

3D Lagrangian Particle Tracking

Andreas Schröder, Daniel Schanz, Matteo Novara,
Sebastian Gesemann, Tobias Jahn

DLR German Aerospace Center
Institute of Aerodynamics and Flow Technology

December 3rd, 2024

3D Lagrangian Particle Tracking (LPT) is an image-based flow measurement technique allowing to determine position, velocity and acceleration along a large number of particle tracks at high temporal resolution within the investigated volume. Providing many long individual particle tracks at statistically significant amounts transport and dispersion properties of the (turbulent) flow can be explored in detail. Furthermore, advanced binning procedures can produce a 3D array of important one- and multi-point statistics of the flow (like averages, Reynolds stresses, two-point-correlations etc.) at very high spatial resolution down to subpixel accuracies. 3D LPT is based on pulsed volumetric illumination of tracer particles inside the flow and imaging of the scattered particle light with temporal resolution onto typically three to six camera projections. The Shake-The-Box (STB) technique is an advanced 3D LPT method which combines the triangulation-based advanced Iterative Particle Reconstruction (IPR) technique with the exploitation of the temporal and spatial coherence of Lagrangian particle tracks in the investigated flow. STB enables the processing of particle image densities up to 0.15 ppp (particles per pixel) under good experimental conditions with an almost complete suppression of ghost particles. Subsequently, the dense scattered particle tracks are temporally filtered for estimating position, velocity and acceleration (material derivative) which can be used in a second step as input for data assimilation approaches using Navier-Stokes-constraints delivering the full time-resolved 3D velocity gradient tensor (VGT)- and pressure fields.

Contents

1	Introduction and Motivation	3
2	Preconditions of 3D LPT using Shake-The-Box techniques.....	5
3	Advanced Iterative Particle Reconstruction.....	7
3.1	Triangulation	9
3.1.1	Camera permutations	10
3.2	Optimization of position and intensity (shaking particles).....	11
3.2.1	Steepest-Descent Method (SDM)	11
3.3	Intensity update	12
3.4	Filtering	13
3.5	Rendering of residual images	13
3.6	Assessments using synthetic test-cases	13
4	3D LPT using (time-resolved) Shake-The-Box.....	17
4.1	Converged phase	19
4.2	Multi-pass processing.....	21
4.3	Variable-time-step STB (VT-STB).....	21
4.4	Examples of STB applications	22
5	3D LPT for high-speed flows: Multi-pulse STB.....	23
6	Conclusions.....	25
7	Acknowledgements	26
8	References.....	26

1 Introduction and Motivation

Most flows we encounter in every-day life are unsteady, turbulent and three-dimensional. The Reynolds numbers reached in nature, aerodynamics (vehicle boundary layers and wakes etc.) or in relevant technical applications (pipe flows, jets, combustion, mixing, heat transfer etc.) are typically far above the onset of turbulence. Literally, human beings are immersed in unsteady fluid flow phenomena from those in our blood vessels over flows around various transport vehicles up to currents in oceans, atmospheric turbulent boundary layer (TBL) flows or mixed thermal convection inside closed rooms. In order to make any technical use of fluid flows one needs to understand their properties in detail in various applications and situations. The main features of turbulent flows are their dynamic energy transfer mechanism in a cascade from large to smaller and smaller vortical flow scales down to dissipation (Richardson 1922) and their increase of scale separation with Reynolds number. During inter-scale energy transfer coherent vortical- and shear- flow structures at multiple scales are developing and decaying in a self-organizing manner driven initially by e.g. large-scale pressure-, velocity- or density gradients. Due to (local) shear instabilities smaller vortical- and other shear flow structures develop which guide the energy, momentum and fluid transport inside the flow. These flow structures can be defined as three-dimensional topologies (stable- and unstable saddle-points or stable- and unstable nodes) which change their shape and orientation in time while convecting (downstream) with the flow. On the other hand, the flow and its structures themselves can be understood and described as a dynamic composition of an infinite number of fluid elements which move with the local flow as Lagrangian trajectories, which are coupled with neighboring elements by viscosity. In a Lagrangian frame of reference these fluid elements are entering and exiting more persistent (vortical) flow structures, thereby keeping them “alive”. Therefore, flow topologies can be defined either in a Eulerian (laboratory) reference system e.g. by the invariants Q and R of the VGT (for incompressible flows) (Chong et al. 1990) or from a Lagrangian perspective (Haller 2015) moving with the fluid elements. Nevertheless, both reference systems describe the same flow, while Lagrangian and Eulerian velocity vectors are identical for each individual time step.

The governing Navier-Stokes-equations are non-linear partial differential equations so that very high computational effort is already needed to predict turbulent flows at moderate Reynolds numbers by Direct Numerical Simulations (DNS). But even with modern High-Performance-Computing (HPC) resources converged results e.g. for the flow around a small passenger aircraft cannot be provided at all and will most probably not be reachable within the next couple of decades. Therefore, various CFD methods aimed at reducing the complexity of the numerical schemes by modelling (small-scale) turbulence and solving the so-called closure problem for the specific equations used in either LES, DES, LBM or (U)RANS codes, which finally allow computing unsteady flows at higher Reynolds numbers, but with intrinsic restrictions. Depending on the code and invested computational effort either only the energy containing scales, large scale (periodic) flow features or mean flow quantities can be delivered reliably for specific adapted flow cases. Nevertheless, the success of resulting numerical flow predictions remains often case critical (e.g. for TBL w/o pressure gradients) and depending on the flow geometry (w/o flow separation). Therefore, (advanced) CFD requires spatially (and

temporally) well resolved experimental validation data to enable tuning of numerical parameters and adapting or further developing turbulence models. In most engineering tasks, the understanding of mean flow topologies and related one-point statistics of the Reynolds stresses are often the focus of interest for validation purposes of (U)RANS methods. These methods are capable of predicting the average flow features of high Reynolds number air vehicles with moderate computational effort and are therefore the main tools in aircraft design processes. As pointed out above, at the borders of the flight envelope where flow separation can occur due to strong adverse pressure gradients and turbulent wake- and shear- flows need to be predicted accurately, RANS methods still require improvements and new (adaptive) turbulence models. A promising field which in future will be closely related to velocity, acceleration and pressure data from 3D LPT measurements are data driven turbulence models (Duraisamy et al. 2019).

During the past decades a tremendous increase of flow field information has been gained from experimental investigations by applying image-based measurement techniques (complementary or alternative to existing probe techniques). For experiments in unsteady and turbulent flows non-intrusive volumetric and time-resolved measurement techniques determining all three components of the velocity- (and acceleration) vectors at many points inside the investigated volume instantaneously are highly desired (Discetti and Coletti 2018). Consequently, particle image-based velocimetry measurement techniques, like PIV, have been extended from 2D to 3D by tomographic methods (tomo-PIV, Elsinga et al. 2006, Scarano 2013) and from snapshot modes to temporal resolution in the recent years (Raffel et al. 2018). Tomo PIV enables flow field estimations on regular 3D velocity vector grids with relatively high spatial resolutions based on particle image densities around ~ 0.05 ppp. For all 3D reconstruction techniques, the particle image density N_I is typically expressed in terms of ‘particles per pixel’ (ppp) on the single camera projections and determines the possible spatial resolution inside the investigated flow. Particle tracking methods, on the other side, have evolved from relatively sparse 3D PTV schemes (e.g. Maas et al. 1993), Dabiri and Pecora 2019) at particle image densities of ~ 0.005 to 0.01 ppp to Lagrangian particle tracking techniques reaching high particle image densities of ~ 0.07 to 0.18 ppp by using the Shake-The-Box (STB) method (Schanz et al. 2016, Jahn et al. 2021, Schröder and Schanz 2023) and particle image or peak- detection schemes based on Convolutional Neural Networks (Godbersen et al 2024a).

Ideally, for many relevant flows time-resolved 3D pressure, temperature and density fields are of high interest as well, which are quite difficult to achieve directly by LPT methods using e.g. temperature- or pressure sensitive particle tracers or with a separate tomo BOS system (Nicolas et al. 2017). Luckily, at least for incompressible flows, 3D pressure fields can be calculated e.g. directly via integration of the Poisson equation from time-resolved 3D velocity vector fields (van Oudheusden 2013, van Gent et al. 2017) or implicitly by data assimilation approaches using dense scattered Lagrangian velocity and acceleration vector fields as input (e.g. Gesemann et al. 2016, Huhn et al. 2018, Raissi 2019, Schneiders 2018, Jeon 2022, Zhou and Grauer 2023, Godbersen et al 2024b).

In the following we will describe the preconditions and the methodology of the advanced 3D LPT method STB and some of its recent developments. STB enables reconstructing long time series of a huge number of individual particle trajectories ($N_P \sim 10^5 - 10^6$) in the investigated

flow volume from only few camera projections. Alternatively, for high speed flows many statistically independent and instantaneous short e.g. 4-, 3- or 2 pulse- tracks can be reconstructed after capturing particle images in a frame-straddling mode when using multi-pulse and multi-illumination STB techniques (Novara et al. 2016, Novara et al. 2019).

2 Preconditions of 3D LPT using Shake-The-Box techniques

For both, PTV/LPT and tomo PIV, three-dimensional reconstruction methods of the particle positions, as well as for the estimation of the underlying flow fields several techniques have been developed. First approaches used manual tracking of singular particles, dating back to the beginning of the 20th century (Nayler and Frazer (1917)). 3D LPT techniques are operating with particle detection-, triangulation- and tracking algorithms generating in a first step only 3D position coordinates (x , y , z) (and intensities (I)) per particle and time step t_i , while with classical 3D PTV techniques particle tracks with a respective ID are extracted by a subsequent tracking step (Malik et al (1993), Ouelette et al. (2006)). With 3D positions and track IDs (w/o intensity) we have the sparsest possible representation of the captured flow information as given by all individually reconstructed particles over time. On the other side, tomo PIV approaches store reconstructed particles as intensity blob distributions in large arrays of voxels ($\sim 10^9$) for each time step and subsequently operate with local 3D cross-correlation schemes of two (or more) temporal subsequently captured and reconstructed intensity distributions (including ghost particles) to determine a regular grid of (low-pass filtered) 3D3C velocity vectors. This procedure has a quite high demand on RAM capacity and computational resources, but has been proven to be robust and reliable, when reducing related known bias errors. However, from a perspective of signal processing the use of 3D LPT techniques is advantageous for several aspects, providing the usable particle concentration is at least comparable to tomo PIV. The key advantages are a) a faster processing with a sparse representation of particles by using only position, intensity and track ID values and b) a maximization of the achievable flow information by tracking each individual particle over time with an improved position, local velocity and acceleration value along the whole track. While the raw position accuracy is already above the one achievable using tomographic methods, the use of a temporal filter function (e.g. 3rd order -B-spline or polynomial) and its derivatives yields further improvements. A third advantage will be demonstrated later when addressing the (advanced) IPR technique, which allows individual single-time-step reconstructions at even higher particle image densities compared to tomo PIV, while producing significantly higher particle position accuracies (Wieneke (2013), Jahn et al. (2019)).

Although the 3D LPT method STB relies on particle- position and tracking based reconstruction techniques, which in a first view seems very similar to classical PTV approaches, important preconditions for STB have been developed in the process of improving the tomographic PIV technique. The multiplicative nature of MART used in tomo PIV (Atkinson (2009), Scarano (2013)) requires subpixel calibration accuracies in order to reduce the ghost particle level: MART is projecting the particle image intensities into the voxel space according to the camera calibration along each pixel's lines-of-sight (LOS). This actually acts as a logical "AND" operation, meaning that, in order to produce finally a 3D intensity blob representing a (possible) particle position at the intersection of all camera LOS in the object space, the pixels

corresponding to the intersecting LOSs on cameras 1,2,3 and 4 must exhibit non-zero grey-value intensities. A slight de-calibration of the LOS of all cameras by only one pixel, would already create a completely different intensity blob distribution inside the volume with much more ghost particles and less true reconstructed particles due to the underlying multiplicative nature of MART.

Therefore, an important step for improving the reconstruction quality was the Volume Self Calibration (VSC) technique (Wieneke (2008)), which enables a higher precision of the intersection of several camera LOS by correcting the initial target-based 3D camera calibration with an error in the order of a few pixels to below ~ 0.1 pixel. VSC is using particle triangulation procedures with sparse particle images (low ppp-values) and calculating their back-projected displacement fields for each sub-volume within the individual cameras for correcting the initial 2D-3D mapping functions. An accurate 3D camera calibration system is mandatory for a successful MART-based particle field reconstruction, which on the PTV side was not identified to be sufficiently critical, mainly because PTV tackles much lower particle image densities at < 0.01 ppp (instead of ~ 0.05 ppp for Tomo PIV) and therefore, the epi-polar line intersection error distance could be treated in a relatively flexible manner (~ 2 px). Although the triangulation idea of nearly intersecting epi-polar lines is already an intrinsic part of the PTV technique the particle based VSC technique was invented within the frame of the voxel based tomo PIV developments.

Furthermore, building upon VSC, a calibration procedure of the optical transfer function (OTF) (Schanz et al. 2013a) of the particles in each sub-volume was developed, accounting for the often varying imaging conditions across the measurement volume between the different cameras and within a single camera frame, thus further reducing the number of reconstructed ghost particle due to more precise projection of the particle image related intensity distribution into the voxel space and the respective subtraction of the found true particle after reprojection into the original images. Note that VSC and OTF are both particle-based in nature and directly use particle imaging properties. This not only makes the gained information ready for use in particle position-based reconstruction schemes, like IPR and STB, but also the necessary algorithmic structures (like peak detection and triangulation) can be directly used within the IPR algorithm, thereby reducing the development effort.

Potentially, the higher the concentration of the seeding the better the spatial resolution of the reconstructed flow field. While high particle numbers are desirable, the reconstruction process gets increasingly difficult and a growing number of non-existing particles (ghost particles) arises in the reconstruction due to the underdetermined nature of the problem. This process heavily depends on the number of imaged particles (N_P), the diameter of the imaged particles (or, more generally, the image source density), the number of cameras (N_C) and the resolution of each camera (Elsinga et al. (2011)).

In following section, the IPR technique and its recent advances will be presented as a major step towards the dense 3D LPT method STB. The advanced IPR technique (Jahn et al. 2021) has been assessed by reconstructions of 3D particle distributions from synthetic single snapshot particle images at up to 0.16 ppp using four virtual camera projections w/o typical imaging noise. Furthermore, results and comparisons of various LPT codes, reconstructing 2- pulse, 4- pulse and time-resolved imaging cases of synthetic flows at various ppp-levels, from the 1st International Challenge on Lagrangian Particle Tracking and Data Assimilation can be found

at Sciacchitano et al. (2021) (including those from the IPR and STB code developments at DLR presented here and the STB code available at LaVision GmbH).

3 Advanced Iterative Particle Reconstruction

The method of Iterative Particle Reconstruction (IPR), introduced by Wieneke 2013, constitutes a major step towards Lagrangian Particle Tracking in densely seeded flows (Schanz et al. (2013b), Schanz et al. (2016)). Here we present novel approaches in several key aspects of the algorithm, which, in combination, triple the working range of IPR in terms of particle image densities. The updated method is proven to be fast, accurate and robust against image noise and other imaging artifacts. Most of the proposed changes to the original processing are easy to implement and come at low computational cost (for more details see Jahn et al. 2021).

IPR introduces an iterative approach to the triangulation procedure, with an intermediate position optimization and a corresponding removal of apparent ghost particles. These measures allowed to reconstruct particle point clouds from images initially with $N_I \leq 0.05$ ppp with reduced occurrence of ghost particles and higher position accuracies, compared to a MART reconstruction. The position optimization (here also termed ‘shaking’ due to the process of adjusting the particles positions in 3D space) is performed for each particle individually using image matching techniques; hence the calibration demands, concerning both the triangulation accuracy and the fidelity of the optical transfer function, are very high. As a consequence, an accurate VSC and OTF calibration is mandatory for IPR.

The underlying problem IPR aims to solve is to find a 3D particle distribution $(x, y, z, I)_i$ for $i = 1..N_p$, that explains a set of measured particle images, simultaneously recorded from multiple cameras from different directions. Here, x, y and z are the three spatial coordinates, I is the particle intensity and N_p the number of particles. To optimally solve this problem, IPR relies on precise knowledge of the viewing directions of all cameras and the imaging properties of the particles on each camera (Optical Transfer Function).

Given a state vector $(x, y, z, I)_i$, a camera mapping function $(u_i, v_i) = cam((x, y, z)_i)$ and a calibrated 2D particle image shape in dependence on the 3D coordinate $shape_i = of((x, y, z)_i)$, the reprojected camera image can be calculated by adding up the images of all particles:

$$rep_{uv} = \sum_i I_i shape_i(u - u_i, v - v_i) \quad (1)$$

with (u, v) being the 2D image coordinates. Due to the finite spatial expansion of the particle’s shapes, it is sufficient to only sum nearby particles for each pixel uv (a typical particle image diameter is 3-6 pixel). To find $(x, y, z, I)_i$, cam and otf in a way that rep_{uv} optimally assembles the measured images img_{uv} , the following concepts are applied:

The IPR method is an iterative procedure for the accurate reconstruction of particle positions. At the center of the IPR is the state vector $(x, y, z, I)_i$, containing the interim estimation for the position x, y, z and intensity I of each of the N_p presumed particles. During the reconstruction this state vector gets iteratively replenished with potential missing particles, purged from potential ghost particles and optimized in terms of position and intensity to best match the observed measurement.

- The IPR workflow is depicted in Figure 1 and the individual steps are described below:
1. All peaks present on the current images (being the original measurements in the first iteration and residuals in all following iterations) are detected using a 2D peak detection scheme.
 2. Potential 3D particle positions are triangulated from the found 2D peaks and added to the state vector. Due to ambiguities, inaccuracies of the peak-detection, noisy images and overlapping particle images, the newly added particles might exhibit a significant fraction of ghost particles and suffer from positional inaccuracies.
 3. The position and intensity of each particle are optimized individually to best match the measured images on all cameras simultaneously using image matching methods.
 4. The state vector is filtered for potential ghost particles, mainly based on the particle intensity (deleting particles that fall below an intensity threshold).
 5. The current interim solution is rendered by projecting all particles onto virtual camera images, as given by the calibration. These projected images I_P are subtracted from the recorded measurement I_{rec} , yielding the residual images I_R .




Figure 1: Scheme of the IPR algorithm. Outer iteration indicated by large arrows. Inner iterations of position/intensity optimization (shaking) and filtering indicated by small arrows (step 3 and 4) (Schröder and Schanz (2023) cited after CC4.0).

The process starts anew with peak detection in I_R . Steps 1-5 are repeated until convergence is reached or for a predefined number of iterations. The key updates to the IPR process introduced in this work are:

1. For the image matching procedure (shaking), the numerical estimation of the derivatives is substituted by analytical derivations of the cost function. This step enables a direct determination of the Gradient (and the Hessian), thereby limiting the number of required evaluations of the cost function (see sections 3.2).
2. It is demonstrated that the completeness of the triangulation process benefits from using several permutations of the camera order, as the triangulation result strongly depends on the first two cameras and their relative angle (see sections 3.1.1)

3. Instead of using a fixed triangulation error ε for all iterations of the IPR procedure, a linearly increasing value of ε is applied, which aids in suppressing the generation of ghost particles, as shown in section 0
4. It is shown that ghost particles can be very effectively suppressed by excluding the camera(s) in which a particle appears the brightest for the optimization of its intensity (see sections 3.3)
5. In addition to the particle-wise optimization procedure (shaking each single particle independently), an ensemble approach is presented, which optimizes the position of all particles at once ('global shake'). Additionally, this procedure can be used to optimize the camera positions, thus representing a mean to perform an in-situ single image calibration correction (see Jahn et al. 2021).

3.1 Triangulation

From the peaks detected on the images of the different cameras, the 3D positions of the underlying particle cloud have to be estimated. To this end, given a certain peak on one camera, the corresponding peaks on the other cameras, belonging to the same particle, have to be identified. This can be achieved by the triangulation procedure, which is based on epi-polar geometry and estimates a set of possible 3D particle positions given detected 2D particle peaks from a set of multiple cameras. The algorithm first selects one peak on Cam 0. The corresponding 3D line of sight is calculated and projected onto Cam 1. On Cam 1 all peaks within a given tolerance ε (allowed triangulation error) around this line-segment are queried. For all found peaks, a 3D position is calculated by intersecting the corresponding Line-of-Sights. These positions correspond to possible 3D positions of the searched particle. To further filter these 3D positions, they are projected onto the remaining cameras. The position is rejected, if no peak can be found on the corresponding image within ε . When a particle is visible on a pre-defined number of cameras, a 3D position is estimated as the point that minimizes the quadratic distance to all corresponding 3D line of sights.

Several effects limit the abilities of the triangulation procedure:

- With increasing particle image densities, the reconstruction of false (ghost) particles rises significantly.
- Overlapping particle images and image noise inhibit a precise localization of 2D peaks, leading to real particles not being triangulated or being displaced, even if only one camera is affected.
- In case the first two cameras are separated by only a small space angle, the projected line of sight becomes very short and the accuracy in the depth direction is compromised, inhibiting a detection of the peak on the other camera images

The IPR procedure as introduced by Wieneke 2013 tries to counteract some of the mentioned drawbacks:

- The iterative approach of triangulating on residual images and shaking the particles to optimal positions lessens the effects of overlapping particle images and noise. Particles that could not be triangulated in the initial iterations might appear later, once other particles, whose images caused an overlap-situation, have been triangulated and their

images have been removed from the residual image. Then all particles can be optimally positioned by the shaking procedure.

- In order to lessen the problem of overlapping particles, Wieneke 2013 introduced some final triangulation iterations, which operate on a reduced set of cameras (a particle needs to be visible only on $N_C - 1$ cameras, instead of the full number of cameras N_C). This way, particles whose image is compromised on a single camera (noise, overlap with other particle images) are still being triangulated. The drawback is that the creation of ghost particles is amplified. Therefore, this measure was only applied during the last IPR iterations, where the residual images already show a reduced particle image density.

Here, we introduce two additional measures to improve the triangulation results: performing multiple triangulations, each working with a permutation of the camera order and a linear increase of the allowed triangulation error ε :

3.1.1 Camera permutations

As indicated before, the first two cameras have a special role in the triangulation process. The solid angle between these cameras determines how accurate the 3D position can be estimated, which is used to query the other cameras. Also, in case the currently investigated peak is overlapping with another one on any of these two cameras, the accuracy of the assumed 3D position can be heavily corrupted. With increasing particle number, a successful identification of the majority of particle images becomes less and less likely due to particle image overlap (Cierpka et al. 2013). Furthermore, for a particle to be recognized, its image needs to be identified in both of the first two cameras. This can be especially problematic, if certain regions - or even the whole image - show reduced intensity on cameras 1 or 2 (due to e.g. an unfavorable viewing angle).

To counteract all these effects, we substitute a typical triangulation, using a fixed camera order (e.g. 1-2-3-4), with a sequence of triangulations with permuted camera order, such that all camera combinations appear once on first and second position (e.g. 1-2-3-4; 2-3-4-1; 3-4-1-2; 4-1-2-3; 1-3-2-4; 2-4-3-1). The order of the remaining cameras is of no relevance. Only such particles are accepted that do not have a direct neighbor in the cloud of already triangulated particles; typically, a minimum distance of 1 px is applied. This way, multiple triangulations of the same particle are avoided. This proximity filter is typically not applied in the case of multi-exposed recordings (see section 5), where the same particle appears multiple times (e.g. two times for double-exposed recordings) within the same 3D point-cloud and the distance between the particle's realizations depends on the local particle displacement. Linear increase of triangulation error

The IPR implementation as introduced by Wieneke 2013 applies the same value for the allowed triangulation error ε for all IPR iterations. Instead, we apply a ramp of ε , linearly increasing within the N_{IPR} outer iterations: $\varepsilon(i) = \varepsilon(1) + i/(N_{IPR} - 1) * \Delta\varepsilon$, with $\Delta\varepsilon = \varepsilon(N_{IPR}) - \varepsilon(1)$. The start and end-values, $\varepsilon(0)$ and $\varepsilon(N_{IPR})$ can be defined according to the experimental conditions. This approach minimizes the creation of ghost particles, especially at the first iterations where the residual image is still very populated and allows to work with a reduced set of cameras for all iterations, reducing the number of parameters.

3.2 Optimization of position and intensity (shaking particles)

The triangulation process is limited by the accuracy of the peak detection and does not suppress ghost particles. Therefore, the positions and intensities of the particles are further optimized using local image matching schemes, which iteratively shift the position and intensity of the particles to minimize the difference between the original camera images and the reprojection of the particles (being the residual images).

The original IPR as introduced in (Wieneke 2013) minimizes the cost function by iteratively optimizing the position x_i, y_i, z_i and intensity I_i of each particle individually by using a procedure roughly equivalent to a Newton-fitting algorithm with numerically determined first and second derivatives for x_i, y_i, z_i and a redetermination of the current optimal intensity I_i in each step. To this end, each single particle is displaced in all three directions of space by e.g. ± 0.1 pixels and the local cost is determined on all of these positions. A 2nd order polynomial fit is applied to the residuals and the particle is moved to the positions of minimal residual (with a maximum step of 0.1 px) in each direction of space.

This approach requires to estimate the local cost around each particle 7 times per camera in each iteration. Depending on the image size and the form of the OTF this can be an expensive operation. Therefore, it is desirable to minimize the number of estimations. Secondly, for best results it is necessary to estimate the numeric derivative with high precision, for which the recommended procedure is to use analytic derivatives whenever possible.

We substitute the optimization using numerical derivatives with analytic derivatives of the cost function. For this we first calculate the derivative with respect to the projected particle positions u_i and v_i and the particle Intensity I_i of particle i on camera j :

$$\frac{d}{du_i} cost_j = 2I_i \sum_{uv} res_{uv} \frac{\delta shape_i}{\delta u} (u - u_i, v - v_i) \quad (2)$$

$$\frac{d}{dv_i} cost_j = 2I_i \sum_{uv} res_{uv} \frac{\delta shape_i}{\delta v} (u - u_i, v - v_i) \quad (3)$$

$$\frac{d}{dI_i} cost_j = -2 \sum_{uv} res_{uv} shape_i (u - u_i, v - v_i) \quad (4)$$

IPR does not optimize the projected particle positions u_i and v_i , but the 3D position x_i, y_i, z_i and the intensities I_i . u_i and v_i are functions of the 3D position of particle i . The exact dependency is defined by the camera-model that maps the 3D positions onto pixel coordinates. It is evident that all derivatives and the local cost can be simultaneously estimated while only iterating once over the immediate proximity of the projected particle position.

3.2.1 Steepest-Descent Method (SDM)

The Steepest Descent Method (Fletcher and Powell 1963) offers the easiest way to use the gradient of the cost function for optimization of the particle position parameters. Let $\vec{x}_i^T = (x_i, y_i, z_i, I_i)$ be the state vector of particle i . We determine the gradient $\vec{\nabla}_i cost(\vec{x}_i^T)$ of the cost function with respect to x_i, y_i, z_i and I_i on each camera. First, we determine the optimal intensity

of particle i on each camera (as discussed in section 3.3) and sum up all derivatives. This summed up gradient gives us the direction, along which we have to move our particle in order to reduce the cost: $\vec{V}_c = \vec{\nabla} cost(\vec{x}_i^T) / |\vec{\nabla} cost(\vec{x}_i^T)|$. In this implementation we apply a simple line search along this line. With a predefined maximum step width d_{max} we query the cost function at three locations

$$1.0 * d_{max} * \vec{V}_c, \quad 0.1 * d_{max} * \vec{V}_c, \quad 0.01 * d_{max} * \vec{V}_c.$$

The particle is simply moved to the location with the lowest value of *cost*. A typical value for d_{max} is 0.4 px in the first IPR iterations. In later iterations d_{max} can be reduced for the sake of accuracy, “shaking” towards an optimal position.

3.3 Intensity update

Along with the optimization of the position, the intensity of each particle is updated to best fit the current residual images. Wieneke (2013) applied a multiplicative corrections scheme, working with the ratio of the particle-augmented residual image to the calibrated particle shape.

In the current implementation, an optimal intensity update is approximated using the cost function. The (local) minimum of *cost* with respect to I_i for a single particle i in camera j can be determined using a one-step Newton method.

The total intensity update $I_{i,new}$ is then calculated as the average of all relevant cameras j . In order to avoid oscillations of the intensity, a dampening factor γ is introduced and applied as follows: $I_{i,new} = \gamma * I_{i,new} + (1 - \gamma) * I_i$. A value of $\gamma = 0.5$ has generally shown good results.

In both Wienekes 2013 approach, as well as earlier in-house DLR implementations, the intensity update was calculated as an average of all cameras, in which the current particle is visible. In the current work we show that by ignoring certain cameras for the update, a much more effective suppression of ghost particles can be achieved. After calculation of the new intensities for all cameras the results are sorted and the N_{Ci} cameras that show the highest intensities are not used for the determination of the average. This results in a slight general underestimation of the particle intensities. However, the impact on the intensity of ghost particles is much higher, since their intensity is often supported by bright peaks on single or few cameras. If these cameras are removed from the intensity calculation, the intensity of the ghost particle drops significantly and the particle is removed during the filtering step using a simple threshold. Another likely scenario is that the brightest camera might show an elevated intensity due to an overlapping situation. In this case, this camera would bias the intensity to higher values, which might result in an aggravated identification of the second peak, with which the image of the current particle is overlapping.

For every outer IPR iteration (starting with peak detection and triangulation), N_{shake} iterations of shaking are performed, each of them optionally followed by a filtering step (see below) – the combination of shaking and filtering we call inner iterations. At the beginning of every inner iteration the residual images res_{uv} are determined anew to reflect the changes in position and intensity from the last iteration.

3.4 Filtering

Each inner iteration of shaking can be followed by a filtering step, where presumable ghost particles are removed from the current interim solution. The most important filtering procedure is removing all particles below a certain intensity threshold I_{min} . This measure is also proposed in the original IPR paper. Ghost particles usually only poorly overlap with measured peaks or draw their intensity mostly from one or two cameras only. Furthermore, the true particles drain intensity from the ghost particles during the optimization-step. For that reason, ghost particles tend to have a lower intensity than true particles, increasingly so with increasing number of iterations (see Figure 2). Opposed to Wieneke 2013, we do not define I_{min} as a percentage of the average intensity, but as a fixed value, which proved more stable in most situations.

Other filtering methods are additionally applied. Particles that are found in close proximity in 3D space (below a threshold of typically 1 pixel) are merged, as it can be assumed that they represent the same true particle. The darker particle is deleted and its intensity is added to the brighter particle. As previously mentioned, the proximity filter is not used in case of multi-exposed recordings (see chapter 5). Also, particles need to be visible within at least two cameras. As soon as they leave this volume, they are deleted.




Figure 2: Particle intensity histograms (broken down into ghost particles (red) and true particles (green)) during the convergence of the IPR algorithm at moderate particle image density. It can be seen how the intensity of ghost particles is distributed on top of the real particles. A separation of the two kinds at a certain threshold becomes more and more effective. The black curve marks the intensity distribution of the ground truth.

3.5 Rendering of residual images

After triangulation and position optimization with filtering, the residual images res_{uv} , are created. These images are the basis for the following outer IPR iteration, as the peak detection will now be carried out on res_{uv} . In order to more effectively erase known particles from the images a factor $f_{pt} \geq 1$ can be used on the particle intensity while creating the projected images prior to the triangulation of new particles on the residual images. Rendering the residual images is not only used in step 1 before the triangulation procedure, but also internally in step 4 within every inner iteration of the shaking process (here $f_{pt} = 1$ is used).

3.6 Assessments using synthetic test-cases

First the performance of the algorithm will be assessed on images free of any defects (case I) and typical measurement noise (case II) (see Figure 3). For case I the particles have a constant intensity of 1000 counts in the volume, which is distributed over all pixels of the particle image

in each camera (leading to typical peaks of 300-400 counts). All pixels of the generated measurement are rounded to the nearest integer. No image pre-processing is applied. The main IPR parameters are indicated in the section above. Specific parameters for this case are: the intensity threshold for identifying peaks on the (residual) images is $P_{min} = 50$ counts, the intensity threshold in the volume, below which a particle is discarded as a ghost particle, is $I_{min} = 200$ counts (i.e. 20 % of the average particle intensity).




Figure 3: Excerpts of the synthetic camera images at different particle image densities (ppp-values) a) without noise and uniform particle intensities (case I) and b) with noise (case II). For both effects of particle overlap become increasingly visible at high ppp-values

For case II (typical noise), the particles have a mean intensity of 1000 counts in the volume, with a Gaussian distribution with $\sigma^2 = 250$ counts, leading to image peaks of different brightness in the range of ~ 200 -500 counts. A constant Gaussian noise floor with a mean of 40 counts and rms width of 25 counts is added to the image. To further simulate photon shot noise, a Poisson process is applied for every pixel, thereby adding noise proportional to the intensity of individual pixels.

Figure 4 shows the convergence behavior of IPR using the Steepest-Decent Method (SDM) for varying particle image densities, when applied to the described noise-free (left) and the typical moderate image noise (right) synthetic dataset. Shown is the ratio of correctly identified particles R_f (reconstructed within a search radius of 1 pixel around the true particles; normalized with the number of true particles and abbreviated as ‘founds’), the ratio of ghost particles R_g (reconstructed particles, for which no true particle could be found in a radius of 1 pixel; also normalized with the number of true particles and abbreviated as ‘ghosts’) and the average accuracy of the particle positioning of the correctly found particles Δ_f (calculated as the average distance of the re-projected points to the original peak position).

For particle image densities ($N_l \leq 0.12$ ppp) the method converges to a stable value of $R_f \approx 1$ in less than 12 to 16 iterations for the clean (Figure 4, left) and noisy cases (Figure 4, right), respectively. Even at $N_l = 0.14$ ppp the reconstruction converges after 15 iterations for the clean case, where only 8 particles out of 136,477 are not found. For the highest seeding density, convergence is foreseeable, but not achieved within the 20 iterations used.

For low values of N_l the number of ghost particles quickly diminishes after a few iterations. At higher particle image densities though, a rise in ghost numbers occurs within the first five iterations, which can go up to $R_g = 1$. This can be understood, as in these cases many particles could not be identified in this early state and the uncertainty of the found ones is high. Therefore, the residual images will show relevant intensities, even in spots where the particles were already triangulated. These residual intensities will be picked up by the peak finder and new false particles will arise.

However, with the identification of more and more true particles, the ghost level is gradually reduced. As soon as all real particles are found, also the ghosts have disappeared. Indeed, at the clean case only for $N_l = 0.16$ *ppp* ghosts are still present after 20 iterations (as this case is not yet converged). As well, for the noisy case (Figure 4, right) the convergence still seems to be possible for $N_l = 0.14$ *ppp* and for all lower *ppp* values the ghost level have been brought down to ~ 0 . However, due to the influence of the image noise, which deteriorates the position optimization scheme by “shaking”, a slight decrease of the position accuracy with *ppp* can be noticed, Δ_f is growing slowly to values slightly bigger than 0.1 px for the converged phases. As the particle image noise for case II is mimicked closely related to typical background optimized experimental conditions a value of 0.1px to 0.15px can be assumed for individual time step reconstruction qualities, when the vast majority of particles have been found and the ghost level is negligible. We will see later, that optimal temporal filtering after STB track evaluation leads to a further reduction of Δ_f by a factor of ~ 2 .




Figure 4: Convergence of single-image IPR reconstructions for various *ppp*-values for clean (left) and typical image noise (right). Shown are the number of correctly identified particles R_f (top), the number of ghost particles R_g (center) (normalized with the number of

true particles), and the position accuracy of the true particles Δ_f (bottom) as a function of the iteration number. After iterations 15 and 18 the maximum step width is reduced.

In order to give a better insight into the working principle within a single IPR iteration, Figure 5 shows the development of ghost- and found true particles for the clean case at $N_I = 0.14 \text{ ppp}$, resolved for the triangulation and all shake iterations (blue and orange lines). Additionally, the end-result of each iteration is connected (black and red lines, which then correspond to the results from Figure 4, left). Looking at the first few outer iterations it becomes obvious that initially high numbers of weak ghost particles are triangulated. These remain in the particle cloud for the first three shake iterations (the short flat parts of the curve at the beginning of each inner iteration). As soon as the filtering of the particles is performed, the majority of these is deleted, as their intensity has been further reduced by the shaking. From there, the number slowly decreases, until the next triangulation drives the number up again. At the same time, the number of found true particles is barely reduced by the shaking. With every outer iteration more true particles are found, reducing residual peaks, which in turn gradually minimizes the number of triangulated (ghost-) particles, until all particles are identified and zero ghost particles remain.




Figure 5: Convergence of found true and ghost particles for case I at $N_I = 0.14 \text{ ppp}$. Particle statistics have been evaluated after triangulation and every inner (shake) iteration (10 for each of the 20 outer iterations). The orange and blue curves show these inner-iteration-resolved results for found true and ghost particles, respectively. The red and black curves connect the end-points of each iteration (corresponding to Figure 4, left). Creation and deletion of particles (resp. ghosts) within each outer IPR iteration become noticeable.

The position accuracy is very much dependent on the presence of ghost particles and the detection of true particles. As soon as these two have converged, high levels of accuracy are reached. The reduction of the maximum step width d_{max} at iteration 15 is very visible. E.g. for the 0.12 ppp case, the position accuracy dips from a plateau at 0.04 px to 0.007 px. The second reduction of d_{max} at iteration 18 yields a further reduction of Δ_f to 0.005 px at this seeding density. At 0.16 ppp it becomes evident that a reduction of d_{max} leads to a reduction in the speed of convergence, when looking at R_f and R_g . Therefore, this measure should only be applied after convergence is reached.

Compared to the results of Wieneke 2013, which were gained using a comparable synthetic experiment, the range of particle image densities has been greatly extended. Wieneke reported a usable particle image density of $N_I = 0.05 \text{ ppp}$ and a sharp decrease of reconstruction quality when going beyond this value using 16 outer and 6 inner iterations. Here we find that images at $N_I = 0.14 \text{ ppp}$ produce results completely void of ghost particles and very few missing real particles (far below Wienekes values at 0.05 ppp). The position accuracy of the advanced IPR is much higher, with values $\Delta_f \leq 0.01 \text{ px}$ for all converged clean cases.

4 3D LPT using (time-resolved) Shake-The-Box

Both previous methods, tomo PIV and 3D PTV, first create a 3D particle reconstruction for each individual time step *separately*¹. The classical triangulation approach of PTV and the MART based tomo PIV reconstruction schemes lead to an increasing number of ghost particles with increasing ppp-values of single-particle images. Then, in a second step, the 3D particle distributions of subsequent time steps or along a whole time-series are connected for estimating particle velocities. PTV is applying tracking schemes along the temporal domain of the 3D particle position distributions (e.g. Malik et al. (1993), Dabiri and Pecora (2019)). Hereby, high ghost particle levels at high ppp-values produces massive ambiguities for the track-connection schemes. Therefore, 3D-PTV has been limited to ppp-values typically below 0.01 ppp even if various predictive steps and decision-tree strategies for the track building have been developed (see e.g. Ouellette et al. (2006), Xu (2008)). For tomo PIV the second step is applying a local cross-correlation scheme on the intensity blob distributions for estimating the velocity vector field from two sub-sequent or even several time-steps (Raffel et al. (2018)). Iterative cross-correlation with volume deformation is robust against noise and provides quite reasonable results even at relatively high ghost levels, but only low-pass filtered velocity vectors can be gained (see Kähler et al. 2012) and correlated ghosts are producing significant bias errors stemming from projections of particles in the bulk flow (Elsinga et al. (2011)).

Different to that, the key idea of STB is, exploit the *temporal coherence* of the relatively sparse 3D particles trajectories following the flow, which is already encoded in the time series of particle images. Although it remains an ill-posed reconstruction problem (2D \rightarrow 3D) we can use properties of the imaging (e.g. OTF) and the flow physics: *sparsity* of particles in the 3D distribution (mean particle distances are larger in 3D than in projected 2D images), *temporal resolution* of particle image shifts (i.e. sampling rate below Kolmogorov time scales) and *smooth 3D trajectories* are given (jerk: D^2U/Dt^2 of the particles in the flow is small for the vast majority of tracks). Additionally, particles cannot disappear over time inside a 3D volume, so that in STB (after a short four-time-step initialization phase and for each further time step) a *time-marching* approach is applied: *First*, a prediction of the 3D particle positions towards time step t_{n+1} by extrapolation of temporal fitting functions along all found tracks in previous time-steps (e.g. t_{n-3} to t_n) is applied, and then, *secondly*, a correction of the estimated 3D positions is done by simultaneous re-projection and image-matching („shaking“) on all cameras for the given time step t_{n+1} . The correction is typically below 1.2 px (less than half the particle image

¹ After a first step, the MTE MART approach made an exemption on the tomo PIV side: The approach iteratively involves two subsequent time-steps in order to reduce the ghost particle problem (Novara et al. (2010)). S-MTE extended this idea to a time-marching scheme, akin to the STB principle (Lynch et al. 2016).

diameter) even for relatively large particle image shifts >10 px, because of the limited jerk for the vast majority of particle tracks, while local position correction by “shaking” is very accurate (~ 0.1 px see chapter 3.6). Ghost tracks which are not coherent over several time steps cannot find respective intensities at the projected image positions at t_{n+1} and are erased when falling below an intensity threshold. In the following, all already found and extended true particle tracks from previous time steps can be subtracted from the original particle images at time step t_{n+1} , before a *first* 3D particle reconstruction with IPR need to be done on the less (and in time lesser) populated residual images. In the initialization (the first few time-steps), as well as during the further runtime, new tracks are added by identifying physically reasonable sequences of reconstructed particles (using IPR) in successive frames (typically four). These are then extended into further time steps. This procedure enables a massive reduction of the reconstruction complexity for each further time-step until convergence in kind of a self-optimizing process. The first implementation of these ideas into an STB algorithm has been presented at the ISPIV 2013 (Schanz et al. 2013b) and produced first 3D LPT results at high ppp-values of a turbulent periodic hill flow (Schröder et al. 2015).

Once again, STB basically *reverses* the typical process of evaluation: after the initialization step (meant to reconstruct a subset of tracks over the first four recordings), 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 and acceleration information along found tracks to create an (estimated) particle distribution at the next time step. The errors introduced by the estimation are small enough to be easily corrected using image matching; no further partner search is required for known particle tracks, i.e. the track (with given ID) is already correctly enlogated after the prediction and correction step. By this the flow physics itself is supporting the reconstruction process of particle trajectories.

Therefore, STB finally transforms the ill-posed global particle reconstruction problem (present for individual time-steps) into a local simultaneous particle position *and* tracking optimization scheme. The prediction-step allows for a severe reduction of iterations and therefore processing time, as well as significant gains in accuracy due to an efficient suppression of ghost particles (which are not temporally coherent). As soon as a sufficient percentage of tracks can be identified, the process is self-stabilizing; one enhancement in a reconstruction property (position accuracy, ghost particle ratio, percentage of found tracks) induces immediately improvements in the others.

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

1. Perform a temporal fit to the last k positions of tracked particles using an optimal Wiener filter based on an analysis of the track position spectrum (Gesemann et al. 2016).
2. Predict the position of the particle in t_{n+1} by evaluating the Wiener filter coefficients.
3. Shake the particles to their correct position and intensity, eliminating the prediction error (see chapter 3.2).
4. Find new particles, entering the measurement domain, on the residual images by IPR.
5. Shake all particles again to correct for residual errors.

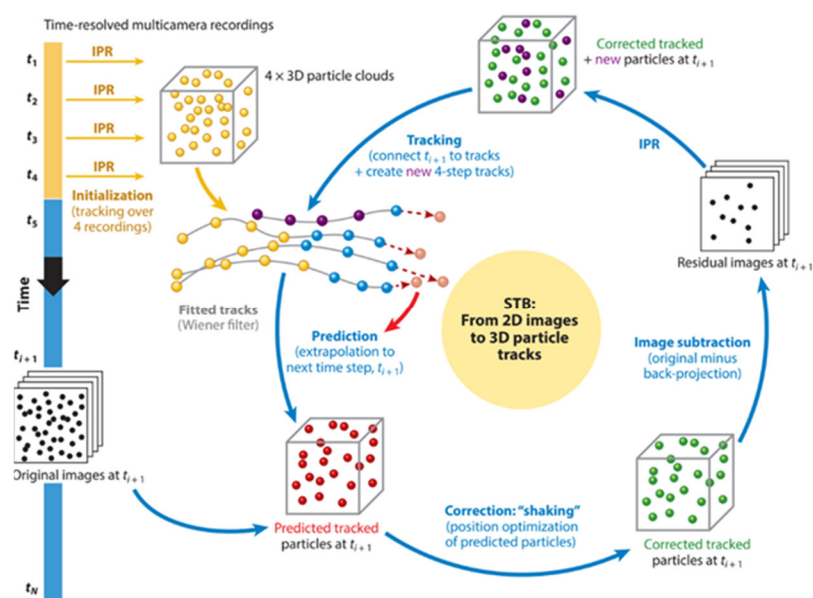
6. Remove particles if leaving the volume or if intensity falls below a certain threshold.
7. Iterate steps 4, 5 and 6, if necessary.
8. Add new tracks for all new particles identified within four consecutive time-steps

After such a processing of a single time-step within the image series 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 pre-solved after the prediction-step and only minor deviations have to be corrected and a few newly entering particles have to be triangulated.

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 dataset 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 provided by IPR), *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, the number of newly found particles is balanced by the number of particles leaving the volume). For the sake of simplicity of the method description we directly step here to the converged phase, which typically is reached after a few time steps building upon the high quality of the advanced IPR reconstruction (Jahn et al. (2021)) and the selective nature of temporally coherent true tracks against incoherent, because arbitrary generated ghost particles and tracks. For more details see Schanz et al. (2016).

4.1 Converged phase

The algorithm needs only a few 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 6: Flow chart of the Shake-The-Box time marching procedure embedding tracking with position prediction and -correction and single time-step iterative loops by IPR (top) (Schröder and Schanz (2023) cited after CC4.0) and of 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) (bottom).

Figure 6 (top) illustrates the flow chart of the Shake-The-Box procedure: The STB scheme is initialized with iterative particle reconstruction (IPR) processing of four subsequent time steps (t_1 to t_4) from multicamera particle images and subsequent temporal connection to 3D particle tracks (yellow). Next, the particle position of each track is predicted for the next time step ($t_i + 1$) by extrapolating suitable temporal fitting functions (red). Then, 3D position and intensity optimizations are applied by the shaking approach and residual images are calculated (green). An IPR process is applied to the residual images, yielding new 3D positions (purple); corrected predictions and new 3D positions are used to finally elongate the existing tracks to the current time step (blue) and identify new tracks from the yet untracked particles of the last four time-steps (purple, akin to the initialization) and so on until convergence.

In Figure 6 (bottom) a single time-step in the converged state and their impact on the residual image of a selected camera is shown. 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 Figure 6 (bottom) 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 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 residuals 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: Polydisperse 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 in space and time). 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.

4.2 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 (relative to the *Initialization* and *Convergence* phases), 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 dataset. 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. 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; more passes can be performed, if necessary.

4.3 Variable-time-step STB (VT-STB)

A recent development of STB is the variable-time-step (VT-) STB method (Schanz et al. 2021) which solves the problem of occurring ghosts in flows with a very high dynamic velocity range. When tracking particles in flows with almost quiescent (or negligible velocity gradient) regions (e.g. around jets or in recirculation bubbles of separated flows etc.) overlapping particle images along several camera LOS can be consistent over many time-steps so that the respective ambiguities of the triangulation-, tracking- and correction process of STB could lead to ghost

particle tracks. VT-STB provides a solution for this specific issue by iteratively reducing the time separation of the STB scheme and first track particles with a large time-separation (e.g. by using every 15th image) within a time series which guarantees a sufficiently large particle image shifts to accurately distinguish individual particles. The tracks of the first iteration are stored and used as input for the next STB iteration with a shorter time separation (e.g. by using every 5th image) and so on down to the full temporal resolution allowing to reconstruct the fastest tracks in the measurement volume.

4.4 Examples of STB applications

In the last few years a large-number of time -resolved and multi-pulse STB experiments have been conducted successfully either by using the DLR in-house code or the code implementation within the DaVis software from LaVision GmbH (a list of references would be quite long and can be explored e.g. at Google Scholar). One example of a successful experimental STB application to a large-scale Rayleigh-Bénard-Convection (RBC) experiment with a volume of 1.05 m³ using Helium-Filled-Soap-Bubbles (HFSB) and pulsed LED illumination at DLR Göttingen is described in Bosbach et al. 2021. The results have been used for an award-winning Gallery of Fluid Motion (GFM) movie contribution at an APS event, created by Godbersen et al. 2020. Three images presenting some Lagrangian and Eulerian flow impressions of the investigated turbulent RBC flow are displayed in Figure 7.




Figure 7: Particle tracks in $\sim 1\text{m}^3$ Rayleigh-Bénard convection: Half-cylinder cut of in total 550k instantaneous particles (left), a single particle trajectory of $\sim 4.4 \cdot 10^4$ timesteps (middle) and full cylinder iso-contours of Q criterion, calculated by FlowFit from the STB tracking result, indicating vortices (right), all color coded by vertical velocity component (movie at <https://gfm.aps.org/meetings/dfd-2020/5f5fe77f199e4c091e67bfe8>)

A second example of a STB and FlowFit application on a flow around a surface mounted cube with laminar and turbulent boundary layers as inflow conditions is described in Schröder et al. 2020. Here, 3D STB and high-resolution 2D TR-PIV have been used for investigating the spatial and temporal behavior of the flow around the cube and the results have been compared with numerical simulations using a Lattice-Boltzmann-Code (LBM). Three images of instantaneous- and mean Lagrangian and Eulerian flow features captured with dense 3D LPT by STB can be found in Figure 8.




Figure 8: Tracking result for $U_\infty = 0.4$ m/s laminar inflow. Left, top: Side view on a 3 mm slice in the middle of the cube. Particle tracks of 60 successive time-steps, color-coded by streamwise velocity. Left, bottom: View from above on a 1.3 mm slice near the top of the cube ($y = 9.5 - 10.8$ mm), color-coded by streamwise acceleration. Middle: Result of 3D mean velocity field for the laminar BL flow case at $U_\infty = 0.2$ m/s based on bin-averaging approach with a total of 9,849,885 volume bins of size 0.25^3 mm³ color coded by u-component of the velocity. Right: Instantaneous snapshot of vorticity iso-surfaces ($\omega = 170$ /s) color coded by u-velocity for $U_\infty = 0.4$ m/s, as calculated by FlowFit from the STB tracking result

5 3D LPT for high-speed flows: Multi-pulse STB

The Multi-Pulse Shake-The-Box (MP-STB) method opened the possibility of extending 3D Lagrangian particle tracking (LPT) to the investigation of high-speed flows, where long time-resolved sequences of recordings are currently not available due to the limited acquisition frequency of high-speed systems. The MP-STB technique makes use of an iterative approach to overcome the limitations posed by the short observation time offered by a multi-pulse recording sequence (see Figure 9).




Figure 9: Iterative processing strategy for MP-STB; the IPR reconstruction of recorded/residual images, the tracking step and the evaluation of back-projected and residual images constitute a single MP-STB iteration. Complete four-pulse tracks are indicated in black, incomplete shorter tracks in gray

Multi-pulse sequences are typically obtained by synchronizing multiple illumination systems to generate bursts of laser pulses where the time separation can be freely adjusted down to less than a microsecond. Several strategies can be adopted for the recording of multi-pulse sequences; a dual camera system can be adopted to separate the single pulses onto the camera frames (either by means of polarization or timing), while the use of multi-exposed frames allows for the employment of a single imaging system, largely reducing the complexity and

cost of the experimental setup. The iterative reconstruction and tracking strategy for MP-STB can be successfully adapted to the case of multi-exposed frames.

Results suggest that, despite the increase in particle image density resulting from the double-exposed particle images, the adoption of multi-exposed recordings has the potential to become the technique of choice for the recording of multi-pulse sequences suitable for Lagrangian particle tracking in high-speed flows (for details of the methodology see Novara et al. (2016a), Novara et al. 2019). The performances of this novel strategy have been demonstrated by means of synthetic experiments (van Gent et al. 2017) and the technique has been successfully applied e.g. to the investigation of turbulent boundary layers (Novara et al. 2019), to data from a subsonic jet experiment at Mach 0.84 or $U_j = 290$ m/s (Manovski et al. 2021) and another jet flow experiments at up to $U_j = 195$ m/s (Sellapan et al. 2020).

An example of a multi-exposed MP-STB application for a high Reynolds number ZPG-TBL flow investigation (see details at Novara et al. 2019) and its basic evaluation steps will be described in the following. After the first IPR reconstruction of the two double-exposed frames, a velocity vector field is obtained by means of Particle Space Correlation (Novara et al 2016b). A 3D cross-correlation with multi-grid window deformation is performed between the two particle objects; unlike for tomographic PIV, where the cross-correlation is applied to the intensity distribution of a voxel space, the PSC makes use only of the particle peak location and intensity as obtained from IPR. The velocity field from PSC is used as a predictor for the particle displacement in the first STB iteration; a tri-linear interpolation is used to evaluate the predictor field at the location of the reconstructed particles (see Figure 10). Approximately 80% of the total number of expected tracks is found at the first MP-STB iteration. These tracks are used as a predictor for the following iteration, delivering additional 3,600 tracks.




Figure 10: Particle tracking scheme for multi-exposed recordings aided by a velocity predictor from Particle Space Correlation (yellow arrows).

As the initial prediction was obtained by means of a low-resolution cross-correlation approach (PSC), particles within regions characterized by strong velocity gradients (e.g. near-wall region) are not tracked in the first two iterations. As a consequence, the third iteration is performed without the aid of a predictor, but following a nearest neighbor approach with a larger search radius for the two-pulse track candidates' identification. Most of the 2,000 tracks identified during the third iteration are low-speed tracks found in the vicinity of the wall.

At the end of the MP-STB iterative strategy, approximately 49,000 instantaneous tracks are identified, corresponding to 92% of the expected number of particles estimated from the

recorded images. These results are confirmed by the similar performances of the reconstruction process when dealing with higher imaged seeding densities. In fact, when the same processing parameters are applied to the 0.046 ppp case, roughly 95,000 instantaneous tracks are successfully reconstructed; the increase in the number of retrieved tracks is consistent with the increase in seeding density.

An instantaneous result from MP-STB is presented in Figure 11 (left), where individual tracks are color-coded by the stream-wise velocity component; the presence of low- and high-speed streaks can be observed along the span-wise direction (Y axis). The scattered results extracted from the track fit at the mid-point of the individual particle tracks are used as the input for the FlowFit data assimilation algorithm (Gesemann et al 2016) in order to interpolate the result onto a grid of continuous 3D B-spline, which enables the visualization of the instantaneous 3D structures by analytical derivatives of the VGT. The typical flow structures of a turbulent boundary layer can be observed in Figure 11-right, with alternating low- and high-speed regions and several horseshoe-shaped vortical structures whose shape and inclination in the flow direction are compatible with the well-known Hairpin vortex model.

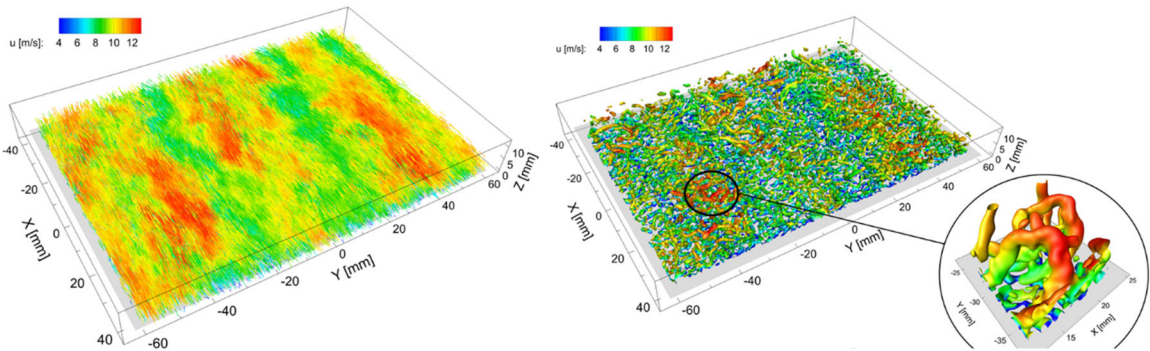


Figure 11: Instantaneous 4-pulse-track result from MP-STB at 0.046 ppp (multi-exposed recordings); approximately 95,000 tracks are shown color-coded with the stream-wise velocity component (along X-axis) (left) and FlowFit interpolation onto regular grid showing instantaneous flow structures visualized by iso-surface of Q-criterion colored by the stream-wise velocity component. Horseshoe-shaped vortices viewed from a different perspective visualized in the zoomed region indicated by the black circle (right)

6 Conclusions

3D Lagrangian Particle Tracking has evolved in the past 30 years from classical 3D-PTV techniques using epipolar line-based particle triangulation for reconstructing 3D particle position distributions per time step and a subsequent application of tracking schemes, which were restricted to ppp-values < 0.01 ppp, towards advanced Lagrangian Particle Tracking methods enabling processing of high particle image densities > 0.1 ppp with Shake-The-Box (Schröder and Schanz 2023). In case of high-speed flows where time-resolved imaging cannot be realized, MP-STB has been made available using simple double-frame image cameras (like for snapshot PIV) with various multiple illumination strategies. As described in the present work, STB made use of improvements developed in the frame of tomo PIV, which are already based on particle image information (e.g. VSC, OTF), the advent of the IPR technique and, as a key function, the temporal resolution of the captured flow physics in a kind of self-optimizing time-marching “prediction and correction” scheme. The processing time for 3D-STB is

significantly shorter ($< 1/20$) than for time-resolved tomo PIV cases, while an almost complete suppression of ghost particles resp. tracks can be provided. The accuracy of the time-resolved mode of STB has been assessed already by two independent PIV and LPT Challenges: At an early stage of the development during the “4th international Challenge on PIV” (see case D in Kähler et al. 2016) and more recently by the “1st International Challenge on LPT and DA” (see Sciacchitano et al. 2021). The two- and four-pulse or MP-STB mode has been additionally assessed by the second event and during the so-called pressure challenge (van Gent et al. (2017)). Further information can be derived from the Lagrangian particle trajectories delivering position, velocity and acceleration along a large number of trajectories. Thus, dense 3D LPT provide ideal data as input to data assimilation (DA) schemes (VIC+ (Schneiders et al. 2016), VIC# (Jeon 2022), FlowFit (Gesemann et al 2021, Godbersen et al. 2024), PINNs (Raissi 2019, Zhou and Grauer 2023) and to high-resolution bin averaging procedures (Agüera et al. 2016, Godbersen and Schröder 2020). All mentioned DA schemes can deliver the full time-resolved VGT and 3D pressure fields at a resolution limited by the mean 3D particle distances, while the high accuracy of the position estimation along the Lagrangian tracks (~ 0.1 px) allows for sub-pixel accurate one- and two-point statistics of the investigated flow volume (when a sufficient amount of LPT data is available).

7 Acknowledgements

This work was partly supported by the Deutsche Forschungsgemeinschaft (DFG) through Grant No. SCHR 1165/5-1/2 as part of the Priority Programme on Turbulent Superstructures (DFG SPP 1881, 1st and 2nd period) and the HOMER project from the European Union’s Horizon 2020 research and innovation program under grant agreement No. 769237.

8 References

- Agüera N, Cafiero G, Astarita T and Discetti S 2016 Ensemble 3D-PTV for high resolution turbulent statistics *Meas. Sci. Technol.* 27 124011
- Atkinson, C. & Soria, J. (2009), An efficient simultaneous reconstruction technique for tomographic particle image velocimetry. *Exp Fluids* 47:563–578
- Bosbach, J., Schanz, D., Godbersen, P. & Schröder, A. (2021), Spatially and temporally resolved measurements of turbulent Rayleigh-Bénard convection by Lagrangian particle tracking of long-lived helium-filled soap bubbles, *Proceedings of 14th International Symposium on Particle Image Velocimetry – ISPIV2021*, August 1–4, 2021, Doi:10.18409/ispiv.v1i1.208
- Chong, M. S., Perry, A. E. & Cantwell, B. J. (1990), A general classification of three-dimensional flow fields, *Phys. Fluids A*2(5), 765–777
- Cierpka, C., Lütke, B. & C Kähler, C. J. (2013), Higher order multi-frame particle tracking velocimetry. *Experiments in Fluids*, 54(5), 1533.
- Dabiri, D. & Pecora, C. (2019), *Particle Tracking Velocimetry*, IOP Publishing, doi: 10.1088/978-0-7503-2203-4
- Discetti, S. & Coletti, F. (2018), Volumetric velocimetry for fluid flows, *Meas. Sci. Tech.* 29, 042001, 26 pp, Doi: 10.1088/1361-6501/aaa571

- Duraisamy, K., Iaccarino G. & Xiao H. (2019), Turbulence Modeling in the Age of Data, Annual Review of Fluid Mechanics, Volume 51, 357-377
- Elsinga, G. E., Scarano, F., Wieneke, B. & van Oudheusden, B. W. (2006), Tomographic particle image velocimetry. *Exp Fluids* 41: 933–947
- Elsinga, G. E., Westerweel, J., Scarano, F., Novara, M. (2011), On the velocity of ghost particles and the bias errors in Tomographic-PIV, *Exp Fluids* 50(4): 825-838
- Fletcher, R., Powell, M. (1963), A rapidly convergent descent method for minimization. *The computer journal*, 6(2):163–168
- Gesemann, S., Huhn, F., Schanz, D., & Schröder, A. (2016) From noisy particle tracks to velocity, acceleration and pressure fields using B-splines and penalties. 18th Int. Symp. on Appl. of Laser and Imaging Tech. to Fluid Mech. Lisbon, Portugal, July 4 – 7 2016
- Godbersen, P. & Schröder, A. (2020), Functional binning: improving convergence of Eulerian statistics from Lagrangian particle tracking, *Measurement Science and Technology*, 31, 9, 095304, IOP Publishing
- Godbersen, P., Bosbach, J., Schanz, D. & Schröder, A. (2020), The beauty of turbulent convection: A particle tracking endeavor. *Gallery of Fluid Motion*, Doi: 10.1103/APS.DFD.2020.GFM.V0074
- Godbersen P, Schanz D & Schröder A. 2024a, Peak-CNN: improved particle image localization using single-stage CNNs. *Exp Fluids* 65, 153 doi:10.1007/s00348-024-03884-z
- Godbersen P, Gesemann S, Schanz D, Schröder A. 2024b. FlowFit3: Efficient data assimilation of LPT measurements, 21st International Symposium on Applications of Laser and Imaging Techniques to Fluid Mechanics, July 8-11, Lisbon, Portugal
- Herman, G. T., Lent, A. (1976), Iterative reconstruction algorithms, *Comput. Biol. Med.* 6 273–94
- Haller, G. (2015), Lagrangian Coherent Structures, *Annu. Rev. Fluid Mech.* 2015. 47:137–62, Doi: 10.1146/annurev-fluid-010313-141322
- Huhn, F., Schanz, D., Manovski, P., Gesemann, S. & Schröder, A. (2018), Time-resolved large-scale volumetric pressure fields of an impinging jet from dense Lagrangian particle tracking, *Exp Fluids*, 59, 5, 1-16
- Jahn, T., Schanz, D. & Schröder, A. (2021), Advanced Iterative Particle Reconstruction for Lagrangian Particle Tracking, *Experiments in Fluids* 62 (8), 1-24, doi: 10.1007/s00348-021-03276-7
- Jeon, Y. J. (2021), Eulerian time-marching in Vortex-In-Cell (VIC) method: reconstruction of multiple time-steps from a single vorticity volume and time-resolved boundary condition, *Proceedings of 14th International Symposium on Particle Image Velocimetry – ISPIV2021*, August 1–4, 2021, Doi: 10.18409/ispiv.v1i1.158
- Jeon YJ, Müller M, Michaelis D (2022) Fine scale reconstruction (VIC#) by implementing additional constraints and coarse-grid approximation into VIC+. *Exp Fluids* 63:70
- Kähler, C. J., Scharnowski, S. & Cierpka, C. (2012), On the resolution limit of digital PIV, *Exp In Fluids*, 52, pp 1629-1639
- Kähler, C. J., Astarita, T., Vlachos, P. P., Sakakibara, J., Hain, R., Discetti, S., La Foy, R. & Cierpka, C (2016), Main results of the 4th International PIV Challenge, *Exp in Fluids*, 57(6), 97

-
- Lynch, K.P., Scarano, F. (2015), An efficient and accurate approach to MTE-MART for time-resolved tomographic PIV. *Exp Fluids* 56, 66. Doi: 10.1007/s00348-015-1934-6
- Maas, H.G., Gruen, A. & Papantoniou, D. (1993), Particle tracking velocimetry in three-dimensional flows - part I: Photogrammetric determination of particle coordinates, *Experiments in Fluids* 15, 133–146, Doi: 10.1007/BF00190953
- Malik, N., Dracos, T. & Papantoniou, D. (1993), Particle tracking velocimetry in three dimensional flows - part II: particle tracking, *Exp Fluids* 15: 279–294, Doi: 10.1007/BF00190953
- Manovski, P., Novara, M., Mohan, N. K. D., Geisler, R., Schanz, D., Agocs, J., Godbersen, P. & Schröder, A. (2021), 3D Lagrangian particle tracking of a subsonic jet using multi-pulse Shake-The-Box, *Experimental Thermal and Fluid Science*, 123, 110346, 2021, Elsevier
- Naylor, J L, Frazer B A (1917) Preliminary report upon an experimental method of investigating, by the aid of kinematographic photography, the history of eddying flow past a model immersed in water. Tech. rep. of the Advisory Committee for Aeronautics for the year 1917-1918, London, 6
- Nicolas, F., Donjat, D., Léon, O. Le Besnerais, G., Champagnat, F. & Micheli F. (2017), 3D reconstruction of a compressible flow by synchronized multi-camera BOS, *Exp Fluids* 58:46, Doi: 10.1007/s00348-017-2325-y
- Novara, M., Batenburg, K. J. & Scarano F (2010), Motion tracking enhanced MART for tomographic PIV. *Meas Sci Technol* 21:035401
- Novara, M., Schanz, D., Reuther, N., Kähler, C. J. & Schröder, A. (2016), Lagrangian 3D particle tracking in high-speed flows: Shake-The-Box for multi-pulse systems, *Exp. Fluids* 57:128
- Novara, M., Schanz, D., Gesemann, S., Lynch, K.P. & Schröder, A. (2016b), Lagrangian 3D particle tracking for multi-pulse systems: performance assessment and application of Shake-The-Box, 18th Lisbon Symposium, July 4-7, Lisbon, Portugal
- Novara, M., Schanz, D., Geisler, R., Gesemann, S., Voss, C. & Schröder A (2019), Multi-exposed recordings for 3D Lagrangian particle tracking with multi-pulse Shake-The-Box. *Experiments in Fluids*, 60(3), 1-19
- Ouellette, N. T., Xu, H. & Bodenschatz, E. (2006), A quantitative study of three-dimensional Lagrangian particle tracking algorithms, *Exp Fluids* 40, 301–313, Doi: 10.1007/s00348-005-0068-7
- Raffel, M., Willert, C. E., Scarano, F., Kähler, C. J., Wereley, S. T., & Kompenhans, J. (2018). *Particle image velocimetry: a practical guide*. Springer.
- Raissi M, Perdikaris P, Karniadakis GE (2019) Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J Comput Phys* 378:686–707
- Richardson, L. F. (1922), *Weather Prediction by Numerical Process*. Cambridge University Press
- Scarano, F. (2013), Tomographic PIV: principles and practice, *Meas. Sci. Technol.* 24 012001, Doi: 10.1088/0957-0233/24/1/012001/
- 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, 10th Symp. on PIV, Delft, The Netherlands
- Schanz, D., Gesemann, S. & Schröder, A. (2016), Shake-The-Box: Lagrangian particle tracking at high particle image densities. *Exp. Fluids*, 57(5), 1-27. Doi: 10.1007/s00348-016-2157-1
- Schanz, D., Novara, M. & Schröder, A. (2021), Shake-The-Box particle tracking with variable time-steps in flows with high velocity range (VT-STB), Proceedings of 14th International Symposium on Particle Image Velocimetry – ISPIV2021, August 1–4, 2021, Doi:10.18409/ispiv.v1i1.65
- Schneiders, J.F.G., Scarano, F. (2016), Dense velocity reconstruction from tomographic PTV with material derivatives. *Exp Fluids* 57, 139 (2016). Doi:10.1007/s00348-016-2225-6
- Schneiders JFG, Avallone F, Pröbsting S, Ragni D, Scarano F (2018) Pressure spectra from single-snapshot tomographic PIV. *Exp Fluids* 59:57
- Schröder, A., Schanz, D., Michaelis, D., Cierpka, C., Scharnowski, S., Kähler, C. J. (2015), Advances of PIV and 4D-PTV” Shake-The-Box” for turbulent flow analysis—the flow over periodic hills, "Flow, Turbulence and Combustion", 95, 2, 193-209, Springer Netherlands
- Schröder, A., Willert, C., Schanz, D., Geisler, R., Jahn, T., Gallas, Q. & Leclaire, B. (2020), The flow around a surface mounted cube: a characterization by time-resolved PIV, 3D Shake-The-Box and LBM simulation, *Exp in Fluids*, 61, 9, 1-22, Springer Berlin Heidelberg
- Schröder A and Schanz D. 2023, 3D Lagrangian Particle Tracking in Fluid Mechanics, Annual Review of Fluid Mechanics, 55:1, pp 511-540, <https://doi.org/10.1146/annurev-fluid-031822-041721>
- Sciacchitano, A., Leclaire, B. & Schröder, A. (2021), Main results of the first Lagrangian Particle Tracking Challenge, Proceedings of 14th International Symposium on Particle Image Velocimetry – ISPIV2021, August 1–4, 2021, Doi: 10.18409/ispiv.v1i1.197
- Sellappan, P., Alvi, F.S. & Cattafesta, L.N. (2020), Lagrangian and Eulerian measurements in high-speed jets using Multi-Pulse Shake-The-Box and fine scale reconstruction (VIC#). *Exp Fluids* 61, 157. Doi: 10.1007/s00348-020-02993-9
- van Gent, P. L., Michaelis, D., van Oudheusden, B. W., Weiss, P.-É., De Kat, R., Laskari, A., Jeon, Y., David, L., Schanz, D., Huhn, F., Gesemann, S., Novara, M., McPhaden, C., Neeteson, N. J., Rival, D. E., Schneiders, J. F. G. & Schrijer, F. F. J. (2017), Comparative assessment of pressure field reconstructions from PIV measurements and Lagrangian particle tracking, *Exp Fluids* 58:33
- van Oudheusden, B. W. (2013), PIV-based pressure measurement, *Meas. Sci. Technol.* 24 032001
- Wieneke, B. (2008), Volume self-calibration for 3D particle image velocimetry, *Exp Fluids* 45:549–556
- Wieneke, B. (2013), Iterative reconstruction of volumetric particle distribution, *Meas Sci Technol* 24:024008
- Xu, H. (2008), Tracking Lagrangian trajectories in position-velocity space, *Meas Sci Technol* 19:075105
- Zhou K, Grauer SJ (2023) Flow reconstruction and particle characterization from inertial Lagrangian tracks. arXiv:2311.09076.