Noname manuscript No.
(will be inserted by the editor)

A scanning particle tracking velocimetry technique for high-Reynolds number turbulent flows

Melissa Kozul · Vipin Koothur · Nicholas A. Worth · James R. Dawson

Received: date / Accepted: date

Abstract We propose a novel robust 3D particle tracking technique based on a scanning laser setup. The method yields Lagrangian statistics in densely-seeded turbulent flows with good spatial and temporal resolution, overcoming some of the inherent difficulty with line-of-sight based volumetric methods. To do this we have developed an effective triangulation method greatly reducing ghost particle reconstruction using images from two cameras. A laser sheet is rapidly traversed ('scanned') across a measurement volume illuminating only a thin slice of the flow at a time. Particle images are taken at closely-spaced, overlapping nominal laser sheet locations giving multiple intensity recordings for each individual particle. The laser-sheet intensity varies as a Gaussian across its thickness, which is here exploited to deduce the particle's probable location along the scan direction to sub-sheet number resolution by fitting a similarly-Gaussian profile to the particle's multiple intensity recordings. The method is presently verified via numerical experiment using a DNS database. Following successful reconstruction of a time series of 3D particle fields, particle tracks are formed from which all components of Lagrangian velocity and acceleration are calculated.

Keywords Particle tracking · Scanning PIV · Lagrangian statistics

M. Kozul · V. Koothur · N. A. Worth · J. R. Dawson
Department of Energy and Process Engineering, Norwegian University of Science and Technology, 7491 Trondheim, Norway
E-mail: melissa.kozul@ntnu.no

1 Introduction

Lagrangian particle tracking (LPT) seeks to uncover the physics of a particular fluid flow by following individual particle tracers over a sequence of time-resolved recordings. The method is usually based on three steps: the identification of illuminated tracer particles from multiple camera views, a triangulation step in which the probable 3D location of the particles is calculated, and the linking of subsequent particle locations to form probable particle trajectories. Commonly implemented as a particle tracking velocimetry (PTV), it is a well-established measurement technique used for the study of 3D Lagrangian particle motion in turbulent flows (Maas et al, 1993; Virant and Dracos, 1997). The goal of establishing such particle tracks is often to calculate Lagrangian velocities (Nishino et al, 1989) and accelerations (Malik et al, 1993).

In order to reconstruct particle trajectories, it must be possible to both accurately triangulate the locations of tracer particles, and then unambiguously link particles to form trajectories. At high particle seeding densities, required for the spatial resolution of the fine scales present in high-Reynolds number flows, both of these operations become challenging. The number of potential particle image matches increases non-linearly with seeding density, resulting in ambiguity during the triangulation procedure giving erroneous particle locations. Difficulties also arise in correctly pairing the same particle with its appearance in subsequent images given many nearest-neighbour candidates. There is therefore an inherent tension between the conditions for accurate Lagrangian particle tracking (LPT) and the high particle densities required to capture the fine scales of high-Reynolds number flow. Implementations of the technique exploiting three-dimensional particle track-

ing velocimetry (3D-PTV) using three (e.g. Maas et al, 1993) or four (e.g. Lüthi et al, 2005) cameras helps resolve ambiguities to an extent, however these 3D-PTV methods are typically limited to relatively low particle image densities N_I of ≈ 0.005 particles per pixel (ppp) (cf. Maas et al, 1993) in order to maintain confidence in the deduced particle tracks.

Recent particle tracking implementations seek to overcome these restrictions, focusing either on accurate tracking in densely-seeded flows, or on the reconstruction of the Eulerian field from sparse Lagrangian data. In the latter category a technique described as 'pouring time into space' was recently introduced by Schneiders and Scarano (2016), making use of both the instantaneous velocity and the velocity material derivative to improve the consistency of the reconstructed instantaneous velocity fields within the framework of sparse tracks. Another approach is the so-called 'FlowFit' method (Gesemann et al, 2016), which employs a system of smooth B-splines and invokes physical constraints during velocity, acceleration and pressure field reconstruction. The method is thus able to increase the spatial and temporal resolution by 'supersampling' the starting scattered data, additionally reducing noise during the process. On the other hand several novel approaches have emerged recently to track particles successfully at higher tracer particle densities. Notable is the 'Shake-The-Box' method (Schanz et al, 2016) extending the Iterative Particle Reconstruction (IPR) methodology of Wieneke (2013). Initialised individual tracks are prolonged to the next time step via extrapolation, where the predicted particle locations are corrected by 'shaking' them in space and image matching in time. Since ghost particles do not typically generate consistent tracks persisting over several successive recordings, essentially ghost-free, and very accurate, particle tracks are generated for particle image densities of up to ≈ 0.125 ppp. Temporal information is thus exploited in addition to the purely spatial information offered by conventional tomographic PIV (Elsinga et al, 2006). Attempts have also been made recently to find 3D particle locations directly from the reconstructed intensity volumes from tomographic PIV, for example the Tomo-3D-PTV of Novara and Scarano (2013), although this approach is susceptible to errors when particle tracking is undertaken since many ghost particles are reconstructed at high seeding densities. The concept of particle prediction as introduced in the 'Shake-The-Box' method (Schanz et al, 2016) is extended into the domain of tomographic PIV by Lynch and Scarano (2015) with an efficient 'motion-tracking-enhanced' intensity field reconstruction. The number of ghost particles is thus significantly reduced, and combined with a suitable tracking scheme, such an approach could increase the particle density at which Tomo-3D-PTV could be applied.

The present LPT methodology is shown to be effective and accurate up to an equivalent particle image density of $N_I \approx 0.125$ ppp, using only two high-speed cameras to record particle images and a scanning laser sheet to illuminate tracer particles. This is similar to the upper particle image density limit tackled by the four-camera setup of Schanz et al (2016), considering this method to be the standard against which comparison is relevant, and builds upon the previous use of scanning PTV at a lower Reynolds number (Hoyer et al, 2005). The present scanning technique thus permits us to extend the seeding density at which a 'classic' LPT methodology such as that of Maas et al (1993), where sequential particle fields are triangulated following which particles are linked to form tracks, can be used. 3D particle field reconstruction is completed for each individual sheet location along the scan, permitting higher particle densities and more accurate triangulation by restricting individual particle locations to the thin illuminated volume where the particle was imaged. The triangulation process is improved by using a fitted sheet number (Knutsen et al, 2017) as opposed to the nominal sheet number for particles identified in each image. Less than 5% of reconstructed particles are ghost particles at a high equivalent particle image density of $N_I \approx 0.05$ ppp, at which point 'classic' volumetric methods as used for PTV (i.e. Nishino et al, 1989; Maas et al, 1993; Malik et al, 1993) typically reconstruct approximately the same number of ghost and true particles (Schanz et al, 2016). Lagrangian particle tracking is performed on the reconstructed particles to calculate velocities and accelerations. At present the method is verified via a synthetic experiment using a database born of direct numerical simulation (DNS). The method is intended for experimental investigation of high-Reynolds number turbulent flows.

1.1 Outline

Sect. 2 describes the present scanning PTV method. How various parameters of the method affect 3D particle triangulation within the measurement volume is considered in Sect. 3, addressing both controllable (Sect. 3.2) and uncontrollable (Sect. 3.3) effects. Results of a synthetic experiment are presented in Sect. 4.

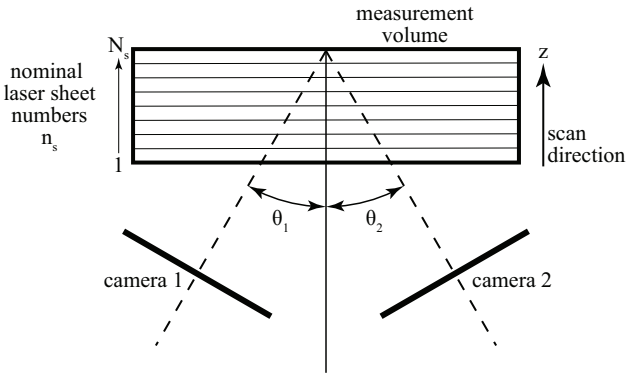

Fig. 1 Schematic of the scanning PTV setup showing two camera views and nominal laser sheet numbers

2 Scanning PTV Methodology

In this section we will introduce and describe the principles of a new method of PTV based on a scanning laser light source. Figure 1 shows a schematic representation of the scanning setup, depicting the two camera views and the subsequent positions of the thin volume illuminated by the scanning laser. Table 1 notes the nomenclature adopted throughout the present work. The experimental configuration for data acquisition is closely related to a standard scanning PIV setup, in which the general concept is that one or more cameras capture particle images as a light sheet is quickly scanned across a measurement volume (e.g. Brücker, 1995; Zhang et al, 2008; Lawson and Dawson, 2014). The speed of the traversing laser sheet ensures the field is approximately 'frozen' as the laser scans through the volume (the error introduced by finite scanning speeds is addressed in Sect. 3.2.4). The main components required to undertake such an experiment are high speed cameras, a pulsed laser light source and a mechanical or optical scanning mechanism. The scanning mechanism deflects the laser beam, expanded to form a laser sheet in order to illuminate particles in thin slices of the domain, such that it traverses across a volume of interest. Over the course of the scan, N_s images are acquired by each camera at nominal sheet numbers $n_s = 1, 2, 3, \ldots, N_s$ which are ideally equispaced along the scanning direction z (figure 1). In scanning PIV, particle images can be used directly to reconstruct velocity fields at different depths z via stereoscopic PIV, and 'stacking' them to recover a volumetric velocity field (i.e. Brücker, 1995), or indirectly via a volumetric intensity reconstruction scheme before application of a cross-correlation algorithm (Lawson and Dawson, 2014). In the present work the particle images are instead used for particle tracking following 3D particle field reconstruction.

PTV algorithms are generally comprised of three steps:

1) detecting the image location of particles from all camera views;
2) determining the probable 3D particle locations, often relying on the intersection of epipolar lines; and
3) linking the particles' locations at successive time instances to form particle tracks

The main contribution of the present work it that of an improved triangulation method for step 2).

2.1 Detection of particle image locations

Laser light scattered from seeding particles is collected via the cameras' optics in the form of particle images. To reduce the impact of noise, image preprocessing is applied. Importantly the preprocessing is chosen in such a way that the intensity information of each particle over the course of a scan is preserved. Dark images are subtracted, and the intensity of each image is rescaled to equalise the image brightness between camera views. Image noise is then reduced by using a Gaussian low-pass filter, with 3×3 pixel (px) window size and a standard deviation of 0.5 px, which also reduces the incidence of false particle detection. Such a filter has a negligible role in the noise-free cases for the range of particle image densities presently investigated, but is important when image noise is present (as an example, using such a filter with the 5% image noise case of Sect. 3.3.1 permits roughly the same number of particles to be reconstructed as for the noise-free case when the equivalent particle image density is $N_I \approx 0.05$ ppp). To identify particle image locations at sub-pixel accuracy, a standard 2D peak-finding algorithm is applied, which first identifies pixels above a prescribed threshold (presently based on the average of the 10 brightest pixels), and then fits two 1D Gaussian functions (Ouellette et al, 2006) to neighbouring pixels.

One of the challenges of higher seeding densities is the effect of overlapping particles within images. Cierpka et al (2013) shows that with increasing particle density, the percentage of particles overlapping also increases, which could lead to a shift in 2D peak identification, possibly resulting in significant positional errors. In scanning PIV, since the whole volume is subdivided into many smaller volumes (i.e. laser sheet thicknesses), the incidence of overlapping particle images is significantly reduced when compared to methods which illuminate the entire volume simultaneously for image taking. The present sub-pixel identification method attempts to avoid erroneous particle identification due to

Table 1 Nomenclature. $[L]$ is the length and $[T]$ the time unit adopted in the physical experiment or DNS database

Parameter	Symbol	Unit
Nominal sheet number	n_s	
Fitted sheet number	f_s	
Number of laser sheets	N_s	
Sheet e^{-1} width	w	$[L]$
Sheet overlap	Δz	$[L]$
Meas. volume dimensions	L_x, L_y, L_z	$[L]$
Particle density (3D)	N_V	ppv
Equiv. particle image density	N_I	ppp
Particle image density, per sheet	$N_{I,s}$	ppp
Particle image diameter	D	$[L]$
Number of particles	N	
Number in a single sheet	N_{ps}	
Number of overlapping particles in a single sheet	N_{os}	
Magnification	M	px/$[L]$
Sheet speed	u_s	$[L]/[T]$
rms turbulent velocity	u'_{rms}	$[L]/[T]$
Kolmogorov lengthscale	η	$[L]$
Kolmogorov timescale	τ_η	$[T]$

the remaining overlapping particles by rejecting particle locations if more than one peak is identified within a radius of 2-3 px.

2.2 Triangulation method

Once the 2D locations in both camera views are identified, starting from an identified particle location in one camera image, an epipolar line in the second camera image can be calculated based on existing knowledge of the camera calibrations. Particle images in the second camera image coinciding with the epipolar line are candidates to be that same particle's image in the second view. The length of this epipolar line is determined by the estimated depth (coordinate z on figure 1) range of the particle in the object space coordinate. In typical triangulation procedures (Maas et al, 1993) this depth is chosen to be that of the illuminated measurement volume, classically the entire volume of interest. At high seeding densities, there can be many possible particle matches from the other view(s), the number being linearly proportional to the length of the epipolar line (Maas et al, 1993). A better estimate of the particle's location along the depth coordinate z in object space will restrict the length of the epipolar line and thus reduce the number of possible matches.

Using a scanning technique (Hoyer et al, 2005) reduces this depth to a fraction of the entire measurement volume. Having restricted the particle's location to the known z position of a nominal laser sheet n_s, Hoyer et al (2005) used the thickness of the laser sheet as the length of the epipolar line for triangulation of individual particles. Such a scheme represents a great improvement in comparison to using the whole measurement volume depth, since the triangulation process begins with a much better estimate of the particle's location along the scan direction z, reducing ambiguous matching to other particles. Yet the chances of incorrect matching between camera images remains an issue, especially when seeding density increases for better spatial resolution of high-Reynolds number flows.

The present technique seeks to improve this scanning setup by further tightening the particle's location along the scan direction. In addition to the nominal sheet number n_s introduced above, a fitted fractional sheet number f_s is central to the current triangulation method. The nominal sheet number corresponds to the sheet number in which the particle is identified, changing value as the particle is illuminated by subsequent (overlapping) laser sheets as the laser sheet scans through the volume. The fractional sheet number f_s indicates the hypothetical sheet number where the particle location would coincide with the centre of the laser sheet (location of peak intensity, assuming some intensity distribution) as the laser sheet scans over it. To find f_s, instances of the same particle being illuminated by successive, overlapped sheets are grouped using a nearest neighbour approach. A search region of only 1-2 px is used as particle displacement during a scan is minimal given high scanning speeds. This is considerably larger than the expected displacement of the particles during image taking for the neighbouring sheets, however a slightly relaxed search radius may account for error during the peak finding step. This search region remains well below the average distance between particles, which can be estimated by $\overline{r}_{2D} = 0.5\, N_{I,s}^{-1/2}$ (Adrian, 1991). We have verified that the search radius for grouping particles in order to calculate f_s remains well below the mean inter-particle distance even at the highest N_V considered here.

Once the particle's appearance in successive scan images is identified, the variation in its intensity during the scan is collated. Given a laser sheet whose intensity varies approximately as a Gaussian across its thickness, the fractional sheet number f_s of the particle is then estimated by fitting a Gaussian curve to these intensities and nominal sheet numbers n_s. The peak of the fitted Gaussian corresponds to the estimated fractional sheet number f_s which we assume corresponds accurately to the particle's location along the scan direction. This technique was originally intended for laser sheet self-calibration (Knutsen et al, 2017), and could still be used thus prior to recording data, but is here repurposed for finding the particle locations themselves along the scan direction.

The triangulation algorithm then follows that used by Maas et al (1993), but now using a reduced search depth giving a shortened epipolar line. The first and last sheet images are excluded since they contain many particles at the edge of the measurement volume not illuminated more than once (obtaining a reasonably accurate f_s requires at least three readings of a particle's intensity). Starting with a particle in camera 1, its f_s is used to fix a projection volume $dz = z(f_s) \pm \Delta_z$, where in practice a small tolerance Δ_z about the z location corresponding to the deduced f_s is used. The estimated 3D locations are projected to both camera views to estimate the re-projection error, which is the difference between the actual location of the particle in the original image to that in its projected image from the candidate triangulation location. The 3D location is only retained as a triangulated particle if the reprojection error is below a certain threshold.

Finally, particles triangulated in a current sheet are removed from particle lists in subsequent sheets where they were identified during the calculation of their f_s. That is, when a particle is triangulated in sheet n_s, it is removed as a potential particle from sheets $n_s + 1$ and so on. Doing so means subsequent sheets have comparatively fewer particles on which a triangulation is attempted, accelerating the procedure along the scan direction z. In addition, an iterative triangulation for the whole volume may be attempted by using residual images. These are formed from subtracting the projected image of triangulated particles from the original image. For example, a 2 px radius about the projected location is identified within the original image, and this intensity distribution is subtracted to form the residual image for that sheet. For the present synthetic experiment, a triangulation on residual images yields an additional modest $\approx 1\%$ of correctly triangulated particles for an equivalent particle image density of $N_I \approx 0.05$ ppp, compared to $\approx 85\%$ of particles being correctly triangulated on the first pass with the original images. If needed, it is possible to use a more sophisticated method for subtracting projected images of triangulated particle locations exploiting a calibrated optical transfer function (Schanz et al, 2013), since overlapping images of particles not yet triangulated are also deleted in such a process. However in this instance, a small fraction of particle images are overlapped due to the division of the volume into thin slices via scanning (being $\approx 2\%$ for the $N_V = 1.8 \times 10^{-4}$ ppv considered in the synthetic experiment, see figure 5 below). Hence such a method is not expected to greatly improve upon a simpler subtraction of all intensity in the proximity of the projected image location. Residual images have not been used in the present work, but could prove useful when the method is used for real experimental data.

2.3 Particle pairing to form tracks

For particle tracking, the scheme of Malik et al (1993) is used. The linking process is guided by three heuristic criteria:

(a) *Nearest neighbour:* particle displacement is limited in all directions
(b) *Minimum acceleration:* Lagrangian acceleration of a particle is limited
(c) *Least change in acceleration:* in cases of multiple choices, the trajectory resulting in the smallest change in acceleration is deemed the most likely one

The scanning PTV method is presently investigated via synthetic experiment. A predictor velocity is required to start the particle tracking algorithm. From a selected particle's location in the first time step, this predictor velocity is then used to estimate the particle's location in the subsequent time step. In an experiment, this would be the corresponding Eulerian velocity from a PIV cross-correlation (i.e. Lawson and Dawson, 2014). For the synthetic experiment as described below, a 'PIV-like' grid of velocity vectors spaced at $\approx 1.53\,\eta$ (Lawson and Dawson, 2014) is sourced at the relevant time step from the DNS database, and the predictor velocity for a particle's location is found via linear interpolation from the gridded data. The velocity fields used are 'ideal' being sourced from spectrally-resolved DNS, where in a real experiment the predictor velocity may be under-resolved or otherwise subject to error. The use of a judiciously chosen search radius based on seeding density about the predicted location means the method is able to tolerate, in practice, some error in the predicted location resulting from a less accurate predictor velocity. A link is created when a particle is found within this specified search region at the subsequent time step. Following the first link, the particle's location in a subsequent time step is predicted using a velocity deduced the first two locations such that the predictor velocity is only required for the first trajectory link. Once possible tracks of length 4 are established, the third criterion (c) is used to further prolong the trajectory. Such particle tracking requires a large number of nearest-neighbour searches, presently undertaken with aid of a GPU and the open-access code of Garcia et al (2010).

3 Study of experimental parameters

The present work aims to improve the triangulation part of the full particle tracking scheme. We here seek to study the influence of various scanning parameters upon the performances of the proposed triangulation method for instantaneous 3D particle fields. Particles triangulated within fields at subsequent times are then linked to form tracks according to the scheme of Malik et al (1993) which we do not alter.

Sect. 3.2 studies the effect of controllable experimental parameters such as the seeding particle density, laser sheet overlap, sheet thickness and scanning laser sheet speed, all being set prior to performing an experiment, where Sect. 3.3 considers uncontrollable factors such as the effect of noise in the images and deviations of the laser sheet intensity profile from an assumed ideal Gaussian one. For the study of controllable experimental parameters, ideal noise-free images were used. The results are quantified in terms of both accuracy (mean error in the triangulated particle locations with respect to 'ground truth' particle locations) and efficacy (fraction of 'ground truth' particles triangulated) of the triangulation procedure. An overview of the relevant scanning parameters, including those to be varied as part of the parametric study, is given in Table 2. The figures in this section result from single realisations of each test, however results appear reasonably converged due to the large number of particles involved. For example, the lowest equivalent particle image density ($N_I \approx 0.005$ ppp) considered involves in excess of 4000 particles; the mean and standard deviation of the error in particle position as shown in figure 3 varies less than 2% over three independent realisations, the fraction of particles found by less than 0.5%. Such variation will be even smaller for higher particle densities with many more particles ensuring confidence in the statistics presented in the parametric studies herein.

3.1 Synthetic particle image generation for the scanning setup

Synthetic images were generated for a simulation of the scanning configuration, which was set up to mimic the previously-realised experiment of Lawson and Dawson (2014). The scanning parameters were chosen based on dimensions of the simulated volume. Two cameras were simulated at angles $\theta = 30°$ and $-30°$, whose sensor sizes were set to 1024×1024 px. Particles were randomly distributed within the volume, based on the required volumetric seeding density N_V in particles per voxel (ppv). The laser sheet positions along the scan were set based on a laser sheet thickness w, sheet spacing Δz and number of sheets N_s, required to scan the depth along the scanning direction L_z. Particles located within the sheet thickness at each sheet position were then projected to image coordinates using a pinhole camera model (Hartley and Zisserman, 2003).

Table 2 Relevant scanning PTV parameters for testing of experimental setup

Parameter	Symbol	Value
Taylor microscale Re	Re_λ	418
Measured volume width	L_x/η	128
Measured volume height	L_y/η	128
Measured volume depth	L_z/η	32
Camera 1 viewing angle	θ_1	30°
Camera 2 viewing angle	θ_2	−30°
Image resolution (px)	–	1024×1024
Particle image diameter	D/η	0.02
Laser sheet thickness	L_z/w	variable
Laser sheet spacing	$M\Delta z$	variable
Laser sheet overlap	$w/\Delta z$	variable
Laser sheet speed	u_s/u'_{rms}	variable

The intensity $I(z)$ of a particle as a function of its z-position within a laser sheet follows the form (Scharnowski and Kähler, 2016):

$$I(z) = I_{max}.\exp\left[-\left|\left(\frac{2z}{\Delta z_0}\right)^s\right|\right], \quad (1)$$

where $\Delta z_0 = w$ is the width at which $I(z)$ drops to I_{max}/e, and s is the shape factor, $s = 2$ being presently used to yield a Gaussian beam profile. The maximum intensity at the light sheet centre, I_{max}, at a scattering angle of $(\pi/2 - \theta)$ is modelled as

$$I_{max} = \frac{4}{\pi k^2 R^2} I_\lambda i(\theta)^2, \quad (2)$$

where I_λ is the input light energy density (set to $2.546 \times 10^4\ J/m^2$), $i(\theta)$ is the Mie scattering coefficient, R is the distance of the particles from the camera sensor and $k = 2\pi/\lambda$ is the wave number for wavelength $\lambda = 532$ nm.

Random background white noise was added to the images at a range of levels (Sect. 3.3.1) with the mean of each noise level set as a percentage of the maximum intensity in the original noise-free image.

3.2 Controllable experimental parameters

3.2.1 Seeding density

To study the effect of increasing seeding density on the efficacy and accuracy of the triangulation method, six seeding densities N_V (in ppv) ranging $N_V = 1.8 \times 10^{-5} - 4.4 \times 10^{-4}$ ppv were considered. An approximate

particle image density N_I (in ppp), representing the equivalent particle image density assuming full-field illumination so as to permit comparison with volumetric 3D-PTV methods, is found by conversion of N_V using an adjusted scanning depth ℓ_z (in px) accounting for the viewing angle:

$$N_I \approx N_V \, \ell_z = N_V \frac{L_z}{\cos\theta}, \tag{3}$$

such that our tested range corresponds to $N_I \approx 0.05 - 0.125$ ppp. The Iterative Particle Reconstruction (IPR) of Wieneke (2013) is able to reconstruct particle fields from single recordings up to around $N_I \approx 0.05$ ppp. Building upon IPR by additionally exploiting time information yielded from a sequence of recordings, the 'Shake-The-Box' method (Schanz et al, 2016) was shown to be highly effective in creating accurate and virtually ghost-free particle tracks over the range $N_I = 0.01 - 0.125$ ppp. With regards to the sheets, $L_z/w = 12$, $w/\Delta z = 5$ were set for this section. Setting $L_z/w = 12$ means the actual particle image densities tackled by the present method corresponds to $\approx N_I/(L_z/w) = N_{I,s} \approx 0.0034 - 0.010$ ppp. That the actual particle image densities faced by the triangulation algorithm are so reduced via partition of the volume using scanning means the 'classic' triangulation of Maas et al (1993), improved using a fitted sheet number f_s, performs satisfactorily. These two sheet parameters are shown in figure 4(b) below to maximise the fraction, for a modest number of sheets, of 'fitted' particles, denoting those particles viewed three or more times such that their position along the scan direction z is estimated accurately by calculating their fitted sheet number f_s via Gaussian curve fitting. In addition $u_s/u'_{rms} = \infty$ was used for this section, such that all images are taken at the same time. In the present method, particle triangulation is aided by tightly fixing the location of each particle along the scan direction z via use of the fitted sheet number f_s. A tolerance of 1/4 th of the sheet thickness w about the estimated z location (from f_s) is used for triangulation.

If a triangulated particle is found within a 1 px radius of a known 'ground truth' particle location then the particle is deemed to be correctly 'detected', this being the same threshold used by Schanz et al (2016). We note that increasing the search range would increase the number of 'found' particles, at the expense of triangulation accuracy. For the purposes of analysis it is however necessary to define a threshold. Figure 2 shows the fraction of particles successfully triangulated over a range of N_V via the present scanning technique. Triangulation where the search depth is the entire volume generally becomes unfeasible (i.e. returns a larger fraction of ghost particles than 'true' particles) at particle

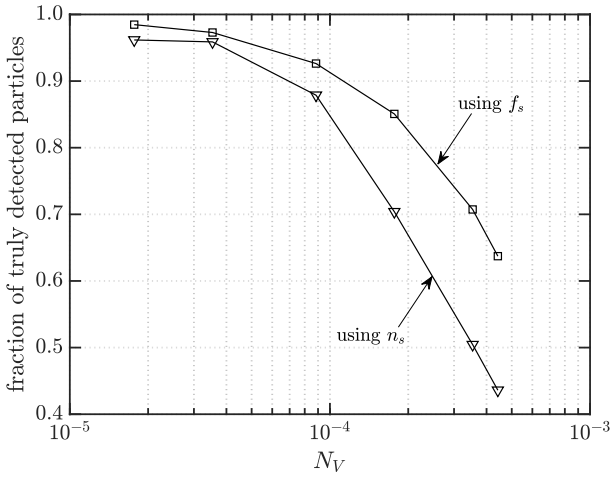

Fig. 2 Fraction of successfully triangulated particles, given the known 'ground truth' locations: using ∇, the nominal sheet number n_s as in Hoyer et al (2005), and □, the fitted sheet number f_s

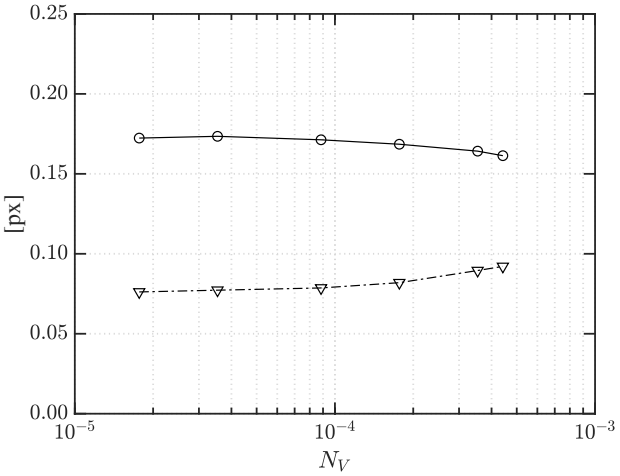

Fig. 3 Positional error in 'truly' detected triangulated particle location for the current scanning PTV: ○, mean error and ▽, standard deviation in error

image densities of $N_I \approx 0.005$ ppp (for the present synthetic setup this corresponds to a volumetric density of $N_V \approx 1.8 \times 10^{-5}$ ppv using (3)). A significant improvement is found over the method of Hoyer et al (2005), which used n_s, as the particle density increases. The vast majority of successfully triangulated particles are fitted with a f_s prior to triangulation; in those cases which they have not the particles' nominal sheet number n_s may be used. The fraction of detected particles decreases rapidly with increasing seeding density for high N_V. At the highest seeding density considered ($N_V = 4.4 \times 10^{-4}$ ppv), the fraction of particles found is only around 0.6.

Figure 3 plots the mean and standard deviation of the positional error between the correctly triangulated ('detected') and ground truth particle locations. We note that as N_V increases, the fraction of successfully detected particles decreases as per figure 2. Interestingly, positional error does not increase with increasing N_V for successfully detected particles, in fact decreasing slightly, however this is at least partially due to the criterion used for deciding on which particles are 'detected', currently set to be those within a 1 px radius of a 'ground truth' particle (i.e. there is an upper bound on the value that the error in figure 3 may take). That the mean positional error is ≈ 0.17 px then suggests that most correctly detected particles are in fact closer than 1 px to the 'ground truth' particle location. However, the percentage of ghost particles increases with increasing N_V. For the tested equivalent particle image densities $N_I \approx \{0.005, 0.01, 0.025, 0.05, 0.1, 0.125\}$ ppp (assuming full-volume illumination) the percentage of ghost particles detected is $\{0.23, 0.38, 1.7, 4.4, 12, 17.3\}$% (as a percentage of the total number of particles triangulated). Although out-performed on positional error by the iterative 'Shake-The-Box' method, which is also able to effectively eliminate ghost particles at the densities considered here, the mean positional error with the current two-camera method compares favourably to that incurred with the iterative MLOS-SMART tomographic reconstruction method, which increases with increasing N_V. Additionally the percentage of ghost particles reconstructed is much smaller, which with MLOS-SMART exceeds 250% at $N_I \approx 0.125$ ppp (Schanz et al, 2016).

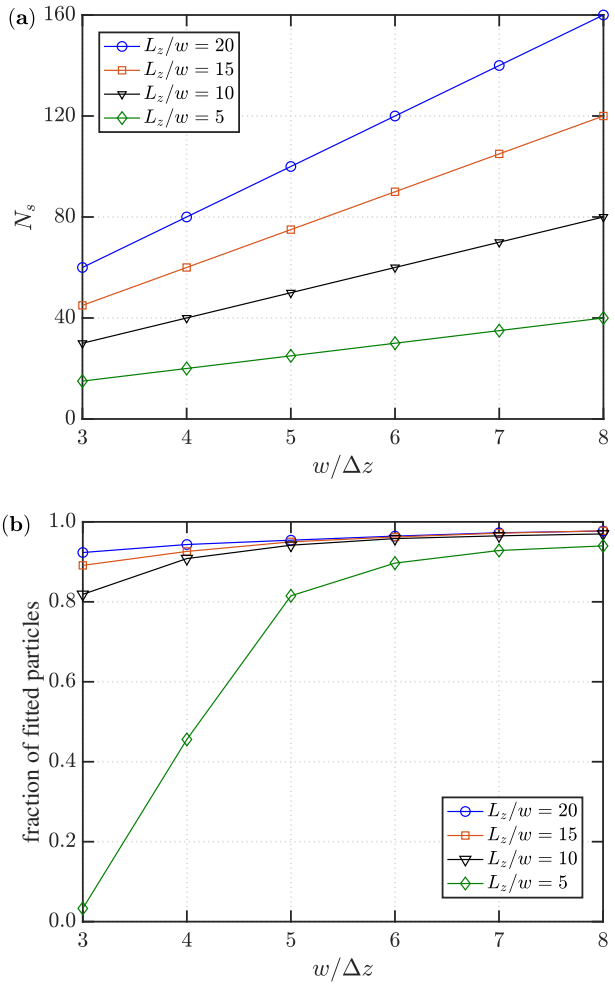

Fig. 4 **a** Number of sheets N_s required as a function of sheet overlap $w/\Delta z$ for different sheet thicknesses L_z/w; **b** fraction of particles successfully fitted

3.2.2 Sheet overlap

Here we consider the effect of sheet overlap $w/\Delta z$ and thickness L_z/w. The triangulation efficacy of our method is directly dependent on the percentage of particles illuminated multiple times over the course of a scan. In order to fit a Gaussian profile, the particle's intensity must be measured at least three times. Hence it is necessary to understand how the fraction of successfully fitted particles depends on these parameters.

Figure 4 shows the result of this parametric study. Sheet overlap ratios of $w/\Delta z = \{3, 4, 5, 6, 7, 8\}$ are considered, corresponding to overlap percentages {66, 75, 80, 83, 85, 87.5}%. For a fixed L_z, sheet thicknesses $L_z/w = \{5, 10, 15, 20\}$ are considered. Figure 4a shows how many sheets (N_s) are required for a fixed L_z as a function of sheet thickness L_z/w and sheet overlap $w/\Delta z$. For higher N_s, the image acquisition rate would need to be increased such that the flow remains 'frozen' as the scan is carried out. For the purposes of this section, sheet speed $u_s/u'_{rms} = \infty$ was used. Figure 4b shows how increasing sheet overlap but decreasing sheet thickness (increasing L_z/w) improves the fraction of fitted particles. The large increase for lower $w/\Delta z$ is expected since more particles will be identified the minimum three times required in order to deduce f_s from a Gaussian fit. The further slow increase at higher $w/\Delta z$ is attributable to more measurements giving a better estimate of f_s (higher N_s for increasing L_z/w and increasing $w/\Delta z$ as shown in figure 4a), and also possibly due to a lower effective seeding density with higher N_s as the scan progresses due to the residual triangulation method (i.e. where a particle, identified over multiple neighbouring sheets during the calculation of f_s, is removed from the particle lists of subsequent sheets where it appears if successfully triangulated in a previous sheet). However gains beyond $w/\Delta z = 5$ and $L_z/w = 10$ are marginal yet fur-

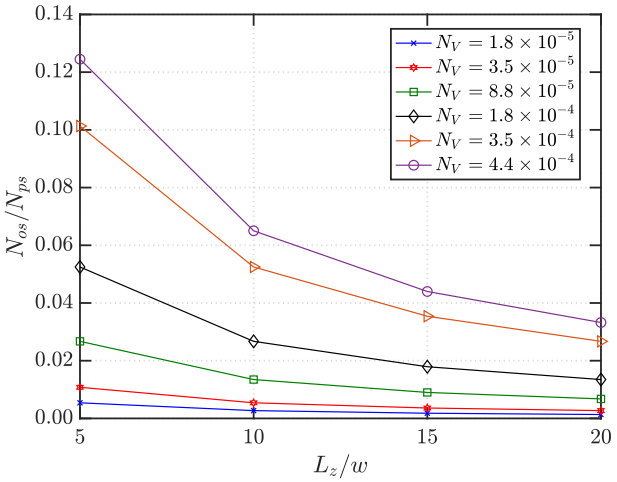

Fig. 5 Fraction of overlapping particles for different sheet thicknesses L_z/w

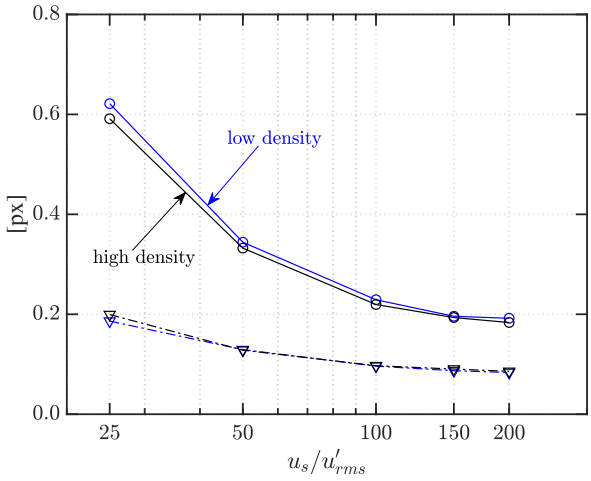

Fig. 6 Error in linearly-corrected triangulated particle locations for different sheet speeds u_s/u'_{rms}; ∘, mean error and ▽, standard deviation in error; curves in black, $N_V = 1.8 \times 10^{-4}$ ppv ($N_I \approx 0.05$ ppp); curves in blue, $N_V = 1.8 \times 10^{-5}$ ppv ($N_I \approx 0.005$ ppp)

ther increasing these requires more sheets. The values $w/\Delta z = 5$ and $L_z/w = 12$ were thus selected for the density study above in Sect. 3.2.1 and for the synthetic particle tracking experiment in Sect. 4.

3.2.3 Overlapping particle images

To study the effect of sheet width and particle density on the incidence of overlapping particle images, all N_V as considered in Sect. 3.2.1 and the four sheet widths (L_z/w) from Sect. 3.2.2 are investigated. The extent of overlapping particles directly affects peak-finding within the camera images and thus the correct identification of individual particles. The number of particles N_{ps} within a single image (i.e. the slice illuminated by the laser sheet) is calculated based on the sheet width as

$$N_{ps} = N \times \frac{w}{L_z}. \tag{4}$$

The number of overlapping particles N_{os} in each sheet image of area $A = L_x \times L_y$ is then given by (Cierpka et al, 2013)

$$N_{os} = (N_{ps} - 1) + \frac{A}{A_{crit}} \left(\exp\left[\frac{-(N_{ps} - 1) A_{crit}}{A} \right] - 1 \right), \tag{5}$$

where $A_{crit} = \pi(D/2)^2$ for particle image diameter D.

Figure 5 shows the percentage of overlapping particles for different sheet widths at the seeding densities studied. A high fraction of overlapping particles will degrade the efficacy of individual particle identification in images for the ensuing triangulation. The fraction of overlapping particle images is essentially zero for the lowest density case $N_V = 1.8 \times 10^{-5}$ ppv ($N_I \approx 0.005$ ppp) at all sheet thicknesses. For the higher density cases, a substantial reduction from around 10% overlapping particle images to around 5% is achieved with $L_z/w = 10 - 15$.

3.2.4 Finite sheet speed

A source of error specific to a scanning setup is the effect of a finite scanning laser sheet speed. Where figure 2 considers an infinite sheet speed, figure 6 considers a range to better understand the effect of finite sheet speed on triangulation accuracy. For both the present parameter assessment and the synthetic experiment presented in Sect. 4, velocity fields from the forced homogeneous isotropic box turbulence case are sourced from the Johns Hopkins Turbulence Database (JHTDB) (Li et al, 2008), with rms velocity u'_{rms}, and are used to advect the particles. Subsequent scans of the measurement volume are separated by a time Δt during which the particles are advected by the local fluid velocity. For a finite sheet speed, requiring a finite time to scan the volume, the particles are also advected *during* the scan. In an experiment, this time is determined by the frame rate f_{scan} and the number of sheets N_s. A decreasing u_s/u'_{rms} means that the particles move increasingly during the scan, giving triangulated locations for a single scan that correspond to an increasing spread of times between particles in the first and last sheets.

Sheet speed ratios of $u_s/u'_{rms} \approx \{25, 50, 100, 150, 200\}$ were considered to study the effect of finite sheet speed.

Following 3D particle reconstruction, a linear (i.e. first-order) positional correction is applied to all particles using their fitted sheet numbers to recover their approximate locations at a simultaneous time corresponding to the end of the scan. That is, particles in the final sheet will be subject to no correction, whereas particles in the first sheet will have a 3D positional correction applied by adding $\boldsymbol{\Delta x} = t_{scan}\boldsymbol{u}_{loc}$ to their triangulated positions, where t_{scan} is the time taken to complete the scan, and $\boldsymbol{u}_{loc}$ is the local velocity vector at the beginning of the scan. Figure 6 shows the mean error in pixels between the 'ground truth' particle locations (recorded at the end of the scan) and the corrected particle locations for the different sheet speed ratios. When the sheet speed is $u_s/u'_{rms} \gtrsim 100$, the mean error in the corrected particle locations is ≈ 0.2 px for both low ($N_V = 1.8 \times 10^{-5}$ ppv, $N_I \approx 0.005$ ppp) and higher ($N_V = 1.8 \times 10^{-4}$ ppv, $N_I \approx 0.05$ ppp) seeding density cases. These values are only ≈ 0.03 px greater than the positional error shown for the infinite-u_s study in figure 3 in Sect. 3.2.1. We note also that all other things being equal, the error will likely increase with a larger L_z (i.e. more sheets) since particles in latter sheets have more time to move during the scan, possibly eventually exceeding that which can be corrected linearly as herein.

3.3 Uncontrollable experimental parameters

3.3.1 Noise

Four levels of random noise are introduced to the ideal images in order to study its impact on the triangulation procedure, being 5, 10, 15 and 20% of the maximum illumination in the noise-free image. This was applied to all seeding densities considered in Sect. 3.2.1. Prior to the triangulation step, 2D intensity peaks must be located within the images, presently carried out using the peak-finding algorithm described in Sect. 2.1. Figure 7a shows the mean error in the 2D peak location estimation when images are subject to these varying levels of noise. A higher intensity threshold for peak-identification was used in this section than for others, so that particles for the 20% noise case could be better distinguished. This in general results in fewer triangulated particles as weaker intensity spots are disregarded. The mean error in 2D intensity peak location was < 0.3 px for all cases and all noise levels investigated. The error appears to grow in a fairly steady manner both with noise level and particle density for the ranges we consider, except for the highest N_V where it increases rather more rapidly. This is presumably due to the intensity peaks becoming more difficult to distinguish with high levels of background noise when there are many closely-spaced particle images. Figure 7b shows fraction of successfully triangulated particles (using the same criterion used for figure 2, i.e. if a triangulated location is within a 1 px radius of a known 'ground truth' particle location, the particle is considered 'found') for the studied range of seeding densities. The performance of the present triangulation procedure is shown to be robust in the presence of mild to moderate noise. However above noise levels of $\approx 15\%$ and particle densities $N_V \approx 3.5 \times 10^{-4}$ ppv, the method triangulates $< 50\%$ of ground-truth particle locations.

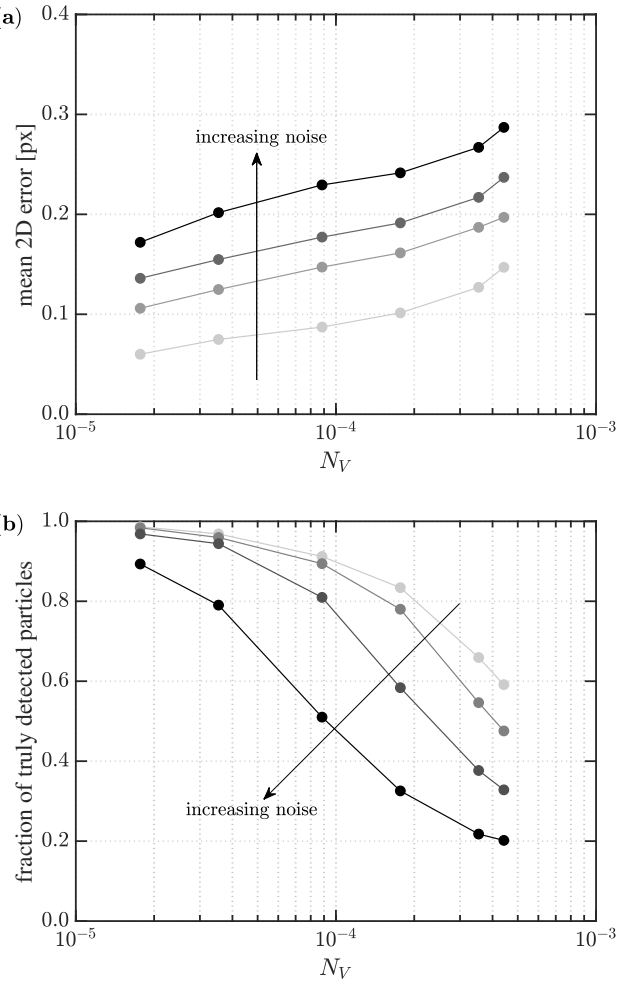

Fig. 7 **a** Mean error in 2D particle image detection as a function of particle density N_V for four different random noise levels $\{5\%, 10\%, 15\%, 20\%\}$, curves darkening with increasing noise level; **b** fraction of successfully triangulated particles, given the known 'ground truth' locations; curves and N_V range same as in **a**

3.3.2 Deviation of the laser sheet profile from an ideal Gaussian

The shape of the laser sheet profile is determined by shape factor s in (1), which when set to $s = 2$ gives a Gaussian profile. Accurate sheet fitting requires some variation of the laser's intensity over the range of nominal sheet number n_s being used to find f_s. We now vary this to $s = \{2, 4, 10\}$ for the seeding density $N_V = 1.8 \times 10^{-4}$ ppv ($N_I \approx 0.05$ ppp). As s increases the profile approaches a top-hat profile. Figure 8a shows the intensity variation of a particle over five sheets, demonstrating how the shape of the laser beam profile changes with s. For a single particle, the fitted sheet numbers (f_s) obtained for $s = \{2, 4, 10\}$ within (1) are essentially exactly the same. The mean error shown in figure 8b also does not vary significantly for the different profile shapes considered presently. Therefore we find the actual profile of the laser sheet doesn't significantly affect the efficacy of triangulation method for the values of s considered here, as long as there is some intensity variation over the laser sheet's thickness, and it is correctly characterized prior to the experiment.

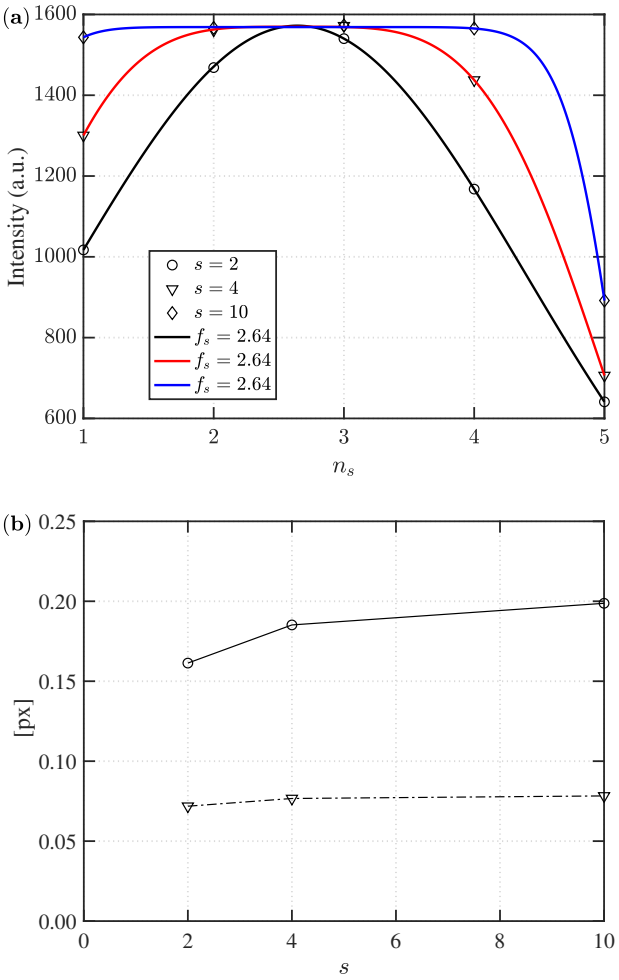

Fig. 8 **a** Intensity variation of a single particle over five sheets within a scan, for three different sheet shape factors s; **b** error in triangulated particle location for the different shape factors for the medium-density case with $N_V = 1.8 \times 10^{-4}$ ppv ($N_I \approx 0.05$ ppp): ∘, mean positional error; ▽, standard deviation of positional error

4 Verification via synthetic experiment

The JHTDB was used for a synthetic experiment on which the present methodology has been tested thus far. The initially random distribution of particles is advected in time subject to the time-evolving velocities from the forced and isotropic DNS fields. For this purpose we made use of the Lagrangian tracking 'GetPosition' function as supplied in the JHTDB toolbox (Yu et al, 2012). Scanning PIV images for two camera angles were created using a code based on the EUROPIV Synthetic Image Generator (Lecordier and Westerweel, 2004). The main details of the synthetic experiment may be found in Table 3. An image resolution of 1024×1024 px was used. Parameters were chosen such as to be dynamically similar to the scanning PIV experiment of Lawson and Dawson (2014). Aside from the sheet overlap, which has been increased slightly for the synthetic experiment in order to deduce the fitted sheet number f_s for individual particles, the remaining parameters are also chosen to be similar, so that our synthetic experiment represents a setup realisable in the laboratory (i.e. feasible sheet speed u_s/u'_{rms}). The present scanning PTV setup places specific demands upon laser sheet speed and sheet overlap for the triangulation to be effective, which may impose practical restrictions upon laboratory experiments. The sheet overlap must be higher than in previous similar setups (Lawson and Dawson, 2014, 2015) such that f_s may be calculated for a majority of particles. Both the sheet overlap and sheet speed requirements of the current methodology are met in the recent experiments in water presented in Lawson et al (2019) for a Taylor microscale Reynolds number of $Re_\lambda \approx 315$. Samples from within the JHTDB are separated in time by the large eddy turnover time $T_L = L/u'_{rms}$, where L is the integral scale, and in space by a distance $\approx 1.5\,L$ within the computational domain, such that the samples may be considered independent.

4.1 Lagrangian statistics

The present experimental scanning method as used by Lawson and Dawson (2014) was first developed to yield Eulerian statistics on a regular grid via PIV cross-correlation.

Table 3 Parameters for the synthetic experiment. Those not listed here are as in Table 1

$N_{samples}$	N_s	N_V	N_I
15	62	1.8×10^{-4} ppv	0.05 ppp
$\Delta t/\tau_\eta$	u_s/u'_{rms}	$w/\Delta z$	L_z/w
0.094	157	5.0	12

Lagrangian statistics are the goal of the present technique. The present method tracks a single particle for over 20 Δt ($\approx 2\,\tau_\eta$) to an accuracy of < 1 px when compared to the final location found using pseudo-tracking directly within the time-evolving DNS. Particle movement during the scan (due to the finite sheet speed) remains small enough such that a linear positional correction using the local velocity is adequate to interpolate their position to that corresponding to a simultaneous time for the whole domain before the tracking algorithm is applied, as was done in Sect. 3.2.4. The search region about a predicted location is presently set to a radius of approximately 1/3 of the mean inter-particle distance. Once tracks of a desired length are established via the particle tracking algorithm, the cubic spline interpolation scheme of Lüthi et al (2005) is used. All components of the particle velocities and accelerations may then be calculated at any point along these trajectories.

4.1.1 Visualisations

Figure 9 shows tracks for two samples of the synthetic experiment. Figure 9a shows a larger number of smaller eddies whereas the sample of figure 9b is seemingly dominated by a single larger structure on the order of the measurement volume itself.

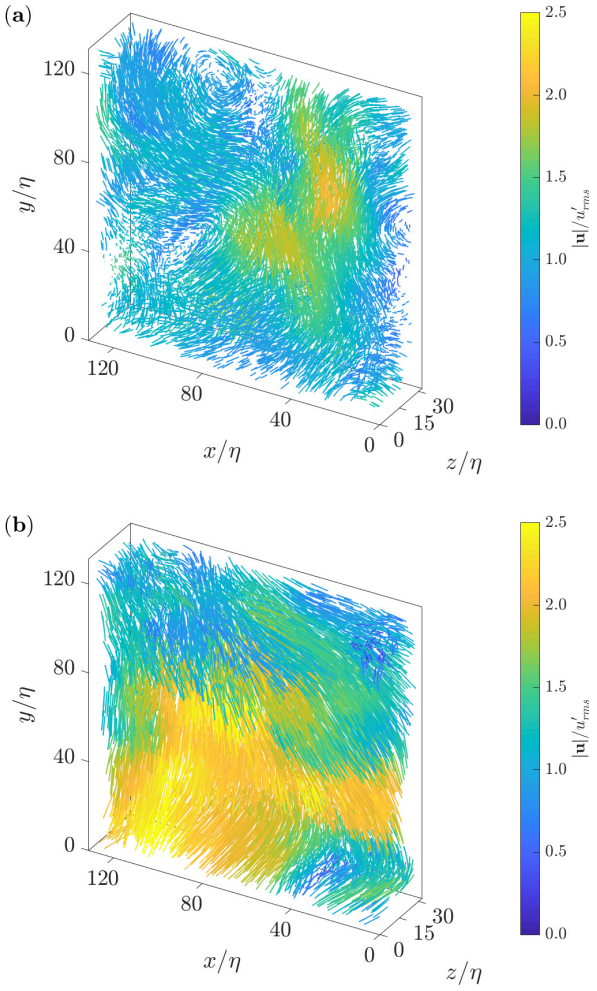

Fig. 9 Particle tracks for two different samples, demonstrating variations in structures for the selected 'slab' of turbulence. Tracks are coloured by the magnitude of the velocity vector $\mathbf{u} = (u, v, w)$ along each track at each point in time

4.1.2 Accelerations

Tracks of length $\approx 0.75\,\tau_\eta$ in time are used to calculate accelerations as in Voth et al (2002). Components of Lagrangian acceleration $\boldsymbol{a} = (a_x, a_y, a_z)$, the distributions of which are shown in figure 10. As demonstrated in Voth et al (2002), the presence of large accelerations signalling extreme events is characteristic of turbulence. For reference we have also plotted the curve fit of Voth et al (2002), although we note their data was taken in a sparsely-seeded flow. For the present high-density synthetic experiment ($N_I \approx 0.05$ ppp), we find the tails to be considerably narrower than that suggested by Voth et al (2002). However wider tails are recovered if the experiment is repeated at a lower density ($N_I \approx$ 0.005 ppp). To avoid ambiguity in the particle matching, the search radius about a predicted location in a subsequent time step is restricted to 1/3 of the inter-particle distance as mentioned above. Hence the sparse case will permit larger accelerations to be recorded, since the algorithm in the dense case is unable to distinguish one potential track from another if the particle's actual location differs from the predicted location by an amount on the order of the inter-particle spacing. The sparse case, with a larger inter-particle spacing, therefore admits larger 'deviations' from a predicted path, and hence more extreme accelerations. The difference in the acceleration PDFs is therefore not the result of a difference in spatial or temporal resolution, but rather exposes an inherent limit on the magnitude of acceleration measurable in dense fields. Presently, the particle tracking algorithm proposed by Malik et al (1993) is used unaltered with a conservative search radius such

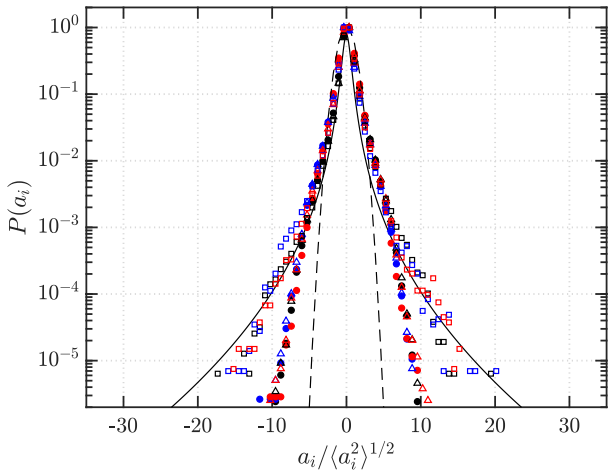

Fig. 10 Probability density function (PDF) of the components of Lagrangian acceleration normalized by the standard deviation for each component: for the present LPT methodology: • a_x, • a_y, • a_z; from pseudo-tracking directly within the DNS fields: ▵ a_x, ▵ a_y, ▵ a_z; repeating the synthetic experiment with a lower equivalent particle image density $N_I \approx 0.005$ ppp: □ a_x, □ a_y, □ a_z. Accelerations are computed from tracks of length $\approx 0.75\,\tau_\eta$ in time from which measurements are bootstrapped along the lengths to increase statistical convergence; ——, fit from Voth et al (2002): $P(a) = C\ \exp(-a^2/((1 + |a\beta/\sigma|^\gamma)\sigma^2))$, where $\beta = 0.539$, $\gamma = 1.588$, $\sigma = 0.508$ and $C = 0.786$; – – – , Gaussian distribution with the same standard deviation as $P(a_x)$

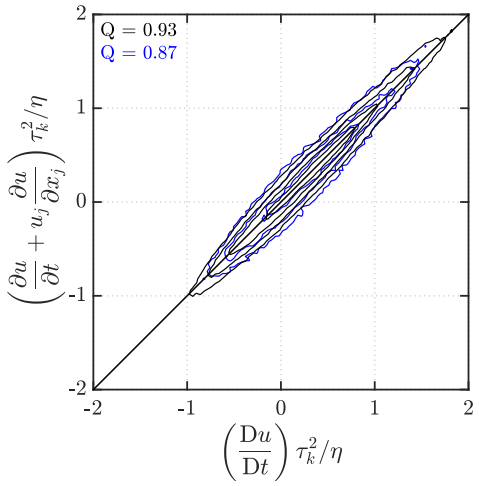

Fig. 11 Acceleration check (joint PDF) relating Lagrangian, local and convective accelerations; in black, for the high density ($N_I \approx 0.05$ ppp) case; in blue, for the low density ($N_I \approx 0.005$ ppp) case. The high correlation coefficient (value of Q) between the two sides of (6) for the high density case demonstrates good spatial resolution

as to avoid ambiguous particle matching. However it is possible that a more sophisticated tracking scheme, perhaps involving a multi-pass approach for 'unclaimed' particles, could admit larger particle accelerations in dense fields.

In general, long tracks formed over a time corresponding to at least $\approx 2\tau_\eta$ are necessary in order to yield smooth Lagrangian statistics (Lüthi et al, 2005), particularly spatial velocity derivatives. Tracks of this length are used to reproduce the acceleration 'test' as in Hoyer et al (2005). The Lagrangian accelerations ($a_i = \mathrm{D}u_i/\mathrm{D}t$) are related to the local accelerations ($a_{l,i} = \partial u_i/\partial t$) and convective accelerations ($a_{c,i} = u_j \partial u_i/\partial x_j$) via the following relation:

$$\frac{\mathrm{D}u_i}{\mathrm{D}t} = \frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j}. \qquad (6)$$

Spatial derivatives for the convective accelerations require information from neighbouring particles. At present the scheme of Lüthi et al (2005) is used with information from the closest 20 particles to the point in space where we seek the spatial derivative, along with information from times t and $t \pm 2\,\Delta t$. The average inter-particle distance in the present fields (at the time of initial seeding) is approximately $1.3\,\eta$, whereas for Lüthi et al (2005) it was around $4.2\,\eta$, and Hoyer et al (2005) used particle fields with average particle separations of $2.4\,\eta$. As detailed by Hoyer et al (2005), relation (6) is a strict test of spatial resolution. In figure 11 we find very high correlation between the left- and right-hand sides of (6) for the a_x component by considering their joint PDFs, confirming the adequate spatial resolution of the present high density ($N_I \approx 0.05$ ppp) synthetic experiment. We find the value of the correlation to be somewhat decreased for the lower density ($N_I \approx 0.005$ ppp) case. Some degradation is to be expected since the fidelity of the velocity gradient tensor, terms of which appear in (6), is reduced at lower seeding density. However we note that for the same number of independent samples, the lower density case represents a smaller number of measurements, meaning the results are less converged. Similar results are found for the other two spatial dimensions.

5 Conclusions

We have introduced a novel method for Lagrangian particle tracking based on a scanning laser technique. Particles are first triangulated in the 3D measurement domain and then linked in subsequent volumes in time to form tracks from which Lagrangian velocities and accelerations can be calculated. The main contribution of the present work is to improve the particle triangulation efficacy. We demonstrated the use of a fitted sheet space number, which serves to further restrict the particle search depth in the scanning direction. The new method accurately triangulates a large fraction of true

particle locations for very high particle densities, up to an equivalent particle image density $N_I \approx 0.125$ ppp, where typical volumetric triangulation typically suffers from large fractions of falsely detected particles when $N_I \approx 0.005$ ppp. The fraction of successfully triangulated particles remains high in the presence of mild to moderate noise, and particle movement during the scan for finite scanning laser sheet speeds can be corrected for if the sheet speed is at least two orders of magnitude larger than the characteristic velocity scale of the flow. The technique was tested via synthetic experiment using a DNS database, mimicking the conditions of a previous scanning setup, for which we were able to calculate Lagrangian velocities and accelerations. Spatial resolution was demonstrated by very good correlation of the Lagrangian acceleration with the local and convective accelerations. An inherent upper limit on the magnitude of accelerations measurable was found for very dense fields using the current tracking algorithm due to a reduced mean inter-particle distance. To avoid ambiguity when forming particle tracks, the more dense a particle field is, the smaller the permissible search radius about a particle's predicted location at a subsequent time step. A more advanced particle tracking algorithm could possibly permit larger accelerations to be detected at higher densities.

The new scanning PTV method presents a robust technique for obtaining Lagrangian statistics in densely-seeded measurement volumes, required for the adequate spatial resolution of flows with high Reynolds number. We intend to use the method introduced herein for high-Reynolds number experimental flow data.

Acknowledgements The authors would like to acknowledge financial support from the Norwegian Research Council, SPRAYICE, project no. 244386.

References

Adrian RJ (1991) Particle-imaging techniques for experimental fluid mechanics. Ann Rev Fluid Mech 23:261–304

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

Cierpka C, Lütke B, Kähler CJ (2013) Higher order multi-frame particle tracking velocimetry. Exp Fluids 54:1533

Elsinga GE, Scarano F, Wieneke B, Van Oudheusden BW (2006) Tomographic particle image velocimetry. Exp Fluids 41:933–947

Garcia V, Debreuve E, Nielsen F, Barlaud M (2010) K-nearest neighbor search: Fast GPU-based implementations and application to high-dimensional feature matching. In: 17th IEEE International Conference on Image Processing (ICIP), Hong Kong

Gesemann S, Huhn F, Schanz D, Schröder A (2016) From noisy particle tracks to velocity, acceleration and pressure fields using B-splines and penalties. In: 18th Int. Symp. on the Application of Laser and Imaging Techniques to Fluid Mechanics, Lisbon, Portugal

Hartley R, Zisserman A (2003) Multiple view geometry in computer vision. Cambridge University Press

Hoyer K, Holzner M, Lüthi B, Guala M, Liberzon A, Kinzelbach W (2005) 3D scanning particle tracking velocimetry. Exp Fluids 39:923

Knutsen AN, Lawson JM, Dawson JR, Worth NA (2017) A laser sheet self-calibration method for scanning PIV. Exp Fluids 58:145

Lawson JM, Dawson JR (2014) A scanning PIV method for fine-scale turbulence measurements. Exp Fluids 55:1857

Lawson JM, Dawson JR (2015) On velocity gradient dynamics and turbulent structure. J Fluid Mech 780:60–98

Lawson JM, Bodenschatz E, Knutsen AN, Dawson JR, Worth NA (2019) Direct assessment of Kolmogorov's first refined similarity hypothesis. Phys Rev Fluids 4:022,601

Lecordier B, Westerweel J (2004) The EUROPIV synthetic image generator (SIG). In: Particle image velocimetry: recent improvements, Springer, pp 145–161

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

Lüthi B, Tsinober A, Kinzelbach W (2005) Lagrangian measurement of vorticity dynamics in turbulent flow. J Fluid Mech 528:87–118

Lynch KP, Scarano F (2015) An efficient and accurate approach to MTE-MART for time-resolved tomographic PIV. Exp Fluids 56:66

Maas HG, Gruen A, Papantoniou D (1993) Particle tracking velocimetry in three-dimensional flows. Exp Fluids 15:133–146

Malik NA, Dracos T, Papantoniou DA (1993) Particle tracking velocimetry in three-dimensional flows. Exp Fluids 15:279–294

Nishino K, Kasagi N, Hirata M (1989) Three-dimensional particle tracking velocimetry based on automated digital image processing. Trans ASME J Fluid Eng 111:384–391

Novara M, Scarano F (2013) A particle-tracking approach for accurate material derivative measurements with tomographic PIV. Exp Fluids 54:1584

Ouellette NT, Xu H, Bodenschatz E (2006) A quantitative study of three-dimensional lagrangian particle tracking algorithms. Exp Fluids 40:301–313

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 24:024,009

Schanz D, Gesemann S, Schröder A (2016) Shake-The-Box: Lagrangian particle tracking at high particle image densities. Exp Fluids 57:70

Scharnowski S, Kähler CJ (2016) Estimation and optimization of loss-of-pair uncertainties based on PIV correlation functions. Exp Fluids 57:23

Schneiders JFG, Scarano F (2016) Dense velocity reconstruction from tomographic PTV with material derivatives. Exp Fluids 57:139

Virant M, Dracos T (1997) 3D PTV and its application on lagrangian motion. Meas Sci Technol 8:1539

Voth GA, La Porta A, Crawford AM, Alexander J, Bodenschatz E (2002) Measurement of particle accelerations in fully developed turbulence. J Fluid Mech 469:121–160

Wieneke B (2013) Iterative reconstruction of volumetric particle distribution. Meas Sci Technol 24:024,008

Yu H, Kanov K, Perlman E, Graham J, Frederix E, Burns R, Szalay A, Eyink G, Meneveau C (2012) Studying lagrangian dynamics of turbulence using on-demand fluid particle tracking in a public turbulence database. J Turbul 13:1–29

Zhang W, Hain R, Kähler CJ (2008) Scanning PIV investigation of the laminar separation bubble on a SD7003 airfoil. Exp Fluids 45:725–743