

LAGRANGIAN STATISTICS FOR DENSELY-SEEDED FLOWS USING SCANNING PARTICLE IMAGE VELOCIMETRY

Melissa Kozul*, Vipin Koothur, Nicholas A. Worth, James R. Dawson

Dept. of Energy and Process Engineering
Norwegian University of Science and Technology
NO-7491 Trondheim, Norway

*melissa.kozul@ntnu.no

ABSTRACT

We propose a novel robust three-dimensional 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 its multiple intensity recordings. Following successful reconstruction of a time series of three-dimensional particle fields, particle tracks are formed from which all components of Lagrangian velocity and acceleration are calculated. The method is presently verified via synthetic experiment using a database born of direct numerical simulation, and is intended for high-Reynolds number experimental flows.

BACKGROUND

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 tracking 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 seeding densities corresponding to ≈ 0.005 particles per pixel (ppp) in the particle images (cf. Maas *et al.*, 1993) in order to maintain confidence in the deduced particle tracks.

METHODOLOGY

We introduce and describe the principles of a new method of LPT based on a scanning laser light source. Figure 1 is a schematic representation of the scanning setup, showing the two camera views and the subsequent positions of the thin volume illuminated by the scanning laser. The experimental configuration for data acquisition is closely related to a standard scanning particle image velocimetry (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; Lawson & Dawson, 2014). The speed of the traversing laser sheet ensures the field is approximately 'frozen' as the laser scans through the volume. 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 successive 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 $n_s = 1, 2, 3, \dots, N_s$ which are ideally equispaced along the scanning direction z (figure 1). In scanning PIV the particle images are used directly, or indirectly following a volumetric intensity reconstruction scheme, with a cross-correlation algorithm. In the present work the particle images are instead used for 3D particle field reconstruction as required for particle tracking.

LPT 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 is an improved triangulation method for step 2).

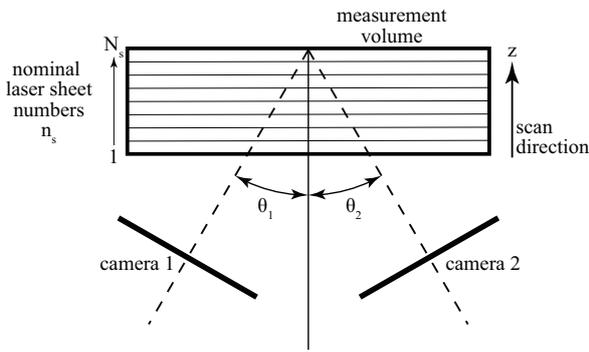


Figure 1. Schematic of the scanning PIV setup to be used for LPT showing two camera views and nominal laser sheet numbers.

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

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 a 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, which is 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 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 the high scanning speeds used. 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 (Knutson *et al.*, 2017), 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. The 3D loca-

tion is only retained as a successfully triangulated particle if the reprojection error is below a certain threshold. Finally, particles triangulated in a current sheet are removed from subsequent sheet images where they were identified during the calculation of their f_s . That is, when a particle is triangulated in sheet n_s , its projected image is subtracted from the image for sheet $n_s + 1$ and so on. Removing already-triangulated particles gives residual images for subsequent sheets with comparatively fewer particles, accelerating the triangulation procedure along the scan direction z .

Figure 2 shows the fraction of particles successfully triangulated 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 image densities of $ppv \approx 0.005$ (for our present synthetic experiment this corresponds to a volumetric density of particles per voxel (ppv) of $ppv \approx 1.8 \times 10^{-5}$, see conversion 1 below). A significant improvement is found over the method of Hoyer *et al.* (2005), which used n_s , as the particle density increases. If a triangulated particle is found within a 1 px radius of a ‘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. For this threshold, the particle positional error is ≈ 0.17 px (with a standard deviation of ≈ 0.08 px), which, for the found particles, does not change significantly over the range of densities spanned by figure 2.

A source of error specific to such a scanning setup is the effect of a finite scanning laser sheet speed. Where figure 2 considers an infinite sheet speed, figure 3 considers a range to better understand the effect of finite sheet speed on triangulation accuracy. Following triangulation of the fields, 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 positional correction, whereas the 3D positions of particles in the first sheet will be ‘corrected’ by adding $\Delta \mathbf{x} = t_{scan} \mathbf{u}_{loc}$ to their triangulated positions, where t_{scan} is the time taken to complete the scan, and \mathbf{u}_{loc} is the local velocity vector at the beginning of the scan. Figure 3 shows the mean error in pixels between the ‘ground truth’ particle location (recorded at the end of the scan) and 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 of the order of 0.2 px for both low ($ppv = 1.8 \times 10^{-5}$, $ppp = 0.005$) and higher ($ppv = 1.8 \times 10^{-4}$, $ppp = 0.05$) seeding density cases. These values are only around 0.03 px greater than the positional error shown for the infinite- u_s study in figure 2. Random noise in the particle images is another obvious source of potential error. For the present modified triangulation procedure an additional ≈ 0.2 px positional error resulted for a random noise level of up to 20% (based on the maximum illumination in the noise-free image), demonstrating the robustness of the method.

Particle pairing to form tracks

For particle tracking using our 3D particle fields, the scheme of Malik *et al.* (1993) is used. The linking process is guided by three heuristic criteria:

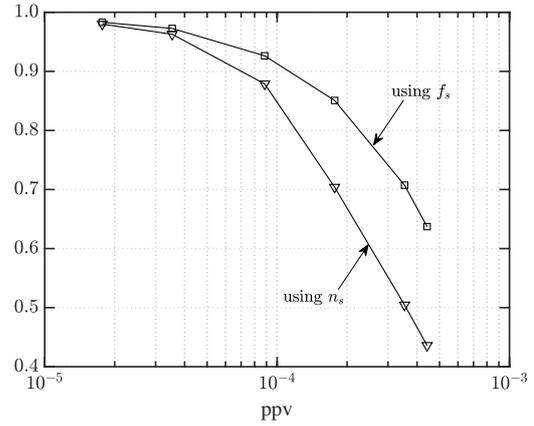


Figure 2. Fraction of correctly triangulated particles, given the known ‘ground truth’ locations: using ∇ , the nominal sheet number n_s , and \square , the fitted sheet number f_s .

- (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

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 & Dawson, 2014). For the synthetic experiment as described below, a ‘PIV-like’ grid of velocity vectors is sourced at the relevant time step from the direct numerical simulation (DNS) database, and the predictor velocity for a particle’s location is found via linear interpolation. A link is created when a particle is found within a specified search region at this subsequent time step. Following the first link, the particle’s location in a subsequent time step is predicted from 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).

VERIFICATION VIA SYNTHETIC EXPERIMENT

The Johns Hopkins Turbulence Database (JHTDB) (Li *et al.*, 2008) was used for a synthetic experiment on which the present methodology has been tested thus far. An initially random distribution of particles at a specified volumetric seeding density ppv is advected in time by evolving velocities from the forced and isotropic DNS fields. Scanning PIV images for two camera angles were created using a code based on the EUROPIV Synthetic Image Generator (Lecordier & Westerweel, 2004). The main details of the synthetic experiment are given in table 1. 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 & Dawson (2014). Aside from the sheet

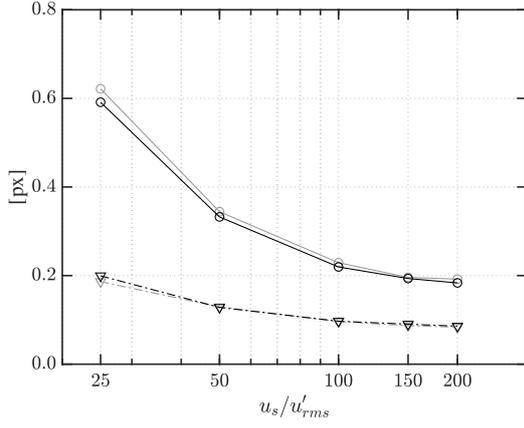


Figure 3. Error in linearly-corrected triangulated particle locations for different sheet speeds u_s/u'_{rms} : \circ , mean error and ∇ , standard deviation in error; curves in black, ppv = 1.8×10^{-4} (ppp = 0.05); curves in grey, ppv = 1.8×10^{-5} (ppp = 0.005).

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} where u'_{rms} is the rms velocity of the isotropic turbulence). 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 $\sim 1.5L$ within the computational domain, such that the samples can be considered independent.

Table 1. Parameters for the synthetic experiment: Re_λ is the Taylor-scale Reynolds number, θ_1 and θ_2 are the camera viewing angles, η is the Kolmogorov lengthscale, τ_η the Kolmogorov timescale, Δt_{PIV} is the time separation between successive scans.

Re_λ	$N_{samples}$	N_s	L_x/η	L_y/η	L_z/η
418	15	62	128	128	32
$\Delta t_{PIV}/\tau_\eta$	u_s/u'_{rms}	ppv	ppp	θ_1	θ_2
0.094	157	1.8×10^{-4}	≈ 0.05	-30°	30°

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 . As noted above, in practice we set the volumetric particle density ppv. The equivalent 2D particle image density ppp, representing the particle density in camera images used for triangulation by volumetric methods illuminating the entire measurement volume simultaneously, is found by conversion of ppv using an adjusted scanning depth ℓ_z (in px) accounting for the viewing angle:

$$ppp = ppv \ell_z = ppv \frac{L_z}{\cos \theta}. \quad (1)$$

Particles located within the sheet thickness at each sheet position were then projected to image coordinates using a pinhole camera model (Hartley & Zisserman, 2003). The intensity $I(z)$ of a particle as a function of its z -position within a laser sheet follows the form (Scharnowski & Kähler, 2016):

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

where $\Delta z_0 = w$ is the width at which $I(z)$ drops to I_{max}/e , and where 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

$$I_{max} = \frac{4}{\pi k^2 R^2} I_\lambda i(\theta)^2, \quad (3)$$

where I_λ is the input light energy density (set to $2.546 \times 10^4 \text{ J/m}^2$), $i(\theta)$ 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 $\lambda = 532 \text{ nm}$.

Lagrangian statistics

Figure 4 shows tracks for two samples for the synthetic experiment. Figure 4(a) shows a larger number of smaller eddies whereas the sample of figure 4(b) is seemingly dominated by a single larger structure on the order of the measurement volume itself.

The present experimental scanning method as used by Lawson & Dawson (2014) was first developed to yield Eulerian statistics on a regular grid via PIV cross-correlation. Lagrangian statistics are the goal of the present technique. The present method tracks a single particle for over $20\Delta t$ to an accuracy of less than a pixel 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 (figure 3). 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.

Accelerations

Tracks of length $\sim 0.75 \tau_\eta$ in time are used to calculate accelerations as in Voth *et al.* (2002). Components of Lagrangian acceleration $\mathbf{a} = (a_x, a_y, a_z)$ are shown in figure 5. 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 that their data was taken in a sparsely-seeded flow. For the present high density synthetic experiment (ppp = 0.05), we find the tails to be considerably narrower than that suggested by Voth *et al.* (2002). However wider tails are recovered if the experiment

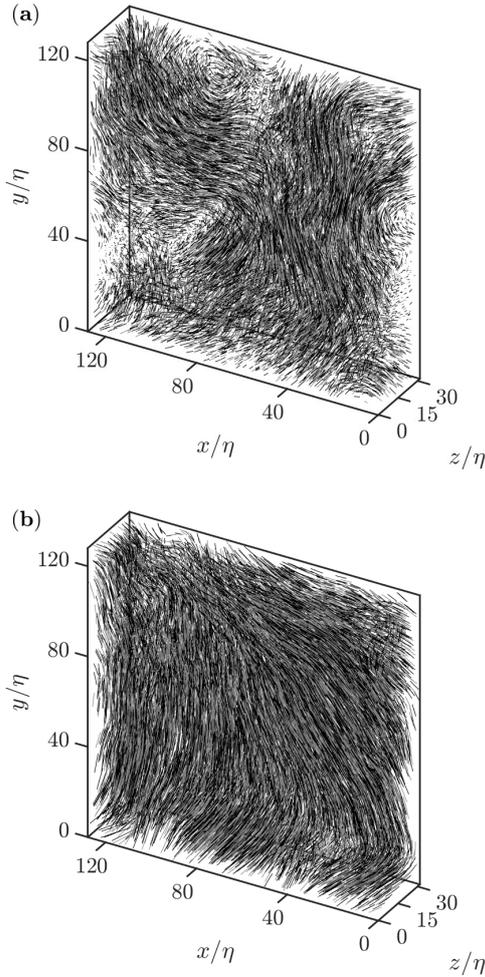


Figure 4. Particle tracks for two independent samples, demonstrating variations in structures for the selected ‘slab’ of turbulence.

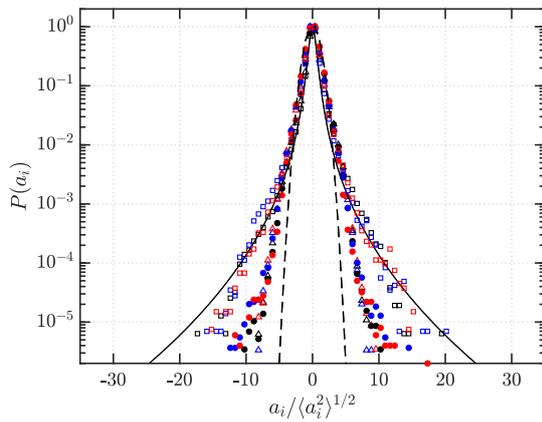


Figure 5. Components of Lagrangian acceleration for the present LPT methodology: \bullet a_x , \bullet a_y , \bullet a_z ; from pseudo-tracking directly within the DNS fields: Δ a_x , Δ a_y , Δ a_z ; repeating the synthetic experiment with a lower seeding density $ppp = 0.005$: \square a_x , \square a_y , \square a_z . Accelerations are computed from tracks of length $\sim 0.75 \tau_\eta$ in time from which measurements are bootstrapped along the lengths to increase statistical convergence; —, fit from Voth *et al.* (2002): $P(a) = C \exp(-a^2 / ((1 + |\alpha\beta/\sigma|^\gamma)\sigma^2))$, where $\beta = 0.539$, $\gamma = 1.588$, and $\sigma = 0.508$; ---, Gaussian distribution with the same standard deviation as $P(a_x)$.

is repeated at a lower density ($ppp = 0.005$). 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, which is much smaller than for the sparse case. The difference in the acceleration PDFs is therefore not the result of a difference in spatial or temporal resolution, but rather exposes an inherent limitation of 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 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. The Kolmogorov constant is given by

$$a_0 = \langle a_i^2 \rangle \frac{v^{1/2}}{\varepsilon^{3/2}}, \quad (4)$$

where ε is the dissipation. Voth *et al.* (2002) examined how this varied with track length. Using the DNS dissipation supplied by authors of the JHTDB, for the present tracks of length $\sim 0.75 \tau_\eta$ we have $(a_0)_x \approx 2.4$ for the high density ($ppp = 0.05$) synthetic experiment and $(a_0)_x \approx 4.8$ for the lower density ($ppp = 0.005$); this lower density value approximately corresponds to that found by Voth *et al.* (2002) for tracks of this length in time at a similar Re_λ .

In general, long tracks formed over a time corresponding to at least $\sim 2 \tau_\eta$ 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 ($a_i = Du_i/Dt$) 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:

$$\frac{Du_i}{Dt} = \frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j}. \quad (5)$$

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 $t \pm 2 \Delta t_{PIV}$. The average inter-particle distance in the present fields (at the time of initial seeding) is approximately 1.3η , whereas for Lüthi *et al.* (2005) it was around 4.2η , and Hoyer *et al.* (2005) used particle fields with average particle separations of 2.4η . As detailed by Hoyer *et al.* (2005), relation (5) is a strict test of spatial resolution. In figure 6 we find very high correlation between the left- and right-hand sides of (5) for the a_x component, confirming the adequate spatial resolution of the present synthetic experiment. Similar results are found for the other two spatial dimensions.

CONCLUSIONS

We have introduced a novel method for Lagrangian particle tracking based on a scanning laser technique. Parti-

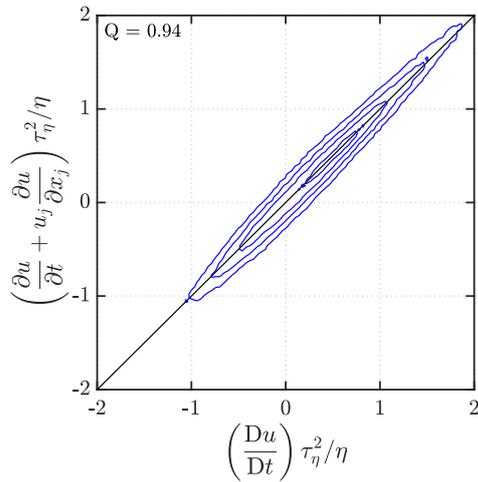


Figure 6. Acceleration check relating Lagrangian, local and convective accelerations. The high correlation coefficient (value of Q) between the two sides of (5) demonstrates good spatial resolution.

cles 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. Combined with the use of residual images, the new method is able to accurately triangulate a large fraction of true particle locations for very high particle densities, up to the equivalent of $\text{ppp} = 0.125$, where typical volumetric triangulation typically suffers from large fractions of falsely detected particles when the particle image density approaches $\text{ppp} = 0.005$. Positional error is only mildly increased in the presence of significant 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 limitation in the magnitude of accelerations measurable was found for very dense fields 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 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.

REFERENCES

- Brücker, Ch. 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, C. J. 2013 Higher order multi-frame particle tracking velocimetry. *Exp. Fluids* **54**, 1533.
- 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*.
- 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, A. N., Lawson, J. M., Dawson, J. R. & Worth, N. A. 2017 A laser sheet self-calibration method for scanning PIV. *Exp. Fluids* **58**, 145.
- Lawson, J. M. & Dawson, J. R. 2014 A scanning PIV method for fine-scale turbulence measurements. *Exp. Fluids* **55**, 1857.
- Lecordier, B. & Westerweel, J. 2004 The EUROPIV synthetic image generator (SIG). In *Particle image velocimetry: recent improvements*, pp. 145–161. Springer.
- 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**, 31.
- Lüthi, B., Tsinober, A. & Kinzelbach, W. 2005 Lagrangian measurement of vorticity dynamics in turbulent flow. *J. Fluid Mech.* **528**, 87–118.
- Maas, H. G., Gruen, A. & Papantoniou, D. 1993 Particle tracking velocimetry in three-dimensional flows. *Exp. Fluids* **15**, 133–146.
- Malik, N. A., Dracos, T. & Papantoniou, D. A. 1993 Particle tracking velocimetry in three-dimensional flows. *Exp. Fluids* **15**, 279–294.
- Ouellette, N. T., 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. 2016 Shake-The-Box: Lagrangian particle tracking at high particle image densities. *Exp. Fluids* **57**, 70.
- Scharnowski, S. & Kähler, C. J. 2016 Estimation and optimization of loss-of-pair uncertainties based on PIV correlation functions. *Exp. Fluids* **57**, 23.
- Voth, G. A., La Porta, A., Crawford, A. M., Alexander, J. & Bodenschatz, E. 2002 Measurement of particle accelerations in fully developed turbulence. *J. Fluid Mech.* **469**, 121–160.