunier and Leweke 2003). For window deformation, the velocity field is interpolated at each pass using a bilinear interpolation and a 10^3 Lanczos interpolation kernel is used for subpixel interpolation. Correlation windows are weighted using a Blackman weighting window (Cameron 2011; Astarita 2007), and peak location is performed using the a 3-point Gaussian fit for the first pass and least-squares Gaussian fit to a 5^3 region near the peak for subsequent passes. False vector rejection using median filtering and

infilling is performed after each pass. To obtain first-order velocity gradients, a “least-squares” finite difference operator is applied over four neighbouring points (Raffel 2007): this choice of operator strikes a balance between noise sensitivity and spatial resolution.

A version of the Lagrangian filtering used by Novara and Scarano (2013) was used to improve the measurement accuracy. In this scheme, tracers are tracked forwards and backwards in time, using the PIV velocity fields, from grid collocation points to yield time-series of tracer positions. A second-order polynomial is fitted to these trajectories, and the coefficients yield velocity and acceleration data over the short interval tracking is performed for (up to $\pm 0.28\tau_k$). This effectively applies a temporal filter to (Lagrangian) time-series data that help to reduce noise, at the expense of some temporal resolution.

Because both tomography problems (single camera or stereo) are split up into multiple, small, independent problems, they can be efficiently solved with parallel computation. On the author's desktop machine (Dell Optiplex 790 with Core i7-2600 CPU), a single-camera reconstruction of the experimental data takes around 2.5 min and a stereo reconstruction takes around 6 min. This is quite good: by comparison, Ponitz et al. (2012) perform scanning reconstructions of a $700 \times 500 \times 519 v_x$ volume in 35 min.

A note on nomenclature: in the subsequent analysis, it is most convenient to use tensor notation. Position is denoted x_i , with subscripts 1, 2 and 3 corresponding to x , y and z , respectively. The velocity field is denoted $U_i = \langle U_i \rangle + u_i$ where u_i is the turbulent fluctuating velocity. For the simulation experiment, $\langle U_i \rangle = 0$ so u_i simply refers to the velocity. Corresponding displacements in the PIV measurement are written δ_i . The (fluctuating) velocity gradient tensor is written as $A_{ij} = \partial u_i / \partial x_j$ and its symmetric part is S_{ij} . The fluctuating vorticity is written ω_i . Error in the i^{th} velocity component, denoted ϵ_i , is the difference between the measurement $\hat{U}_i(\mathbf{x}, t)$ taken at time t and the ground truth $U_i(\mathbf{x}, t)$ and is always presented in voxel units.

3 Results

3.1 Baseline error and convergence

In this paper, the RMS velocity error will be used extensively to assess measurement accuracy. It is therefore useful to get an idea of the size of error associated with the PIV algorithm alone. To this end, ideal reconstructions were simulated from the particle tracking data, by placing gaussian blobs of characteristic radius 1.4 v_x at the location of each particle. Thus, the method is identical to the scanning simulation, but the reconstruction process has been replaced with an idealised operation. The RMS

velocity errors in the x , y and z components were 0.036, 0.038 and 0.038 v_x . The RMS errors in longitudinal and transverse velocity gradients were around 2.0×10^{-3} and 2.4×10^{-3} v_x/v_x , respectively. For reference, the RMS velocity gradient is 16.5×10^{-3} v_x/v_x for longitudinal gradients and 23.3×10^{-3} v_x/v_x for transverse gradients.

Given the small number of volumes simulated, there is some uncertainty in the reported statistics of velocity error, because the error has a dependence upon the flow. This is mostly due to the PIV algorithm. To obtain estimates of statistical uncertainty in the RMS errors, the bootstrap procedure was applied using 1,000 bootstrap samples to all of the simulated data (Efron and Tibshirani 1994). Both non-parametric case resampling and the block bootstrap were applied. In case resampling, error vectors are randomly sampled with replacement, to generate a set of bootstrap samples. For the block bootstrap, independent velocity fields constitute separate blocks, which are then sampled with replacement to make up the bootstrap sample. Using these bootstrap samples, 95 % confidence intervals on RMS velocity errors can be calculated for the population in question.

Naturally, the block bootstrap sample estimates larger confidence intervals. For example, in the aforementioned ideal reconstruction, the 95 % confidence intervals (obtained from the bootstrap sample) of the RMS x velocity error are approximately $\pm 6.6 \times 10^{-3}$ v_x , with similar levels for y and z components. Across the entire survey, the block bootstrap confidence intervals for RMS velocity error were typically around ± 11 % and no larger than ± 20 % of the reported value. The regular, case resampling population shows much lower uncertainty in RMS error, with typical confidence intervals around ± 0.2 % and no larger than ± 2.5 % of the reported value. Thus, whilst there is moderate uncertainty in the absolute value of RMS error, the trends identified in the following sections are distinguishable from statistical uncertainty.

3.2 Single-camera parametric investigation

The problem of single-view reconstruction is considered first. A preliminary exploration of the parameter space is made, before incorporating additional complications such as the effect of finite sheet speed and noise. For a camera angle of -30° (camera 1), simulations with sheet thicknesses $w/\Delta z = 2, 4, 6$ and spacing $M\Delta z = 2, 4, 8, 16$ voxels ($N_S = 129, 65, 33, 17$) were performed. For the plane-normal camera ($\theta = 0^\circ$, camera 2), additional simulations were performed at $w/\Delta z = 3, 5$. Error in velocity was found to be anisotropic: it is worst in the viewing direction and best in the x and y directions.

For the plane-normal camera, the effect of sheet thickness and spacing on accuracy are shown in Fig. 5.

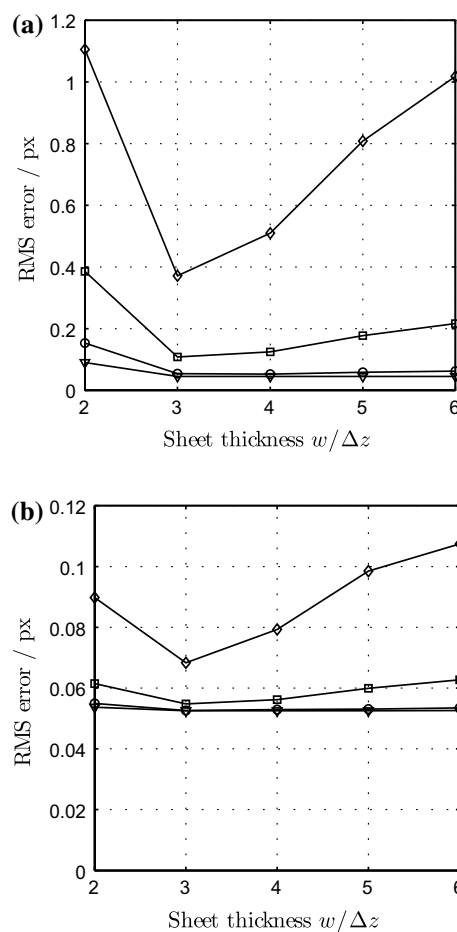


Fig. 5 RMS velocity error in **a** z direction (viewing direction) and **b** y direction, as a function of sheet spacing and thickness for plane-normal camera. Inverted triangle $M\Delta z = 2$ v_x , circle 4 v_x , square 8 v_x , diamond 16 v_x

Only RMS errors in the y and z directions are shown. Unsurprisingly, the best accuracy is obtained at the lowest spacing (highest sampling rate in z). However, error remains comparable using half the number of laser sheets. There is an optimum laser sheet thickness: too small ($w/\Delta z \simeq 2$) and particles become undersampled in depth, effectively introducing peak-locking. Too large ($w/\Delta z \simeq 6$) and the particle image is spread over too large a depth, increasing uncertainty in its position. There exists a happy medium somewhere between $w/\Delta z = 3 - 4$ where error in the z component is minimised. Error in the component orthogonal to the viewing direction is largely unaffected by sheet spacing or thickness. When considering the $M\Delta z = 2$ v_x case, error is in fact marginally higher in y (and x , not shown) compared to the z direction. This is because the resampling of particle images in planes of constant z , which occurs during the reconstruction process, introduces some error in x and y but not in z .




Fig. 6 RMS velocity error in **a** viewing direction and **b** y direction, as a function of sheet spacing and thickness for off-axis camera (1). Inverted triangle $M\Delta z = 2$ vx, circle 4 vx, square 8 vx, diamond 16 vx

The same analysis is repeated for camera 1 which is off-axis, see Fig. 6. The trends are the same, but error in the viewing direction is larger. This is partly because the effective laser sheet spacing (as measured along a ray from the camera pinhole) is larger and scales as $M\Delta z / \cos(\theta_j)$. Since any trends observed for the plane-normal reconstruction are expected to be similarly true for the off-axis reconstruction, the investigation now focuses on a set-up with a camera normal to the laser sheet.

It is preferable to use as large a sheet spacing as possible because, for a fixed reconstruction depth, this minimises the number of sheets per reconstruction. However, the accuracy in the viewing direction depends jointly on the sheet spacing and noise. To simulate this, for $w/\Delta z = 3$ (around the optimal sheet thickness), poisson noise is added to synthetic images before reconstruction. Six noise levels were chosen: the mean number of counts corresponded to 0, 1, 2, 5, 10 and 25 % of the maximum greyscale image value (255 counts). For example, a noise level of 1 % corresponds to poisson noise with a mean of 2.55 counts. This




Fig. 7 Sensitivity of measurement accuracy to additive poisson noise in source images, for camera 2 with $w/\Delta z = 3$ in **a** the z direction and **b** the y direction. Noise levels circle 0 %, square 1 %, diamond 2 %, triangle 5 %, right pointing triangle 10 %, inverted triangle 25 %

qualitatively approximates sensor noise and weakly illuminated particles. However, it fails to represent other potential sources of noise, such as shot-to-shot variation in pulse energy, deviation from a gaussian laser sheet profile or jitter in sheet position.

Figure 7 shows the joint dependence of error in the y and z directions upon the sheet spacing and image noise level. The strongest trend is in the view-normal (z) velocity component. There is a simple trade-off: with fewer sheets and therefore a larger sheet spacing, sensitivity to image noise is greater and the initial error is larger when no noise is added. There is large relative difference in error between the zero-noise case and lowest level of noise: 17 % greater at $M\Delta z = 2$ vx growing to 50 % greater by $M\Delta z = 6$ vx. For a moderate level of image noise (2 % of peak brightness), a sheet spacing below around 5 voxels is necessary to maintain RMS error below 0.1 vx in the z direction.

When noise is present, a weak trend appears in the error in the view-orthogonal (y) velocity component: it is

marginally reduced at larger sheet spacings. This curious effect can be understood thus: the smaller the sheet spacing, noise artefacts of smaller dimension (in depth) can be reconstructed. The closer these artefacts in size to a true particle image, the more they may influence error in the view-orthogonal component through the nonlinear peak fitting operation during PIV cross-correlation, which is (by design) sensitive to small wavelength features of the reconstruction.

3.3 Stereo camera parametric investigation

A more accurate velocimetry may be obtained using two cameras for stereo scanning reconstructions. This allows information about a particle's position in depth (z) to be encoded not only by its variation in illumination during a scan, but also by its triangulation between views. From a tomography perspective, the inclusion of how a true particle's brightness varies as it is recorded provides extra information that enables it to be distinguished from ghost particles.

To identify the optimal sheet thickness and spacing for the stereo case, a similar test was performed for two cameras at $\pm 30^\circ$ and $u_s/u = \infty$. The result is plotted in Fig. 8, which shows that the measurement error is lower than the single-camera case (for equivalent thickness and spacing) and is also less sensitive to the choice of sheet thickness and spacing. As expected, the greatest error lies in the scanning direction, z . There is an optimal $w/\Delta z$ at around 3 where the error in the z direction is similar to both the x and y components and is close to the best achievable by the PIV algorithm.

To test sensitivity to noise, as in Sect. 3.2, varying levels of poisson noise were added to source images before reconstruction. This was performed for $M\Delta z = 2, 3, 4, 5, 6, 8$ vx and $w/\Delta z = 3$. Samples of the reconstructions are shown in Fig. 9 which indicate the algorithm is quite robust to noise: the particle field can be faithfully reconstructed even when the mean noise level reaches 25 % of the dynamic range. Relative to the particle brightness, reconstruction artefacts become increasingly apparent in the background as the noise in the source images increases.

Figure 10 shows the dependence of error in the y and z directions upon the image noise level and sheet spacing. If one contrasts this against Fig. 7, which shows noise sensitivity for the single-camera case, one immediately sees that the stereo technique is more robust to noise. Two separate trends are apparent in the error in y and z components. For the z component, RMS error tends to increase as more image noise is added. There is a significant increase in error beyond $M\Delta z \simeq 8$ vx.

For the y component, when noise is present, the velocimetry actually becomes more accurate when fewer laser

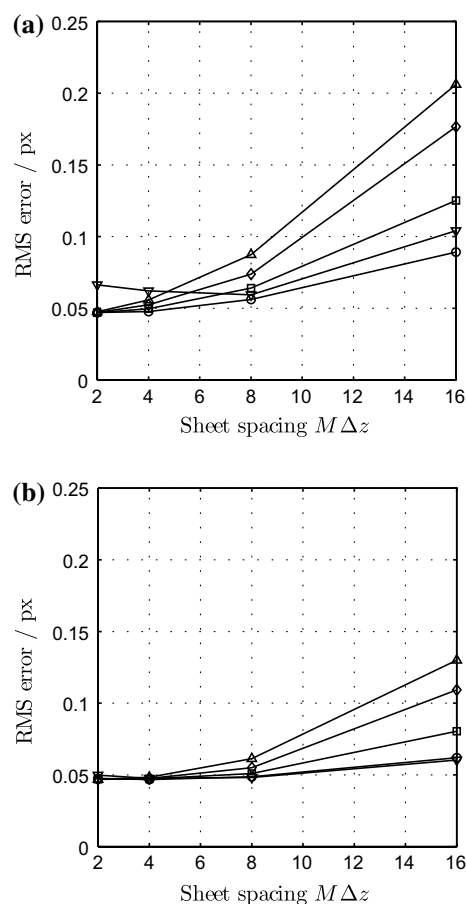


Fig. 8 RMS velocity error in **a** z direction and **b** y direction as a function of sheet spacing and thickness for stereo configuration. Inverted triangle $w/\Delta z = 2$, circle 3, square 4, diamond 5, triangle 6

sheets are used. This trend does not continue indefinitely: accuracy tends to get worse for $M\Delta z \geq 8$ vx. This is similar to the weak trend observed for the single-camera case. Likewise, this is attributed to the presence of reconstruction artefacts, such as those in Fig. 9, which may become smaller and more particle-like at smaller sheet spacings. For a moderate level of image noise (2 %) the RMS error in both y and z components can be maintained below 0.1 voxel with a sheet separation as large as $M\Delta z \simeq 18$ vx.

3.4 The effect of finite sheet speed

Until now, the effect of finite sheet speed has not been addressed. Aside from interfering with the reconstruction of the particle field, finite sheet speed can have two additional consequences for PIV. As the acquisition of an image volume takes a finite amount of time, the reconstruction of a pair of particles separated by some distance δ_3 in the scanning direction will correspond to their positions at different times, separated by an interval $\delta t = \delta_3/u_s$. Thus, if the same particle is recorded in consecutive volumes, during which it

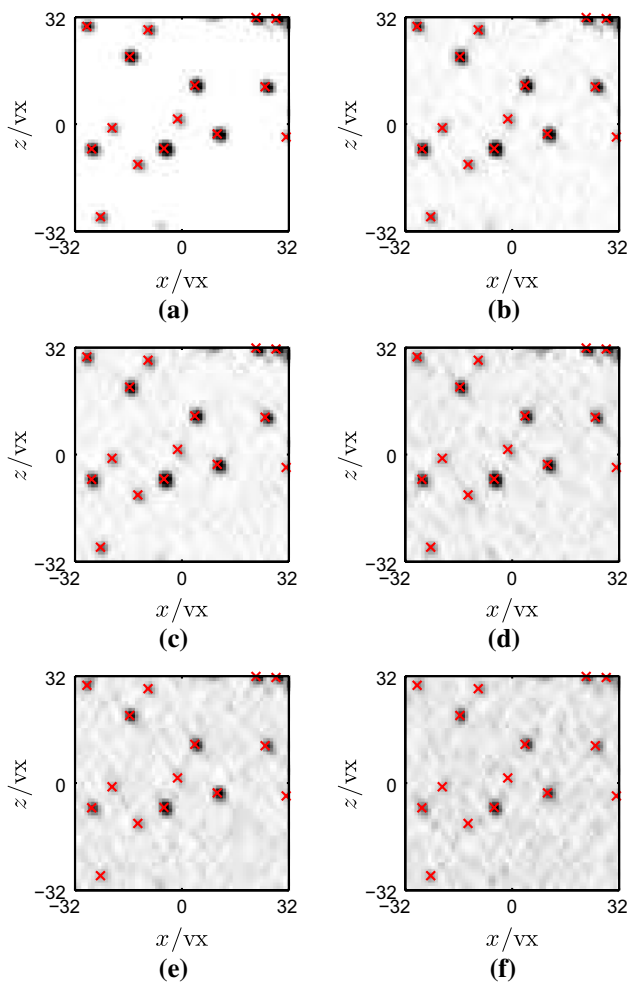


Fig. 9 Single-voxel-thick slices through stereo reconstruction at different levels of additive Poisson noise. Markers indicate true particle positions. **a–f** 0, 1, 2, 5, 10 and 25 % noise

has been observed to undergo a displacement δ_3 , there is an effective time separation $\Delta t + \delta_3/u_s$ between consecutive particle images. To first order, the particle velocity is

$$u_i = \delta_i / (\Delta t + \delta_3/u_s) \tag{7}$$

where δ_i is the particle’s displacement as measured by the PIV algorithm. If one does not make this “delta-t correction”, then there will be an error, which to first order is $u_i \delta_3/u_s$. This effect is referred to as “delta-t dilation”, as the PIV time separation has been effectively extended/shortened by the particle motion across the volume.

Secondly, the aforementioned time gradient in the scan direction means that the velocity measured at one side of the volume corresponds to the motion at a slightly different time compared with the other side of the velocity field. If unsteady effects are large, this may contribute to error. In a worst-case scenario, the time separation between opposite sides of the volume is around the




Fig. 10 RMS velocity error in **a** z direction and **b** y direction as a function of sheet spacing and noise level for stereo configuration with $w/\Delta z = 3$. Noise levels of *circle* 0 %, *square* 1 %, *diamond* 2 %, *triangle* 5 %, *right pointing triangle* 10 %, *inverted triangle* 25 %

size of the PIV time separation, and this unsteady effect would be comparable to error introduced by the finite PIV sampling rate.

These hypothesised effects have been tested by simulating a stereo measurement with finite sheet speed. The configuration was chosen to match the laboratory experiment: $w/\Delta z = 4$, $M\Delta z = 4$ vx and $\theta \pm 30^\circ$. Eight different sheet speeds were tested, corresponding to $u_s/u = 10, 20, 30, 60, 100, 160, 250$ and ∞ . This is equivalent to varying the time taken to scan a complete volume, which is determined by the frame rate and the number of sheets. This varied between $N_S/f_{scan}\tau_k = 0$ and 0.307 (at $u_s/u = 10$). In comparison, the PIV time separation is $0.056\tau_k$. In practice, the time taken to scan the volume would be less than the PIV time separation. This unphysical scenario was simulated to test the effects of very low sheet speeds.

Figure 11 shows the mean error in the z component of velocity u_3 , conditioned on u_3 . When the sheet speed is very fast, the bias error is small. At lower sheet speeds, there is a

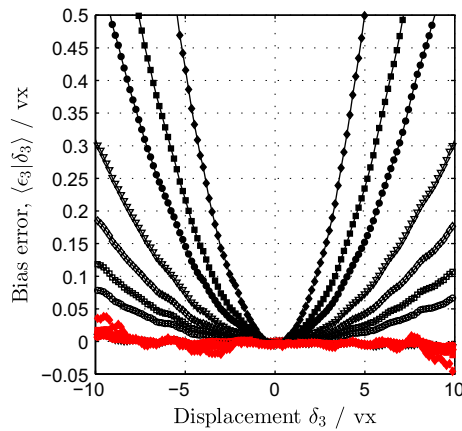


Fig. 11 Conditional mean bias error in z component of velocity, before (black) and after (red) delta-t correction, at different sheet speeds triangle $u_s/u = \infty$, circle 250, square 160, diamond 100, inverted triangle 60, filled circle 30, filled square 20, filled diamond 10




Fig. 12 RMS error in x (circle), y (square) and z (diamond) velocity components before (unfilled, black) and after (filled, black) delta-t correction. Red markers show RMS error after delta-t correction and accounting for flow-dependent unsteady effects

significant bias error due to delta-t dilation which is quadratic in u_3 , a consequence which may be inferred from (7). Red dotted lines show the conditional mean error after the delta-t correction has been applied. It is immediately apparent that the bias error has been removed.

Figure 12 plots the RMS error in each (uncorrected) velocity component at different sheet speeds. As sheet speed is decreased, error increases. On the same figure, the filled markers show the RMS error in each velocity component after the correction for finite sheet speed has been applied. Whilst the correction makes little difference when the sheet travels very quickly, there is a substantial reduction in error at slower sheet speeds. Even at very low sheet speeds ($u_s/u \simeq 25$), the RMS error in the corrected velocity




Fig. 13 RMS error in velocity gradients before (unfilled) and after (filled) delta-t correction circle $\partial u_1/\partial x_1$, square $\partial u_3/\partial x_1$, diamond $\partial u_1/\partial x_3$ and triangle $\partial u_3/\partial x_3$

remains below 0.1 vx. At $u_s/u = 100$ ($N_S/f_{scan} \tau_k = 0.031$), there is virtually no additional error beyond the delta-t dilation.

Figure 13 plots the RMS error in selected velocity gradients at different sheet speeds. The velocity gradients are markedly less sensitive to finite-speed effects than the velocity. This counter-intuitive result can be understood by looking at the error term arising from the delta-t dilation. To first order, the error in $\partial u_i/\partial x_j$ is

$$\frac{\partial \epsilon_i}{\partial x_j} = \frac{\partial}{\partial x_j} \frac{u_i u_3}{u_s} = \frac{\partial u_i}{\partial x_j} \frac{u_3}{u_s} + \frac{u_i}{u_s} \frac{\partial u_3}{\partial x_j} \tag{8}$$

Thus, large errors in velocity gradients require large velocity gradients in regions of high velocity. After the delta-t correction has been applied, the effect upon the RMS error in velocity gradients remains modest down to quite low sheet speeds ($u_s/u = 30$).

The increase in error at lower sheet speed is only partly due to the different time separation. There is also an effect associated with unsteadiness. The simulation has been arranged so that the sheet passes $z = 0$ at $t = \pm \Delta t/2$ in the first and second scans, respectively, i.e. centred around $t = 0$. Thus, vectors at depth z are effectively measured at time $t = z/u_s$: there is a time gradient across the measurement volume. Figure 14 shows the RMS error in the corrected velocity conditioned on the z coordinate. Two cases are presented: a reference ($u_s/u = \infty$) and a case where unsteady effects are significant ($u_s/u = 30$). When the sheet speed is very high, the conditional average only reflects the fact that vectors near the edges of the volume are less reliable. At low sheet speeds, unsteady effects are apparent, as error becomes larger further from the middle of the volume. Near the middle of the volume, where unsteady effects are negligible, error is comparable to the infinite sheet speed case.




Fig. 14 RMS error in each corrected velocity component, conditioned on z coordinate, for a reference case ($u_s/u = \infty, \dots$) and a case where unsteady effects are significant ($u_s/u = 30, -$)

It is clear that the errors reported here, after delta- t correction has been applied, depend upon the degree of flow unsteadiness and the time taken to record the measurement volume. To distinguish the unsteady effect from other sources of error, the RMS error has been re-calculated by comparing the measurement $\hat{u}_i(\mathbf{x}, t)$ to the ground truth $u_i(\mathbf{x}, t + z/u_s)$ after the time gradient has been compensated for. This is shown in red in Fig. 12. This residual error has a very weak dependence upon sheet speed. Thus, the increased error at low sheet speed (after delta- t correction) is due to the time gradient across the measurement volume. More generally, the error associated with unsteadiness will depend upon the flow in question, but the error associated with remaining factors is expected to scale independently (i.e. with the other parameters identified in this investigation).

3.5 Self-calibration

Just as in tomographic PIV, a small misalignment between cameras can cause large errors in reconstruction. Figure 15 shows a 9-vx thick slice through part of a simulated stereo reconstruction made with and without a small misalignment of the camera set-up. Markers indicate true particle positions. For the aligned reconstruction, particles are reconstructed at their correct position and there are only a few, weak ghost artefacts. In contrast, the misaligned reconstruction contains many ghost particles and exhibits poor recovery of the true particle field even though the misalignment is quite small (mean disparity between stereo cameras was 4.56 and 2.13 vx in x and y , respectively). This highlights the need for self-calibration: slight misalignment in even the most carefully conducted of experiments can lead to grossly erroneous reconstructions.




Fig. 15 A sample of a scanning stereo reconstruction **a** with and **b** without a small misalignment of the set-up

A typical scanning stereo set-up was simulated in the absence of flow ($\theta = \pm 30^\circ, M\Delta z = 4\text{ vx}, w/\Delta z = 4, u_s/u = \infty$). If the reconstruction set-up is known perfectly, then the disparity should be zero. To measure the disparity, a 2D cross-correlation was applied in planes of constant z in a manner similar to Brückner et al. (2013). Whilst the mean disparity was found to be close to zero, the RMS disparity is 0.064 and 0.026 vx in x and y , respectively. In this case, there is no disparity error due to misalignment and the fluctuating error is due to PIV uncertainty caused by elongation of particle shape in the viewing direction of either camera. In this instance, re-sliced particle images are elongated in the x direction, and this is the reason for the larger fluctuating disparity error in x compared to y .

966 Monte Carlo trials were conducted where a small perturbation was applied to the camera set-up and the resulting disparity between single-camera back-projections was measured. Excellent correlation was found between the measured disparity and that predicted by the self-calibration model, with an R^2 correlation coefficient of 1.000 in both x and y (to 4 s.f.). The RMS error between measurement and prediction is 0.070 and 0.029 vx in x and y , respectively, which is close to the level of the PIV uncertainty quantified by the zero-disparity case.

Whilst the self-calibration model has been demonstrated to correctly predict the disparity error for a known misalignment, the objective of self-calibration is to infer the misalignment given the disparity and subsequently correct for it. To check the validity of the disparity correction, one can calculate the residual disparity vectors after correction has been made to the back-projection model. The RMS residual is 0.057 and 0.016 vx in x and y . The residual disparity error has actually been reduced below the level of the unperturbed case. This is a subtle consequence of the self-calibration “correcting” an error introduced by cross-correlating reconstructions from different cameras.

Successful minimisation of disparity is no fluke since the unperturbed back-projection model can be recovered with some confidence. Before disparity correction, the RMS error in laser sheet position is 7.37 vx , the RMS error

in sheet orientation is 1.16 mrad, and in (absolute) camera orientation, it is 1.14 mrad. After disparity correction, the RMS error in laser sheet position is 0.15 vx, its orientation 0.54 mrad, and in camera orientation 0.54 mrad. However, if one considers the relative orientation of cameras, one finds that the disparity correction reduces the RMS error from 1.58 to 0.05 mrad. For this configuration of cameras, it turns out to be quite difficult to identify a joint tilt of both cameras about the x axis using disparity vectors alone. Indeed, if one evaluates the PDFs of alignment cosines (not shown), they show that error in absolute camera orientation is strongly aligned with the x direction. Thus, disparity correction can reliably detect relative rotation between cameras, but not the absolute rotation of both. On the other hand, sheet position and orientation errors can be corrected for quite well.

3.6 Large mixing tank experiment

An example of a typical reconstruction from the laboratory experiment after intensity normalisation is shown in Fig. 16, which has been truncated to 48^3 vx (the size of an interrogation window). Clearly the seeding density is adequate and that particle images are well defined, with few artefacts from reconstruction. The seeding concentration was measured to be around 25 particles/ 32 vx^3 , which is quite high. To achieve this seeding concentration, a tomographic PIV measurement spanning the same size volume (without scanning) would have a source density of 0.9,




Fig. 16 An example of an experimental stereo scanning PIV reconstruction, showing a $48 \times 48 \times 48$ vx subvolume from different angles

which would be far too high to yield valid reconstructions (Westerweel et al. 2013).

Making accurate and well-resolved measurements of velocity gradients of fine-scale turbulence is fraught with difficulties, so it is very important to be able perform checks on the data. Buxton et al. (2011) studied the effects of spatial resolution and noise on the measurement of fine-scale statistics of turbulent flows using DNS and experimental data. They highlighted a number of measurements related to velocity gradients which have well-known behaviours in turbulent flows, which are also sensitive to measurement noise and spatial resolution. As there is no simultaneous, superior reference measurement available for the laboratory experiment, the tests they suggest are an extremely valuable tool in the assessment of measurement accuracy.

One of the simplest measures to test the accuracy of velocity gradients is to test the divergence-free condition ($\partial U_i / \partial x_i = 0$) of the incompressible flow. One may calculate the joint PDF of quantities in the divergence sum, e.g. $\partial U_1 / \partial x_1$ and $-(\partial U_2 / \partial x_2 + \partial U_3 / \partial x_3)$. In an ideal measurement of an incompressible flow, the correlation coefficient between these terms is unity. However, truncation errors associated with applying a finite difference method to obtain gradients would reduce this figure below unity even in an ideal measurement. Figure 17a shows logarithmically spaced contours of probability density showing that large departures from incompressibility being rare. The correlation coefficient Q is 0.978, which compares favourably to Ganapathisubramani et al. (2007) (0.82, using stereo PIV and Taylor's hypothesis) and Tsinober et al. (1992) (0.7, using multi-probe hot-wires). Without the Lagrangian filter (Fig. 17b), the data are slightly noisier, but the correlation coefficient is still very good (0.943). For a comparison to other experimental data sets concerning turbulent flows, see Ganapathisubramani et al. (2007) and Casey et al. (2013).

Zhang et al. (1997) introduced a dimensionless quantity ξ that can quantify the level of error in the velocity gradients (9). $\langle \xi \rangle = 0$ implies mass continuity is always satisfied and $\langle \xi \rangle = 1$ if the velocity gradients in the sum are independent (no continuity constraint). Over 1,003 independent turbulent velocity fields, $\langle \xi \rangle = 0.156$ before Lagrangian filtering and $\langle \xi \rangle = 0.045$ after. It is worth noting that the PIV algorithm used incorporates no information about mass continuity. This compares favourably to Mullin and Dahm (2006), Ganapathisubramani et al. (2007), Casey et al. (2013) and Zhang et al. (1997), who report values of $\langle \xi \rangle$ of 0.12, 0.18, 0.36, and 0.5, respectively. However, $\langle \xi \rangle$ cannot be used as a marker of accuracy alone as Zhang et al. (1997) show that one can improve the reported value substantially by low-pass spatial filtering, at the sake of resolution. The rest of this section will present data after




Fig. 17 Joint PDF of $\partial U_3/\partial x_3$ and $-(\partial U_1/\partial x_1 + \partial U_2/\partial x_2)$ **a** with and **b** without Lagrangian filtering. Contours of probability density are logarithmically spaced from 10^{-3} to 10^1

Table 2 selected ratios of second moments of velocity gradients

K_1	K_2	K_3	K_4	H_1
$2 \frac{(\partial_2 u_2^2)^{1/2}}{(\partial_1 u_1^2)^{1/2}}$	$2 \frac{(\partial_2 u_2^2)^{1/2}}{(\partial_3 u_3^2)^{1/2}}$	$2 \frac{(\partial_2 u_2^2)^{1/2}}{(\partial_1 u_1^2)^{1/2}}$	$2 \frac{(\partial_2 u_2^2)^{1/2}}{(\partial_3 u_3^2)^{1/2}}$	$2 \frac{\langle S_{ij} S_{ij} \rangle}{\langle \omega_i \omega_i \rangle}$
1.47	1.48	1.43	1.46	1.01

Lagrangian filtering has been applied, with the intent of demonstrating both the high fidelity and spatial resolution of this data set.

$$\xi = \frac{(\partial U_i/\partial x_i)^2}{(\partial U_1/\partial x_1)^2 + (\partial U_2/\partial x_2)^2 + (\partial U_3/\partial x_3)^2} \quad (9)$$

At this juncture, it is worth noting that for this flow, the anisotropy presents at the large scale persists down to the fine scales. Consider Table 2, which lists a select number of ratios of moments of the VGT. For homogeneity, a prerequisite of axisymmetry or isotropy, $\langle \omega_i \omega_i \rangle = 2 \langle S_{ij} S_{ij} \rangle$. For isotropy, the ratios K_i should all be equal to unity, but

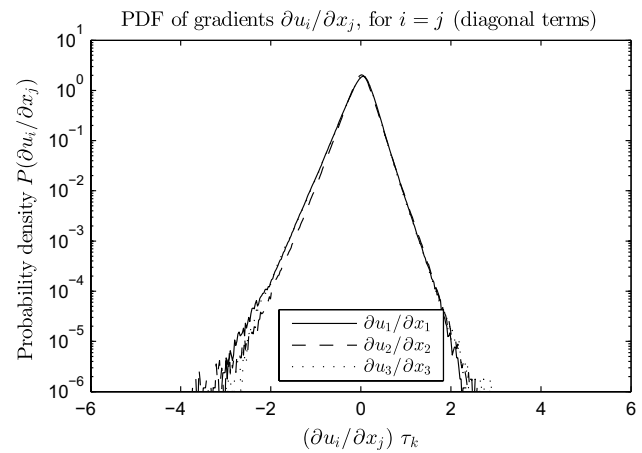


Fig. 18 PDF of longitudinal velocity gradients ($\partial u_i/\partial x_j$ for $i = j$)

under axisymmetry about the 2 direction, only $K_1 = K_2$ and $K_3 = K_4$ with no restriction on their values (George and Hussein 1991). The measurement is therefore consistent with homogeneous, axisymmetric turbulence, but not isotropy. One must therefore be careful about what checks are appropriate for the data on the basis of turbulence statistics, which are often derived on the basis of isotropy. Further checks on axisymmetry are presented later in this section.

Figure 18 shows PDFs of the longitudinal velocity gradients (on-diagonal components of the VGT), which shows good qualitative agreement with that of homogeneous isotropic turbulence at similar R_λ (see Ishihara et al. 2007, Fig. 5 or Buxton et al. 2011, Figs. 1, 2 for a comparison). The PDFs exhibit the well-known negative skewness and wide, non-gaussian tails, which Buxton et al. (2011) note are features that are particularly sensitive to spatial resolution, with skewness S decreasing with decreasing spatial resolution. For this flow, $S = -0.54, -0.29$ and -0.57 for velocity gradients in the 1, 2 and 3 directions which (despite the obvious anisotropy) is still in approximate agreement with the power law trend identified by Ishihara et al. (2007) for isotropic turbulence, which predicts $S \simeq -0.62$. The collapse is quite good for the positive tails (given the normalisation), and it is apparent that the negative tail is the reason for the less negative skewness in the two directions. Thus, departure from isotropy is likely the explanation for the more positive skewness of $\partial u_2/\partial x_2$, rather than a lack of spatial resolution. Furthermore, the PDFs of longitudinal gradients are consistent with axisymmetry.

Figure 19 shows PDFs of the transverse velocity gradients (off-diagonal components of the VGT), which do not exhibit the same degree of skewness as the longitudinal gradients. The tails of the PDF of the transverse velocity gradients exhibit the non-exponential, fat tails more strongly than the longitudinal gradients in agreement with

Ishihara et al. (2007). Buxton et al. (2011) noted that the fat tails were a feature that disappeared as resolution coarsened and are almost completely lost by $\Delta x = 5\eta_k$. Axisymmetry requires that $\partial u_2/\partial x_1$ and $\partial u_2/\partial x_3$, $\partial u_1/\partial x_3$ and $\partial u_3/\partial x_1$, and $\partial u_1/\partial x_2$ and $\partial u_3/\partial x_2$ exhibit the same collapse, which approximately holds.

The joint PDF of invariants $Q_A = -A_{ij}A_{ji}/2$ and $R_A = -\det(A_{ij})$ of the VGT has been shown to take a characteristic tearing drop shape in a wide variety of turbulent flows (Ooi et al. 1999; Soria et al. 1994; Chacin and Cantwell 2000; Andreopoulos and Honkan 2001; Elsinga and Marusic 2010) and has been hypothesised to be universal features of turbulence (Tsinober 2009; Elsinga and Marusic 2010). Whilst the shape is relatively insensitive to spatial resolution when normalised by small-scale

quantities, it is very sensitive to noise in velocity gradient measurements, as Q_A and R_A are second- and third-order products of velocity gradients, respectively. Figure 20a shows the joint PDF for the present experimental data, whilst Fig. 20b provides a comparison with data taken from the JHU turbulence database. Despite the differences in Reynolds number, very good agreement is observed with DNS. The JPDF takes the characteristic tearing drop shape, exhibiting sharp changes in the angle of contours around the line of zero discriminant (denoted by the black, tent-like curve), which tends to become more rounded at higher levels of noise. The PDF is strongly skewed towards the so-called Vieillefosse tail which is another feature that disappears in the presence of noise.

Finally, as a check on the spatial resolution, one can compare the mean dissipation to that predicted by Kolmogorov's second similarity hypothesis using measurements of the modified third-order longitudinal velocity structure function, following Nie and Tanveer (1999). They show that the well-known Kolmogorov 4/5th Law can be derived for the more general case where isotropy does not hold. They define a modified third-order structure function $\tilde{S}_3(r)$ which reaches an asymptotic value of $-\frac{4}{3}\langle\epsilon\rangle r$ in the inertial regime $\eta_k \ll r \ll L_{\text{turb}}$. Their modified third-order structure function is defined

$$\tilde{S}_3(r) = \frac{1}{4\pi} \int \langle (\delta u_i \delta u_i)^{3/2} \cos(\theta_{\delta \mathbf{u}, \mathbf{r}}) \rangle d\Omega_{\mathbf{r}} \quad (10)$$

where $r = (r_j r_j)^{1/2}$ and the integration represents taking an average over all solid angle $\Omega_{\mathbf{r}}$. The angle $\theta_{\delta \mathbf{u}, \mathbf{r}}$ is the angle between the velocity increment $\delta u_i = u_i(x_j + r_j) - u_i(x_j)$ and the displacement r_j . The integrand in (10) is the

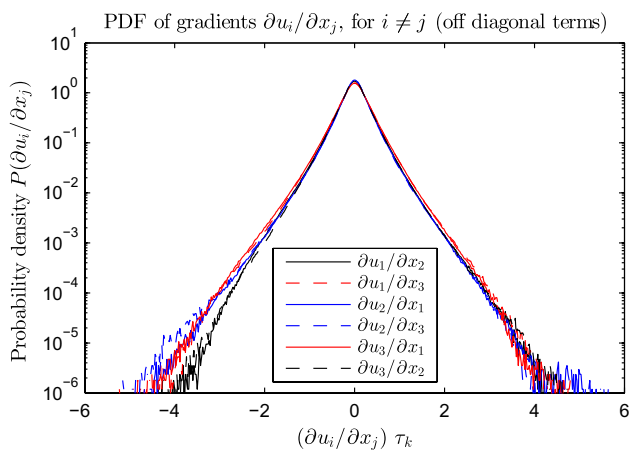


Fig. 19 PDF of transverse velocity gradients ($\partial u_i/\partial x_j$ for $i \neq j$)




Fig. 20 Joint PDFs of invariants Q_A and R_A of the VGT for a the laboratory experiment and b DNS of homogeneous isotropic turbulence at $R_\lambda \simeq 443$ (Li et al. 2008)




Fig. 21 Modified third-order structure function, taken in x (dot), y (times) and z (circle) direction and also solid angle averaged form (square)

third-order longitudinal structure function (11) without solid angle averaging and is an axisymmetric function of r in this case. When pre-scaled by $-\frac{4}{3}\langle\epsilon\rangle r$, it becomes a dimensionless measure of the average kinetic energy flux across the boundary of a spherical control volume, of radius r , in a frame of reference moving with the fluid.

$$\tilde{S}_3^c(r_j) = \langle (\delta u_i \delta u_i)^{3/2} \cos(\theta_{\delta \mathbf{u}, \mathbf{r}}) \rangle \quad (11)$$

This third-order structure function is plotted in Fig. 21 without solid angle averaging in the three principal directions, as well as the solid angle average which was obtained with recourse to axisymmetry. The structure functions have been scaled by $-\frac{4}{3}\langle\epsilon\rangle r$ where $\langle\epsilon\rangle = 2\nu\langle S_{ij}S_{ij}\rangle = 0.833\text{ mm}^2\text{ s}^{-3}$ comes directly from velocity gradient measurements. There is excellent agreement between $\tilde{S}_3(r)$ and its asymptotic form in the inertial range (within 3 %), indicating that the dissipation has been well resolved. The structure function in the 1 and 3 directions is in agreement and is substantially different to the two directions. This is consistent with axisymmetry, but not isotropy. That they are so markedly different is not necessarily erroneous: it indicates that eddies in the inertial range are largely fed kinetic energy by fluctuations in velocity orthogonal to the symmetry axis. Velocity fluctuations parallel to the symmetry axis exhibit a much weaker net flux of kinetic energy.

4 Conclusion

This paper has introduced, tested and applied a hybrid variant of scanning PIV, which combines aspects of tomographic PIV and conventional scanning PIV. The objective has been to develop a technique applicable to big, slow,

turbulent flows that can yield accurate measurements of the full VGT in 3D with good space and time resolution. This requires a technique that can handle high seeding densities, in order to capture the contribution of the smallest scales of motion to the velocity gradients.

Velocimetries are made by cross-correlating volumetric reconstructions of the particle field obtained from particle images corresponding to different depths through a measurement volume. Only one or two cameras are necessary to make reconstructions, which can be completed in just a few minutes on a high-end desktop computer. The details of the reconstruction method have been exposted, and a self-calibration scheme has been introduced to correct misalignment when using the stereo reconstruction method.

A parametric investigation using synthetic particle images and velocity fields from numerical simulation of isotropic turbulence has been conducted. The effect of six dimensionless parameters on velocimetry accuracy has been tested: the number of cameras and their orientation, the sheet spacing, sheet thickness, sheet speed and image noise level.

The single-camera method produces less accurate velocimetries and error is anisotropic: it is largest in the viewing direction. The most accurate single-camera velocimetries are made when the camera is oriented normal to the laser sheet. The stereo method is able to extract depth information about particles as encoded by their brightness and their triangulation between views. As such, it is more accurate and error is less anisotropic, although remains largest in the scanning direction.

For both single-camera and stereo configurations, there is an optimal laser sheet thickness of $w/\Delta z \simeq 3$. For the single-camera configuration, there is a trade-off to be made between accuracy and the sheet separation used in the presence of image noise. Using a smaller sheet separation (more laser sheets) tends to improve accuracy for a given level of image noise. The stereo configuration is less sensitive to image noise than the single-camera configuration. At a moderate level of image noise (2 % of the image's dynamic range) and in the absence of finite sheet speed effects, the single-camera configuration achieved an RMS error in the scanning direction below 0.1 vx for sheet separations $M\Delta z$ up to 5 voxels. For the stereo configuration, the same level of error was achieved with a sheet spacing as large as 18 voxels.

The time taken to record an acquisition in comparison with the PIV timescale is characterised by the dimensionless sheet speed u_s/u . Two sources of error are associated with finite sheet speed: one is due to an effective modulation of the PIV time separation ("delta-t dilation") and the other is associated with flow unsteadiness. The former can be readily corrected for using information available from the measurement itself and can make a substantial reduction in error at low sheet speed. For this study, after making the "delta-t

correction”, the effect of finite sheet speed was found to be negligible for $u_s/u \geq 100$ and RMS errors below 0.1 voxel were achieved for sheet speeds as low as $u_s/u = 30$.

The error associated with flow unsteadiness is more difficult to correct for and its size will naturally depend upon the unsteady behaviour of the flow in question. After accounting for this unsteady effect using information available from the numerical simulation, and applying the delta-t correction, the error has only a weak dependence upon sheet speed over the range tested. The implication is that surprisingly low sheet speeds can be used, provided error associated with flow unsteadiness can be tolerated.

The self-calibration scheme for stereo reconstruction has been validated using a Monte Carlo approach to apply small perturbations to the camera set-up and subsequently test how well the perturbation can be corrected for. The measured disparity vectors, after applying each perturbation, agree with the self-calibration model to within the PIV uncertainty. A least-squares fitting operation of the self-calibration model to the disparity vectors has been shown to correct the disparity to within the PIV uncertainty. This is no fluke: the exact details of the applied perturbation can be recovered with good confidence.

Stereo scanning PIV has been applied to a laboratory investigation of homogeneous turbulence in a von Kármán swirling water flow at $R_\lambda \simeq 219$. Statistics of velocity gradients and two-point correlation functions indicate a turbulence consistent with homogeneity and local axisymmetry, rather than local isotropy.

A number of tests upon accuracy and spatial resolution of velocity gradients have been made. Errors in velocity gradients, as tested through mass continuity, are substantially lower than those reported by previous authors ($\langle \xi \rangle = 0.045$). The PDFs of longitudinal and transverse velocity gradients exhibit the well-known fat, and non-gaussian tails and the longitudinal velocity gradients exhibit levels of skewness approximately consistent with isotropic turbulence at similar Reynolds number. The well-known tearing drop shape of the joint PDF of invariants Q_A and R_A has been reproduced and is in good agreement with data from isotropic turbulence. Using measurements of the modified third-order structure function, the presence of an inertial range has been demonstrated, with a level of dissipation consistent with that directly measured from velocity gradients to within 3 %. These results confirm that stereo scanning PIV is capable of making accurate and spatially resolved measurements of turbulent flows.

References

- Andreopoulos Y, Honkan A (2001) An experimental study of the dissipative and vortical motion in turbulent boundary layers. *J Fluid Mech* 439:131–163
- Astarita T (2007) Analysis of weighting windows for image deformation methods in PIV. *Exp Fluids* 43(6):859–872
- Bodenschatz E, Malinowski SP, Shaw RA, Stratmann F (2010) Can we understand clouds without turbulence? *Science* 327(5968):970–971
- Brücker C (1995) Digital-particle-image-velocimetry (DPIV) in a scanning light-sheet: 3D starting flow around a short cylinder. *Exp Fluids* 19:255–263
- Brücker C (1996) 3-D scanning-particle-image-velocimetry: technique and application to a spherical cap wake flow. *Appl Sci Res* 56:157–179
- Brücker C (1997) 3D scanning PIV applied to an air flow in a motored engine using digital high-speed video. *Meas Sci Technol* 8(12):1480–1492
- Brücker C, Hess D, J K (2013) Single-view volumetric PIV via high-resolution scanning, isotropic voxel restructuring and 3D least-squares matching (3D-LSM). *Meas Sci Technol* 24(2):024001
- Buxton ORH, Laizet S, Ganapathisubramani B (2011) The effects of resolution and noise on kinematic features of fine-scale turbulence. *Exp Fluids* 51(5):1417–1437
- Cameron SM (2011) PIV algorithms for open-channel turbulence research: accuracy, resolution and limitations. *J Hydroenviron Res* 5(4):247–262
- Casey TA, Sakakibara J, Thoroddsen ST (2013) Scanning tomographic particle image velocimetry applied to a turbulent jet. *Phys Fluids* (1994-present) 25(2):025102
- Chacin JM, Cantwell BJ (2000) Dynamics of a low Reynolds number turbulent boundary layer. *J Fluid Mech* 404:87–115
- David L, Jardin T, Braud P, Farcy A (2012) Time-resolved scanning tomography PIV measurements around a flapping wing. *Exp Fluids* 52(4):857–864
- Diez FJ, Cheng Y, Villegas A (2011) Time resolved visualization of structures of velocity gradients measured with near Kolmogorov-scale resolution in turbulent free-shear flows. *Exp Therm Fluid Sci* 35(6):1223–1229
- Douady S, Couder Y, Brachet ME (1991) Direct observation of the intermittency of intense vorticity filaments in turbulence. *Phys Rev Lett* 67:983–986
- Efron B, Tibshirani R (1994) An introduction to the bootstrap. Chapman & Hall, New York
- Elsinga GE, Marusic I (2010) Universal aspects of small-scale motions in turbulence. *J Fluid Mech* 662:514–539
- Elsinga GE, Scarano F, Wieneke B, van Oudheusden BW (2006) Tomographic particle image velocimetry. *Exp Fluids* 41(6):933–947
- 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(4):825–838
- Ganapathisubramani B, Lakshminarasimhan K, Clemens NT (2007) Determination of complete velocity gradient tensor by using cinematographic stereoscopic PIV in a turbulent jet. *Exp Fluids* 42(6):923–939
- George WK, Hussein HJ (1991) Locally axisymmetric turbulence. *J Fluid Mech* 233:1–23
- Hartley R (2003) Multiple view geometry in computer vision. Cambridge University Press, Cambridge, UK
- Herman GT, Lent A (1976) Iterative reconstruction algorithms. *Comput Biol Med* 6(4):273–294
- Hori T, Sakakibara J (2004) High-speed scanning stereoscopic PIV for 3D vorticity measurement in liquids. *Meas Sci Technol* 15(6):1067–1078
- Ishihara T, Kaneda Y, Yokokawa M, Itakura K, Uno A (2007) Small-scale statistics in high-resolution direct numerical simulation of turbulence: Reynolds number dependence of one-point velocity gradient statistics. *J Fluid Mech* 592:335–366
- Lecordier B, Westerweel J (2004) The EUROPIV synthetic image generator (S.I.G.). In: Stanislas M, Westerweel J, Kompenhans J

- (eds) Particle image velocimetry: recent improvements, pp 145–161, EUROPIV 2 Workshop, Zaragoza, Spain, Mar 31–Apr 01, 2003
- Li Y, Perlman E, Wan M, Yang Y, Meneveau C, Burns R, Chen S, Szalay A, Eyink G (2008) A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence. *J Turbul* 9:1–29
- Meunier P, Leweke T (2003) Analysis and treatment of errors due to high velocity gradients in particle image velocimetry. *Exp Fluids* 35(5):408–421
- Mullin JA, Dahm WJA (2006) Dual-plane stereo particle image velocimetry measurements of velocity gradient tensor fields in turbulent shear flow. I. Accuracy assessments. *Phys Fluids* 18(3):035101
- Nie Q, Tanveer S (1999) A note on third order structure functions in turbulence. *Proc R Soc Lond Ser A Math Phys Eng Sci* 455(1985):1615–1635
- Nobach H (2011) Influence of individual variations of particle image intensities on high-resolution PIV. *Exp Fluids* 50(4):919–927
- Novara M, Scarano F (2013) A particle-tracking approach for accurate material derivative measurements with tomographic PIV. *Exp Fluids* 54(8):1584
- Ooi A, Martin J, Soria J, Chong MS (1999) A study of the evolution and characteristics of the invariants of the velocity-gradient tensor in isotropic turbulence. *J Fluid Mech* 381:141–174
- Ponitz B, Satsuba M, Brücker C, Kitzhofer J (2012) Volumetric velocimetry via scanning back-projection and least-squares-matching algorithms of a vortex ring. In: Proceedings of the 16th international symposium on applications of laser techniques to fluid mechanics
- Raffel M (2007) Particle image velocimetry: a practical guide. Springer, New York
- Scarano F (2013) Tomographic PIV: principles and practice. *Meas Sci Technol* 24(1):012001
- Schröder A, Geisler R, Elsinga G, Scarano F, Dierksheide U (2008) Investigation of a turbulent spot and a tripped turbulent boundary layer flow using time-resolved tomographic PIV. *Exp Fluids* 44(2):305–316
- Soodt T, Schröder F, Klaas M, van Overbrüggen T, Schröder W (2011) Experimental investigation of the transitional bronchial velocity distribution using stereo scanning PIV. *Experiments in Fluids* 1–10
- Soria J, Sondergaard R, Cantwell BJ, Chong MS, Perry AE (1994) A study of the fine-scale motions of incompressible time-developing mixing layers. *Phys Fluids* 6:871–884
- Sreenivasan KR (2004) Possible effects of small-scale intermittency in turbulent reacting flows. *Flow Turbul Combust* 72(2–4):115–131
- Tsinober A (2009) An informal conceptual introduction to turbulence. Springer, Dordrecht
- Tsinober A, Kit E, Dracos T (1992) Experimental investigation of the field of velocity gradients in turbulent flows. *J Fluid Mech* 242:169–192
- Wallace JM (2009) Twenty years of experimental and direct numerical simulation access to the velocity gradient tensor: what have we learned about turbulence? *Phys Fluids* 21(2):021301
- Westerweel J, Elsinga GE, Adrian RJ (2013) Particle image velocimetry for complex and turbulent flows. *Annu Rev Fluid Mech* 45(1):409–436
- Wieneke B (2005) Stereo-PIV using self-calibration on particle images. *Exp Fluids* 39(2):267–280
- Worth N (2010) Tomographic-PIV measurement of coherent dissipation scale structures. PhD thesis, University of Cambridge
- Worth NA, Nickels TB (2008) Acceleration of Tomo-PIV by estimating the initial volume intensity distribution. *Exp Fluids* 45(5):847–856
- Worth NA, Nickels TB (2011) Time-resolved volumetric measurement of fine-scale coherent structures in turbulence. *Phys Rev E* 84:025301
- Zhang J, Tao B, Katz J (1997) Turbulent flow measurement in a square duct with hybrid holographic PIV. *Exp Fluids* 23:373–381