

Particle position optimization of incompletely reconstructed particle clouds

Philipp Godbersen^{1,*}, Daniel Schanz¹, Andreas Schröder^{1,2}

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

2: Brandenburgisch Technische Universität (BTU) Cottbus-Senftenberg, Germany

*Corresponding author: Philipp.Godbersen@dlr.de

Keywords: Lagrangian Particle Tracking, particle reconstruction, position optimization, shaking.

ABSTRACT

A key aspect in Lagrangian particle tracking with Iterative Particle Reconstruction is the position optimization of triangulated particles. Such an optimization is conducted by slightly displacing the individual particles to better match their back projected images with those captured on the cameras. However, the quality of this optimization can be degraded in the early stages of particle reconstruction when only a small number of particles has been found yet. We investigate this effect, motivated by worse than expected performance of DLR multi-pulse STB code in our initial submission to the 2nd LPT Challenge, and provide some initial fixes to this problem together with the introduction of a novel placeholder approach, capable of preventing such issues in the future. The proposed approach is useful in high seeding density situations, as well as measurement setups that image a not insignificant amount of particles outside of the common camera volume.

1. Introduction

For 3D Lagrangian Particle Tracking (LPT) approaches the accurate determination of particle positions from the images of several cameras is an important aspect. We will consider this in the context of the Shake-the-Box (STB, Schanz et al. (2016)) approach, where an image-matching scheme is used to optimize the 3D positions of triangulated or tracked particles. The image matching scheme relies on minimizing the local residual (summed over all cameras) and was first introduced as part of the Iterative Particle Reconstruction method (IPR, Wieneke (2012)). In the remainder we will refer to this scheme as "shaking" of the particles. Our investigation focuses on the performance of the shaking stage within STB in situations where the particle clouds have only been partially recovered, which can negatively impact the performance of the shaking optimization. Such an effect has been described by Cornic et al. (2021) but has not been seen as an issue in our evaluations before. This submission is motivated by our investigation into the worse than expected particle position accuracy in some cases of our initial submission to the two-pulse LPT_Case02 Dataset




Figure 1. Position error anomaly observed in our initial submission to the 2nd LPT Challenge Case 02 TP for the high ppp cases with our submission highlighted: (left) initial result showing greatly increased error for the high ppp cases with 0.2 ppp case at over 0.3 voxels rms position error. (right) drastically reduced particle positioning errors achieved as a result of this investigation where the 0.2ppp case now shows below 0.1 voxels rms position error.

of the 2nd LPT&DA Challenge (Leclaire et al., 2025) as presented in June 2025 on a workshop in Tokyo, connected to the International Symposium on Particle Image Velocimetry. Understanding of this effect has allowed us to significantly reduce particle position errors in those affected cases as show in figure 1.

For clarity of presentation we will discuss this purely on the single time step Iterative Particle Reconstruction (IPR, Wieneke (2012)) sub-component of STB instead of the full predictor-corrector STB approach without loss of generality. When IPR is utilized on images from a calibrated multi-camera setup the first step consists of 2d particle position detection on the individual images, which is termed peak detection. The detected peaks on the cameras are then used in a triangulation approach in order to obtain corresponding particle positions in 3d space. Especially for high particle image densities this is a challenging reconstruction problem where only some of the correct particles are recovered and false detections can occur. These positions and intensities of these particles are then optimized by the shaking step: Due to an empirically calibrated optical transfer function (OTF) together with the camera model it is possible to compute a projected particle image on each of the cameras for a given 3d particle. Subtracting all of these back-projections from the original camera image forms the so called residual images. The goal of the optimization is now to slightly modify the 3d particle positions in a way that the sum of the residual images is minimized, which results in a more accurate particle position. This approach can be implemented via gradient descent and is a key factor in the performance of IPR based techniques.

However, a high position accuracy for the particles via this method can only be achieved if all particles in the cloud have been successfully discovered, as illustrated in figure 2 for a very simplified one-camera 1d case, consisting of two particles. In this scenario the shaking becomes a simple multi Gaussian fitting problem but the core issue remains the same. Fitting with two Gaussians leads to a correct outcome but if only one is used for the fit the result is degraded. The attempt to cover both peaks with a single Gaussian results in a positioning in-between the peaks. Due to this




Figure 2. Fitting [red] of multi-Gaussian data [blue] using: (left) correct number of Gaussians. (right) just one Gaussian used which results in the fit belonging to the right peak being dragged away into the center of the two.

effect the optimizer can actually decrease the position accuracy of a particle. During standard IPR processing with conventional peak detection utilized so far this has not necessarily been a problem. As conventional peak detection is not particularly accurate in overlap situations anyway, the IPR processing pipeline is configured in a robust way that allows to iteratively recover most of the particles even if they are not well positioned yet at which point the optimization works well again and accurately positioned particles are obtained.

This changes with the switch to PeakCNN (Godbersen et al., 2024) for peak detection, utilizing a machine learning approach. Peaks are found on the images much more accurately even in high particle density images allowing for tighter tolerances in the triangulation. This is necessary to push the capabilities of the IPR further towards very high particle image densities. As we discovered in our initial submission of the challenge test case greater attention must be directed to the shaking procedure in these high density configurations.

1.1. Shaking procedure

The way of optimization of particle positions has changed somewhat over the life of the STB approach; initially starting with simple variation of a fixed displacement and soon turning towards a more systematic gradient descent based optimization (Jahn et al., 2021). A detailed explanation and analytic computation of the needed gradients is given in that paper but we'll provide an overview here. The purpose of the shaking is to refine the 3d positions of the triangulated particles by back-projection onto the camera images. 3D coordinates of a given particle are projected into a cameras coordinate system by use of the calibrated camera model and an particle image can then be rendered by use of the empirically determined optical transfer function (Schanz et al., 2012) (OTF) which models how the diffraction limited point source image appears on the camera image. In an ideal situation these projected particle images match up perfectly with the ones visible in the measurement image so that subtraction leads to an empty image. Any remaining intensity in this residual image is indicative of a mismatch between the actual tracer particles and the reconstructed particle positions (ignoring the aspects of image noise or OTF and camera calibration errors for now). The mean absolute residual value across the images of all cameras can thus serve

as loss value for a particle position optimization. The loss is expressed in terms of 3d particle position and its gradient is analytically determinable by chaining the subtraction of the Gaussian OTF in the camera coordinate system with the camera models Jacobian. Particle positions are then optimized with regards to this loss by gradient descent.

This optimization problem could be performed globally as one large optimization across all reconstructed particles simultaneously using widely available optimization libraries but for the specific requirements of the IPR a particle individual sequential local optimization has been found to be advantageous (Jahn et al., 2021). Each particle is thus treated sequentially and a gradient descent with a simple line search is used for optimization of its position. This line search is performed by selecting a maximum displacement (maximum shake-step, in the following), querying at that position and several lower values along the line. The particle displacement corresponding to the lowest loss is then selected and the procedure repeated for the next particle. Multiple iterations of this shaking procedure are then performed, resulting in particles with improved position accuracy. The approach using the fixed line search as described above can be implemented very performantly.

2. Investigation

In order to investigate the behaviour of our shaking procedure for yet incompletely reconstructed particle clouds we created two synthetic experiments. In the first one we generate a random particle cloud and use it to simulate images on a known camera setup. We then mimic a partial IPR evaluation by selecting a subset of the ground truth particles, adding some noise to their position and designating them as recovered particles. We then run our shaking procedure on these particles and then record the their displacement with respect to the ground truth. A perfectly operating shaking procedure should fully remove the added position noise and drag all recovered particles back to their ground truth positions.

As long as all particles were designated as recovered even large initial errors could be fully corrected in this setup. However, for simulated partial recovery the result degraded, with the position error being in some cases significantly worse than prior to optimization. Figure 3 shows the results of this analysis in a 2d heatmap with the two axes as the ratio of particles declared as discovered and the added position error. The colormap and numerical labels on the heatmap correspond to the difference in position rms error after and before the shaking. As seen in the figure the position optimization is actively worsening the result in a large region. These are the cases where the initial errors are rather small but not all particles recovered which results in the optimizer dragging them off their initial positions even if they exhibited low error initially. It is exactly this operating region the high ppp cases from figure 1 fall into as for such high particle densities the first triangulation does not find all particles anymore. In synthetic tests similar to the challenge data we observed the




Figure 3. Examination of positioning error difference after and before shaking procedure showing negative values as improvement. Varied along the axes are the added particle position error and the ratio of kept particles from the ground truth: (left) using our original parameterization actively worsening error in some cases. (right) Maximum optimizer step in gradient descent severely decreased to 0.01 px

first triangulation to yield less than 60% of the total particles at a positioning error before shaking of about 0.1 px thanks to the PeakCNN input. The shaking step therefore worsened the particle position error by more than a factor of two, massively damaging the result. In the same figure we also show the same analysis with a modified shaking procedure where we limit the maximum step size in the gradient descent to a much lower value. This limits the maximum extent on the line search to 0.01 px and prevents the optimizer from pulling particles too far into the wrong directions while still offering some minor improvement. Limiting the optimiser this way then yielded a 10% improvement of the error as opposed to simply not optimizing at all.

2.1. Single particle pair tests

This effect can be more precisely investigated in a second experiment where only two particles of equal brightness are considered in isolation. Only one of them is provided as recovered in a similar investigation as above but now the influence of the distance of the two particles is investigated in more detail. Figure 4 shows the results of such an investigation as the displacement of the single particle from its placement at its ground truth position a function of separation between the two particles. For our standard shaking procedure the single particle is moved to a position exactly in the middle between the two particles for increasing separations until its displacement is clipped by the maximum possible stepping distance visible as a plateau in the figure which is bounded by the product between the maximum step size and the number of shake iterations. This directly due to the effect shown in figure 2. For large distances between the particles, the error after optimization is decreased again as the particles overlap less and less at this point culminating in zero error once no overlap is achieved. The effect of the reduced max step size in the Small variant is immediately apparent as a much sooner acting clip on the displacement resulting in a much larger plateau




Figure 4. Two particle investigation: (left) Image of the two particles for a very large separation without overlap for illustration purposes. Ground truth position labeled in black circles, position of the used particle before and after optimization in yellow and red markers. (right) Position error after optimization depending on separation of the two particle images. Standard and Small refer to a max step size of 0.1 px and 0.01 px respectively, both applying 10 steps. The dashed black line represents a distance exactly in the middle of the two particles

thereby preventing the particle to be dragged to far away from its position into the middle between the two particle images.

2.2. Placeholder approach

As discussed above the shaking procedure can be negatively affected if not all particles visible on the measurement images have been found yet. This is the case for the reconstruction of dense particle clouds where initial IPR iterations are not able to recover all particles yet or even for sparsely seeded images in situations where the camera setup is in such that a non-negligible fraction of imaged particles is outside of the common camera volume. Such particles may only be visible in one of the cameras and could thus never be reconstructed in 3d, but still contaminate the residual image during shaking. Our proposed solution to this problem is based on advances in particle image localization (peak detection). The PeakCNN approach (Godbersen et al., 2024) allows to accurately detect particle image locations even in very densely seeded measurements with high position accuracy by use of a convolutional neural network. The traditional peak detection approach struggled with densely overlapping particles in which cases the position accuracy suffered as well. This was sufficient for the IPR using the triangulation and subsequent shaking and filtering together with multiple IPR iterations. But the much more trustworthy peak positions obtained from the PeakCNN allow to use them in a new approach:

Particle reconstruction is conducted as normal at first with a peak detection on the images and a subsequent triangulation. Instead of directly proceeding with the particle position optimization




Figure 5. Fitting [red] of multi-Gaussian data [blue] using a passive placeholder on the possibly slightly off peak detection of one peak [black]: (left) situation before shaking (right) The placeholder allows the second peak to be fitted correctly instead of being drawn to the middle like shown in figure 2

via shaking we analyze the triangulated particles with respect to the detected peaks first. Generally only a limited number of the detected peaks has been successfully triangulated to form a 3d particle at this stage. Immediately shaking these would be disadvantageous as described in the first section of the paper. However, if we trust the results of the peak detection reasonably well, then nearly all of the detected ones correspond to actual 3d particles that simply have not been triangulated yet and we can take this into account during the shaking. We select only those detected peaks which do not correspond to a already triangulated particle and simply temporarily subtract them from the residual image during shaking. This effectively creates a placeholder on those particle images which have not been triangulated yet and prevent them from negatively affecting the shaking process. These placeholders are only temporally used in the shaking stage and any further particle reconstruction steps proceed using the standard residual images. The next IPR iteration then discovers additional 3d particles so less placeholders are needed, a process that continues over the iterations and ideally ends in a fully recovered particle cloud where no placeholders are needed anymore. Or, in the case of a limited common camera volume, the only placeholders remain are those that belong to particles outside the common volume. A sketch of the application of a placeholder during position optimization is shown in figure 5.

2.3. Placeholder tests

We demonstrate this shaking with placeholders approach using a similar test as before for figure 4. To better show another effect of the Small shake case we now introduce a variable amount of initial displacement error on the particle before shaking. Whereas in the previous test it was sufficient for the shaking to just keep the particle in its initial location as it was placed exactly, now the shaking procedure must actually correct for an error. A downside of reducing the maximum step-size of the optimizer to limit the initial errors in particle displacement is its reduced capacity to correct large particle position errors as not enough movement might be allowed to correct this. Figure 6 shows the same evaluation as in figure 4 but now with this additional position noise added to the particle before optimization. As expected the small variant with a step size of 0.01 px and 10 steps is only able to correct at most 0.1px of displacement error. The placeholder approach shows good results, being able to fully correct the error for most of the particle distances. Only for the




Figure 6. The same two-particle investigation as before but now with a variable initial position error added to the particle which the shake needs to correct. The black dotted line corresponds to a position halfway between the two particles which the standard shake at a maximum shake step of 0.1px locks to. Using the small shake step of 0.01px prevents this but since the particle movement is limited the initial error can not fully be corrected for. For the placeholder approach the lines for the different errors collapse into one fully correcting the error for most distances




Figure 7. The same analysis as conducted in fig. 6 but now using ten times as much shake iterations. The standard and small case now become very similar showing how the advantage of the small case was solely due to the limitation of total maximum displacement. The 100 iterations now allow the particle in the small case to be displaced as far as the standard case at 10 iterations and this is reflected in the outcome in the figure.

very close placements does it behave like the standard case, displacing the particle to the midpoint between the two of them. This is simply due to the fact that no placeholder was used for these distances since the PeakCNN did not detect two separate peaks there. As soon as the separation has increased enough that two separate particles are detected by the network does the placeholder kick in allowing the shake to the correct position.

The plateau for the Small case in figure 4 already indicated that the small maximum shake step was primarily limiting the maximum possible displacement of the particle, thereby limiting the occurring errors. A similar effect can be seen in figure 6, but we can additionally see that this then limits the ability to correct larger errors. The particle is simply not allowed to move far enough so even when the distance between the particles is high enough that they no longer overlap the error can not be corrected anymore. This can be seen even more clearly in figure 7 where the same study has been conducted but now for a total number of 100 shake iterations. Such a large number of iterations allows the Small case to displace the particle as far as the Standard case using the 10 iterations and this is now reflected in the very similar results between the two. The performance of the placeholder approach is essentially unaffected and remains superior.

This analysis, using only two particles, allows for insight into the core problem causing the increasing position error during the standard shaking approach, but is not sufficient to demonstrate the suitability of the placeholder technique in real world applications as these are essentially perfect conditions. In actual measurements a multitude of particles, some detected, some not, are involved, so further investigation is needed.

We therefore test this on a synthetic four camera experiment at 0.2 ppp and designate a certain percentage of the known particles as found to be used for shaking. Their positions are displaced from the ground truth positions by random position noise and then subjected to the shaking procedure. The quality of the shaking process can thus be determined by how well the particles are shaken back to their ground truth positions. Figure 8 shows the results of this analysis in terms of the rms particle position error after shaking. One can see that the performance of the standard shake approach is highly sensitive to the number of found particles. For the 30% case the position accuracy even decreases after the shaking. But if more particles have been found it becomes more and more successful. Using a small shake step decreases the position error in all cases but the optimizer can not sufficiently correct the position errors in the cases where more particles have been found. The placeholder approach is highly successful in optimizing the particle positions for the cases where not all particles have been found yet. Only for the 100% case where no placeholders would be needed it performs worse as the standard processing. This could be improved upon in the future by better managing when placeholders are needed and when not.

The setup as presented corresponds to the shaking in a single IPR iteration. If the position accuracy is actively worsened such as with the standard shaking approach then subsequent IPR iterations may have trouble improving the particle yield or progressing in such a way that the particle er-




Figure 8. Analysis of particle position errors of the three shaking variants on a dense synthetic test case at 0.2 ppp for various amounts of the percentage of found particles in the measurement. All variants are initialized with the same amount of simulated measurement error but performance of the shaking approaches depends strongly on the ratio of found particles with the standard shake approach exhibiting the greatest variability.

ror accumulated at the beginning may never successfully be reduced again. Since the multiple IPR iteration mean the shaking is repeated additional times this leaves an application of the small shaking approach still viable. As seen in figure 8 the positioning error is reduced only by a small amount, but at least there is no active worsening. This then allows subsequent IPR iterations to acquire more particles which then improve shaking conditions in the later iterations. The total number of shake iterations over all IPR iterations then leads to a similar behaviour as shown in figure 7 with 100 shake iterations (similar to 10 shake iterations each for 10 IPR iterations) allowing to correct even large errors successfully. The major benefit of spreading the 100 shake iterations over the IPR iterations is that due to the IPR additional particles keep being added allowing the algorithm to escape the detrimental region of few yet found particles. With the standard shaking approach this was not possible since the early shake iterations actively worsen the particle positions thus preventing the convergence behavior.

While this is a viable alternative, the placeholder approach promises better results as it is able to quickly reduce position errors which makes subsequent IPR iterations less challenging. The small shake step strategy can also only work if it is even possible to ultimately converge to a state where almost all particle images in the camera image have been successfully reconstructed. If the common camera volume (where triangulation is possible) is smaller than the imaged region, i.e. some particles seen in one camera are not visible in the others, then this approach will fail. The images of the non-triangulatable particles will always be present in the measurement images and contaminate the shaking without any chance of them being eventually found. Since the placeholder approach relies only on the peak detection it is not affected and would keep such particles

as permanent placeholders, thus allowing for correct shaking of the found particles in the common volume.

3. Conclusion

Due to time reasons we were unfortunately not able to further integrate the placeholder concept into our full IPR and STB processing yet but we believe the preliminary investigations show good potential. As a consequence of this investigation we have used the small shaking technique for a corrected submission to the LPT Challenge case as no additional implementation work was required. Using this approach we were indeed able to reduce the large position errors present in the initial submission, as shown in figure 1, obtaining very good reconstruction quality.

The current implementation of the placeholder approach can be further optimized. At the moment the placeholders are purely passive components in the shaking procedure as they are temporally "burned" into the residual images. This fully relies on the position accuracy of the peak detector which - while good - may still be improved. This could be improved by a joint global optimization of the 2d placeholder positions together with the 3d particles. Such an approach would need significant changes to the existing shaking architecture. The basic implementation of the placeholder concept already shows promising results. We would have liked to further demonstrate the effect of the placeholder concept within a full IPR processing or application to the LPT challenge data but were unable to implement the changes into the IPR and STB algorithms in time. Especially the I/U variants of Case 2 would be a good testing ground due to the reduced common volume. As discussed above, the small shake approach is limited in such cases and our higher position error for this case seems to confirm this.

References

- Cornic, P., Champagnat, F., & Leclaire, B. (2021, August). Lagrangian particle tracking: a link between localization error and fraction of missed particles. *14th International Symposium on Particle Image Velocimetry*, 1(1). doi: 10.18409/ispiv.v1i1.186
- Godbersen, P., Schanz, D., & Schröder, A. (2024, October). Peak-CNN: improved particle image localization using single-stage CNNs. *Experiments in Fluids*, 65(10). doi: 10.1007/s00348-024-03884-z
- Jahn, T., Schanz, D., & Schröder, A. (2021). Advanced iterative particle reconstruction for lagrangian particle tracking. *Experiments in Fluids*, 62(8), 1–24.

- Leclaire, B., Cornic, F., P. and Champagnat, Fabre, E., & Calmels, F. (2025). *A web portal for automatic performance evaluation of lagrangian particle tracking and data assimilation algorithms*. Retrieved from "https://onera-flow-benchmarks.fr/en/Second_LPT_DA_Challenge"
- Schanz, D., Gesemann, S., & Schröder, A. (2016, Apr 27). Shake-the-box: Lagrangian particle tracking at high particle image densities. *Experiments in Fluids*, 57(5), 70. doi: 10.1007/s00348-016-2157-1
- Schanz, D., Gesemann, S., Schröder, A., Wieneke, B., & Novara, M. (2012, dec). Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. *Measurement Science and Technology*, 24(2), 024009. doi: 10.1088/0957-0233/24/2/024009
- Wieneke, B. (2012). Iterative reconstruction of volumetric particle distribution. *Measurement Science and Technology*, 24(2), 024008. doi: 10.1088/0957-0233/24/2/024008