

3D particle position determination and correction at high particle densities

D. Schanz¹, T. Jahn¹, A. Schröder¹

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

* Correspondent author: daniel.schanz@dlr.de

Keywords: Lagrangian Particle Tracking, Shake-The-Box, Iterative Particle Reconstruction, Triangulation

ABSTRACT

The method of Iterative Particle Reconstruction (IPR), introduced by Wieneke in 2013, constitutes a major step towards Lagrangian Particle Tracking of densely seeded flows (Schanz et al. 2016). This work presents the advanced Iterative Particle Reconstruction method that was recently published in *Experiments in Fluids* (Jahn et al. 2021), detailing numerous approaches to enhance the original IPR working principle. Most of them are easy to implement, come at low computational cost and, in combination, nearly triple the processable particle image densities. The updated method is proven to be fast, accurate and robust against image noise and other imaging artifacts.

Additionally, the abilities of the method to correct for misplaced particle positions e.g. by biased triangulation or the predictive step of the Shake-The-Box scheme are quantified and a method to account for large mis-predictions (in the range of 2-10 px) is introduced. The results can be used to derive guidelines for the layout of experiments in terms of the required acquisition frequency as a function of the expected Lagrangian accelerations in order to guarantee reliable particle tracking.

1. Introduction

The introduction of the iterative particle reconstruction (IPR, Wieneke 2013) overcame the limitations of 3D-PTV and was a first major step towards performing Lagrangian Particle Tracking (LPT) with high particle image densities (Schanz et al. 2013b, Schanz et al. 2016), comparable to Tomographic PIV (Elsinga et al. 2006). IPR represents 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 at particle image density $N_I \leq 0.05$ ppp with reduced occurrence of ghost particles and higher position accuracies, compared to a MART reconstruction. When further increasing N_I the results deteriorated rapidly, when processing single images. The optimization of the positions (here also termed ‘shaking’ due to the process of moving the particles around in 3D space) is performed for each particle individually using image matching techniques; hence the demands on the calibration

are quite high. An accurate Volume-Self-Calibration (VSC, Wieneke 2007) and calibration of the optical transfer function (OTF, Schanz et al. 2013) is mandatory for IPR.

IPR is actively used in the LPT method ‘Shake-The-Box’ (STB, Schanz et al. 2016), which uses an elaborate predictor/corrector scheme to integrate the temporal domain into the reconstruction process. It has been shown that this scheme converges to a solution where the vast majority of particles is tracked, ghost particles are almost non-existent and – due to the absence of ghost particles – a high position accuracy is attained. This process has been shown to work at particle image densities up to $N_I = 0.2$ ppp for synthetic data (Raffel et al. 2018, Sciacchitano et al. 2021). Recent international competitions (Kähler et al. 2016, van Gent et al. 2017) show distinct advantages over Tomo-PIV in terms of accuracy, spatial resolution, reconstruction completeness and computer resources. Furthermore, Lagrangian properties (such as accelerations / material derivatives) needed for e.g. pressure estimation schemes are directly available (van Gent et al. 2017).

With the IPR algorithm as one key element of the most recent advances in LPT, it is evident that an improved IPR performance has a direct positive impact onto the reconstruction and tracking performances of several recent methods. This work introduces a number of advancements of the IPR idea in several key domains (see also Jahn et al. 2021).

2. Advanced Iterative Particle Reconstruction

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 Fig. 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 ('shaking' of the particles).
4. The state vector is filtered for potential ghost particles, mainly based on the particle intensity (deleting particles that fall below an intensity threshold). Points 3 and 4 are iterated several times ('inner iterations' N_{shake})
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 .

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 'outer' iterations (an outer iteration constituting of a full processing of points 1 to 5).

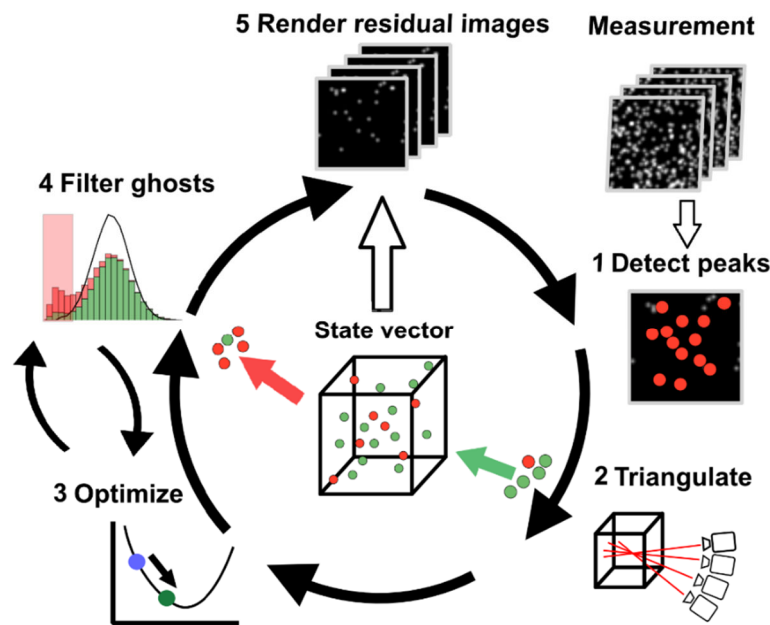


Fig. 1: Scheme of the IPR algorithm. Outer iteration indicated by large arrows. Inner iterations of position/intensity optimization (shaking) and filtering indicated by small arrows (step 3 and 4).

The key updates to the IPR process presented in this work are:

1. 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

2. 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
3. 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

Jahn et al. (2021) goes beyond these points by further introducing an analytic determination of the cost function gradient used within the position optimization (Steepest-Descend Method, SDM), as well as documenting the effects of varying the number outer and inner iterations.

3. Synthetic Test-Case

The performance of the extended IPR processing method when applied to synthetic data is analyzed below. To this end, a synthetic test case was created, which replicates the one used in (Wieneke 2013): Four virtual pin-hole cameras of 1300×1300 pixels each are oriented in a pyramidal viewing configuration and image randomly created particles in a domain of $50 \times 50 \times 16$ length units. All cameras view the full volume. The used particle shape is a 2D Gaussian ($otf(u, v) = I_{a,r} e^{-0.5(u^2 a + 2buv + cv^2)}$) with diameter of approx. 2 pixels ($I_{a,r} = 1.0, a = 2, b = 0, c = 2$). No variations of the OTF were included here. The pixel intensity values are always integers. Three different cases were considered: perfect imaging conditions (case I), moderate noise (case II) and heavier noise (case III). Particle image densities ranging from 0.02 ppp to 0.16 ppp were investigated.




Fig. 2: Excerpts of the synthetic camera images at different particle image densities (ppp-values) without noise and uniform particle intensities. Effects of particle overlap become increasingly visible.

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, see Fig. 2). 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.

For **case II** (moderate noise), the particles have a mean intensity of 1000 counts in the volume, with a Gaussian distribution with $\sigma^2 = 250$ counts, leading to particle peaks on the images 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. Instead of just rounding to the closest integer, the final simulated intensity of each pixel gets randomly chosen from a Poisson-distribution of the respective pixel value as mean parameter. Thereby we simulate photon shot noise and add noise proportional to the intensity of the individual pixel.

Case III (heavier noise) applies a Gaussian distribution with $\sigma^2 = 500$ counts to the particle intensities, leading to image peaks in a broad range of ~ 40 -800 counts. A constant Gaussian noise floor with a mean of 80 counts and rms width of 50 counts is added to the image. Again, photon shot noise is additionally applied. See Fig. 3 for examples of case II and III.




Fig. 3: Excerpts of the synthetic camera images at different noise levels (case II: upper row, case III: lower row) and particle image densities

Unless otherwise noted, the following parameters were used for the IPR reconstructions of all cases: 20 outer IPR iterations N_{IPR} were conducted. For the sake of clarity and visualization of the results, this value remained constant, even if convergence is usually reached earlier. The allowed triangulation error ε starts at $\varepsilon(1) = 0.4$ px and is linearly increased to $\varepsilon(N_{IPR}) = 0.8$ px (see section 5.2). Due to the low initial values of ε the triangulation can immediately start with using a reduced set of cameras (only $N_c - 1 = 3$ cameras are required to have a peak in order to register the particle). Each outer IPR iteration (as depicted in Fig. 1) used $N_{shake} = 10$ inner iterations of

shaking. From those, particle filtering (intensity, duplicates, FOV) is performed for the last seven iterations. The number of brightest cameras left out in the intensity update (see section 5.3) is set to $N_{Ci} = 2$. The camera order is permuted 4 times (1-2-3-4, 4-1-2-3, 3-4-1-2, 2-3-4-1) and a triangulation is carried out for each permutation (see section 5.1). The maximum step width in shaking is initially set to $d_{max} = 0.4 px$. For iterations 15-17 d_{max} is reduced to $0.1 px$; for iterations 18-20 a further reduction to $d_{max} = 0.02 px$ is applied.

4. IPR Reconstructions of the synthetic test case

The created simulated images are reconstructed and the results compared to the original particle distribution. This section shows results for the full IPR processing, incorporating all extension, while the effects of the different new approaches are documented in section 5.

Fig. 4 (left column) shows the convergence behavior of IPR for varying particle image densities, when applied to the described noise-free synthetic dataset (case I). 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 reprojected points to the original peak position).

For lower particle image densities ($N_l \leq 0.08 ppp$) the method converges to a stable value of $R_f \approx 1$ in less than 5 iterations. At $N_l = 0.14 ppp$ the reconstruction converges after 15 iterations. Even for this high seeding density 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, only for $N_l = 0.16 ppp$ ghosts are still present after 20 iterations (as this case is not yet converged). For all other cases the absolute number of detected ghosts is 0.

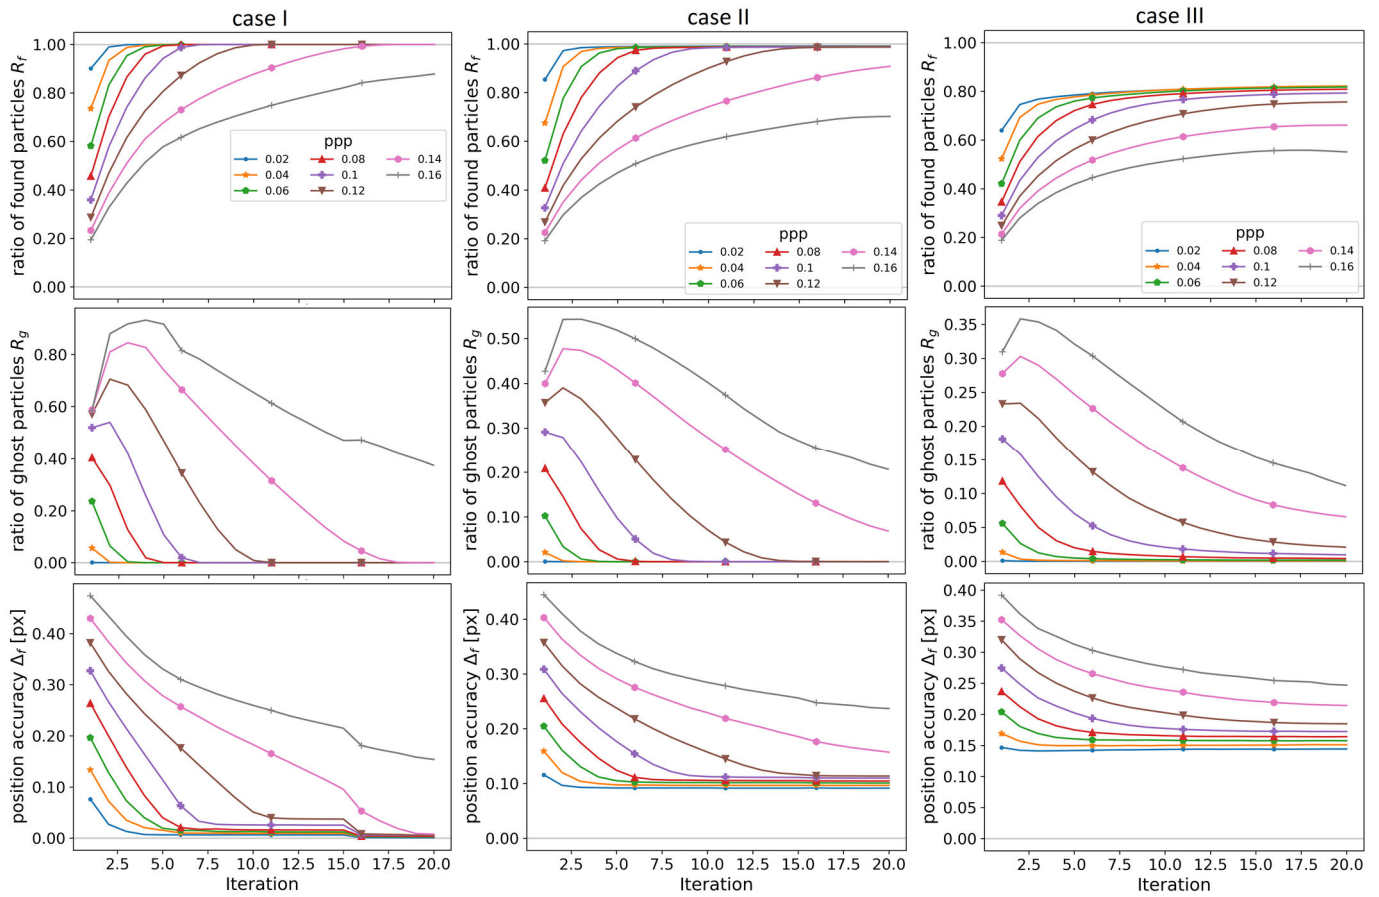


Fig. 4: Convergence of single-image particle IPR reconstructions for different particle image densities and noise cases. Shown are the number of correctly identified particles R_f (top), the number of ghost particles R_g (center) (both normalized with the number of true particles), and the position accuracy of the true particles Δ_f (bottom) as a function of the iteration number. At iterations 15 and 18 the maximum step width is reduced.

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. The convergence of the advanced IPR processing when applied to images with medium and high noise are given in Fig. 4, middle and right columns. For the medium noise level (case II), it can be seen that the general convergence is slowed down compared to the noise-free case I. However, all reconstructions with a particle image density below $N_l = 0.14$ ppp converge to $R_f \approx 0.99$ within at most 13 iterations. For $N_l = 0.14$ ppp a similar convergence would likely be achieved when applying more than 20 iterations. For all but the highest particle image densities the number of ghost particles is virtually zero, once convergence is reached. The position accuracy shows a

constant level, determined by the noise level, slightly increases with growing N_I due to increasing effects of particle overlap. For the converged cases values between $\Delta_f = 0.09 px$ and $\Delta_f = 0.12 px$ are attained.

For the higher noise levels and intensity variation of case III, the ratio of found particles drops to values of $R_f \approx 0.8$, which are achieved up to $N_I = 0.1 ppp$. For higher particle image densities lower values are seen; still, approx. 66 % of the real particles are correctly identified at $N_I = 0.14 ppp$ for this relatively high noise level. In general, the convergence is much slower, as the peak detection accuracy suffers from the image noise, requiring higher values of ε for a successful triangulation, which are only attained within later IPR iterations. Despite a high number of noise peaks that reach above the peak detection threshold, the occurrence of ghost particles is limited ($R_f < 0.03$ for $N_I < 0.14 ppp$). The position accuracy is more affected by the stronger noise, with most seeding densities ranging between $\Delta_f = 0.14 px$ and $\Delta_f = 0.18 px$. For both noisy cases, the close proximity of Δ_f for a wide range of particle image densities points to the image noise as the dominant source for the position uncertainty.

5. Details of advancements

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 $N_I = 0.05 ppp$ 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 $N_I = 0.14 ppp$ 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 $\Delta_f \leq 0.01 px$ for all converged cases. These quite profound advances in reconstruction quality result from the combination of the several enhancements to the IPR algorithm mentioned above, which are detailed and quantified in the sections below.

5.1 Camera permutations

The first two cameras have a special role in the triangulation process. The space 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 (Michaelis et al. 2010, 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 row 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 further 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. In case of experimental setups using many cameras (e.g. 6 or more), the number of triangulations could be limited by manually selecting a subset of permutations, based on the space angle between the first and the second camera in each respective permutation.

Fig. 5 compares results from case II, which were reconstructed once using four successive triangulations with camera orders (1-2-3-4, 4-1-2-3, 3-4-1-2, 2-3-4-1) and once using a single triangulation with order (1-2-3-4).




Fig. 5: Convergence of IPR reconstructions with triangulations using a single (fixed) camera order compared to using several camera order permutations at $N_l = 0.06 \text{ ppp}$ (top) and $N_l = 0.12 \text{ ppp}$ (bottom) of case II (moderate noise)

The permutation-approach clearly enhanced convergence. For both particle image densities, this case yields higher numbers of true particles, starting from the first iteration. From there, the reconstructions converge within 5 ($N_l = 0.06 \text{ ppp}$) and 14 ($N_l = 0.12 \text{ ppp}$) iterations, while the

standard approach fails to fully converge within twenty iterations in both cases. At the beginning, the permuted cases also show higher ghost levels, which can be understood, given that more triangulations are involved. However, the ghost particles are quickly suppressed as more and more true particles are found. For the standard case at $N_l = 0.12 \text{ ppp}$ the ghost level rises within the first five iterations and is only slowly reduced in the following.

5.2 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.

Fig. 6 compares the results of reconstructions using fixed values of $\varepsilon = 0.6 \text{ px}$ and $\varepsilon = 0.8 \text{ px}$ to the linear ramp approach for two particle image densities from case III. At low particle image density ($N_l = 0.04 \text{ ppp}$), the results do not differ much, with the reconstruction using $\varepsilon = 0.6 \text{ px}$ converging quickest, followed by the ramp approach. Using $\varepsilon = 0.8 \text{ px}$ results in a slightly slower convergence, likely due to an increased appearance of ghost particles, caused by noise peaks and the relatively high allowed triangulation error. Eventually all approaches converge to a comparable level of approx. 82 percent of found particles.

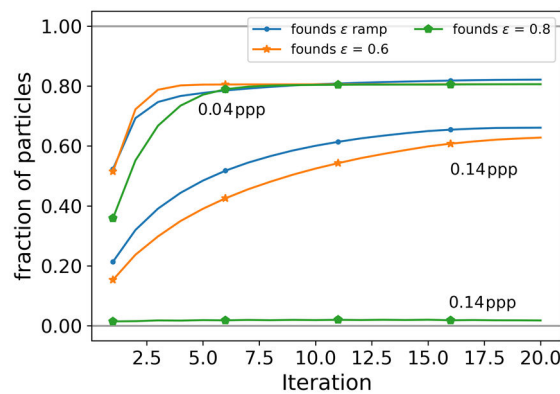


Fig. 6 Convergence of IPR reconstructions of case III (heavier noise) using either a fixed triangulation error of $\varepsilon = 0.6 \text{ px}$ and $\varepsilon = 0.8 \text{ px}$ or a linear ramp between $\varepsilon = 0.4 \text{ px}$ and $\varepsilon = 0.8 \text{ px}$, increasing with the IPR iterations.

At high particle image density ($N_l = 0.14 \text{ ppp}$), the results are different. Here, the ramp approach is consistently yielding the best results. Using a fixed $\varepsilon = 0.6 \text{ px}$ shows a slower convergence and

a lower end value, while the fixed $\varepsilon = 0.8 \text{ px}$ case completely fails to yield a significant number of true particles.

5.3 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 both Wieneke's (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. 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.

Fig. 7 shows selected results for case I and III. For noise-free images at moderate seeding density ($N_l = 0.08 \text{ ppp}$), all versions start at the same fraction of found true particles, however the more cameras are ignored, the better ghost particles can be suppressed. In case all cameras are used for the update ($N_{Ci} = 0$), the ghost fraction is too high to be quickly corrected for and convergence is delayed. At lower particle image densities all versions perform approx. the same. In these noise-free conditions, the intensity information from peaks on a single camera is sufficient for an accurate determination, as can be seen when looking at the result with higher particle image density ($N_l = 0.14 \text{ ppp}$). Again, the fraction of found true particles starts at the same value for all values of N_{Ci} and the ghost levels vary. The case with $N_{Ci} = 3$ converges within 12 iterations. Including two cameras in the intensity update ($N_{Ci} = 2$) leads to a slower convergence, which is reached after 15 iterations. For the other cases convergence is not attained; at $N_{Ci} = 0$ only approx. 60 % of the true particles can be identified. The reason for these differences is the strong difference in ghost suppression capability. A ghost particle is typically supported by strong peaks on one or two cameras only; the other cameras show low-intensity residuals or only partial overlap with a particle image. Therefore, the chances of ghost particles retaining a high intensity are greatly reduced when using $N_{Ci} = 2$ or $N_{Ci} = 3$.

The situation changes to a certain degree when noise and other imaging defects are present. The result of case III at $N_I = 0.12 \text{ ppp}$ shows that here the reconstruction using $N_{Ci} = 3$ performs worse than $N_{Ci} = 2$ and is even becoming slightly unstable (R_g going down in the last iterations), such that the reconstruction using $N_{Ci} = 1$ is reaching the same level, despite converging slower initially. Again, including all cameras in the intensity update performs worst. The reason for the worse performance of $N_{Ci} = 3$ in this case is likely a stronger deletion of true particles by the intensity filtering due to the presence of weak particles and the influence of the noise on the intensity registered in each camera.

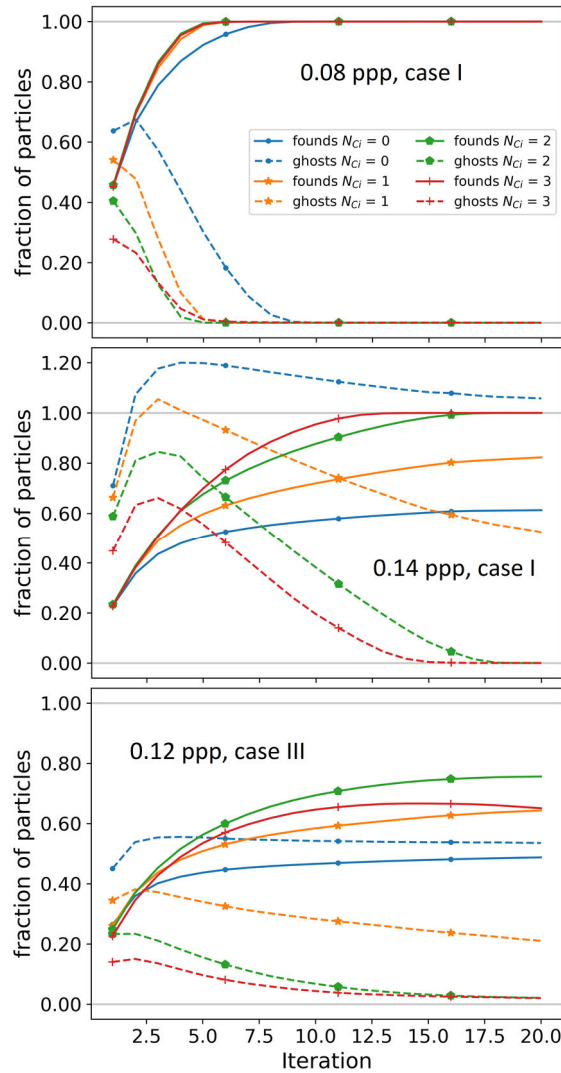


Fig. 7: Convergence of IPR reconstructions using several value for the number of locally brightest cameras being ignored during the intensity update (N_{Ci}). Shown are $N_I = 0.06 \text{ ppp}$ and $N_I = 0.14 \text{ ppp}$ for case I (no noise, top and middle) and $N_I = 0.12 \text{ ppp}$ for case III (heavier noise, bottom).

Due to the high relevance of noise in real experimental data $N_{Ci} = 2$ was chosen as the standard value for a four-camera setup, as regarded here. When using more cameras, a higher value of N_{Ci} might also be chosen.

6. Working range of the position optimization

In this section, we analyze the working range of the Steepest-Descent method (SDM) for particle position optimization (“shaking”), described in Jahn et al. (2021), for various image noise levels and prediction error (ε_p) distributions. Synthetic Lagrangian particle trajectories, extracted from a Direct Numerical Simulation (DNS) of a highly turbulent flow around and in the wake of a smooth cylinder at $Re = 3,900$ (Sub-domain 1 from Khojasteh et al. 2022; open data under the CC BY license), were used to characterize occurring ε_p in dependence on the temporal sampling. The downloadable set contains the positional information of approx. 197,000 particles at 350 time-steps, each spaced 10 DNS time-steps (t_{DNS}) apart. For this evaluation we temporally subsample the data of the full field to tracks with spacings of $dt = 80, 160, 240$ and $320 t_{DNS}$. Three different imaging conditions (noise-free, moderate and high noise), corresponding to case I, II and III from above, were considered. The track positions were scaled to pixel values and noise was added to each spatial component according to probability density functions (PDFs) of position accuracy (P_{pos}), gained by IPR reconstructions at $N_I = 0.1 ppp$ for each noise-case. Successive positions at times t_0 to $t_{9,dt}$ were used as input to a predictive Wiener filter (Wiener 1949), as described in Schanz et al. (2016) to predict the position at time-step $t_{10,dt}$.

The occurring prediction errors gathered in PDFs (P_{pred}) for each noise case, temporal sampling and component. At $dt = 320 t_{DNS}$ the particles exhibit an average and maximum shift of $20.8 \bar{p}\bar{x}/dt$ and $39.5 \bar{p}\bar{x}/dt$, respectively. Fig. 8 a) shows the x- and z-components of P_{pred} for this case at moderate image noise (grey curves). The conformity of the distributions documents that in this highly turbulent case the prediction errors are dominated by the occurring high accelerations (and jerks). The higher positional noise in the z-component has a negligible effect.

Even at this large dt and particle image shifts, the vast majority of the particles can be predicted with good accuracy (89 % of the particles show $\varepsilon_p < 1 \bar{p}\bar{x}$). Events with high acceleration do occur, but they are rare ($\varepsilon_p > 10 \bar{p}\bar{x}$ for 0.6 % of the particles). This shows that the vast majority of particles in turbulent flows can be tracked reliably at relatively low sampling frequencies, however the gained statistical properties will be biased towards lower acceleration values, as the rare events are lost. Also, the amplitudes of extreme vorticity events and corresponding strong pressure minima, which are mainly causing these rare high acceleration events, will be modulated when applying subsequent data assimilation approaches.

In order to characterize the ability of SDM to correct such prediction errors, random particle fields were imaged on four cameras under the three considered noise conditions, recreating the test setups from section 3. The particle image shape is Gaussian with 2 pixels diameter. The particles were displaced according to the different P_{pred} and 25 iterations of the SDM position optimization

scheme were applied. A particle is regarded as successfully corrected if its location after optimization is within $1 \bar{p}\bar{x}$ of the true position. The solid curves in Fig. 8 a) show the fraction of corrected particles as a function of their initial displacement in x- and z-direction for the case with moderate noise and $dt = 320 t_{DNS}$ at $0.1 ppp$. It can be seen that for the depth direction (z), initial ε_p of up to ± 3.1 pixels can be reliably corrected. For the in-plane direction x (y is similar), ± 2.3 pixels can be corrected. This discrepancy is explained by the geometry of the camera setup. In order to determine a meaningful gradient of the cost function, the reprojected particle image has to overlap at least partly with the original on at least one camera. A displacement along the depth direction results in a larger remaining overlap compared to an in-plane displacement. The smaller the total aperture of the camera system, the larger this effect will be. At the same time, the depth position accuracy is inversely dependent on the system aperture. A reduced positional accuracy in the depth direction is therefore counteracted by an enhanced correction ability. The extent of the correctable range scales with the particle image diameter. The resulting mean positional root-mean-square error of the corrected particles for the three noise cases is $0.0005 \bar{p}\bar{x}$, $0.095 \bar{p}\bar{x}$ and $0.172 \bar{p}\bar{x}$.

It is obvious that for this large particle image shift case not all predicted particles could be matched to their correct position using SDM. The total fraction of corrected particles for this case is $f_{ct} = 0.962$. Fig. 8 b) (black curves) documents f_{ct} for the different dt and noise cases at $0.1 ppp$. At low dt virtually all particle predictions are correctable, only the heavy noise case shows around 0.1 % of lost particles. For all noise levels, the number of lost particles increases steadily with dt , reaching around 4.7 % under heavy noise and $dt = 320 t_{DNS}$.

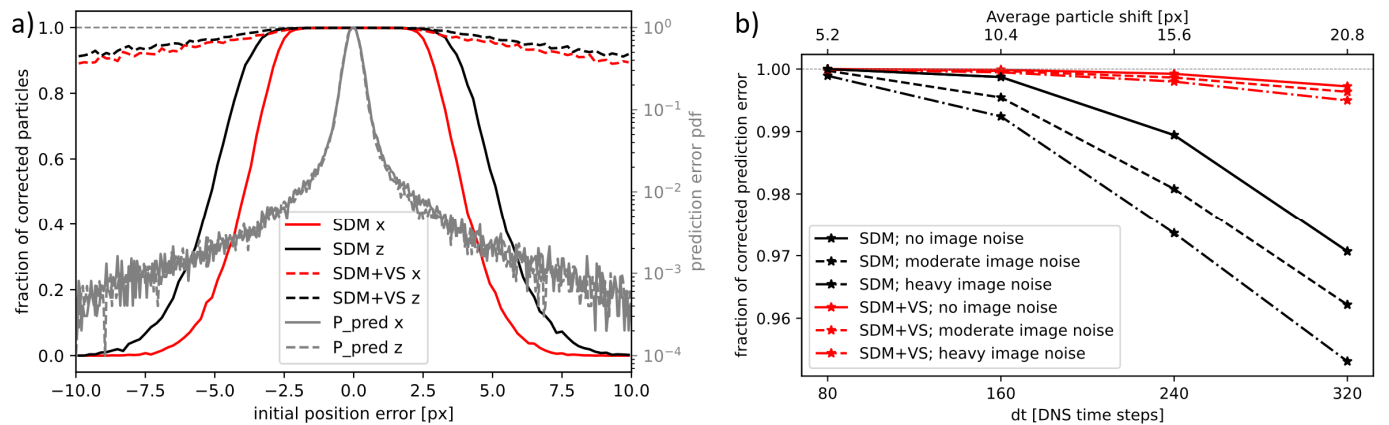


Fig. 8: a) fraction of particles, whose position was correctly optimized after the prediction process as a function of their initial misplacement in x- and z-component for SDM optimization and additional Variable Space step (VS) application; the misprediction PDF P_{pred} is given for both components (moderate image noise, $0.1 ppp$, $dt = 320 t_{DNS}$);

b) fraction of corrected particle for SDM and SDM + VS as a function of the temporal sampling.

7. Correcting large prediction errors

An experimental investigation under no technical constraints typically aims to not exceed a maximum particle shift of around $10 \bar{p}\bar{x}$ - falling in-between the $dt = 80 t_{DNS}$ and $dt = 160 t_{DNS}$ - cases. According to our results, the SDM method provides a sufficient range of correctable prediction errors for such conditions. However, depending on the flow and the available cameras and illumination, the desired temporal sampling cannot always be reached and an extended range of correctable errors would be needed.

Here we propose an extension to SDM to enable the correction of large - but rare - prediction errors using an iterative grid search for a small subset of particles, that we call the 'Variable Space step' (VS) approach. The vast majority of particles are correctly positioned after the predicted particle field has been treated by position optimization. The (few) particles that exhibited a non-correctable prediction error, will be placed somewhere, typically close to the prediction point within a local minimum of the cost function, formed by image noise or residuals. To identify particles whose position is not well supported by the images after the SDM optimization, the gradient of the cost function of SDM (basically the sum of the local residual images) is evaluated at the current position for all particles. Correctly positioned particles will show a high gradient of the cost in all directions; for incorrectly positioned particles the gradient is much flatter (at least in most directions), allowing a reliable discrimination of particles by setting a threshold.

For all misplaced particles identified this way, the surrounding space is scanned within a grid of $(2n + 1)^3$ points and a spacing of d_g , centered in the current position. Each of these 3D coordinates is projected onto all cameras and the intensity of the 3x3 pixels surrounding the projection points is summed over all cameras. The particle is then placed on the location with the highest intensity sum. Using a multiplicative instead of an additive approach proved to be less stable. As small prediction errors are more likely than large ones, the method is applied iteratively. Here, the grid size is increased linearly in 8 steps from $n = 2$ to $n = 16$, always retaining a grid spacing of $d_g = 2 \bar{p}\bar{x}$. Therefore, a maximum error of $32 \bar{p}\bar{x}$ in every direction can be corrected in the last iteration. d_g is chosen according to the particle image size - ensuring an overlap of the recorded image with the 3x3 kernel, while maximizing the step size for computational efficacy. After each VS iteration, all reconstructed particles are treated by four iterations of SDM optimization and the detection of misplaced particles is performed anew. This way, only few particles are left during the later VS-iterations, where the search space becomes so large that ambiguity might become a problem. The computational demand is reasonable due to the low number of interrogated particles. VS was inspired by the 'Initial Shake' method, introduced in Schanz et al. (2016), but proves to be far more capable due to the iterative approach and the restriction to salient particles.

The results of a combined application of SDM and VS are given in Fig. 8 a) (dashed curves) and Fig. 8 b) (red curves). It can be seen that the region of successful correction has been increased significantly for all spatial directions. Particles with a (very rare) $10 \bar{p}\bar{x}$ displacement now exhibit a 90% chance of being corrected. The total fraction of corrected particles f_{ct} increases distinctly. For the $dt = 320 t_{DNS}$ - case at high noise level, only around 0.5 % of the particles cannot be corrected, thereby greatly improving the tracking of highly accelerated particles. The continuation of tracks even under such unfavorable circumstances further stabilizes the prediction and correction scheme of the STB technique and only a very limited number of tracks need to be terminated. Consequently, this procedure enhances the DAR of the STB measurement scheme significantly. Although the track can be continued with high curvature effects, the used temporal fitting scheme might introduce truncation errors related to such extreme events, which can only be avoided by increasing the temporal sampling rate.

Acknowledgements

This work was 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).

References

- Cierpka C, Lütke B, Kähler CJ (2013) Higher order multi-frame particle tracking velocimetry. *Experiments in Fluids*, 54(5), 1533.
- Elsinga GE, Scarano F, Wieneke B, van Oudheusden BW (2006) Tomographic particle image velocimetry. *Exp Fluids* 41:933–947
- Gesemann S, Huhn F, Schanz D, Schroder 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
- Jahn, T., Schanz, D. & Schröder, A. Advanced iterative particle reconstruction for Lagrangian particle tracking. *Exp Fluids* **62**, 179 (2021). <https://doi.org/10.1007/s00348-021-03276-7>
- 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. *Experiments in Fluids*, 57(6), 97
- Michaelis D, Novara M, Scarano F, Wieneke B. (2010). Comparison of volume reconstruction techniques at different particle densities. 15th int. symp. on appl. of laser tech. to fluid mech.
- Raffel M, Willert C E, Scarano F, Kähler C J, Wereley S T, Kompenhans J (2018) Particle image velocimetry: a practical guide (Digital Content). Springer

- Sciacchitano, A., Leclaire, B., & Schroeder, A. (2021). Main results of the first lagrangian particle tracking challenge. In *14th International Symposium on Particle Image Velocimetry*, Doi: 10.18409/ispiv.v1i1.197.
- 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 and 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. <https://doi.org/10.1007/s00348-016-2157-1>
- van Gent PL, Michaelis D, Van Oudheusden BW, Weiss P-É, De Kat R, Laskari A, Jeon Yj, David L, Schanz D, Huhn F, Gesemann S, Novara M, McPhaden C, Neeteson NJ, Rival DE, Schneiders JFG, Schrijer FFJ (2017) Comparative assessment of pressure field reconstructions from PIV measurements and Lagrangian particle tracking, *Exp Fluids* 58:33
- Wieneke B (2007) 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