

Shake-The-Box: Lagrangian particle tracking at high particle image densities

Daniel Schanz¹ · Sebastian Gesemann¹ · Andreas Schröder¹

Received: 2 November 2015 / Revised: 31 March 2016 / Accepted: 2 April 2016 / Published online: 27 April 2016
© Springer-Verlag Berlin Heidelberg 2016

Abstract A Lagrangian tracking method is introduced, which uses a prediction of the particle distribution for the subsequent time-step as a mean to seize the temporal domain. Errors introduced by the prediction process are corrected by an image matching technique (‘shaking’ the particle in space), followed by an iterative triangulation of particles newly entering the measurement domain. The scheme was termed ‘Shake-The-Box’ and previously characterized as ‘4D-PTV’ due to the strong interaction with the temporal dimension. Trajectories of tracer particles are identified at high spatial accuracy due to a nearly complete suppression of ghost particles; a temporal filtering scheme further improves on accuracy and allows for the extraction of local velocity and acceleration as derivatives of a continuous function. Exploiting the temporal information enables the processing of densely seeded flows (beyond 0.1 particles per pixel, ppp), which were previously reserved for tomographic PIV evaluations. While TOMO-PIV uses statistical means to evaluate the flow (building an ‘anonymous’ voxel space with subsequent spatial averaging of the velocity information using correlation), the Shake-The-Box approach is able to identify and track individual particles at numbers of tens or even hundreds of thousands per time-step. The method is outlined in detail, followed by descriptions of applications to synthetic and experimental data. The synthetic data evaluation reveals that STB is able to capture virtually all true particles, while effectively suppressing the formation of ghost particles. For the examined four-camera set-up particle image densities N_1 up to

0.125 ppp could be processed. For noise-free images, the attained accuracy is very high. The addition of synthetic noise reduces usable particle image density ($N_1 \leq 0.075$ ppp for highly noisy images) and accuracy (still being significantly higher compared to tomographic reconstruction). The solutions remain virtually free of ghost particles. Processing an experimental data set on a transitional jet in water demonstrates the benefits of advanced Lagrangian evaluation in describing flow details—both on small scales (by the individual tracks) and on larger structures (using an interpolation onto an Eulerian grid). Comparisons to standard TOMO-PIV processing for synthetic and experimental evaluations show distinct benefits in local accuracy, completeness of the solution, ghost particle occurrence, spatial resolution, temporal coherence and computational effort.

1 Introduction and motivation

Lagrangian particle tracking (LPT) signifies the tracking of individual tracer particles in a three-dimensional volume, typically following a flow which is being sampled by the seeding particles. Depending on the methods used for extraction of the particle tracks, it is often referred to as 3D particle tracking velocimetry (3D PTV) or, more recently, tomographic PTV.

The Shake-The-Box (STB) method represents an advanced particle tracking scheme that incorporates the recent advancements of both 3D PTV and tomographic particle image velocimetry (TOMO-PIV, Elsinga et al. 2006a; Scarano 2013) and complements these methods with an extensive use of the temporal information contained in time-resolved data sets. STB was first introduced in 2013 (Schanz et al. 2013b) and gradually improved in performance since then. It features a considerable improvement

✉ Daniel Schanz
daniel.schanz@dlr.de

¹ German Aerospace Center (DLR), Institute of Aerodynamics and Flow Technology, Göttingen, Germany

compared to previous methods in both accuracy (in relation to both 3D PTV and TOMO-PIV) and the applicable particle image densities N_I (especially in relation to 3D PTV). In order to classify the methods used in STB, both TOMO-PIV and 3D PTV will be briefly introduced in this section; their advantages and drawbacks are discussed, followed by an overview of the methods adapted by STB.

1.1 Tomographic PIV

Since its introduction by Elsinga et al. (2006a), TOMO-PIV has been rapidly accepted as a reliable tool for 3D flow measurements. Applications range from (often spatially highly resolved) two-pulse measurements in water (e.g. Hain et al. 2008; Scarano and Poelma 2009) and air (e.g. Elsinga et al. 2006a; Schanz et al. 2012; Henningsson et al. 2015) to time-resolved measurements in water (e.g. Schröder et al. 2011; Violato et al. 2011; Schröder et al. 2015a) and air (e.g. Schröder et al. 2008; Ghaemi and Scarano 2011), just naming a few. Like nearly all three-dimensional measurement techniques, TOMO-PIV deduces the spatial position of particle tracers from projections on multiple two-dimensional camera images. An iterative approach to the reconstruction allows for relatively high particle image densities (typically around 0.05 particles per pixel, ppp), using algorithms like MART (Herman and Lent 1976) or SMART (Atkinson and Soria 2009). Particles are reconstructed as intensity peaks in a voxel space. 3D cross-correlation is applied after the reconstruction process of subsequent time-steps, ensuring a robust deduction of velocity information. Performing a particle-based correction of the calibration function (volume self-calibration, Wieneke 2007) reduces calibration errors from typically 1–2 pixels (px) down to below 0.1 px. Improved particle reconstruction and accuracy are attained and higher seeding concentrations can be processed. Calibrating the 3D position-dependent particle image shape (optical transfer function, OTF, Schanz et al. 2013a) further increases accuracy and reduces the occurrence of ghost particles (ambiguities in the reconstruction problem).

1.2 TOMO-PIV and temporal information

When dealing with time-resolved data, information within the temporal domain can be used in order to improve the quality of each single time-step by seizing the different views on the (virtually) identical flow, provided by successive time-steps (Elsinga and Tokgoz 2014). Motion tracking enhanced MART (MTE-MART, Novara et al. 2010) applies such a concept in the reconstruction step by combing reconstruction results from multiple time-steps, using the velocity field to deform the particle fields, i.e. voxel spaces; fluid trajectory correlation (FTC, Lynch and Scarano 2013)

employs advanced nonlinear multi-frame window deformation in combination with a scheme to extract pseudo-Lagrangian trajectories of fluid parcels. Both methods work successfully in enhancing the quality of the reconstruction (MTE) and the correlation (FTC)—albeit each at a high computational cost due to the need of repeatedly processing each time-step within several iterations.

Very recently, an extension to the MTE method was introduced, termed sequential MTE (SMTE, Lynch and Scarano 2015). This method adapts the general ideas of MTE and combines it with an STB like approach of predicting particle positions in a future time-step (by computing an enhanced guess of the voxel intensity distribution for the subsequent time-step, followed by some iterations of MART to refine this guess). The authors demonstrate an effective suppression of ghost particles and a reduction in computation time compared to MTE as no iterative feedback between reconstruction and correlation is needed. Reconstruction speed and memory requirements are comparable to normal TOMO-PIV processing.

1.3 Limitations of TOMO-PIV

Processing robustness, along with the developed measures to increase calibration accuracy, explains the huge success of TOMO-PIV since its introduction; however, some drawbacks are associated with the technique: ghost particles have an influence on the velocity vector result, especially when using high particle image densities. Cross-correlation applies spatial averages over interrogation volumes and therefore smooths out velocity gradients and fine flow structures. As an extreme example, Atkinson et al. (2011) investigated a turbulent boundary layer and found average velocity errors of up to 1.5 px/time-step within the region of high shear close to the wall. This effect can be slightly lessened by the use of Gaussian (Discetti and Astarita 2012) or adaptive (Novara et al. 2013) weighting in the correlation process, but cannot be fully compensated.

The discretization of the particle representation using a voxel space introduces unavoidable errors in the particle position accuracy, which are found to be in the order of 0.15–0.2 px even for perfect synthetic data (e.g. Wieneke 2013). Under experimental circumstances, other factors—such as image noise, illuminations problems or unfavourable optical access, can influence the error in particle positioning. The correlation process reduces random errors to a certain degree (by averaging over multiple particles). From various experimental investigations, average velocity errors ranging from 0.2 to 0.5 px (Elsinga et al. 2006b; Atkinson et al. 2011; Lynch and Scarano 2014) have been reported. With a typical maximum particle displacement of 10–20 px, the dynamic velocity range (DVR, Adrian 1997) gets reduced to 20–40 for a typical experiment. Other, more

general, downsides are the computational time and the large amounts of data that need to be kept in main memory and (at least temporarily) saved to hard disk.

1.4 Lagrangian particle tracking

The limitations in accuracy, as well as the computational considerations of TOMO-PIV, show that it is desirable to gain direct knowledge of particle positions in space. Tracking such particles in time and applying a temporal filter enables locally precise velocity and acceleration determination without the need of spatial averaging. Lagrangian particle statistics, such as pair dispersion studies (e.g. Bourgoin et al. 2006), become possible. As the number of variables is dramatically reduced (number of particles vs. number of voxels), the computational cost and memory requirements decrease. Three-dimensional particle tracking velocimetry (3D PTV, Nishino et al. 1989; Maas et al. 1993; Malik et al. 1993) has been applied for over 25 years. Particle positions are deduced from projections on few images by triangulation (typically using epipolar lines) for each time-step; matching particles in successive time-steps are searched either in image or on world space. This technique has been successfully used for fundamental research, e.g. to investigate acceleration statistics (La Porta et al. 2001), velocity statistics (Xu et al. 2006) or vorticity dynamics (Lüthi et al. 2005) by examining (few) particle tracks over very long examination times. The downside of the approach is the limitation of the triangulation process in particle image density. Only particle numbers an order of magnitude lower compared to TOMO-PIV can be processed (around 0.005 ppp), as for higher seeding concentrations the occurrence of ghost particles becomes dominant. These are interfering with particle position accuracy and the following tracking process. Overlapping particle images (for 0.05 ppp and a particle image diameter of 2.5 px approximately twenty per cent of the particles are overlapping, Cierpka et al. 2013) tend to shift the 2D peak identification, resulting in high positional errors.

1.5 Alternative particle detection methods

The seeding concentration limitations of the 3D PTV technique led to attempts of extracting particle tracks from tomographic reconstructions, using a 3D Gaussian peak fitter (Schröder et al. 2011; Novara and Scarano 2013) and very recently the introduction of hybrid algorithms (Cornic et al. 2015), relying on an initial tomographic reconstruction and a subsequent particle discretization using heuristical considerations (such as the number of expected particles).

The method of ‘Iterative reconstruction of Volumetric Particle Distribution’ (IPR), introduced by Wieneke (2013)

was the first purely particle position-based method to alleviate the problem of limited particle image density (up to 0.05 ppp). An iterative approach of particle triangulation is applied, combined with an image matching technique to enhance accuracy (realized by moving—‘shaking’—the particle around in 3D space until the local residual is minimized. Wieneke (2013) demonstrates increased position accuracy compared to TOMO-PIV, but still the problem of ghost particles remains—whose number is rapidly increasing as soon as the particle image density approaches 0.05 ppp.

Another approach of directly determining particle distributions, using a marked point process (Ben-Salah et al. 2015), recently surfaced and shows promising first results.

1.6 Shake-The-Box (STB)

Most mentioned methods of reconstructing 3D particle distributions rely on an individual treatment of every single snapshot of the particle distribution (with the exception of SMTE and, to a lesser extent, MTE). For TOMO-PIV, a tomographic reconstruction of every time-step is performed, with a subsequent correlation of two consecutive voxel spaces. The 3D PTV and IPR methods compute particle distributions from scratch for every snapshot, with IPR requiring many iterations until converging to the solution. If the data at hand are sufficiently time resolved though, such approaches neglect the possibility of utilizing already processed data to extract a priori knowledge on the currently processed step.

Shake-The-Box was specifically designed to incorporate as much temporal and spatial information as possible. It combines the progresses in the different fields—taking the calibration methods (volume self-calibration and OTF calibration) of TOMO-PIV development and the iterative triangulation, as well as the image matching (shaking), introduced by IPR. As an additional key step, the temporal domain is exploited by predicting the particle distribution in each subsequent time-step via extrapolation of known trajectories. This predicted particle distribution is used as an initialization to an extended IPR process, which first corrects for errors in the prediction and only in a second step identifies new particles that are not tracked as of now. The result is a method, which allows fast processing of three-dimensional data with high particle concentrations, while capturing the vast majority of true particles and creating virtually no ghost particles.

STB basically reverses the typical process of evaluation: the tracking process precedes the reconstruction process. Instead of first determining particle distributions, followed by a deduction of the velocity (by means of correlation or pair identification), STB uses the available velocity information to create an (estimated) particle distribution. The

errors introduced by the estimation are small enough to be easily corrected using image matching; no further partner search is required for known particles. By this, the flow physics itself is supporting the reconstruction process of particle trajectories.

Extrapolation of known trajectories has been applied in 2D PTV (e.g. Dalziel 1992; Bastiaans et al. 2002) and 3D PTV (e.g. Willneff 2003) before, however, mainly to reduce the search radius for finding connected particles either in image or in object space. Willneff (2003) demonstrated the potential of prediction of particle positions to improve tracking results in terms of completeness of the found tracks. However, this implementation still showed severe limitations in particle concentration ($N_1 \leq 0.005$ ppp for experimental data).

1.7 Features of STB

The prediction step allows for a severe reduction in iterations and therefore processing time, as well as significant gains in accuracy due to an efficient suppression of ghost particles. As soon as a sufficient percentage of tracks can be identified, the process is self-stabilizing; an advance in one reconstruction property (position accuracy, ghost particle ratio, percentage of found tracks) induces advances in the others.

The dense particle trajectory fields obtained by STB can be evaluated in different ways, producing, e.g. accurate Lagrangian velocity and acceleration statistics (similar to 3D PTV, but at much higher particle image densities, see, for example, Schröder et al. 2015a) or highly resolved profiles by slicing the measurement volume into small bins and averaging over all particles of the data set located within such a bin (Kasagi and Nishino 1990; Schröder et al. 2015b, Discetti et al. 2015). The superiority of this approach to correlation-based techniques has been shown in detail for the 2D case (Kähler et al. 2012a, b). To allow for the evaluation of spatial derivatives, a B-spline-based interpolation method ('FlowFit') is introduced in Sect. 4.3. This scheme allows for the interpolation of velocity (and acceleration) values given at discrete particle positions on arbitrary Eulerian grids under several physical constraints. The FlowFit method is designed to translate a maximum of information from the locally very accurate particle data to the interpolated volumes.

1.8 Outline of the manuscript

The general working principle of STB is laid out in Sect. 2, while supplementary topics are given in Sect. 4, where the application to experimental data is illustrated. Section 3 gives quantitative values of the achievable accuracy at different particle image densities and noise levels, extracted

from the results of applying STB on synthetic particle image data.

2 The Shake-The-Box method

The basic concept of the STB method relies on two assumptions: (1) particles within the measurement volume do not disappear, and (2) the knowledge of a particle trajectory enables a fairly accurate estimation of the particles 3D position in the next time-step. The latter can be experimentally ensured by balancing the sampling rate with the Kolmogorov timescale, or the maximum expected acceleration values.

Assuming that the trajectories of (nearly) all particles within the system are known for a certain number of time-steps t_n , the STB scheme for the single time-step t_{n+1} is as follows:

1. Perform a Fit to the last k positions of tracked particles using an optimal Wiener filter—see Sect. 2.2.1.
2. Predict the position of the particle in t_{n+1} by evaluating the Wiener filter coefficients—see Sect. 2.2.1.
3. Shake the particles to their correct position and intensity, eliminating the prediction error—see Sects. 2.2.2 and 2.2.3.
4. Find new particles, entering the measurement domain, on the residual images—see Sect. 2.2.5.
5. Shake all particles again to correct for residual errors.
6. Remove particles if leaving the volume or if intensity falls below a certain threshold—see Sect. 2.2.4.
7. Iterate steps 4, 5 and 6, if necessary.
8. Add new tracks for all new particles identified within four consecutive time-steps—see Sect. 2.2.6.

After such a processing of a single time-step, the known particle tracks have been accurately extended to the current time-step and new particle tracks have been added, capturing particles entering the volume. The entirety of these tracks can now be predicted for t_{n+2} and the process starts anew. This way, STB can work its way through an entire time series, consisting of possibly thousands of images. The effort needed for every single time-step is low, as the system is largely presolved after the prediction step and only minor deviations have to be corrected.

However, as the knowledge of a vast majority of particle tracks is not a given (at the beginning the method has to start from scratch), the evaluation of a data set has to converge to such a stable solution. The progress of the algorithm can be described in three main phases: *Initialization* (trying to find as many particle tracks as possible within a few time-steps), *Convergence* (the complexity of the reconstruction problem is gradually reduced by identifying more

and more true tracks) and *Converged State* (all true tracks are known, and the number of newly found particles is balanced by the number of particles leaving the volume). The three phases will be discussed in detail, while simultaneously describing the main elements of the STB algorithm within the appropriate context.

2.1 Initialization phase

For the first images of every data set, no track information is known a priori. Therefore, a system has to be established that is able to identify enough correct particle tracks within these first images in order to allow the method to reach the convergence phase.

To this end, two measures are conducted for the first n_{init} time-steps: (1) the effort in identifying particles within these initialization time-steps is extended, and (2) optionally, a predictor field is used to limit the amount of falsely detected tracks. Particle identification is typically tackled using iterative triangulation (extended IPR—see Sect. 2.2 for a basic description of the method). Other approaches such as normal triangulation (for low particle image density) or peak determination in tomographic reconstructions (for high particle image density, Schröder et al. 2011) are also feasible.

The identified particle positions will be called *particle candidates*, based on the idea that only such particles for which a track can be identified are considered as true, reconstructed particles. All other—currently untracked—particle candidates are potential ghost particles. Typically, the initialization is applied to the first four time-steps ($n_{\text{init}} = 4$); this number was chosen by experience and may be varied.

Following the particle identification, coherent trajectories need to be extracted from the distributions of particle candidates for the first n_{init} time-steps. A wide variety of methods can be applied for this purpose. The approach used in the current implementation is a relatively simple system, suitable for the identification of short tracks (as only tracks of length 4 need to be identified). It relies on finding matches in consecutive frames by applying a search radius around either the particle position or a predictor location. The details of this method will not be discussed further in this context.

If the flow at the investigated time-steps is—at least roughly—known, a predictor can be constructed, helping with the tracking process. An obvious way to gain a priori knowledge of the flow is to perform a TOMO-PIV evaluation of the first time-step(s). These results can be used to create a predictor for every point in space—allowing for a much smaller search radius, significantly reducing the number of falsely detected tracks. In case the flow is at least partly predictable (e.g. a turbulent boundary layer

with a roughly known velocity profile), an averaged result of previous measurements or even CFD results can be used to derive predictors for the search of particle tracks. In this case, even the search radius could be parameterized, using, for example, the rms value gained from a previous evaluation.

All tracks of length n_{init} that are found by the tracking system are checked for sanity (velocity and acceleration below certain thresholds) and the ones that pass are added to the system of tracked particles.

2.2 Convergence phase

Following the initialization, the gained tracks are used to predict particle positions for the next time-step. Errors induced by the prediction are compensated using image matching techniques, facilitating the reconstruction problem for the yet untracked particles. This way, additional tracks can be found with every time-step, until convergence is reached.

2.2.1 Predicting next position of a tracked particle

The N_T particles being tracked at time-step t_n are extended to time-step t_{n+1} by applying a Wiener filter (Wiener 1949) for extrapolation. The filter parameters are determined based on an estimation of the signal and noise spectra of the particle location signals. That way, the sensitivity of the filter can be adjusted to the experimental conditions, i.e. the influence of noise on the position accuracy. The filter parameters are optimized independently for different track lengths, thereby considering the increased accuracy when more information is available. The found filter coefficients are evaluated at t_{n+1} , and the new (temporary) positions of all tracked particles are set to these extrapolated coordinates.

In earlier versions of the code, producing the results presented in (Schanz et al. 2013b, 2014), polynomials of different length, determined by means of a Savitzky–Golay filter (Savitzky and Golay 1964), were used to fit the data points and extrapolated for the prediction. While this approach works well if the data quality is suited for the chosen polynomial order and length, an adaptation to varying data quality was difficult. Therefore, the Wiener filter approach provides more flexibility, easier application and, in general, a better reliability of the prediction.

2.2.2 Position refinement ('shaking')

The predicted particle positions will be close to the real ones—not more than one or two pixels off, typically only a fraction of a pixel (depending on the flow, the noise level and the temporal sampling). A mean to correct this

error in particle position is to use image matching techniques, which try to (locally) minimize the residual image I_R . One such method was introduced by Wieneke (2013) and relies on moving ('shaking') the particle around in space in small steps, while simultaneously determining the local residual. A brief outline of the (slightly altered) method follows, for a detailed description see (Wieneke 2013).

2.2.2.1 Normal shake The shake-method treats all particles independently and successively. Every particle is moved in small steps in all directions of space, starting from its initial position $[x_I, y_I, z_I]$. The particle is first moved in x -direction to the three positions $x'_1 = x_I - \delta_S$, $x'_2 = x_I$ and $x'_3 = x_I + \delta_S$, with a typical shake width $\delta_S = 0.1$ px. For all positions, the local residual R is calculated using the recorded image I_{rec} , the projected image I_{proj} and the projection of the currently treated particle $I_{\text{Part}[x,y,z,I_P]}$, being localized at position $[x, y, z]$ with intensity I_P :

$$R[x', y', z', I_P] = \left(I_{\text{res}+p} - I_{\text{Part}[x', y', z', I_P]} \right)^2, \quad \text{where} \quad (1)$$

$$I_{\text{res}+p} = I_{\text{res}} + I_{\text{Part}[x, y, z, I_P]} \quad \text{and} \quad I_{\text{res}} = I_{\text{rec}} - I_{\text{proj}}.$$

R is calculated as the sum over all cameras and in a window of selectable size around the projection points. $I_{\text{res}+p}$ expresses the local residual without projecting the currently investigated particle; therefore, $R[x', y', z', I_P]$ describes the residual resulting from repositioning the current particle to $[x', y', z']$. Optimized particle reprojection by application of a calibrated optical transfer function (OTF, Schanz et al. 2013a) is recommended to achieve high accuracy.

The size of the evaluation window can be varied. Using a small window (e.g. 4×4 px) in combination with a central sampling of the OTF (e.g. 2×2 px) for the particle projection yields the most stable shaking results, as noise effects on the low intensity particle image tails are avoided. For large particle images, these values could be increased.

Evaluating R at the positions x'_1 , x'_2 and x'_3 yields three residual values, which are fitted using a polynomial of second order. The extremum $x_{R,\text{ex}}$ of this parabola is determined, and the new particle position $x_{I,\text{new}}$ is set to:

$$x_{I,\text{new}} = \begin{cases} x_I - \delta_S & \text{if } |x_{R,\text{ex}}| > \delta_S \text{ and } R[x'_1, y, z, I_P] > R[x'_3, y, z, I_P] \\ x_{R,\text{min}} & \text{if } x_I - \delta_S < x_{R,\text{min}} < x_I + \delta_S \\ x_I + \delta_S & \text{if } |x_{R,\text{ex}}| > \delta_S \text{ and } R[x'_1, y, z, I_P] < R[x'_3, y, z, I_P] \end{cases} \quad (2)$$

This process is repeated for y - and z -direction, using $[x_{I,\text{new}}, y_I, z_I]$ and $[x_{I,\text{new}}, y_{I,\text{new}}, z_I]$ as initial positions, respectively. If good image data are available, it is possible to reduce δ_S in order to gain even more accuracy (e.g. $\delta_S = 0.025$ px for the last three shake iterations).

2.2.2.2 Initial shake Depending on the temporal resolution of the data sets, the distance covered by the particles can become relatively far. For large separations in the range of 10–30 px, the relative error introduced by the prediction scheme can be 1–2 px (depending on accelerations). In such a case, the normal shake process might converge to a local minimum of the residual that does not coincide with the true particle position. To prevent such a situation, an additional step to roughly place the particle within the neighbourhood of the particle-based 'global' minimum of the residuum is introduced: The particle is moved on a coarse grid (with 0.4–0.8 px separation between the grid points) around the predicted point. The particle is finally put to the point where the smallest residual was found.

2.2.3 Intensity correction

Following each shake iteration, the particle intensity I_P is updated by applying

$$I_{P,\text{new}} = I_P * \sqrt{\frac{\sum_{\text{px}} (I_{\text{res}+p})}{\sum_{\text{px}} (I_{\text{Part}[x_{I,\text{new}}, y_{I,\text{new}}, z_{I,\text{new}}, I_P])}},$$

with both sums running over all pixels of the considered cameras within the chosen window around the projection point.

Taking the root of the intensity ratio proved to dampen intensity oscillations, as does capping the intensity ratio at $3/2$ and $2/3$, respectively. By omitting the camera showing the highest (local) intensity for the current particle, the occurrence of ghost particles (and, especially, ghost tracks) can be reduced: ghost particles often take their energy mainly from a particle peak on one single camera, while the other cameras only show residual peaks or noise. By excluding the (locally) brightest camera from the intensity correction, the intensity of such ghost particles is reduced considerably, leading to a rapid deletion due to the intensity threshold (see next paragraph).

2.2.4 Deleting particles

In case the intensity falls below a specified threshold (e.g. 5 % of the average particle intensity), it is assumed that the particle was lost and the shake process is not converging to the true particle position. In that case, the particle is deleted and the track ends.

Using the steps described above, the positions of all tracked particles are predicted, then roughly repositioned using initial shake and finally iteratively shifted in all directions of space, while constantly updating I_P . Applying an initial shake, followed by five to ten normal shake iterations, is typically sufficient to correct for the prediction errors.

2.2.5 Identifying new particle candidates

Following the prediction and correction of tracked particles, new particles are identified on the residual images of the currently investigated time-step. During the convergence phase, not all particle tracks have been found yet; therefore, the residual images will still show a significant amount of particle images. However, the perceived particle image density will be lower compared to the original image, as particles that are already correctly tracked are removed. The complexity of the reconstruction problem is reduced in relation to the initialization phase, allowing for a reduction in the applied IPR iterations.

New particle candidates are triangulated from the residual images using a low allowed triangulation error (0.5–1.5 px, depending on image quality and seeding concentration). All such that are located within a radius of one pixel around existing particles are discarded. Both the new particle candidates and the tracked particles are then treated by a few iterations of shaking in order to optimize the relative positioning and intensity. New residual images with even less particle images are created. These can be used to start another iteration of triangulation and shaking.

With every iteration, new particle candidates are detected; false particle candidates (ghosts) are taken out by the intensity threshold and the residual is reduced, until the process settles at a stable number of particle candidates for the current image. At this point, it can be beneficial—depending on the particle image density and the number of cameras—to introduce triangulations using a reduced set of cameras. This way particles can be identified whose image is shifted on one camera, e.g. by an overlap situation.

When calculating the residual image before a triangulation iteration, as many traces as possible of the known particle should be subtracted. Therefore, a wide sampling of the calibrated OTF (Schanz et al. 2013a) is applied (opposed to the shake step, where a tight central sampling of the OTF increases stability); a 6×6 to 10×10 pixel grid has proven effective. As laid out in Sects. 2.3 and 4.1, it can be of benefit to multiply the particle intensity with a constant factor $f_{pt} > 1$ in order to ensure a sufficient reduction in the residual on all cameras prior to triangulation iterations.

2.2.6 Adding and deleting tracks

After the complete processing of time-step t_n , a mixture of particle candidates and tracked particles is available. The particle candidates may still contain a significant amount of ghost particles, while the tracked particles should represent true particles. Using the particle candidates of time-steps t_{n-3} to t_n , additional tracks of length 4 are searched. One difference to the initialization phase is the origin of the predictor for the track searching. Here, no predictor field from

cross-correlation is available, but a predictor can be constructed from neighbouring tracked particles. If at least two tracked particles are found within a neighbourhood (e.g. three times the average particle distance), the predictor is calculated as a Gaussian-weighted average of the velocities of these particles. If not enough neighbouring particles are found, a general, larger search radius (e.g. corresponding to the largest expected particle shift) is applied to the position of the particle.

Approved tracks of length 4 are spotted, and the corresponding particles are added to the list of tracked and predicted particles. Tracks leaving the measurement domain are terminated.

The algorithm continues with time-step t_{n+2} , which will again be easier to reconstruct. This process of finding tracks, which in turn facilitate the identification of new ones, will continue until (nearly) all true particles are tracked. At this point, convergence is reached.


2.3 Converged phase

As seen in the previous paragraph, the algorithm needs some time-steps to converge to a stable state, where the number of tracked particles does not change significantly. In this stage, the vast majority of the particles is known and tracked. From there on, most tracks end only when the corresponding particles leave the measurement volume and new tracks are found when particles are entering the interrogation volume. The general processing remains the same as in the convergence phase.

Figure 1 illustrates the main steps of Shake-The-Box for a single time-step in the converged state and their impact on the residual image of a selected camera. At the beginning of the processing of this time-step, the residual image matches the recorded camera image. After predicting the positions of the tracked particles, residuals reflecting the errors of the prediction are visible. New particles, entering the measurement domain from the left and the bottom, appear unaltered in the residual images. After performing some iterations of shaking, the residuals of the tracked particles vanish (nearly) completely—only the new particles remain. These are then tackled by the triangulation/IPR process. Due to the low effective particle image density at this point, only particles with overlapping images remain undetected by the triangulation. Most of these situations can be resolved by successively leaving out single cameras during the triangulation. The end result is a nearly completely blank residual image.

Following the example given in Fig. 1, it can be seen how much the prediction step—thus the inclusion of the temporal information—simplifies the problem of particle position detection. At the point where the first triangulations are performed, the perceived particle image density has drastically

Fig. 1 Schematic description of the Shake-The-Box procedure for one time-step in the converged state by illustrating the effects of the different computation steps on the residual image of one single camera (out of multiple)



decreased, enabling a fast and reliable determination of the previously undetected particles. In the converged state, these are mostly new particles that have entered the measurement domain within the last four time-steps. However, it can also happen the track of a particle is lost. Such events are usually caused by overlapping particle images in more than one camera. The involved particles can be pulled to wrong locations during the shaking process. In this case, the prediction for the next time-step will be compromised and the particle will likely get deleted due to low intensity. Such particles will show up on the residual images of the next time-steps, and the track of the particle has to be picked up again.

A complete elimination of residual will only be seen when using synthetic data. With experimental data, the intensities of the images of a single particle on the different cameras will not be as balanced as in a synthetic case. Even a thorough calibration of the OTF will not be able to fully compensate such effects, as particles do not behave equal: poly-disperse particles have different scattering properties, which will lead to varying intensities in the different cameras; the same is true for ‘potato’-shaped particles (e.g. polyamide seeding particles, which rotate and, depending on orientation, scatter differently). However, even if the residuals of such particles do not vanish within all cameras, they will do so for at least some cameras—as the particle intensity is an average over the image intensities—therefore preventing the particle from being picked up again by the next triangulation.

To counteract the effects of non-vanishing residual, it can be useful to multiply the particle intensity by a constant factor f_{pt} when projecting the particles (only for the triangulations, not for the shaking). By that measure, residual peaks in cameras in which the current particle is imaged brighter than average can be avoided; especially for experimental data, a clear reduction in triangulation effort and ghost particle creation can be achieved, while only slightly reducing convergence speed (see Sect. 4.1.1 for an example). Values of up to $f_{pt} = 2.0$ have been used.

2.4 Outlier removal

It can happen that tracked particles get lost during the tracking process (the reconstructed particle deviates from

the true particle trajectory due to, for example, overlapping particle images and/or noise), but still survive without being deleted due to low intensity (caused, for example, by image noise). However, these particles will most likely show velocities that deviate significantly from the surrounding particles. Hence, an outlier validation of the particle velocity with its neighbours can be carried out: particles are searched within a radius (e.g. four times the average particle distance) around each tracked particle. The 10–50 closest ones are chosen, their velocities averaged (v_{avg}), and the root mean square of the velocities difference in relation to v_{avg} is computed (rms_v), giving a coarse measure of the velocity gradients present. The velocity difference Δ_v of the current particle to v_{avg} is determined, and if $\Delta_v > T_o \cdot rms_v$, the particle is deleted. The outlier threshold factor T_o can be chosen according to the flow and the spatial sampling of the scales. Typical values are 5–15. If not enough particles are found in the vicinity, the particle is left as is.

Another sign of a lost particle can be an erratic trajectory. To detect such behaviour of a particle, a linear fit is applied to the last four time-steps. If the particle shows an average deviation Δ_f above a certain threshold $T_{\Delta f}$, the particle is deleted. $T_{\Delta f}$ should depend on the temporal sampling of the flow scales and the noise present on the images. Typical values for experimental data are 0.7–1.5 px.

2.5 Multi-pass processing

Even though the tracking process of STB is rather robust compared to traditional 3D PTV methods, interrupted or incomplete tracks occur; especially, the first time-steps, when not all particle tracks have been identified yet, exhibit a lot of missing particles. One easy method to improve on this situation is to perform a second pass of STB, running temporally backward through the data set. By this, track fragments might be connected and particle tracks are extended backwards to the time-step when the particle first occurred within the measurement volume. To this end, the tracks identified by the first pass are temporally filtered (see Sect. 4.2). For each time-step, the filtered particles from the previous pass are taken as a predefined particle distribution. In addition, new particle candidates are triangulated,

enabling the search for new tracks. If a track ends (i.e. its starting point in the first pass is reached), the track is extended backwards in time by the usual prediction scheme as long as the particle stays within the measurement domain or it is lost due to either the intensity threshold or the other validations discussed above. Section 4.1 demonstrates the effects of a second pass on an experimental data set; more passes can be performed, if necessary.

3 Synthetic data validation

In order to quantify the features of the STB algorithm, the concept is applied to synthetic data sets of varying particle image density and noise level. The current section will first describe the creation of the synthetic data, followed by the results of STB reconstructions and a comparison to tomographic PIV.

3.1 Creation of synthetic tracks

A process to extract synthetic tracks from a known velocity volume was designed. For the first image, particles are randomly distributed in the selected domain; the velocity of the particles is calculated as the Gaussian-weighted average of the eight neighbouring velocity vectors of the source vector volume. In the next step, the particles are moved, using a simple Euler scheme, according to the determined velocity and the chosen time separation. To ensure smooth acceleration along the trajectory, a temporal filter was applied to the tracks. Thus, smooth particle tracks following the source vectors with first-order accuracy can be created.

To serve as source vector volume, a result of an experimental data set of the flow behind a series of periodic hills (Schröder et al. 2015a) was used. A sub-volume of $1000 \times 1000 \times 400$ voxels was taken from the middle of the volume and the vector results originating from this volume were used as source for the track creation. For the placement of reinserted particles, a buffer of 30 non-imaged

voxels was left at each interface of the volume, so that an imaged volume of $940 \times 940 \times 340$ voxels remained.

As time separation, the original sampling rate was chosen, resulting in a mean 3D particle displacement of around 6 px and a maximum displacement of around 11 px. The particle positions determined from the track creation scheme are projected (parallel projection) onto four virtual cameras in pyramidal configuration (with a square basis and an angle of $\pm 30^\circ$ in x - and y -direction, 1200×1200 pixels each, 16-bit integer). For particle imaging (OTF), a two-dimensional Gaussian peak is used with a particle diameter of around 2.4 px (intensity fall-off to e^{-2}). The average peak height is 6500 counts. Particle image densities ranging from $N_1 = 0.01$ – 0.125 ppp (calculated with respect to the imaged volume—in this case $N_1 = N_p/(940 \times 940)$ with N_p : number of particles) were realized. Figure 2 shows excerpts of a virtual camera image for three particle image densities.

3.2 Application of STB to synthetic images

The Shake-The-Box scheme was applied to the created image time series. The used parameters are given in Table 1. Time series of 50 images were processed for each particle image density. Figure 3, 4, 5 and 6 illustrate the temporal development over these images of four parameters, describing the quality of the reconstructions:

F_u , the fraction of undetected true particles; the reconstructed tracks are compared to the original ones by searching for reconstructed particles in a radius of 1 px around every true particle. If no reconstructed particle is found within this radius, the particle is registered as undetected (see Fig. 3).

$F_{g(tr)}$, the fraction of *tracked* ghost particles, calculated by searching around all reconstructed particles that are part of a track in a radius of 1 px within the source volumes; if no true particle is found, the reconstructed one is counted as a ghost particle (see Fig. 4).

$F_{g(tot)}$, the fraction of *total* ghost particles, calculated by searching around all reconstructed particles (tracked ones, as well as all particle candidates, reconstructed by

Fig. 2 Details of virtual camera image for ppp = 0.01, ppp = 0.05 and ppp = 0.125 (from left to right)




Table 1 STB parameters applied to synthetic image data from time-step 5 onwards

Triangulation iterations using N_{cam}	$m = 3$ (6)
Triangulation iterations using N_{cam}^{-1}	$n = 2$ (4)
Number of shake iterations	$k = 8$
Shake width	$\delta_s = 0.1-0.025$ px
Number of Initial shake iterations	1
Allowed triangulation error [px]	$\varepsilon = 0.5$ px
Search radius without predictor [px]	12 px
Search radius with predictor [px]	4 px
Particle intensity threshold	$T_{int} = 0.1 \cdot I_{avg}$
Smoothness threshold	$T_{\Delta f} = 0.4$ px (0.8 px)
Outlier detection threshold (T_O)	–
Projection factor prior to triangulation	$f_{pt} = 1.0$
Predictor for initialization	TOMO-PIV result
Number of passes	1

Values for time-step 1–4 (initialization phase) are given in brackets

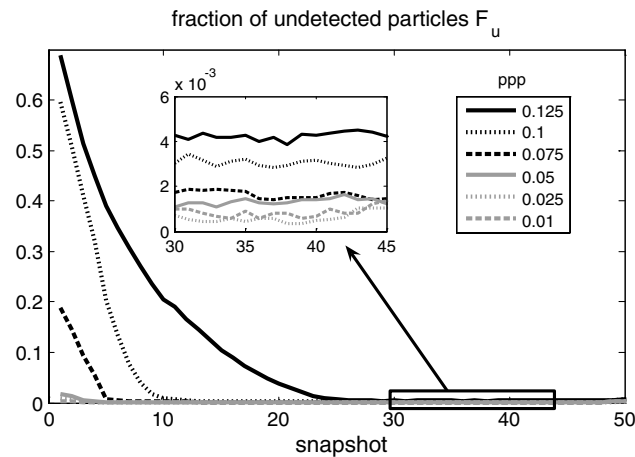


Fig. 3 Number of *undetected particles* relative to the true particle number

triangulation/IPR) in a radius of 1 px within the source volumes; if no true particle is found, the reconstructed one is counted as a ghost particle (see Fig. 5).

Δ_p , the arithmetic mean error of the detected true particles; the magnitude of deviation of the position of the reconstructed particles \vec{x}_R to the original position \vec{x}_S is determined in pixel units and averaged: $\Delta_p = \frac{1}{N} \sum_1^N |\vec{x}_R - \vec{x}_S|$ (see Fig. 6).

All four plots document the temporal convergence of the STB method (especially for high particle image densities). A strong reduction in undetected particles, tracked and untracked ghost particles and positional error with progression of time can be seen. In order to document the three main stages of the STB method—introduced in Sect. 2—the plots are first examined for the initialization phase

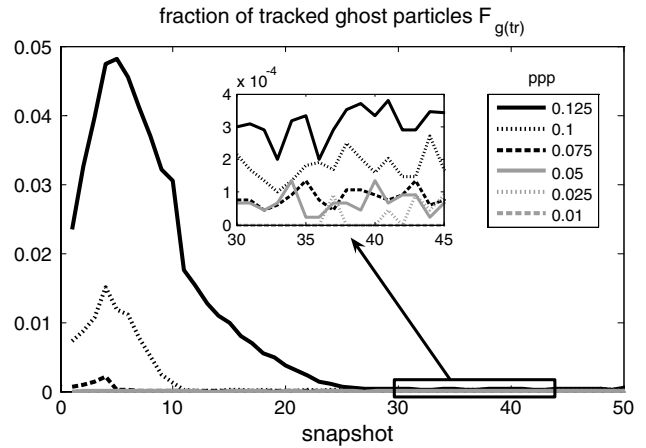


Fig. 4 Number of *ghost particles* within the system of *tracked particles* relative to the number of true particles




Fig. 5 *Total number of reconstructed ghost particles* (tracked and untracked) relative to the number of true particles

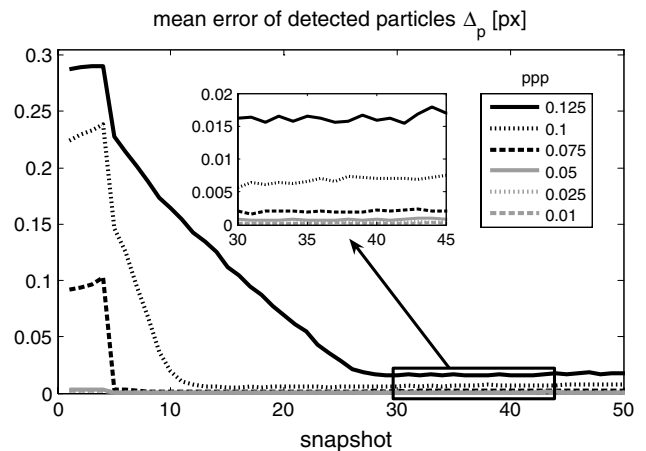


Fig. 6 *3D positional error* in pixel, averaged over all tracked particles relative to the true particle position

Table 2 Comparison of tomographic reconstruction (MLOS-SMART) and subsequent particle peak identification to tracking results of STB

Particle image density N_1	[ppp]	0.125	0.1	0.075	0.05	0.025	0.01
Real particles		110,868	88,584	66,439	44,233	22,060	8904
Undetected particles	SMART	9387 8.48 %	3400 3.83 %	1119 1.68 %	405 0.91 %	135 0.61 %	47 0.53 %
	STB	483 0.43 %	268 0.30 %	104 0.16 %	63 0.14 %	18 0.06 %	9 0.10 %
Ghost particles	SMART	282,090 254.8 %	206,990 233.7 %	121,680 183.1 %	42,229 95.5 %	3721 16.8 %	134 1.5 %
	STB	36.5 0.033 %	16.1 0.018 %	5.8 0.008 %	3.5 0.008 %	1 0.005 %	0 0.0 %
Avg. position error Δ_p [px]	SMART	0.308	0.278	0.243	0.201	0.155	0.135
	STB	0.0177	0.0076	0.0023	0.0008	0.0005	0.0002

Values averaged over images 40–44 of the time series discussed in Sect. 3.1

(time-steps one to four), followed by a description of the convergence phase and finally the converged state. Table 2 summarizes the results for the converged state.

Please note that the results for the initialization can be identified by looking at the very first data point of each plot (t_1). The following data points (t_2, t_3) already contain additional tracks that were found at t_5 and t_6 , respectively, as the track identification always reaches four time-steps in the past.

Looking at the fraction of undetected particles (Fig. 3), it can be seen that for lower particle image densities ($N_1 \leq 0.05$ ppp), the track initialization is very effective in finding nearly all true particles: For 0.05 ppp, only 2 % of the true particles are not found after the first four time-steps. When using higher particle image densities, this value quickly rises (19 % for 0.075 ppp; 69 % for 0.125 ppp). More and more particle images overlap, preventing an accurate determination of the peak position on the 2D images. Therefore, many particles cannot be successfully triangulated due to the small allowed triangulation error of $\varepsilon = 0.5$ px. As an additional constraint, only those particles that were successfully triangulated in all four time-steps of the initialization can be successfully tracked.

The fraction of ghost particles within the tracked particles $F_{g(tr)}$ (Fig. 4) is very low already immediately after the initialization, even for the higher particle image densities (around 3–5 % for 0.125 ppp). However, the absolute number of ghost particles (see Fig. 5) is high for the initialization time-steps and particle image densities above 0.075 ppp (around four out of five particle candidates are ghost particles at 0.125 ppp within the initialization, albeit most of low intensity). These post-initialization results basically constitute pure IPR results and coincide well with the findings of Wieneke (2013).

The ratio of tracked versus total ghost particles clearly shows that (for the investigated case) the occurrence of ghost particles is quickly decorrelating with time. As shown by Novara et al. (2010) and Elsinga et al. (2011),

the decorrelation of ghost particles is strongly dependent on the flow situation. For uniform flow, ghost particles do not decorrelate at all. For a given experiment, the flow can be assessed beforehand using the formulas given in (Elsinga et al. 2011) to what extent a pairing of ghost particle is expected. However, in general the temporal domain allows an efficient separation of real and ghost particles, even for short time series of four images. For $N_1 \leq 0.05$ ppp, ghost particles are virtually non-existent (<0.1 % tracked ghosts, 6.3 % absolute ghosts after the initialization). At these particle image densities, four time-steps are sufficient to eliminate virtually all ambiguities.

For low particle image densities ($N_1 \leq 0.05$ ppp), the positional error Δ_p (see Fig. 6) is very low already after the initialization ($\Delta_p \approx 0.0035$ px for 0.05 ppp), while for higher particle image densities much higher errors are seen ($\Delta_p \approx 0.3$ px for 0.125 ppp). These values document the effect of ghost particles on the accuracy of the true particles. Ghost particles draw energy from the particle images, making the projected image a sum of the projections of real and ghost particles—which causes the true particle to shift in space to better fit the deformed image. For these accuracy considerations, the total number of ghost particles $F_{g(tot)}$ is crucial. In case of low particle image densities, ghost particles are rare—consequently, the accuracy of the true particles is high. When looking at the curve of Δ_p for $N_1 = 0.075$ ppp, this effect is plainly visible as the sharp increase in accuracy after time-step 4. For the initialization steps, the temporal information is not yet used in the reconstruction; therefore, the results resemble those of non-time-resolved evaluations (around 33 % of the triangulated particles are ghost particles at $N_1 = 0.075$ ppp, see Fig. 5). After the fourth time-step, the identified tracks (around 81 % of the real particles) are extracted and predicted to the next time-step. After shaking these to their correct positions, the system is largely presolved. The remaining undetected particles can be triangulated from the residuals without

occurrence of ghost particles. The accuracy of the true particles subsequently increases significantly.

When going to higher particle image densities, this process requires more time-steps, as the total number of ghost particles present in the initialization rapidly increases (up to a fraction of 4.3 for $N_I = 0.125$ ppp). The less accurate placement of the true particles leads to a reduced identification of particle tracks. The following time-steps, representing the convergence phase, gradually identify more tracks, as the ones that are already known help to reduce the complexity of the reconstruction problem. Falsely or inaccurately detected tracks from the initialization are thrown out of the tracking system. This explains the notable rise in tracked ghost particles during the first images after the initialization (see Fig. 4). Both $F_{g(\text{tot})}$ and $F_{g(\text{tr})}$ sink rapidly.

The lower particle image densities converge instantaneously after the initialization; convergence takes two time-steps for $N_I = 0.075$ ppp, 7 time-steps for 0.1 ppp and 24 time-steps for 0.125 ppp. It has to be noted that for the 0.125 ppp case, the number of normal triangulation iterations was increased to 4, followed by 3 iterations with reduced camera numbers. Each triangulation is followed by 10 shake iterations. Using the iteration numbers applied to the other cases leads to a slower convergence (around 60 images). Above 0.125 ppp, the algorithm fails to converge with the chosen (simple) initialization: The IPR reconstruction of the first images is highly erroneous with respect to found true particles, number of ghost particles and particle accuracy, such that the algorithm is not able to compensate for these errors in the following steps. For such high particle image densities, different approaches to the initialization have to be attempted—for instance taking the particle peaks from tomographic reconstruction (ideally from advanced methods like SMTE) and/or increasing the number of initialization time-steps.

By time-step 30, all evaluations have reached the converged phase. The insets in Fig. 3, 4, 5 and 6 show an enlarged view of the results within this phase. It can be seen that, except for minor oscillations, the values remain on a constant level, with very low numbers of undetected particles and tracked ghost particles (0.4 and 0.03 %, respectively, for $N_I = 0.125$ ppp). Also the untracked ghost particles introduced by the IPR process are reduced to a minimum, as documented by the total number of ghost particles (0.4 % at 0.125 ppp). Particle peak accuracy is very high, with $\Delta_p < 0.003$ px for $N_I \leq 0.075$ ppp and $\Delta_p \approx 0.016$ px for $N_I = 0.125$ ppp.

The results presented here are slightly different to the ones shown in Schanz et al. (2014), mainly due to some added processing steps: excluding the brightest camera in the intensity update led to a more effective suppression of ghost particles; reducing the shake width to $\delta_s = 0.025$ px

for the last iterations led to a noticeable increase in position accuracy.

3.3 Comparison to tomographic reconstruction

In order to compare the results to the technique commonly used for processing three-dimensional data at high particle concentrations, tomographic reconstructions of the synthetic images were performed. An MLOS-SMART algorithm (Atkinson and Soria 2009) with a 2D B-spline weighting function (OTF) (Schanz et al. 2013a) was applied. As this technique does not utilize temporal information, only five volumes (time-steps 40–44, enabling a direct comparison to STB data in converged state) per seeding concentration were reconstructed. Following the MLOS initialization, five iterations of SMART, each with subsequent volume smoothing (Discetti et al. 2013) and contrast elevation to accelerate convergence, were performed. The used voxel-to-pixel ratio was 1.0, resulting in voxel spaces of dimensions $1000 \times 1000 \times 400$.

In order to compare position accuracy, as well as the fraction of undetected and ghost particles, a 3D Gaussian peak finder from LaVision Davis 8 was used to identify particle positions within the reconstructed volume. The accuracy determination and particle/ghost identification from the original track data were conducted analogous to the STB data. For both MLOS-SMART and STB, the results from steps 40–44 were averaged and are given in Fig. 7 and Table 2, respectively.

Looking at the MLOS-SMART tomographic reconstructions, most true particles are correctly reconstructed (>99 % for low N_I , 92 % for 0.125 ppp) with a positional error that rises from $\Delta_p = 0.13$ px for 0.01 ppp to $\Delta_p = 0.31$ px for 0.125 ppp. The fraction of the summed ghost particle intensity to the summed true particle intensity is low for $N_I \leq 0.025$ ppp, but rises quickly with increasing N_I . For 0.125 ppp, the ghost particles contain more energy than the reconstructed true particles. As given in Table 2, the number of detected ghost particles surpasses the number of true particles at 0.075 ppp; however, the average intensity of a ghost particle is lower than for a true particle; therefore, the intensity fraction is lower (around 0.5 at 0.075 ppp). For 0.125 ppp, over 280.000 ghost particles are found—a factor of 2.5 to the true particle number of around 110.000.

The peak accuracy results for MLOS-SMART are in good agreement with the values given for MART by Wienieke (2013). However, for high N_I both the number of true particles and the number of ghost particles are significantly higher for the MLOS-SMART case. As MART and SMART should produce comparable results, this difference is most likely explained by different thresholds used for the 3D peak detection: A higher value leads to a reduction




Fig. 7 Comparison of results gained by tomographic reconstruction with subsequent particle peak identification to tracking results by STB for varying particle image densities. Values averaged over images 40–44 of the time series discussed in Sect. 3.1

in detected ghost particles, but also reduces the number of correctly found real particles.

Comparing to the results gained by Shake-The-Box, it becomes obvious the inclusion of the temporal information opens a door to results of much higher quality. As discussed in the previous paragraph, the number of correctly identified particles remains above 99.5 %, even for 0.125 ppp. The ghost particle problem is nearly completely resolved—for 0.125 ppp, a ghost particle proportion of 0.03 % is detected. At particle image densities below 0.1 ppp, ghost levels of <0.01 % can be reached. For low particle image densities, the positional error is practically zero, and for 0.075 ppp, it is still below 0.003 px. The error rises for higher particle image densities, but remains below 0.02 px for 0.125 ppp—being an order of magnitude lower compared to SMART (0.31 px).

Wieneke already demonstrated the very high accuracy achievable by the image matching process used in IPR—however in case of single images, this holds only for very low particle image densities (Wieneke 2013). Starting at 0.005 ppp, the error rises and reaches $\Delta_p = 0.1$ px by around 0.03 ppp. For 0.1 ppp, Wieneke finds average errors of around 0.6 px for single-image IPR. He writes: ‘Convergence starts to fail above 0.05 ppp when the solution is no longer unique’. This problem is solved by the inclusion of the temporal domain, as each snapshot provides a new view on the system—essentially adding a new system of cameras, as argued by Novara et al. (2010)—leaving the whole spatio-temporal system only one solution—the real one—to converge to.

The fact that for STB the system is already close to the real solution after the prediction step allows for a low number of used triangulation iterations. Each of these is fast, as the residual images are sparse. This combined effect leads to the computational efficiency documented in Fig. 7d. Reconstruction times for a single snapshot on an eight-core Xeon server (2× Xeon E5520 quad-core CPUs, 24 GB Ram) are compared for the different seeding

concentrations. It can be seen that STB is 4–6 times faster compared to MLOS-SMART. Computation time rises with N_I because of the increasing number of peaks detected on the images (leading to a more complex triangulation process) and the increase in tracked (and shaken) particles. The rise for 0.125 ppp is caused by the increase of used triangulation iterations. For SMART, the percentage on nonzero voxels increases, leading to rising computational effort. Additionally, the reconstructed voxel spaces still need to be processed further (3D cross-correlation or particle peak fitting and partner search), while STB directly yields velocity (and acceleration) data.

Comparing the results to conventional 3D PTV evaluations of synthetic, noise-free data, Ouellette et al. (2006) report position accuracies of around 0.025 px for particle image densities $N_I \leq 0.01$ ppp (STB: 0.0002 px), reflecting the accuracy of the shake process supported by the use of a calibrated optical transfer function. The ability to retain particle tracks quickly diminishes with particle image density for standard PTV approaches. At 0.01 ppp, the different algorithms already show between 6 and 25 % of undetected particles (Ouellette et al. 2006), clearly documenting the limits in seeding concentration for this technique.

3.4 Influence of image noise

As shown in the previous paragraphs, the STB concept yields very accurate results for a wide range of particle image densities when looking at perfect imaging conditions. However, image noise will have an influence on several parts of the algorithm. The triangulation process for identifying new particles will be affected, as noise tends to shift 2D particle position identification. This will directly influence the triangulation error—therefore, the allowed value ϵ has to be altered in order to find a sufficient number of particles. A higher value of ϵ will lead to a higher probability of ghost particle formation. Secondly, the image matching process will not be able to find a perfect match

for the particle position, as the particle image is altered by the noise in relation to the calibrated OTF used for residual determination.

Data sets have been created for three different noise levels. The noise was introduced after imaging the synthetic particle distribution by adding a randomized intensity to every pixel, taken from a normal distribution with variance σ , derived from the average peak intensity of a particle image $I_{p,avg}$ ($\sigma = 0.03 \cdot I_{p,avg}$, $\sigma = 0.1 \cdot I_{p,avg}$ and $\sigma = 0.2 \cdot I_{p,avg}$). Figure 8 shows exemplary excerpts of one camera image for 0.01 ppp. The first two cases can be seen as representative for good to normal experimental circumstances (considering noise levels), while the high-noise case is approaching experiments with poorly controlled conditions (sparse illumination, small tracer particles). Application of any kind of image preprocessing was omitted in order to not introduce further parameters. For the sake of clarity, only two particle image densities (0.01 and 0.05 ppp) were considered, for which all STB runs converged within a time series of 50 images. At the lowest noise level, convergence is still reached up to 0.125 ppp, for $\sigma = 0.1 \cdot I_{p,avg}$ until 0.1 ppp and for the highest noise level until 0.075 ppp (not within the 50 images though).

The created image time series were reconstructed both by STB and by MLOS-SMART. Concerning STB,

the allowed triangulation error was set $\varepsilon = 0.85$ px for the two cases with lower noise and $\varepsilon = 1.1$ px for the high-noise case. Results of the STB track reconstruction are shown in Fig. 9. It can be seen that the convergence time of the algorithm rises with the noise level; especially, the high-noise case with 0.05 ppp illustrates that the system has to work much harder in order to identify true particles and to get rid of ghost particle tracks. For this case, convergence is reached around 20 time-steps after the initialization, with a then constant ratio of undetected particles of around 1.5 %. The 0.05 ppp case with medium noise level converges much quicker (5 time-steps after initialization), while the low-noise case converges instantly. For 0.01 ppp, the lower noise levels converge instantly, but around 5 iterations are needed for the high noise level.

Looking at the mean displacement of the particles, it is obvious that the very high accuracies seen for perfect imaging cannot be reached. For $\sigma = 0.03 \cdot I_{p,avg}$, displacement errors of around 0.03 px are found. The error rises to 0.1 and 0.24 px for the higher noise cases. As soon as the system is converged, the error is largely dependent on the noise level and less on the particle image density—clearly indicating the reduced accuracy of the image matching process as source for the position error.

Fig. 8 Detail view from camera image for $N_1 = 0.01$ ppp for different levels of artificially added noise: $\sigma = 0.03 \cdot I_{p,avg}$, $\sigma = 0.1 \cdot I_{p,avg}$, $\sigma = 0.2 \cdot I_{p,avg}$ (from left to right). σ is the variance of the normal distribution used for the random noise generation; $I_{p,avg}$ denotes the average peak intensity of a particle image





Fig. 9 Temporal development of STB runs for 0.01 and 0.05 ppp at different image noise (parameterized by σ in units of $I_{p,avg}$). **a** Fraction of non-detected particles; **b** particle position accuracy; and **c** fraction of ghost particles within the tracked particles

The level of tracked ghost particles (Fig. 9c) stays very low for the two low-noise cases. For the high-noise case, the ghost level rises slightly in the non-converged state. At this noise level, single noise peaks can be of the same intensity as the particle images. These false particle peaks can lead to random ghost particles in the pure IPR reconstructions, which is plainly visible when looking at the total number (tracked + untracked) of ghost particles, given in Fig. 10. For the high-noise case, fractions of 0.2 (at 0.01 ppp) and around 1.4 (at 0.05 ppp) are seen. The necessity of using a larger search radius of $\epsilon = 1.1$ px (due to peaks being shifted by the noise), combined with the occurrence of high-intensity noise peaks, causes the significant rise in registered ghost particles. While the number of ghost particles is high in the set of particle candidates, it is still unlikely that they are incorporated into a

track, which is demonstrated by the still very low tracked ghost particle fraction—even in the first time-steps of the evaluation. When the tracking system is converging, also the occurrence of triangulated ghost particles is significantly reduced; however, a fraction of 0.08 remains for the 0.05 ppp case, which is notably up from the no-noise case (<0.002). The triangulated (untracked) ghost particles will have no influence on the final result, as they are discarded. However, they can induce inaccuracies in the positioning of the true particles.

Figure 11 compares the converged STB results to reconstructions using MLOS-SMART and a subsequent particle peak identification. The fraction of undetected particles is largely similar between the two techniques for the lower noise cases, while STB shows better results for high-noise images.

Comparing the total ghost particle intensity $F_{Ig(tot)}$ shows—as before—distinct advantages for the STB technique. Maximum values of $F_{Ig(tot)} = 0.08$ for STB are opposed by values of around 2.2 for SMART. For the SMART case, the fraction of ghost particles is actually higher for the low particle image density, as the number of noise peaks does not change, while the number of true particles is reduced.

The mean positional error (Fig. 11b) rises for both STB and SMART at approximately the same rate, when not filtering the STB results (black curves, empty circles). As already discussed, STB does not reach the very high accuracy seen in the previous paragraph, but always holds the accuracy advantage compared to tomographic reconstruction. SMART exhibits an error of around $\Delta_p = 0.39$ px for the high-noise case, while STB shows around 0.24 px.

In order to improve the positional accuracy and to extract velocity and acceleration information from the tracks, temporal filtering is applied. To this end, a third-order B-spline




Fig. 10 Total (tracked and untracked) number of ghost particles relative to the true particle number for different noise levels (σ in units of $I_{p,avg}$) and particle image densities

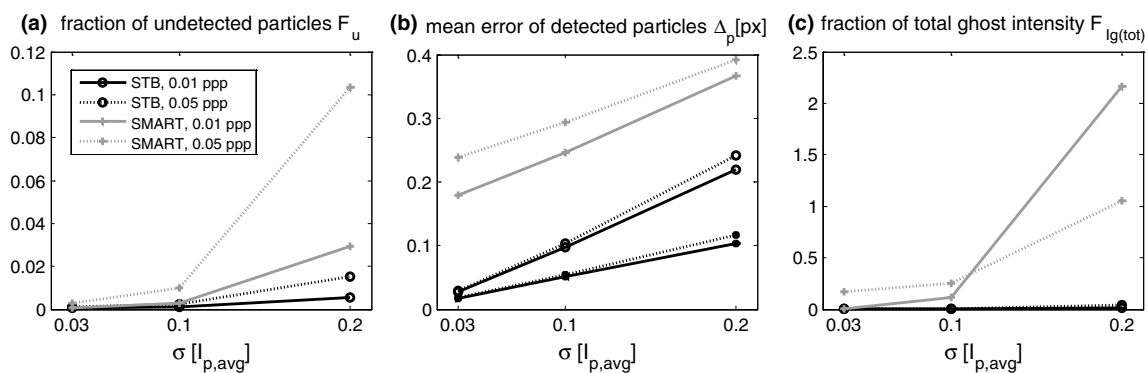


Fig. 11 Comparison of results gained by tomographic reconstruction with subsequent particle peak identification and tracking results by STB for varied amounts of image noise. **a** Fraction of non-detected particles; **b** mean error of detected particles; for STB, two curves are supplied: mean positional error of the raw data (black curves, empty

circle) and mean positional error after applying a temporal fit (black curves, filled circle); and **c** fraction of intensity sum over all ghost particles in relation to the intensity sum of all true particles. Values averaged over images 40–44 of the time series analysed in Fig. 9

function (see Sect. 4.2) is fitted to the particle positions. A basic outlier filter is applied to remove obviously deviating values of velocity and acceleration (occurring either at the end/beginning of tracks or induced by false pairings of source/reconstructed particles; around 0.5 % removal rate). Average positional errors between $\Delta_p = 0.018$ px for $\sigma = 0.03 \cdot I_{p,avg}$ and $\Delta_p = 0.117$ px for $\sigma = 0.2 \cdot I_{p,avg}$ are attained for the filtered tracks at 0.05 ppp (black curves, filled circles).

Table 3 summarizes the results on position accuracy and gives additional values for the accuracy of velocity and acceleration. The velocity error, determined as the arithmetic mean of the difference between reconstructed and original velocity for each particle, is between $\Delta_v = 0.014$ px/time-step for $\sigma = 0.03 \cdot I_{p,avg}$ and $\Delta_v = 0.052$ px/time-step for $\sigma = 0.2 \cdot I_{p,avg}$. Knowing the velocity error allows determining the dynamic velocity range (DVR), as introduced by Adrian (1997). To this end, the errors are expressed by means of root mean square, $rms_v = \sqrt{1/N \sum_1^N (\vec{v}_R - \vec{v}_S)^2}$ and set in relation to the maximum velocity present in the data set (~ 11 px). DVR values ranging between 680 and 170 are found for the different noise levels at 0.05 ppp, while the noise-free case shows DVR in excess of 10,000.

Table 3 Arithmetic mean (Δ) and root mean square (rms) of magnitude of position error (in px), velocity error (in px/time-step) and acceleration error (in px/time-step²) for STB and SMART—reconstructions at 0.05 ppp for different noise levels

Noise level σ	0	0.03 $I_{p,avg}$	0.1 $I_{p,avg}$	0.2 $I_{p,avg}$
SMART pos.				
Δ_p	0.201	0.234	0.290	0.387
rms_p	0.229	0.261	0.322	0.429
STB pos.				
Δ_p	0.0008	0.029	0.103	0.239
rms_p	0.0012	0.033	0.116	0.271
STB filter				
Δ_p	0.0004	0.018	0.054	0.117
rms_p	0.0007	0.020	0.061	0.133
STB vel.				
Δ_v	0.0003	0.014	0.030	0.052
rms_v	0.0006	0.016	0.034	0.065
DVR	17,900	684	319	170
STB acc.				
Δ_a	0.0004	0.019	0.024	0.031
rms_a	0.0008	0.021	0.027	0.039
DAR	1060	43	33	23

Dynamic velocity range (DVR) and dynamic acceleration range (DAR) calculated as ratio of the maximum flow velocity (~ 11 px/time-step) and acceleration (~ 0.9 px/time-step) to the respective rms

Comparing accuracy and DVR to planar PIV, it can be stated that STB delivers *average* velocity accuracies that are at least on par with the ones gained by planar PIV (a typical velocity error of 0.1 px was repeatedly found for planar PIV, see, for example, Nobach and Bodenschatz 2009; Sciacchitano et al. 2013; Wieneke 2015). However, planar PIV suffers from much higher errors in regions of shear or turbulence. Sciacchitano et al. (2013) investigate a transitional jet in water and report errors of around 0.3 px in the shear layer (caused by shear within the correlation windows that even window deformation techniques cannot fully compensate) and up to 0.5 px in the turbulent regions (caused by severe out-of-plane motion inducing loss of pairs). STB does not suffer from both of the observed problems, being a volumetric, particle-based method and therefore yields an increased depiction of highly dynamic regions.

The temporal fit to the particle trajectory additionally yields values for Lagrangian acceleration (material derivative). The accuracy of the gained results was assessed similarly to velocity by computing the root mean square of the acceleration error (rms_a). The different noise levels show values of rms_a between 0.021 and 0.039 px/time-step². Following Adrians definition of DVR (1997), a *dynamic acceleration range* (DAR) is calculated as the ratio of rms_a and the maximum value of acceleration (~ 0.9 px/time-step², taken from the ground truth tracks). This yields values for DAR between 43 and 23 for the different noise cases (see Table 3).

As observed in the previous paragraph, the data reported for STB show improvements over the values given in (Schanz et al. 2014) due to improved processing.

4 Application to experimental data

The STB evaluation scheme has been applied to a time-resolved data set of a transitional jet in a water tank. This set of images originates from a time-resolved TOMO-PIV experiment, which was carried out in 2010 at the water jet facility at TU Delft. The same set-up was used for a number of different analysis, both on the flow phenomena (e.g. Violato et al. 2011, 2012) and on the characterization of methods and tools for 3D measurements (e.g. Novara et al. 2010; Lynch and Scarano 2013). The data set discussed here was recorded within the scope of the work on calibrating and applying the optical transfer function (OTF) on 3D reconstruction problems (Schanz et al. 2013a). For a detailed description of the experimental apparatus, see (Violato et al. 2011).

Figure 12a shows the experimental set-up, which consists of a water jet created by a nozzle with a diameter of $d = 10$ mm and an exit velocity of $v_s = 0.43$ m/s. The jet




Fig. 12 **a** Set-up of the TOMO-PIV experiment at the water jet facility at TU Delft, the Netherlands. Three high-speed cameras image a water jet, using circular laser illumination; **b** full image of one exemplary camera; and **c** camera image detail

is illuminated from above using a light volume of cylindrical shape. As light source, a Quantronix Darwin-Duo Nd:YLF high-speed laser with a repetition rate of 1 kHz and a light output of 2.25 mJ is used. The water is seeded using polyamide particles with an average diameter of 56 μm . Imaging is realized by three LaVision Imager pro HS high-speed cameras, observing an interrogation volume extending from $y = 1.4\text{--}54$ mm above the nozzle. The illuminated diameter is approximately 28 mm (at the top) to 22 mm (at the bottom). The octagonal shape of the water tank allows all cameras to image the volume perpendicular to the air/glass/water interfaces. The cameras are equipped with 105 mm Nikon lenses, using $f_{\#} = 22$. A cropped resolution of 672×1024 pixels is used. Figure 12b, c shows particle images, as seen by one camera.

The recorded time series consist of 500 images per run, one of which was reconstructed by the STB algorithm. The active image area is approximately 0.43 MPix per camera, and the particle image density averaged over the whole image is around 0.035 ppp. Due to the cylindrical illumination, the central part of the images shows significantly more particles, compared to the borders. The perceived particle image density varies between around 0.01 and 0.05 ppp within the camera images.

4.1 Algorithm parameters and tracking evaluation

The parameters for STB were chosen as given in Table 4. In this case, no triangulations using a reduced set of cameras were performed, as only three cameras are available. Two passes were conducted, going once forwards and backwards through the time series.

Table 4 STB parameters as applied to experimental image data from a transitional jet from time-step 5 onwards

Triangulation iterations using N_{cam}	$m = 3$ (6)
Triangulation iterations using N_{cam}^{-1}	$n = 0$ (0)
Number of shake iterations	$k = 7$
Shake width	$\delta_s = 0.1\text{--}0.025$ px
Initial shake iterations	0
Allowed triangulation error [px]	$\varepsilon = 1.0$ px
Search radius without predictor [px]	18 px
Search radius with predictor [px]	4 px
Particle intensity threshold	$T_{\text{int}} = 0.05 \cdot I_{\text{avg}}$
Smoothness threshold	$T_{\Delta f} = 0.8$ px (1.6 px)
Outlier detection threshold (T_O)	–
Projection factor prior to triangulation	$f_{\text{pt}} = 2.0$ (1.1)
Predictor for initialization	TOMO-PIV result
Number of passes	2

Values for time-step 1–4 (initialization phase) are given in brackets

Figure 13a shows the number of tracked particles and the total number of particles (tracked + particle candidates) for the 500 images and both passes. It can be seen that for the first few images, when the algorithm is not yet converged, over 30,000 particle candidates per time-step are triangulated. Out of those, around 6000 4-step tracks are found after the initialization phase. In the following, the algorithm quickly converges: after 11 time-steps around 10,000 tracks are found, and finally, from time-step 30 on, around 11,300 particles are tracked in the first pass. When time is reversed at the end of the data set, the number of tracked particles climbs to around 12,600. These additional




Fig. 13 **a** Number of tracked particles (*grey*) and total number of particles (*black*) within each time-step for two passes of STB and **b** distribution of track lengths

particles originate mostly from tracks that were not identified immediately when entering the measurement domain in pass 1 and are now extended to the edge of the volume by walking backwards along the track in pass 2. The number of additionally triangulated particle candidates quickly falls after the initialization and levels around 1000 for pass 1 and around 450 for pass 2. These numbers document how effective the prediction system reduces the complexity of the system: in the converged state, the number of newly triangulated particle candidates is $<5\%$ of the number of tracked particles in pass 2. The reconstruction process is very effective, and high accuracy is attained due to the (nearly) completely solved positioning problem. Most of the triangulated particle candidates are ghost particles, as only around 100 new 4-step tracks per time-step are identified for pass 1 (particles entering the domain, balanced with the number of particles leaving the domain) and only around 10 for pass 2 (tracks that were previously missed).




Fig. 14 Excerpt of inverted single camera image of the jet experiment and reprojected—as well as residual image, for two different the non-vanishing residual, in spite of a (nearly) complete reconstruction

4.1.1 Remarks on residual images

The reason for the occurrence of ghost particle candidates becomes apparent by looking at the residual images (see Fig. 14). When using a reprojection factor before each triangulation $f_{Pt} = 1.0$ (Fig. 14, left side), the residual image is still populated by a significant amount of peaks. Judging only by the residual image it appears as if many particles are yet undetected. Comparing the reprojected image with the original camera image reveals, however, that virtually all particles are present in the reconstruction. The residual peaks are caused by a discrepancy in particle intensity from original to reconstruction. One reason for this phenomenon is the fact that the camera showing the brightest image is left out for the intensity update (in order to avoid ghost tracks, see Sect. 2.2.3), which reduces average particle intensity relative to the recording. Another reason is the scattering behaviour of the particles, of which only the average can be calibrated by the OTF. Individual particles can show very different scattering—depending on their size and shape—giving rise to intensity ratios between

the images of a particle on the various cameras that differ significantly to the ones as calibrated by the OTF. Therefore, one particle, having a certain reconstructed intensity, can be brighter than the recording on one camera, while it is considerably less bright on other cameras. On each of these other cameras, a residual peak will remain when subtracting the reprojected image from the original image. These residual peaks are often high enough to be picked up by the next triangulation iteration, causing superfluous calculation effort and, possibly, the occurrence of ghost particles.

As reasoned in Sect. 2.3, a method to counteract these effects is to apply a reprojection factor $f_{pt} > 1$. The intensity of the particles is multiplied by f_{pt} prior to each triangulation iteration (it is not applied prior to shake iterations). Figure 14 shows camera image, reprojected image and residual image for one time-step for a small region of one camera for two STB runs—one using $f_{pt} = 1.0$, the other $f_{pt} = 2.0$. It can be seen that virtually the same particle distribution was reconstructed, matching well to the camera image. The effect on the residual is easily visible, with a significant reduction in remaining peaks. The number of additional particle candidates triangulated in each step decreases from around 4200 for $f_{pt} = 1.0$ to around 450 for $f_{pt} = 2.0$. Due to the increased efficiency and reliability of the tracking system, $f_{pt} = 2.0$ was used for the processing in the jet case and all presented plots originate from this version. This value for f_{pt} is quite high—in most experimental cases, a value of 1.2–1.4 was found to be sufficient. The reason for the high value used here is most possibly the used seeding material, which produced particle images of quite different intensities and sizes (as seen in Figs. 12c, 14).

The thoughts on the residual image show that STB is—in contrast to tomographic reconstruction algorithms like MART or SMART—not a technique that minimizes the residual at all cost. As a matter of fact, in a situation where single particles do not scatter the light as given by the calibration it is not possible to realize a disappearing residual on all cameras without adding additional particles (which are ghost particles). This is exactly what MART and SMART do when encountering imaging situations as described: as it is not possible to completely solve the reconstruction problem by giving intensity to voxels that represent the true particle, additional voxels have to be used to further reduce the residual. These voxels either deform the 3D particle shape or are not connected to the true particle at all (and are therefore ghost particles/ghost energy). Therefore, by not claiming a completely disappearing residual, STB (and all particle-based methods) actually operates closer to reality, as the reconstruction problem is not completely solvable due to inconsistent particle imaging.

4.1.2 Track length statistics

The circa 12,600 particles that are tracked for every time-step in pass 2 are part of tracks with very different lengths: some particles are tracked over the whole sequence of 500 images (these are slow particles in the outer limits of the measurement volume that are not swallowed by the jet at some point), while many other tracks are much shorter. Figure 13b shows the distribution of track lengths for both passes. The insert zooms into the region of very short tracks. Such short tracks originate either from (small) particles with very weak intensity, which are lost and found multiple times during their stay in the volume or are made up by false tracks that might occur during and briefly after the initialization phase. For both passes, a maximum in the track length is seen at 6 time-steps; however, this maximum is clearly reduced for the second pass. Another peak can be seen at track lengths of around 125 time-steps. This peak originates from the fast particles in the centre of the jet that move around 8 pixels per time-step and remain visible on the cameras for around for 120–130 time-steps. The track length distribution slowly decreases, reaching a last maximum at exactly 500 time-steps. For the first pass, 777 particles are tracked over all 500 images, and for pass 2, this number is 1261. The average track length rises from 75.3 time-steps in pass 1 to 82.1 time-steps for pass 2.

4.2 Track filtering

STB identifies long tracks, comprised of particle positions. Compared to two-frame recordings, which are limited to velocity estimations up to second-order accuracy (Wereley and Meinhart 2001), the use of multiple frames allows for higher-order accuracy in the velocity determination (Cierpka et al. 2013). In order to accurately extract velocity (and acceleration) information from the data, a suitable fitting function needs to be applied to the time series of coordinates. A first version of the code used third-order polynomials of predefined length being gradually moved through the track, thus filtering each time-step in relation to its temporal neighbours (Savitzky and Golay 1964). In order to utilize a fit that is more adapted to the source data, this method was exchanged for a third-order B-spline fit, which approximates the optimal Wiener filter for our model of particle motion. This filter is a low-pass filter with a selectable cut-off frequency, which is determined using the spectral distribution of the unfitted tracks (much like the Wiener filter used for particle position prediction). As a further measure, the filter could be adapted to local areas of the measurement volume, taking specific imaging or flow conditions into account. More details on the method, which is based on (Eilers and Marx 1996), can be found in (Gesemann 2015).

For the jet case, a cut-off frequency of 0.25 times the Nyquist frequency was chosen and all tracks were fitted, relocating the particles to a new—corrected—position on the trajectory as given by the continuous B-spline function. Velocity and acceleration are extracted as the first and second derivative of that function. On average, the particles are moved by 0.061 px in x -direction, 0.050 px in y -direction and 0.081 px in z -direction, making for a total average correction of 0.113 px. For both passes, the corrections are very similar. These numbers reflect the positioning of the cameras (x : only middle camera views perpendicular; y : all cameras are viewing perpendicular; z : no camera views perpendicular, middle camera views directly from top). The correction of the particle position may be used as a rough measure of the particle position accuracy. However, care has to be taken in order to use a filtering scheme attuned to the expected noise level.

Figure 15 shows results from the track reconstruction, as given by STB with subsequent B-spline fitting; (a) and (b) display the tracks of 100 successive time-steps out of 500, colour coded by streamwise velocity (v). It can be seen how

the fast particles coming from the nozzle are surrounded by a field of low-velocity particles that describe a slow circular motion around the jet. The detail plot in Fig. 15b shows how the particles are excited to curling trajectories by passing vortex rings that originate from Kelvin–Helmholtz (KH) instabilities at the shear layer (Violato et al. 2011). When using window methods like correlation, such details are lost. The shear layer is thin near the nozzle, but quickly widens with increasing distance. Figure 16 gives a better impression of the steepness of the velocity drop-off, by displaying radial bin-averaged (ensemble-averaged) profiles of streamwise velocity v at different heights. To create this plot, the velocity of all particles from the 500 time-steps located in a height interval of ± 2 mm of the given y -value (e.g. for $y = 1.0$ d, all particles within $y = 8$ –12 mm) were averaged in bins, according to their distance to the jet axis. The bins have similar volume; therefore, the spacing is closer in the outer regions of the jet. 100 bins are distributed over $r = 0$ –15 mm, leading to particle numbers of around 5000–10,000 per bin in each of the six heights. Very close to the nozzle ($y = 0.2$ d), a steep velocity falloff

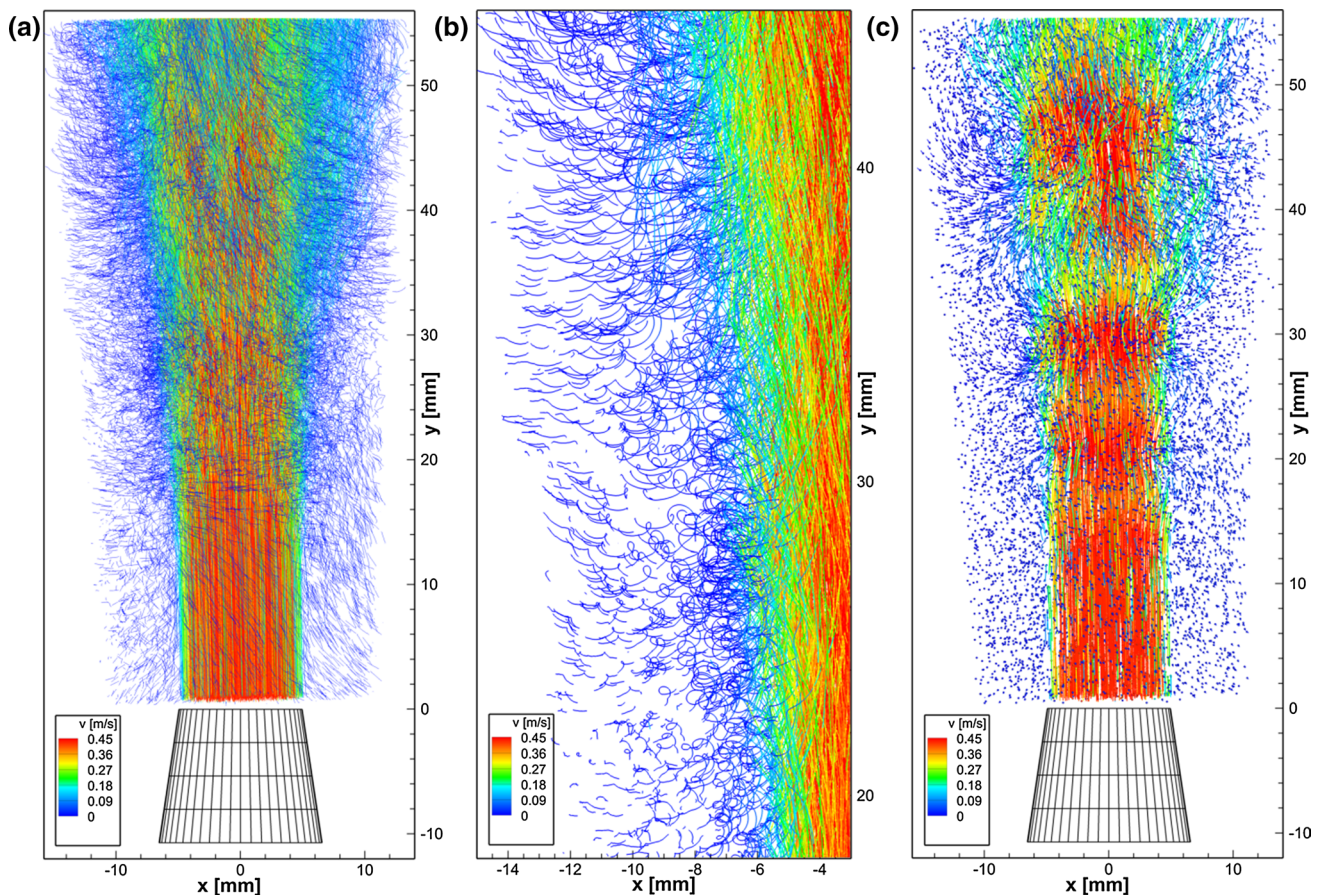


Fig. 15 **a** Particle tracks of water jet, reconstructed by STB. Overlay of 100 time-steps, colour coded by streamwise velocity; **b** detail of (a), showing the entrainment of particles; and **c** tracked particles for

single time-step (t_n , given by dots) with a tail of 15 time-steps (reaching back to t_{n-14})

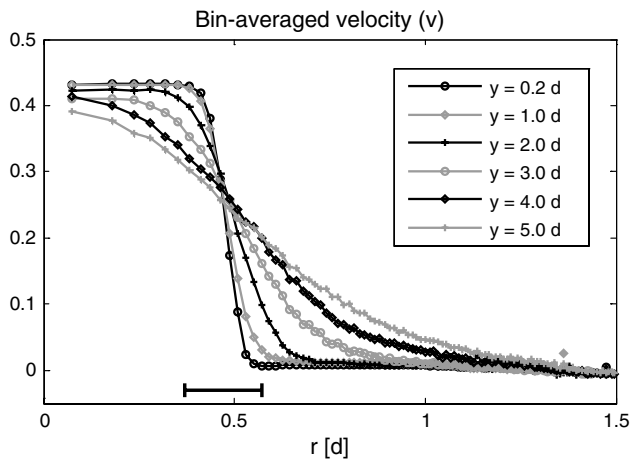


Fig. 16 Radial profiles of streamwise velocity (v) at different heights of the jet. The size of a 36-pixel wide correlation window is given as a reference

can be seen; moving upstream, the profiles flatten, as the shear layer breaks down gradually and turbulent structures begin to spread the high-velocity fluid over a larger area. Comparing these profiles to the ones given in (Violato et al. 2011)—which were created from averaging 4000 images of a Stereo-PIV measurement using ensemble-averaged 6×6 px correlation windows—it can be seen that, even for the relatively low number of images used here (500), the bin-averaging approach allows for the generation of reasonably converged, high-resolution profiles from three-dimensional track data. Opposed to correlation-based methods—which are limited to a certain window size—the resolution of the bin-averaging approach is solely dependent on the total number of samples within each volume unit. This number can easily be increased via the number of images, thereby allowing the use of smaller bins. See, for example, (Schröder et al. 2015b) for an application of the concept on a turbulent boundary layer, yielding profiles with a resolution below $1/10$ of a viscous unit.

As a by-product of the track filtering process, probability density functions (PDFs) of velocity and acceleration (see Schröder et al. 2015a for an example) or other flow properties, like wall-shear stress fluctuations (see Schröder et al. 2015b), can efficiently be created on a particle basis.

Moving on from averaged flow properties, Fig. 15c visualizes a more instant flow situation by plotting the particles of one time-step as dots, followed by a tail of 15 time-steps extending in the past.

4.3 Interpolation to Eulerian grid ('FlowFit')

To allow for a better identification of flow structures by known variables—such as vorticity, Q -criterion or λ_2 swirling strength—an accurate interpolation of the discrete

particle-based information onto an Eulerian grid is desirable. A scheme to interpolate the reconstructed tracks on such a grid, while maintaining as much of the high local accuracy by avoiding spatial smoothing, was developed and termed 'FlowFit' (Gesemann 2015): Each component of the flow field is modelled as a weighted sum of three-dimensional and evenly spaced quadratic B-splines. In order to evaluate this flow field on arbitrary coordinates, the weights have to be determined according to the known flow velocities at certain locations (being the particles with their velocity and acceleration). This results in a system of linear equations, where for each known flow velocity at some particular position, three equations are created. In addition to these equations based on the measurements, other equations are used to regularize the equation system by penalizing nonzero curvatures and optionally (when the flow can be regarded as incompressible) by penalizing nonzero divergences on a regular grid. This results in an overdetermined system where measurements and different kinds of regularizations can be weighted differently. This equation system is solved iteratively via the conjugate-gradient algorithm. The resulting flow field is then sampled on a regular grid including its spatial derivatives so that the derived values, such as vorticity or Q -criterion, can be computed without numerical differentiation. In order to retain accuracy at the particle positions, the underlying B-spline systems oversamples the particle field (the number of B-spline cells is typically chosen to be 5–20 times the number of particles). The splines within empty cells have to only fulfil smoothness and divergence criteria, whereas the ones in non-empty cells additionally have to describe the velocity data of the contained particle(s), thereby defining the shape of the whole system.

A related approach to reconstructing a velocity field from LPT/PTV data was very recently introduced by Schneiders et al. (2015), using the vortex-in-cell method (VIC+). While the methods are differing quite substantially (VIC+ operates on a fixed grid, in which it computes vorticity as the only variable and derives all other quantities from there), the basic concepts are comparable. VIC+ adds the inclusion of the acceleration field into the regularization, improving upon results only using the incompressibility constraint. Schneiders et al. (2015) demonstrate the general superiority of flow field reconstructions from Lagrangian tracks over correlation-like methods.

Evaluating the jet case, the measurement domain was divided into a cell system of quadratic B-splines, such that on average every eighth cell contains a particle (0.125 ppc, 'particles per cell'). The pitch of the resulting cell system is 0.62 mm; each cell represents a volume of approximately 0.25 mm^3 ; each particle represents approximately 2.0 mm^3 . For reference, a 36^3 px correlation volume applied for TOMO-PIV processing comprises approximately 8.4 mm^3 . The solution of the resulting equation system is sampled

on a 3D grid with 0.25 mm pitch—thus oversampling the cell system by a factor of 2.4—resulting in volumes of $124 \times 240 \times 124$ vectors.

Figure 17a displays the result of FlowFit applied to the particle field shown in Fig. 15c. Isosurfaces of vorticity are complemented by the particle tracks, enabling a joint visualization of the data on an Eulerian grid with the underlying discrete Lagrangian information. Vortices can now be easily identified as ring-shaped Kelvin–Helmholtz instabilities, originating from a circular shear layer around the laminar flow emanating from the nozzle. The vorticity isosurfaces show a spatially consistent depiction of both the large-scale vortices and smaller structures occurring after the breakdown of the KH instabilities, as well as an even representation of the shear layer close to the nozzle.

4.4 Lagrangian accelerations

Applying a temporal fit to the particle tracks not only yields Lagrangian velocities, but also accelerations. These values

are of special interest, as it has proven to be difficult to extract accurate acceleration data from spatially smoothed velocity fields gained from TOMO-PIV. Shake-The-Box allows for the accurate extraction of Lagrangian accelerations, while remaining spatially well resolved. Figure 17b displays the acceleration in streamwise direction for the same time-steps as given in Fig. 15c, visualizing how the particles are accelerated when drawn into one of the large ring vortices and decelerated when they are ejected again. The same effect is visible for the larger vortices that occur after the KH breakdown.

One application of acceleration (or material derivative) data is the extraction of pressure distribution, which is actively been worked on recently. While some methods directly work on discrete particle tracks (Neeteson and Rival 2015), most of the developed methods require acceleration data on an Eulerian grid (Violato et al. 2011; Novara et al. 2013; Huhn et al. 2015). The FlowFit interpolation scheme can be used to create Eulerian gridded data from the discrete acceleration values. The same basic

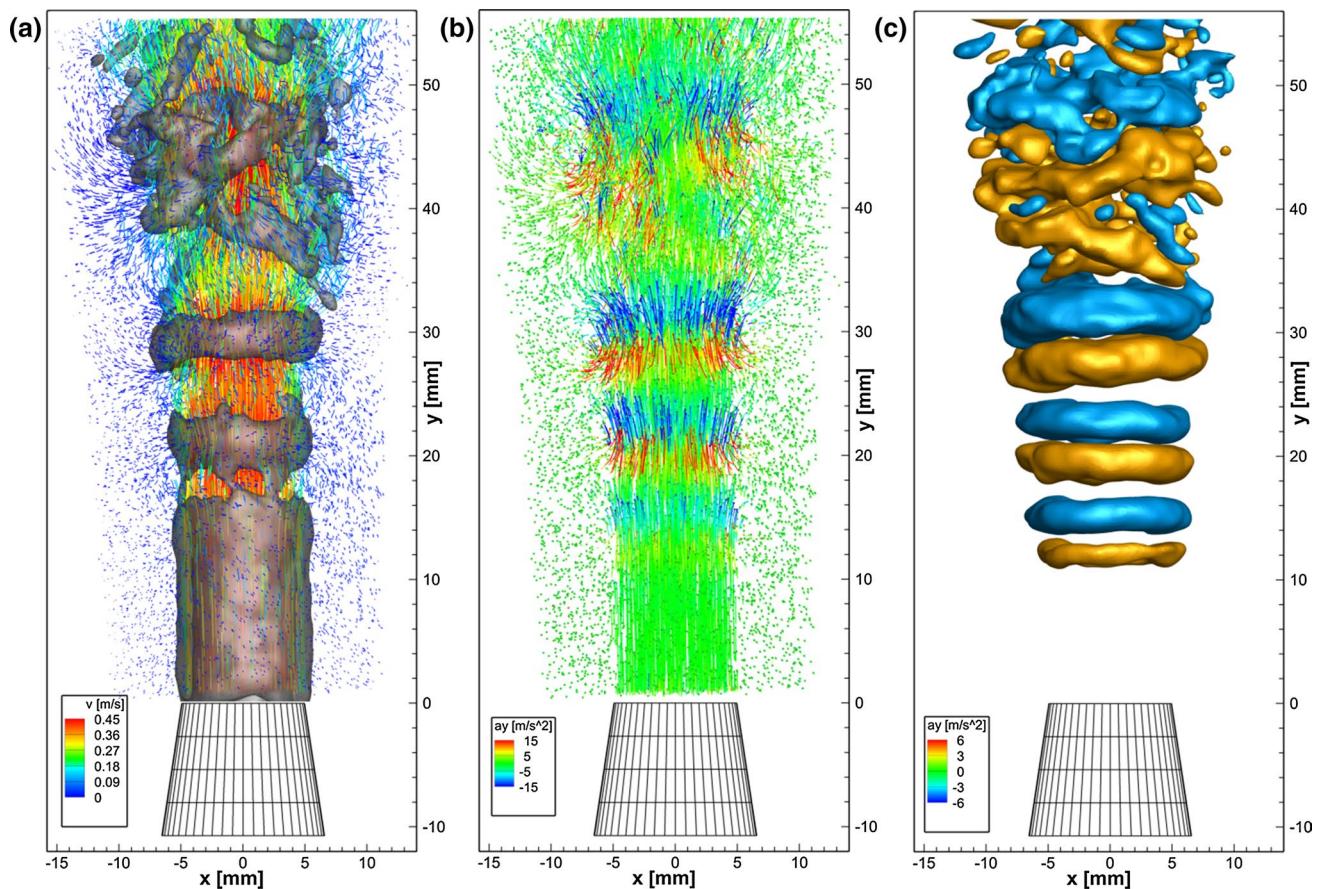


Fig. 17 **a** Isosurfaces of vorticity ($\omega = 175/s$) as calculated by FlowFit at single time-step (t_n) with superimposed particle tracks, extending 5 time-steps back in forth in time; **b** tracked particles for single time-step (t_n , given by spheres) with a tail of ten time-steps, colour

coded by streamwise acceleration; and **c** isosurfaces of streamwise acceleration ($a_y = 4 m/s^2$ and $a_y = -4 m/s^2$) at single time-step t_n as calculated by FlowFit with penalization of rotation

principle as with velocity is applied, only the penalization of divergence is substituted by a penalization of rotation (the acceleration field should be rotation-free in incompressible flows when neglecting viscosity). Figure 17c shows the result of interpolating the particle acceleration values via FlowFit onto a grid with 0.25 mm pitch. A value of 0.2 ppc was used. Isosurfaces of streamwise acceleration are displayed, colour coded by negative and positive orientation. Surrounding the KH-type ring vortices, clearly defined rings of uniform positive and negative streamwise acceleration are identified; the vortex is located in-between these rings, with the acceleration vectors pointing towards the vortex centre axis.

4.5 Comparison to TOMO-PIV

In order to judge the quality of the STB results, the particle images were used to perform a tomographic reconstruction using five iterations of the SMART algorithm (Atkinson and Soria 2009), including volume smoothing (Discetti et al. 2013) and contrast enhancement after each iteration in order to reduce ghost particles and accelerate convergence. The resulting voxel spaces were correlated using 3D direct correlation (Discetti and Astarita 2012), as implemented in LaVision Davis 8.2. A window size of 36^3 voxels with 75 % overlap was applied; using smaller windows resulted in a notable increase in noise. Figure 18 compares isosurfaces of vorticity for TOMO-PIV and STB + FlowFit. The spatial sampling of the FlowFit B-spline system was reduced to 0.5 mm to match the resolution of the 3D Correlation. The time instant is the same as in Fig. 17, albeit rotated by 180 degree for better visualization. The results look quite different, with a more uneven representation of large structures for TOMO-PIV, combined with a lack of many small structures that are rendered by STB + FlowFit. The first ring vortex clearly shows secondary vortices in streamwise orientation, reaching out to the second ring vortex, that are visible in the STB calculation, but are missing for TOMO-PIV. The shear layer displays a lot of patches when using standard TOMO-PIV processing, but is represented evenly by the FlowFit interpolation of STB tracks.

The improvements on spatial coherence observed for the STB + FlowFit evaluation stem from a combination of ghost-free reconstructions, position accuracy, a lack of spatial filtering induced by a correlation window and the possibilities of pouring physical constraints, like divergence penalization, into the cost function of the interpolation. These features are the effect of STB being purely based on the reconstruction of fluid element trajectories (which is true for low Stokes number tracers) and FlowFit being able to translate the local accuracy into a volumetric representation. Similar results were found in (Schneiders et al. 2015).

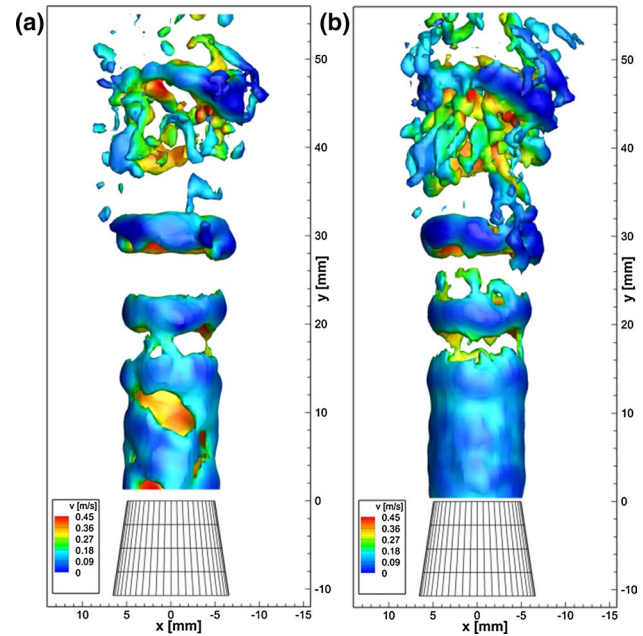


Fig. 18 **a** Isosurfaces of vorticity ($\omega = 175/s$), colour coded by streamwise velocity, gained by TOMO-PIV processing; **b** result gained from STB and grid interpolation using FlowFit at the same spatial resolution

Temporal coherence is a strong point of tracking techniques, as accurate temporal fitting is intrinsic to the methods. This is supplemented by Fig. 19, which shows the flow field at three different time instances, each separated by ten time-steps (10 ms). For better clarity, the vorticity threshold has been increased compared to previous images. While TOMO-PIV shows severe variations in the flow structure representation—even for the largest of these—STB retains a high temporal consistency even for smaller structures. The shapes and distortions of the ring vortices are retained over long periods of time, clearly separating the influence of noise from the structure development. Post-processing TOMO-PIV results, doing, for example, spatio-temporal filtering like polynomial least squares regression—which was applied in (Violato and Scarano 2011) to a very similar data set—would presumably improve the temporal coherence; however, the influences of spatial filtering would still persist. Applying more advanced TOMO-PIV methods, like SMTE (Lynch and Scarano 2015), would most likely result in a much higher similarity of the results between STB and TOMO, as both methods heavily incorporate temporal information. Upcoming investigations shall give in-depth comparisons of the two techniques.

4.6 Computational effort

Looking at the calculation time for the different methods, the rapid processing by STB becomes apparent: the

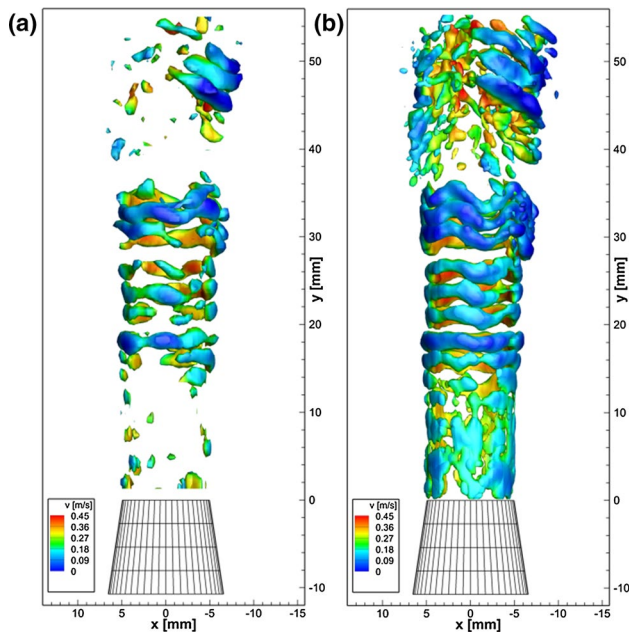


Fig. 19 **a** Isosurfaces of vorticity ($\omega = 280/s$), colour coded by streamwise velocity for three snapshots spaced by 10 time-steps, gained by TOMO-PIV; **b** result from STB + grid interpolation using FlowFit at same spatial resolution

tomographic reconstruction of 500 images, using SMART with a $568 \times 1064 \times 568$ voxel space, takes around 510 min on a state-of-the-art high-end server (dual Intel Xeon E5-2680 ten-core processors at 2.8 GHz each). The direct correlation adds another 470 min, for a total of 980 min. On the same machine, STB needs 69 min for the first pass and 63 min for the second pass, for a total of 131 min. What is more, due to the relatively low number of particles the server is not working to full capacity: it is possible to start three to five STB processes simultaneously without any noticeable increase in calculation time. Alternatively, keeping it at one process, the same reconstruction times are achievable using a modern laptop instead of a server. The (optional) calculation of Eulerian data using FlowFit requires between 23 min (using moderate quality settings) and 265 min (using highest quality settings) for the processing of all 500 images.

An interesting view on the development of evaluation codes for 3D flow experiments gives the numbers provided in (Violato and Scarano 2011), where a very similar case (same jet, 4 cameras, $600 \times 1000 \times 600$ voxels) is discussed. The algorithms available at that time needed around 20 min for reconstruction and 60 min for 3D cross-correlation per time-step, making for a total of 40,000 min for 500 time-steps. These details explain why at that time only 500 images were recorded. In the circa 5 years, the evaluation time for such a time series has dropped from over 650 h on a double quad-core server to around 2 h on a good laptop.

5 Conclusions

A novel method to extract Lagrangian particle tracks from time-resolved series of images with high particle image density has been introduced and termed ‘Shake-The-Box’. By effectively seizing, the temporal domain particle tracking at particle image densities typically reserved for TOMO-PIV becomes feasible. The claim for temporal coherence successfully reduces ghost particle problems to a negligible extent. Combining this approach with methods from both TOMO-PIV development (volume Self-calibration, OTF calibration) and advanced 3D PTV (iterative particle reconstruction) results in highly accurate determination of position, velocity and acceleration.

By applying the Shake-The-Box method to time series of synthetic images, the ability to solve the tempo-spatial reconstruction problem for a broad range of seeding concentrations and imaging conditions was demonstrated. Noise-free images with a particle image density of 0.125 ppp could be successfully evaluated. After a convergence phase, nearly full completeness (>99.5 % of identified particles at 0.125 ppp), virtually no ghost particles (<0.04 % false particles) and high accuracy (average position error of 0.018 px) were attained. Using lower particle image densities further improved on these numbers. Compared to conventional 3D PTV techniques, the applicable seeding concentration was enhanced by at least an order of magnitude. Tomographic reconstruction on the other hand is able to process such high particle image densities, albeit at much higher ghost particle occurrence (>250 % false particles at 0.125 ppp) and a significantly lower accuracy (average position error of 0.3 px).

Adding noise to the particle images showed a reduction in accuracy for both STB and tomographic reconstruction. Noise interferes with the convergence behaviour of STB; for highly noisy images, convergence could only be reached until 0.075 ppp. Operating below this particle concentration is advised when dealing with low-quality images. While the position error grows, STB retains a nearly ghost-free status even with increased noise. Applying a temporal filter to the reconstructed tracks (realized by a third-order smoothing B-spline) increases the accuracy of particle placement (an average position error of 0.018–0.12 px was found for the different noise levels at 0.05 ppp). Lagrangian velocity and acceleration are extracted using the derivations of the continuous B-spline curve. A track-based dynamic velocity range (DVR) of 300–670 was found for average to low noise levels; even for highly noisy images, the DVR remained above 150. The dynamic acceleration range (DAR) varied between 23 and 43, depending on the noise level, proving that STB is able to generate acceleration data with accuracy suitable for, for example, the extraction of 3D pressure distributions (Neeteson et al. 2015; Huhn et al. 2015).

Time-resolved images from a volumetric three-camera measurement on a transitional jet in water were used to demonstrate the application of STB to experimental data. A two-pass approach of STB was realized, going once forwards and backwards through the time series. The extracted particle tracks reveal flow details that are lost in correlation-based evaluations. Applying a third-order smoothing B-spline to the reconstructed tracks yields accurate velocity and acceleration values, as demonstrated by the extraction of highly resolved axial velocity profiles using a bin-averaging technique. While the particle number (12,600) is rather low for the presented case (due to only three cameras being used at an effective resolution below 0.5 MP per camera), STB has been shown to be able to process particle numbers at least an order of magnitude higher (Schröder et al. 2015a).

In order to interpolate the locally highly accurate information of particle tracks onto an Eulerian grid, the ‘Flow-Fit’ algorithm was introduced. The resulting 3D velocity volumes show detailed flow structures, surpassing the results of TOMO-PIV processing both in spatial resolution and in temporal consistence. Applying the FlowFit scheme to discrete particle accelerations allows for the extraction of Eulerian gridded acceleration fields.

The computational effort of Shake-The-Box was shown to be low due to the mostly presolved status of the reconstruction problem after the prediction step. Reconstruction time is mainly depending on the number of particles and was found to be a factor of 5–8 lower compared to TOMO-PIV processing. Overnight processing of extended time series becomes possible using standard PC hardware.

The results from both the synthetic test cases and the experimental time series demonstrate that the precise knowledge of trajectories of tracer particles—which can be regarded as fluid elements—allows for optimal exploitation of data. All ways of data extraction, e.g. bin averaging, determination of derivations or interpolation to a regular grid, can be performed as post-processing on the extracted track data. These steps can be specifically designed to retain as much of the available information as possible.

The details of the STB algorithm are still under development. Recently, an updated version of the code—featuring a more effective 4-step track identification scheme—was applied to the synthetic data presented Sect. 3. Using this code, noise-free images with particle image densities up to 0.2 ppp and up to 0.1 ppp at the highest noise level were successfully processed.

The algorithm may be improved in several ways: single-image volume self-calibration (Michaelis and Wolf 2011; Earl et al. 2015), a global approach to the particle correction scheme (complementing the local shake scheme, Cornic et al. 2014) or an advanced system of connecting

track fragments (see, for example, Willneff 2003; Xu 2008) come to mind.

An ongoing process is the adaptation of the ideas of STB on the processing of short time series (4 or even only 2 pulses). As demonstrated by Schröder et al. (2013), the use of a dual-volume-set-up facilitates the task of separating true particles from ghost particles using the temporal domain. Novara et al. (2015) carried this idea forward by extracting four-step tracks from experimental data on a turbulent boundary layer using an iterative STB approach. Though the temporal information is limited in such set-ups, the technique translates the advantages of particle tracking methods to flow regimes of much higher Reynolds numbers by overcoming the limitations of camera repetition rate and laser power for time-resolved measurements.

Acknowledgments The work has been conducted in the scope of the DFG-project ‘Analyse turbulenter Grenzschichten mit Druckgradient bei großen Reynoldszahlen mit hochauflösenden Vielkammeressverfahren’ (Grant KA 1808/14-1 and SCHR 1165/3-1). The authors thank Dr. Daniele Violato and Dr. Matteo Novara from TU Delft for the experimental set-up and the collaboration on conducting the experiment used for the evaluation of the algorithm. Furthermore, the authors thank Bernhard Wieneke for fruitful discussions on IPR and for providing an algorithm to remove divergence from a given vector volume which served as a source for the synthetic track generation.

References

- Adrian RJ (1997) Dynamic ranges of velocity and spatial resolution of particle image velocimetry. *Meas Sci Technol* 8:1393–1398
- Atkinson C, Soria J (2009) An efficient simultaneous reconstruction technique for tomographic particle image velocimetry. *Exp Fluids* 47:563–578
- Atkinson C, Coudert S, Foucaut JM, Stanislas M, Soria J (2011) The accuracy of tomographic particle image velocimetry for measurements of a turbulent boundary layer. *Exp Fluids* 50:1031–1056
- Bastiaans RJM, van der Plas GAJ, Kieft RN (2002) The performance of a new PTV algorithm applied in super-resolution PIV. *Exp Fluids* 32:345–356
- Ben-Salah R, Alata O, Tremblais B, Thomas L, David L (2015) Particle volume reconstruction based on a marked point process and application to Tomo-PIV. EUSIPCO 2015. Nice, France, 31 Aug–3 June 2015
- Bourgoin M, Ouellette NT, Xu H, Berg J, Bodenschatz E (2006) The role of pair dispersion in turbulent flow. *Science* 311(5762):835–838
- Cierpka C, Lütke B, Kähler CJ (2013) Higher order multi-frame particle tracking velocimetry. *Exp Fluids* 54:1533
- Cornic P, Champagnat F, Plyer A, LeClaire B, Cheminet A, Le Besnerais G (2014) Tomo-PTV with sparse tomographic reconstruction and optical flow. In: 17th international symposium on applications of laser techniques to fluid mechanics. Lisbon, Portugal, 07–10 July 2014
- Cornic P, Champagnat F, Cheminet A, LeClaire B, Le Besnerais G (2015) Fast and efficient particle reconstruction on a 3D grid using sparsity. *Exp Fluids* 56:62

- Dalziel SB (1992) Decay of rotating turbulence: some particle tracking experiments. *Appl Sci Res* 49:217–244
- Discetti S, Astarita T (2011) A fast multi-resolution approach to tomographic PIV. *Exp Fluids* 52:765–777
- Discetti S, Astarita T (2012) Fast 3D PIV with direct sparse cross-correlations. *Exp Fluids* 53:1437–1451
- Discetti S, Natale A, Astarita T (2013) Spatial filtering improved tomographic PIV. *Exp Fluids* 54:1505
- Discetti S, Agüera N, Cafiero G, Astarita T (2015) Ensemble 3D-PTV for high resolution turbulent statistics. In: 11th international symposium on PIV–PIV15. Santa Barbara, USA, 14–16 Sept 2015
- Earl A, Cochard S, Thomas L, Tremblais B, David L (2015) Implementation of vibration correction schemes to the evaluation of a turbulent flow in an open channel by tomographic particle image velocimetry. *Meas Sci Technol* 26:015303
- Eilers PHC, Marx BD (1996) Flexible smoothing with B-splines and penalties. *Stat Sci* 11(2):89–121
- Elsinga GE, Tokgoz S (2014) Ghost hunting—an assessment of ghost particle detection and removal methods for tomographic-PIV. *Meas Sci Technol* 25:084004
- Elsinga GE, van Oudheusden BW, Scarano F (2006b) Experimental assessment of tomographic-PIV accuracy. In: 13th international symposium on applications of laser techniques to fluid mechanics. Lisbon, Portugal, 26–29 June 2006
- Elsinga GE, Scarano F, Wieneke B, van Oudheusden BW (2006a) Tomographic particle image velocimetry. *Exp Fluids* 41:933–947
- Elsinga GE, Westerweel J, Scarano F, Novara M (2011) On the velocity of ghost particles and the bias errors in Tomographic-PIV. *Exp Fluids* 50:825–838
- Gesemann S (2015) From particle tracks to velocity and acceleration fields using B-splines and penalties. arXiv 1510.09034
- Ghaemi S, Scarano F (2011) Counter-hairpin vortices in the turbulent wake of a sharp trailing edge. *J Fluid Mech* 689:317–356
- Hain R, Kähler CJ, Michaelis D (2008) Tomographic and time resolved PIV measurements on a finite cylinder mounted on a flat plate. *Exp Fluids* 45:715–724
- Henningsson P, Michaelis D, Nakata T, Schanz D, Geisler R, Schröder A, Bomphrey RJ (2015) The complex aerodynamic footprint of desert locusts revealed by large-volume tomographic particle image velocimetry. *J R Soc Interface* 12:108
- Herman GT, Lent A (1976) Iterative reconstruction algorithms. *Comput Biol Med* 6:273–294
- Huhn F, Schanz D, Gesemann S, Schröder A (2015) Pressure fields from high-resolution time-resolved particle tracking velocimetry in 3D turbulent flows. In: Proceedings of NIM2015 Workshop, Poitiers, France
- Kähler CJ, Scharnowski S, Cierpka C (2012a) On the resolution limit of digital particle image velocimetry. *Exp Fluids* 52:1629–1639
- Kähler CJ, Scharnowski S, Cierpka C (2012b) On the uncertainty of digital PIV and PTV near walls. *Exp Fluids* 52:1641–1656
- Kasagi N, Nishino K (1990) Probing turbulence with three-dimensional particle tracking velocimetry. *Exp Therm Fluid Sci* 4:601–612
- La Porta A, Voth GA, Crawford AM, Alexander J, Bodenschatz E (2001) Fluid particle accelerations in fully developed turbulence. *Nature* 409:1017–1019
- Lüthi B, Tsinober A, Kinzelbach W (2005) Lagrangian measurement of vorticity dynamics in turbulent flow. *J Fluid Mech* 528:87–118
- Lynch K, Scarano F (2013) A high-order time-accurate interrogation method for time-resolved PIV. *Meas Sci Technol* 24:035305
- Lynch K, Scarano F (2014) Experimental determination of tomographic PIV accuracy by a 12-camera system. *Meas Sci Technol* 25:084003
- Lynch K, Scarano F (2015) An efficient and accurate approach to MTE-MART for time-resolved tomographic PIV. *Exp Fluids* 56:66
- Maas HG, Grün A, Papantoniou D (1993) Particle tracking in three dimensional turbulent flows—part I: photogrammetric determination of particle coordinates. *Exp Fluids* 15:133–146
- Malik N, Dracos T, Papantoniou D (1993) Particle tracking in three dimensional turbulent flows—part II: particle tracking. *Exp Fluids* 15:279–294
- Michaelis D, Wolf CC (2011) Vibration compensation for tomographic PIV using single image volume self calibration. In: 9th international symposium on particle image velocimetry—PIV11. Kobe, Japan, 21–23 July 2011
- Neeteson NJ, Rival D (2015) Pressure-field extraction on unstructured flow data using a Voronoi tessellation-based networking algorithm: a proof-of-principle study. *Exp Fluids* 56:44
- Neeteson NJ, Bhattacharya S, Rival DE, Michaelis D, Schanz D, Schröder A (2015) Pressure-field extraction from lagrangian flow measurements. In: 11th international symposium on PIV–PIV15. Santa Barbara, USA, 14–16 Sept 2015
- Nishino K, Kasagi N, Hirata M (1989) Three-dimensional particle tracking velocimetry based on automated digital image processing. *Trans ASME J Fluid Eng* 111:384–390
- Nobach H, Bodenschatz E (2009) Limitations of accuracy in PIV due to individual variations of particle image intensities. *Exp Fluids* 24:045302
- Novara M, Scarano F (2013) A particle-tracking approach for accurate material derivative measurements with tomographic PIV. *Exp Fluids* 54:1584
- Novara M, Batenburg KJ, Scarano F (2010) Motion tracking-enhanced MART for tomographic PIV. *Meas Sci Technol* 21:035401
- Novara M, Ianiro A, Scarano F (2013) Adaptive interrogation for 3D-PIV. *Meas Sci Tech* 24:024012
- Novara M, Schanz D, Kähler CJ, Schröder A (2015) Shake-The-Box for multi-pulse tomographic systems: towards high seeding density particle tracking in high speed flows. In: 11th international symposium on PIV–IV15. Santa Barbara, USA, 14–16 Sept 2015
- Ouellette NT, Xu H, Bodenschatz E (2006) A quantitative study of three-dimensional Lagrangian particle tracking algorithms. *Exp Fluids* 40(2):301–313
- Savitzky A, Golay MJE (1964) Smoothing and differentiation of data by simplified least squares procedures. *Anal Chem* 36(8):1627–1639. doi:10.1021/ac60214a047
- Scarano F (2013) Tomographic PIV: principles and practice. *Meas Sci Technol* 24:012001
- Scarano F, Poelma C (2009) Three-dimensional vorticity patterns of cylinder wakes. *Exp Fluids* 47:69–83
- Schanz D, Schröder A, Heine B, Dierksheide U (2012) Flow structure identification in a high-resolution tomographic PIV data set of the flow behind a backward facing step. In: 16th international symposium on applications of laser techniques to fluid mechanics, Lisbon
- Schanz D, Gesemann S, Schröder A, Wieneke B, Novara M (2013a) Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. *Meas Sci Technol* 24:024009
- Schanz D, Schröder A, Gesemann S, Michaelis D, Wieneke B (2013b) Shake-the-Box: a highly efficient and accurate tomographic particle tracking velocimetry (TOMO-PTV) method using prediction of particle position. In: 10th international symposium on particle image velocimetry—PIV13. Delft, The Netherlands, 1–3 July 2013

- Schanz D, Schröder A, Gesemann S (2014) Shake-the-Box—a 4D PTV algorithm: accurate and ghostless reconstruction of Lagrangian tracks in densely seeded flows. In: 17th international symposium on applications of laser techniques to fluid mechanics. Lisbon, Portugal, 07–10 July 2014
- Schneiders JFG, Azijli I, Scarano F, Dwight RP (2015) Pouring time into space. In: 11th international symposium on particle image velocimetry—PIV15. Santa Barbara, California, 14–16 Sept 2015
- Schröder A, Geisler R, Elsinga GE, Scarano F, Dierksheide U (2008) Investigation of a turbulent spot and a tripped turbulent boundary layer flow using time-resolved tomographic PIV. *Exp Fluids* 44:305–316
- Schröder A, Geisler R, Staak K, Wieneke B, Elsinga G, Scarano F, Henning A (2011) Lagrangian and Eulerian views into a turbulent boundary layer flow using time-resolved tomographic PIV. *Exp Fluids* 50:1071–1091
- Schröder A, Schanz D, Geisler R, Willert C, Michaelis D (2013) Dual-volume and four-pulse tomo PIV using polarized laser light. In: 10th international symposium on particle image velocimetry—PIV13. Delft, The Netherlands, 1–3 July 2013
- Schröder A, Schanz D, Michaelis D, Cierpka C, Scharnovski S, Kähler CJ (2015a) Advances of PIV and 4D-PTV “Shake-The-Box” for turbulent flow analysis—the flow over periodic hills. *Flow Turb Comb* 95(2–3):193–209
- Schröder A, Schanz D, Geisler R, Gesemann S, Willert C (2015b) Near-wall turbulence characterization using 4D-PTV Shake-The-Box. In: 11th international symposium on particle image velocimetry—PIV15. Santa Barbara, California, 14–16 Sept 2015
- Sciacchitano A, Wieneke B, Scarano F (2013) PIV uncertainty quantification by image matching. *Meas Sci Technol* 24:045302
- Violato D, Scarano F (2011) Three-dimensional evolution of flow structures in transitional circular and chevron jets. *Phys Fluids* 23:124104
- Violato D, Moore P, Scarano F (2011) Lagrangian and Eulerian pressure field evaluation of rod-airfoil flow from time-resolved tomographic PIV. *Exp Fluids* 50:1057–1070
- Violato D, Ianiro A, Cardone G, Scarano F (2012) Three-dimensional vortex dynamics and convective heat transfer in circular and chevron impinging jets. *Int J Heat Fluid Flow* 37:22–36
- Wereley ST, Meinhart CD (2001) Second-order accurate particle image velocimetry. *Exp Fluids* 31:258–268
- Wieneke B (2007) Volume self-calibration for stereo PIV and tomographic PIV. *Exp Fluids* 45:549–556
- Wieneke B (2013) Iterative reconstruction of volumetric particle distribution. *Meas Sci Technol* 24:024008
- Wieneke B (2015) PIV uncertainty quantification from correlation statistics. *Meas Sci Technol* 26:074002
- Wiener N (1949) *Extrapolation, interpolation, and smoothing of stationary time series*, vol 2. MIT Press, Cambridge
- Willneff J (2003) *A spatio-temporal matching algorithm for 3D particle tracking velocimetry*. Swiss Federal Institute of Technology Zurich Diss. ETH No. 15276
- Xu H (2008) Tracking Lagrangian trajectories in position-velocity space. *Meas Sci Technol* 19:075105
- Xu H, Bourgoin M, Ouellette NT, Bodenschatz E (2006) High order lagrangian velocity statistics in turbulence. *Phys Rev Lett* 96:024503