Advanced iterative particle reconstruction for Lagrangian particle tracking

Tobias Jahn¹ · Daniel Schanz¹ · Andreas Schröder¹

Received: 16 February 2021 / Revised: 21 May 2021 / Accepted: 21 July 2021 / Published online: 4 August 2021
© The Author(s) 2021

Abstract

The method of iterative particle reconstruction (IPR), introduced by Wieneke (Meas Sci Technol 24:024008, 2013), constitutes a major step toward Lagrangian particle tracking in densely seeded flows (Schanz et al. in Exp Fluids 57:1–27, 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. Furthermore, a bundle adjustment scheme that simultaneously updates the 3D locations of all particles and the camera calibrations is introduced. While the particle position optimization proved to be more effective using localized ‘shake’ schemes, this so-called global shake scheme constitutes an effective measure to correct for decalibrations and vibrations, acting as an in-situ single-image volume-self-calibration. Further optimization strategies using such approaches are conceivable.

Graphic abstract

This paper presents an in-depth analysis of the working scheme and the performance of advanced IPR in different imaging conditions.

✉ Daniel Schanz
daniel.schanz@dlr.de

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

1 Introduction

Determining the three-dimensional position of small, point-like objects from their projection on several (few) cameras has been the object of many developments in the last decades. Particularly in the field of volumetric optical flow measurements, various schemes for the accurate

reconstruction of particle positions were explored. Here, the flow of a fluid—otherwise invisible—is revealed by adding small tracer particles ('seeding the flow'). The projections of the particles are recorded by their scattered light over several time-steps with multiple (typically 3–8) cameras. Finally, the flow field is estimated from their imaged motion. Depending on the examined flow and the equipment, the number of recorded consecutive time-steps varies from two (double-pulse recordings), over 4–8 (multi-pulse recordings) to sometimes hundreds of thousands (fully time-resolved recordings).

Potentially, the higher the concentration of the seeding the better the spatial resolution of the reconstructed flow field. The particle image density N_I is typically expressed in terms of 'particles per pixel' (ppp) on a single camera. 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).

For the three-dimensional reconstruction of the particle positions, as well as the estimation of the underlying flow field, several techniques have been developed. First approaches used manual tracking of singular particles, dating back to the beginning and mid of the twentieth century (Nayler and Frazer 1917, Chiu and Rib 1956). The first attempts of three-dimensional particle tracking velocimetry (3D-PTV, e.g., Nishino et al. 1989) applied triangulation using epipolar lines (Maas et al. 1993) and connected the reconstructed particle clouds of each time step to Lagrangian particle trajectories by several approaches, e.g., minimizing the acceleration over four consecutive frames (Malik et al. 1993). The triangulation process is susceptible to the reconstruction of ghost particles. In order to minimize false particle tracks, 3D PTV was typically applied on images with $N_I < 0.005$ ppp (at $N_C = 4$, which is used as the standard case for the rest of manuscript), which for most flows does not yield instantaneous results at sufficient spatial resolution, therefore limiting the attainable results to Lagrangian and average flow properties.

The desire to resolve instantaneous flow structures volumetrically led to the development of tomographic particle image velocimetry (Tomo PIV, Elsinga et al. 2006). The approach to tomographically reconstruct the volume in a so-called voxel-space (Atkinson and Soria 2009), followed by a 3D cross-correlation, was later extended by methods using the temporal information, like MTE (Novara et al. 2010) or SMTE (Lynch and Scarano 2015). Hybrid approaches, combining aspects of tomographic and particle-based

reconstruction, emerged (Schröder et al. 2011; Novara and Scarano 2013; Cornic et al. 2020).

Due to the multiplicative nature of the algorithm, the accuracy of the camera system calibration with a (multi-plane) calibration target proved to be not sufficient for the tomographic reconstruction. Volume self-calibration (VSC, Wieneke 2008) uses triangulation of a sparse ensemble of particles to correct global and local decalibrations of the camera system, decreasing the error from $N(1 \text{ pixel})$ to below 0.1 pixel (px). Building upon VSC, a calibration procedure of the optical transfer function (OTF) (Schanz et al. 2013a) of the particles was developed, accounting for the often-varying imaging conditions across the measurement volume between the different cameras and within a single camera frame.

The introduction of the iterative particle reconstruction (IPR, Wieneke 2013) overcame the limitations of 3D-PTV and was a first major step toward performing Lagrangian particle tracking (LPT) with high particle image densities (Schanz et al. 2013b) comparable to Tomo PIV. 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 with $N_I \leq 0.05$ ppp with reduced occurrence of ghost particles and higher position accuracies, compared to a MART reconstruction. When further increasing N_I , the results deteriorate rapidly, when processing single images. The position optimization (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 calibration demands are even higher, compared to Tomo PIV. An accurate VSC and OTF calibration is mandatory for IPR.

Both VSC and OTF calibration rely on determining 3D particle positions, for which they apply triangulation. As the particle position information is intrinsically available within IPR, it is conceivable to substitute the separate determination of VSC and OTFs with an integrated process, where the calibration and the particle imaging are determined along the optimal particle positions within a bundle adjustment process. This solution allows to dynamically adapt to changing imaging conditions, induced by, e.g., vibrations or varying density gradients. A first step in this direction is introduced in this publication, applying an LBFGS algorithm to simultaneously optimize particle positions and correct for camera vibrations.

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 $N_1 = 0.2$ ppp for synthetic data (Raf-fel et al. 2018). 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).

Experimentally, STB has been applied at particle image densities $N_1 > 0.1$ ppp (Huhn et al. 2017; Bosbach et al. 2019). Recently, an open-source version of the algorithm was developed (Tan et al. 2020). For high-speed flows, where time-resolved recordings cannot be conducted, the method was adapted to a multi-pulse reconstruction strategy (Novara et al. 2016a, 2019a, b). A joint optimization of two time-steps proved to improve the reconstruction quality for both instances (Kähler et al. 2016; Lasinger et al. 2019).

The availability of highly accurate and sufficiently dense particle trajectories spawned the development of data assimilation schemes, which incorporate constraints from flow physics into the interpolation procedures, thus enhancing the spatial resolution beyond the sampling provided by the particles. Two examples for such data assimilation methods are the FlowFit algorithm (Gesemann et al. 2016) and the VIC+—scheme (Schneiders and Scarano 2016).

With the IPR algorithm as one key element of the most recent advances in LPT, it is evident that an improved IPR performance will have a direct impact onto the reconstruction and tracking performances of several recent methods. The present publication introduces a number of advancements of the IPR idea in several key domains.

Section 2 discusses some prerequisites needed by the IPR processing. Section 3 describes the working principle of IPR, gives in-depth introductions to all main steps and highlights the novel aspects introduced in this work. Section 4 introduces a synthetic test case, presents the results of our advanced IPR when reconstructing the generated images with varying imaging conditions, quantifies the effect of each updated approach and discusses the influence of major parameters. Section 5 demonstrates the capabilities of the novel ‘global shake’ approach to instantaneously correct decalibrations of the camera system as part of the reconstruction. Section 6 lists various experimental investigations, which applied the advanced IPR processing. Finally, Sect. 7 summarizes the findings and gives an outlook onto further ideas to improve the particle reconstruction.

2 Prerequisites

The underlying problem IPR aims to solve is to find a 3D particle distribution $(x, y, z, I)_i$ for $i = 1 \dots N_p$, 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 $(u_i, v_i) = \text{cam}((x, y, z)_i)$ and a calibrated 2D particle image shape in dependence on the 3d coordinate $\text{shape}_i = \text{otf}((x, y, z)_i)$, the reprojected camera image can be calculated by adding up the images of all particles:

$$\text{rep}_{uv} = \sum_i I_i \text{shape}_i(u - u_i, v - v_i) \quad (1)$$

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:

2.1 Camera model and calibration

IPR only reconstructs the particle distribution $(x, y, z, I)_i$ and therefore relies on prior knowledge of the camera system. Typically, we aim at reconstructing the particle positions with accuracies below 0.1 px (much smaller for noise-free synthetic data). Therefore, we have high requirements on the calibration of the camera system, which are typically met by performing a volume-self-calibration (VSC, Wieneke 2008) on top of the geometrical calibration using a 3D calibration plate.

The gained correspondences of world coordinates to image coordinates are used to optimize the parameters of a camera model. This model represents the camera cam as a function that projects from 3D world-coordinates onto the 2D image coordinates, as well as a function cam^{-1} that projects from 2D image coordinates onto 3D line-of-sights (represented by two points lying on said line). Further, the Jacobian matrix J with the derivatives of the camera function is required for the optimization of the IPR particles and the calibration of the cameras.

The actual implementation of the camera model can be freely chosen, as long as the aforementioned properties are fulfilled. For the synthetic test cases below, we use a perspective camera (akin to a pin-hole model). For experimental cases, such a model can be amended with second- and/or fourth-order correction terms.

2.2 Optical transfer function (OTF)

The images of illuminated particles typically vary in shape and intensity, depending on their position within the measurement volume and the different viewing angles of the cameras. The attainable position accuracy of an image matching process highly depends on the conformity of the two involved images—the re-projected—and the original particle image. Therefore, the imaging properties of the particles are calibrated using the methods described in (Schanz et al. 2013a). The resulting average shapes of the particle images within several sub-volumes are modeled as a 2D Gaussian with 4 parameters for each camera:

$$\text{otf}(u, v) = I_{a,r} e^{-0.5(u^2a+2buv+cv^2)} \quad (2)$$

The parameters a , b and c encode the circular or elliptical shape and orientation of the Gaussian and the parameter $I_{a,r}$ the average brightness of the particles on the current camera, relative to the other cameras. Since the particle images are assumed to be diffraction limited, the four parameters are solely dependent on the particles 3D position and the local OTF of the respective camera.

Within the IPR processing, the calibrated shape parameters are used as kernels around the projection points of the reconstructed particles while creating the projected images [as per Eq. (1)].

2.3 Rendering of reprojected and residual images

The task to optimally reproduce the measured image img_{uv} by a reprojection of a 3D particle cloud is tackled by reducing the pixel-wise squared distance between img_{uv} and rep_{uv} for each camera with as few and as bright particles as possible. To that end, we define the residual image res_{uv} as the pixel-wise difference between img_{uv} and rep_{uv} :

$$\text{res}_{uv} = \text{img}_{uv} - \text{rep}_{uv} \quad (3)$$

With this residual image, the cost function on camera j becomes

$$\text{cost}_j = \sum_{uv} |\text{img}_{uv} - \text{rep}_{uv}|^2 \quad (4)$$

and the total cost of the reconstruction is the sum of the costs of each camera

$$\text{cost} = \sum_j \text{cost}_j. \quad (5)$$

In an ideal situation, res_{uv} would be completely void of residuals of particles already contained in the state vector. However, due to particle imaging issues (noise, deviation of a particle image from the calibrated OTF, non-Gaussian

particle images, particles in outer regions only visible by a subset of cameras, etc.) and due to the existence of ghost particles (at least in the first iterations of IPR), residuals of the images of already reconstructed particles are common.

3 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:

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.

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.

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 position and intensity of each particle are optimized individually to best match the measured images on all cameras simultaneously using image matching methods. The state vector is filtered for potential ghost particles, mainly based on the particle intensity (deleting particles that fall below an intensity threshold).

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 pre-defined 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 Sects. 3.3 and 4.9).
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 Sects. 3.2.1 and 4.6).
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 Sects. 3.2.2 and 4.8.
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 Sects. 3.3.2 and 4.7).
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 Sects. 3.6 and 5).

The five main steps of the IPR method and the changes introduced in this publication are described in-depth in Sects. 3.1 to 3.5. Section 3.6 introduces the global shake method.

3.1 Peak detection

The detection of particle images on the camera images is carried out using a custom procedure that understands all pixels of the image as the parameters of 2D B-splines and therefore effectively smooths the image. While the smoothing reduces noise the splines also provide an interpolation between the pixels of the image and spatial derivatives. The peak-detection iterates over the B-spline grid looking for bright spots above a certain threshold P_{min} . For every bright spot, a Levenberg–Marquardt algorithm (Moré 1978) is used to find the associated maximum with subpixel accuracy (approximately on the order of 0.1 px). Further, it estimates the shape of each found peak, suitable for calibrating the OTF.

3.2 Triangulation

From the peaks detected on the (residual) 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 epipolar geometry and estimates a set of possible 3D particle positions given detected 2D particle peaks from a set of multiple cameras (see Fig. 2 for a scheme).

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

Fig. 2 Illustration of the triangulation procedure. Red dots depict selected peaks, green depicts the queried surface. The algorithm first selects one peak on Cam 0. The corresponding 3D line of sight (blue) is calculated and projected onto Cam 1. On Cam 1 all peaks within a given tolerance ε (allowed triangulation error) around this line-segment (green area) 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 parti-

cle. 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 (green circles). In the example, four possible 3D particle positions remain. When a particle was 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

- 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 heavily 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 :

3.2.1 Camera permutations

As indicated before, 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.

The effect of the permutation procedure is documented in Sect. 4.6

3.2.2 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: $\varepsilon(i) = \varepsilon(1) + i / (N_{\text{IPR}} - 1) * \Delta\varepsilon$, with $\Delta\varepsilon = \varepsilon(N_{\text{IPR}}) - \varepsilon(1)$. The start and end-values, $\varepsilon(0)$ and $\varepsilon(N_{\text{IPR}})$ 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. The effect of this procedure is documented in Sect. 4.8

3.3 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 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 ± 0.1 pixels and the local cost is determined on all of these positions. A second-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 :

$$\frac{d}{du_i} \text{cost}_j = 2I_i \sum_{uv} \text{res}_{uv} \frac{\delta \text{shape}_i}{\delta u} (u - u_i, v - v_i) \quad (6)$$

$$\frac{d}{dv_i} \text{cost}_j = 2I_i \sum_{uv} \text{res}_{uv} \frac{\delta \text{shape}_i}{\delta v} (u - u_i, v - v_i) \quad (7)$$

$$\frac{d}{dI_i} \text{cost}_j = -2 \sum_{uv} \text{res}_{uv} \text{shape}_i (u - u_i, v - v_i) \quad (8)$$

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. Given the Jacobian of the camera function, the chain-rule can be applied to Eqs. (6) and (7), yielding

$$\begin{aligned} \frac{\delta}{\delta x_i} \text{cost} &= \frac{\delta u_i}{\delta x_i} \frac{\delta \text{cost}}{\delta u_i} + \frac{\delta v_i}{\delta x_i} \frac{\delta \text{cost}}{\delta v_i} \\ \frac{\delta}{\delta y_i} \text{cost} &= \frac{\delta u_i}{\delta y_i} \frac{\delta \text{cost}}{\delta u_i} + \frac{\delta v_i}{\delta y_i} \frac{\delta \text{cost}}{\delta v_i} \\ \frac{\delta}{\delta z_i} \text{cost} &= \frac{\delta u_i}{\delta z_i} \frac{\delta \text{cost}}{\delta u_i} + \frac{\delta v_i}{\delta z_i} \frac{\delta \text{cost}}{\delta v_i} \end{aligned} \quad (9)$$

Using the same schemes, also the second derivate of the cost function with respect to the spatial coordinates can be determined analytically. 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.

The classical approach to use these cost derivatives to fit the interim solution to the measurement is to apply common fitting algorithms, like LBFGS, to solve the problem for all particles simultaneously. Unfortunately, standard libraries often do not allow adding and removing of particles at runtime (we expect it to be possible to adapt the LBFGS to these needs, however that is a work in progress). Furthermore, special features, like ignoring a selection of cameras for each particle to actively suppress ghost particles are not available. Due to these reasons, localized methods that treat each particle individually have shown a better reconstruction performance. Still, global solvers have a potential of being effective tools for the position optimization, once specialized versions are developed. A first useful development and its application are presented in Sects. 3.6 and 5.

Using the derivatives introduced above, several optimization procedures working on individual particles have been tested, including ones that query the second derivative to estimate the optimal step size, thereby requiring ideally only one sampling of the cost function.

However, only the technique that proved to be the most stable is presented here, which is the steepest-descent method, requiring only the first derivative.

3.3.1 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 parameters. Let $\vec{x}_i^T = (x_i, y_i, z_i, I_i)$ be the state vector of particle i . Using Eq. (9), we determine the gradient $\vec{\nabla}_i \text{cost}(\vec{x}_i^T)$ 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 Sect. 3.3.2) 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: $\vec{\nabla}_c = \vec{\nabla} \text{cost}(\vec{x}_i^T) / |\vec{\nabla} \text{cost}(\vec{x}_i^T)|$. 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

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

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. Several other approaches have been tried (fitting a function and identifying the minimum, using more queries, applying an adaptive step size); however, the described method showed the best compromise between computational efficiency (only three queries of the cost function) and stability (low tendency of oscillations).

3.3.2 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. To this end, in addition to the first derivative $\partial_{I_i} \text{cost}_j$ (8), the second derivative

$$\partial_{I_i}^2 \text{cost}_j = 2 \sum_{uv} \text{shape}_i^2 (u - u_i, v - v_i) \quad (10)$$

is determined, which can, for sufficiently wide shapes, be approximated as $\partial_{I_i}^2 \text{cost}_j = 2 \|\text{shape}_i\|^2$. The new intensity can then be expressed as

$$I_{i,j,\text{new}} = I_{i,j} - \frac{\partial_{I_i} \text{cost}_j}{\partial_{I_i}^2 \text{cost}_j} \approx I_{i,j} - \frac{\partial_{I_i} \text{cost}_j}{2 \|\text{shape}_i\|^2} \quad (11)$$

The total intensity update $I_{i,\text{new}}$ 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: $I_{i,\text{new}} = \gamma * I_{i,\text{new}} + (1 - \gamma) * I_i$. A value of $\gamma = 0.5$ has generally shown good results.

In both Wieneke's (2013) approach and 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 according to Eq. (11), 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. The characteristics of the SDM are compared to the original shake approach by Wieneke in Sect. 4.9.

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 Fig. 3). 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. Also, particles need to be visible within at least two cameras. As soon as they leave this volume, they are deleted.

Fig. 3 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.

3.5 Rendering of residual images

After triangulation and position optimization with filtering, the residual images res_{uv} , as introduced in Sect. 2.3, 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 $f_{\text{pt}} \geq 1$ 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 $f_{\text{pt}} = 1$ is used).

3.6 Position optimization using ensemble-based image matching ('global shake')

In Sect. 3.3, we showed how to efficiently determine the derivatives of the cost function with respect to the $4N$ parameters x_i, y_i, z_i, I_i for all N particles (by estimating the derivatives for each particle on each camera and looping over its immediate proximity in the residual image).

As an alternative approach, the cost function together with this $4N$ -dimensional gradient vector allows us to apply common fitting algorithms, like the LBFGS algorithm for a global optimization of the problem. This way, all particle positions will be optimized within the same fitting step, rather than being treated consecutively.

As indicated in Sect. 3.3, these algorithms currently cannot compete with approaches that optimize each particle individually. Another feature of this approach, however, proves to be very useful. Namely, during the estimation of the derivatives with respect to the particle parameters it is possible to also estimate the derivatives of the cost function with respect to internal camera parameters with little additional cost. To take advantage of this, we introduced an additional global 2D shift

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

on each camera to model small shifts induced due to vibrations during the measurement.

$$\begin{aligned} u_i &= u_i(x_i, y_i, z_i) + \text{shift}_u \\ v_i &= v_i(x_i, y_i, z_i) + \text{shift}_v \end{aligned} \quad (12)$$

These two extra parameters per camera can simply be added to the gradient vector, resulting in a total of $4N + 2N_{\text{cams}}$ number of degrees of freedom to be optimized. This way, the global shake procedure can be used to correct for vibrations in an experimental setup or for small initial deviations from the camera calibration, which might occur between subsequent realizations of an experiment.

The capability of the global shake, which can be regarded as a bundle adjustment method, to account for decalibrations is demonstrated in chapter 6.

4 Application of advanced IPR to synthetic data

The IPR algorithm, including all introduced extensions, has been implemented as a C++ library that allows easy scripting via a Python interface. In order to quantify the reconstruction capabilities of extended IPR, it is applied to several synthetic test cases.

Section 4.1 describes the creation of synthetic data and specifies the general parameters of the IPR used for the reconstructions. Section 4.2 gives results for perfect imaging conditions, Sect. 4.3 quantifies the effects of noise and other imaging artifacts, Sect. 4.5 describes the influence of the number of IPR iterations, while Sect. 4.4 does the same for the number of shake and filtering iterations. Section 4.6 demonstrates the effect of permutating the camera order, Sect. 4.7 quantifies the influence of the intensity update and Sect. 4.8 covers the choice of the allowed triangulation error. Finally, Sect. 4.9 compares two approaches of calculating the derivative of the cost function.

4.1 Synthetic test-cases

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 particle shape is Gaussian with a diameter of 2 pixels ($I_{a,r} = 1.0$, $a = 2$, $b = 0$, $c = 2$, see Sect. 2.2). No variations of the OTF were included here. The pixel 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.

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 $\varepsilon(N_{\text{IPR}}) = 0.8$ px (see Sect. 3.2). 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 $N_{\text{shake}} = 10$ 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 Sect. 3.3) is set to $N_{\text{Ci}} = 2$. 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 Sect. 3.2)). The maximum step width is initially set to $d_{\text{max}} = 0.4$ px. For iterations, 15–17 d_{max} is reduced to 0.1 px; for iterations, 18–20 a further reduction to $d_{\text{max}} = 0.02$ px is applied. The factor applied to the reprojected image was $f_{\text{pt}} = 1.0$.

The created virtual images are reconstructed and the results compared to the original particle distribution.

4.2 Perfect imaging conditions

First the performance of the algorithm will be assessed on images free of any defects (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. 4). 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 $P_{\text{min}} = 50$ counts, the intensity threshold in the volume, below which a particle is discarded as a ghost particle, is $I_{\text{min}} = 200$ counts (i.e., 20% of the average particle intensity).

Figure 5 shows the convergence behavior of IPR using the steepest-descent method (SDM) for varying particle image densities, when applied to the described noise-free 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 reprojected points to the original peak position).

For lower particle image densities ($N_1 \leq 0.08$ ppp), the method converges to a stable value of $R_f \approx 1$ in less than 5 iterations. At $N_1 = 0.14$ 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 (see also Sect. 4.4).

For low values of N_1 , 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 $N_1 = 0.16$ ppp ghosts are still present after 20 iterations (as this case is not yet converged). For all other cases, the number of detected ghosts is 0.

In order to give a better insight into the workings within a single IPR iteration, Fig. 6 shows the development of ghost and found true particles at $N_1 = 0.14$ 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 Fig. 5). 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

Fig. 4 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

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.

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 0.04–0.007 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 $N_I=0.05$ 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 Wieneke's 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 introduced above, whose effect is quantified in Sects. 4.6 to 4.9.

The next section will discuss the influence of camera noise and particle intensity variations on the reconstruction results.

4.3 Effect of image noise and intensity variations

Reconstructions of noise-free synthetic images provide a mean to characterize the theoretic performance of algorithms; however, images gathered within experiments will show numerous deviations from perfect conditions. Camera noise can play an important role, as the light budget is typically limited for volumetric investigations, leading to low signal-to-noise ratios. Seeding particles are typically not perfectly monodisperse, leading to—sometimes broad—distributions of peak intensities. A single particle might even show variations in imaged intensity for the different camera projections due to effects of scattering angles (after Mie theory) (Manovski et al. 2021) or uneven particle shape. In order to characterize the effects of some of these experimental shortcomings on the reconstruction quality, synthetic test cases with two different levels of simulated camera noise and a distribution of particle size (i.e., intensity) were created (cases II and III).

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 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. 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, pixel shot noise is additionally applied. Case II was chosen to represent an experimental case with very good imaging conditions, e.g., recordings of Helium-filled soap bubbles illuminated by LED light (see e.g., Huhn et al. 2017). Case III represents tougher experimental conditions, like laser-illuminated droplets in air (e.g., Manovski et al. 2021) or particles in water (e.g., Neeteson et al. 2016; Schröder et al. 2020). We are aware

Fig. 5 Convergence of single-image particle IPR reconstructions for different particle image densities. 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

that other effects (like speckles and scattering laws) are as well relevant in this context; however, these topics are beyond the scope of this work. Figure 7 shows exemplary excerpts from both cases at low- and high-particle image densities.

Image pre-processing consisted of subtracting a constant of 40 counts (case II) and 80 counts (case III), which roughly corresponds to subtracting the minimum or average image, as often performed for experimental cases. The main IPR

Fig. 6 Convergence of found true and ghost particles for case I at $N_1=0.14$ 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 Fig. 5). Creation and deletion of particles within each outer IPR iteration become noticeable

parameters are the same as indicated in Sect. 4.1. Specific parameter for this case: the intensity threshold for identifying peaks on the (residual) images had to be increased to $P_{\min}=70$ counts (case II) and $P_{\min}=90$ counts (case III), in order to limit the influence of noise peaks. In order to more effectively remove false particles, triangulated from the remaining noise peaks, the intensity threshold in the volume, below which a particle is discarded as a ghost particle, is increased to $I_{\min}=350$ counts (case II) and $I_{\min}=450$ counts (case III), respectively.

Figure 8 shows the convergence of the advanced IPR processing for varying particle image densities, when applied to the noisy synthetic datasets. For the lower noise level (left column), it can be seen that the general convergence is slowed down compared to the noise-free case. However, all cases below $N_1=0.14$ ppp converge to $R_f \approx 0.99$ within at most 13 iterations. For $N_1=0.14$ 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, which is slightly offset by 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 $N_1=0.1$ ppp. For higher particle image densities, lower values are seen; still, approx. 66% of the real particles are correctly identified at $N_1=0.14$ 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. Consequently, the reduction of the maximum shake width to 0.1 and 0.02 px in the last IPR iterations is of no consequence here. It is evident that especially for case III not all particles could be reconstructed, which is likely caused by the large range of particle intensities, leading to peaks that are within the noise level and below the peak detection threshold for most or all cameras.

Figure 9 (top) shows the inner-iteration-resolved convergence of case III at $N_I = 0.12$ ppp. Opposed to the processing

Fig. 7 Excerpts of the synthetic camera images at different noise levels (case II: upper row, case III: lower row) and particle image densities

of case I (see Fig. 6), the number of found true particles is now affected by the filtering process. Even at the later iterations, a certain number of true particles is deleted by the intensity threshold filter, which is likely caused by the presence of low-intensity particles, which are bright enough to be picked up by the peak search, but fall below the relatively

high volume intensity threshold of $I_{\min} = 450$ counts. The number of ghost particles fluctuates strongly within each IPR iteration, illustrating that the noise peaks are indeed strong enough to be detected by the peak finder, but do not form coherent 3D particles, which are easily removed by the intensity filter. The increase of triangulated ghost particles with convergence can be attributed to the increasing allowed triangulation error. The lower absolute number of ghost particles at the beginning in relation to the noise-free case I can be explained by the increased peak detection threshold (90 counts vs. 50 counts).

The combination of high noise levels and a significant spread of particle intensities inhibits a reconstruction of all particles for case III. The majority of the missing real particles are of so low intensity that they fall below the specified intensity threshold. Figure 9 (bottom) documents this by plotting the ratio of found particles R_f^* , this time normalized by the number of true particles that have been assigned an intensity above $I_{\min}$ (450 counts). It can be seen that approx. 95% of these particles are correctly identified up to $N_I = 0.1$ ppp. The remaining undetected particles are missed due to the heavy noise on the particle images, inhibiting a correct detection of the particle peak.

If performing STB evaluations of an experiment with multiple time-steps (ideally fully time-resolved), a full reconstruction is still possible. Due to the predictive step within STB, particles whose trajectory has once been found are typically not lost before leaving the measurement volume. Therefore, it is sufficient to identify a particle only once in four consecutive time-steps within its residence time, thereby greatly enhancing the chances of a successful detection. Two- (Jahn et al. 2017; Lasinger et al. 2019) and multi-pulse particle tracking approaches (Novara et al. 2019a), using, e.g., particle space correlation (PSC) (Novara et al. 2016b) as a local flow predictor, could apply iteratively reduced thresholds in order to track as many particles as possible.

Depending on the type of experiment, several other imaging defects might occur (especially when using laser illumination) that are not integrated into the current simulation. For such experimental data, even lower thresholds might have to be chosen, resulting in high amounts of ghost particles in single time-steps. Again, the dimension of time has to be used to separate them from the true particles.

A useful tool is a histogram of the number of detected peaks in relation to the peak finder threshold, which can help finding a suited threshold to separate noise from signal. Another measure that would be of benefit in such cases is to develop/apply a peak detection method that relies less on the intensity as the sole detection measure. If, e.g., searching for a specific particle image shape instead of a bright peak, the detection would be less susceptible to noise.

Fig. 8 Convergence of single-image particle reconstruction with advanced IPR algorithm using the *Steepest-Decent Method* for different particle image densities. Left column: case II; right column: case III

4.4 Number of IPR iterations

The effects of changing the number of outer (IPR) iterations N_{IPR} is investigated by comparing results of reconstructions using $N_{\text{IPR}} = 5$, $N_{\text{IPR}} = 20$ and $N_{\text{IPR}} = 100$. The linear increase of the triangulation radius from $\varepsilon = 0.4$ to $\varepsilon = 0.8$ is retained, only the spacing changes according to the number of iterations.

All other parameters remain unchanged compared to Sects. 4.2 and 4.3. Figure 10 (top) shows the convergence of these reconstructions for case I at three particle image densities. It can be seen that for the two lower values ($N_{\text{IPR}} = 0.04$ ppp and $N_{\text{IPR}} = 0.1$ ppp) the differences are minor, however the reconstruction using $N_{\text{IPR}} = 5$ actually converges the fastest, due to the quicker increase in ε . At

very high-particle image densities ($N_1 = 0.16$ ppp), the differences are profound: For $N_{\text{IPR}} = 5$ only approx. 60% of the true particles can be identified, using $N_{\text{IPR}} = 20$ increases this value to approx. 85%. When applying 100 iterations, the reconstruction converges to a full solution after 40 iterations.

A slightly different trend can be seen when using the noisy images from case III [Fig. 10 (bottom)]. At low-particle image densities, all reconstructions converge to $R_g \approx 0.82$; however, the less iterations are used, the faster the convergence due to the quick increase of ε . At medium values of N_1 however, this reconstruction using $N_{\text{IPR}} = 5$ is not able to converge to the same value of R_g as the other two and at $N_1 = 0.14$ the differences between the three reconstructions are again profound ($R_g \approx 0.42$ at $N_{\text{IPR}} = 5$, $R_g \approx 0.66$ at $N_{\text{IPR}} = 20$ and $R_g \approx 0.74$ at $N_{\text{IPR}} = 100$).

Fig. 9 Top: convergence of found true and ghost particles for case III at $N_I=0.12$ ppp. Particle statistics evaluated as in Fig. 6. Bottom: Convergence of found particles, normalized with the number of real particles above the volume intensity threshold (450 counts) for case III

To summarize, increasing the number of IPR iterations is beneficial at very high particle image densities. For typical experimental conditions, using low values for N_{IPR} (between three and ten) usually works equally well and further iterations might even not improve the values due to usually more complex experimental imaging conditions than those mimicked in the present synthetic test cases.

4.5 Number of shake iterations

In order to quantify the effect of the number of inner (shake) iterations I_i on the results, the created images have been evaluated by IPR using $N_{shake} = 5$, $N_{shake} = 10$ and $N_{shake} = 20$ shake iterations within each outer iteration. Filtering was always performed starting from the fourth inner iteration. All other parameters remain unchanged compared to Sects. 4.2 and 4.3.

Figure 11 (top) shows the results for three different particle image densities of case II. It can be seen that only for the highest particle image densities noticeable differences can be seen. At $N_I=0.06$ ppp, all cases perform equally. Going

to $N_I=0.1$ ppp, the reconstructions using more iterations show a slightly increased convergence. At $N_I=0.14$ ppp, the reconstruction using $N_{shake} = 5$ does not converge, opposed to the other two cases. The reason for the instability of this case is further analyzed in. When comparing the inner-iteration-resolved convergence (Fig. 11, bottom) to the case using $N_{shake} = 10$ (Fig. 11, middle), it can be seen how at the first iterations the additional iterations in the latter case lead to an ongoing reduction in ghost particles within every outer iteration. This, combined with a better positioning and intensity reconstruction of the particles, leads to a less populated residual image, which in turn leads to a reduced ghost particle triangulation in the next iteration. For the case using $N_{shake} = 5$ on the other hand, the number of ghost particles triangulated within every outer iteration is increasing more and more as the triangulation error is growing with the iteration number. The high number of total particles leads to a reduction of average particle intensity, explaining the decrease of reconstructed true particles in the last iterations, as more and more true particles are deleted by the intensity threshold. Using a reprojection factor $f_{pt} > 1$ might stabilize the process, at the expense of some unidentified true particles.

All in all, it can be stated that for particle image densities that are typically encountered in experimental conditions, using low numbers of shake iterations in order to limit processing time comes with little penalty.

4.6 Camera order while triangulating

Conducting several triangulations with permuted camera order, as introduced in Sect. 3.2.1, has a strong influence on the convergence behavior of IPR. Figure 12 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). All other parameters remain unchanged compared to Sect. 4.3.

It is apparent that the permutation approach greatly 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 ($N_I=0.06$ ppp) and 14 ($N_I=0.12$ 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 $N_I=0.12$ ppp, the ghost level rises within the first five iterations and is only slowly reduced in the following.

For cases I and III, similar results were attained. The slight increase in computational effort by conducting

Fig. 10 Comparison of IPR reconstructions of case I (no noise, top) and case III (heavier noise, bottom) using different numbers of outer IPR iterations for three particle image densities at a logarithmic scale

multiple triangulations is easily compensated by the much faster convergence.

4.7 Intensity update variations

As described in Sect. 3.2.2, all reconstructions presented until now use an intensity update, which ignores the two brightest camera ($N_{Ci}=2$), i.e., uses the average of the two (locally) weakest cameras to compute the updated intensity. In this section, we analyze the impact of changing N_{Ci} . To this end, reconstructions of all test cases and particle image densities were conducted, varying N_{Ci} from 0 to 3.

Figure 13 shows selected results for case I and III. For noise-free images at moderate seeding density ($N_I=0.08$ 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

Fig. 11 Top: comparison of IPR reconstructions of case II (moderate noise) using different numbers of inner iterations for three particle image densities. Middle and bottom: Inner-iteration resolved convergence of found true and ghost particles for case II at $N_I=0.14$ ppp with 10 and 5 inner iterations. Particle statistics evaluated as in Fig. 6

when looking at the result with higher particle image density ($N_I=0.14$ 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 $N_I = 0.12$ 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.

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.

4.8 Choice of triangulation error ϵ

As introduced in Sect. 3.2.2, all previously presented IPR reconstructions apply a triangulation radius, linearly increasing within the N_{IPR} outer iterations:

Fig. 12 Convergence of IPR reconstructions with triangulations using a single camera order compared to using several camera order permutations at $N_I = 0.06$ ppp (top) and $N_I = 0.12$ ppp (bottom) of case II (moderate noise)

$\epsilon(i) = \epsilon(0) + i / (N_{IPR} - 1) * \Delta\epsilon$, with $\Delta\epsilon = \epsilon(N_{IPR}) - \epsilon(0)$ and $\epsilon(0) = 0.4$ px, $\epsilon(N_{IPR}) = 0.8$ px. In this section, the differences to using a fixed value of ϵ , as applied, e.g., in Wieneke (2013), are discussed, as well as other functions of increasing ϵ .

Figure 14 compares the results of reconstructions using fixed values of $\epsilon = 0.6$ px and $\epsilon = 0.8$ px to the linear ramp approach for two particle image densities from case III. At low-particle image density ($N_I = 0.04$ ppp), the results do not differ much, with the reconstruction using $\epsilon = 0.6$ px converging quickest, followed by the ramp approach. Using $\epsilon = 0.8$ 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% of found particles.

At high-particle image density ($N_I = 0.14$ ppp), the results are different. Here, the ramp approach is consistently yielding the best results. Using a fixed $\epsilon = 0.6$ px shows a slower convergence and a lower end value, while the fixed $\epsilon = 0.8$ px case completely fails to yield a significant number of true particles. To better understand these results, Fig. 15 shows the inner-iteration-resolved convergence of the two fixed cases at $N_I = 0.12$ ppp, which can be directly compared to the linear-ramp result, given above in Fig. 9. It is obvious that both cases using a fixed value of ϵ exhibit a high level of triangulated ghost particles at the beginning of the reconstruction ($R_g \approx 2.5$ at $\epsilon = 0.6$ px and $R_g \approx 5.0$ at $\epsilon = 0.8$ px), which can be avoided using a low value of ϵ , as the ramp approach does (here R_g always remains below 0.8). At a fixed value of $\epsilon = 0.6$ px, the reconstruction is able to recover from the initially high ghost fraction, as enough true particles are above the intensity threshold to allow convergence; at $\epsilon = 0.8$ px however, the intensity available from the images is distributed over so many triangulated particles that nearly all of them fall below the threshold of $I_{min} = 450$ counts. From the initially triangulated 55% of true particles, less than 3% remain after filtering. The ghost particles are nearly completely deleted. Therefore, the residual image used for the next triangulation iteration represents an almost unaltered original image and the reconstruction shows no signs of convergence over all 20 iterations.

Of course, these results will vary with the chosen parameters. If the triangulations of the first IPR iterations were conducted using the full set of cameras (i.e., requiring a peak on all 4 cameras, instead of just 3), the number of triangulated particles would be reduced; however, this would require more parameters and proved to be less flexible.

To summarize, using a fixed triangulation error can yield competitive results, especially at lower seeding densities and low-noise conditions, as long as it is chosen small enough. However, this requires a very accurate calibration that is not

Fig. 13 Convergence of IPR reconstructions using several value for the number of locally brightest cameras being ignored during the intensity update (N_{CI}). Shown are $N_I=0.06$ ppp and $N_I=0.14$ ppp for case I (no noise, top and middle) and $N_I=0.12$ ppp for case III (heavier noise, bottom)

always attainable in real-life experiments. Therefore, we recommend using a triangulation error that is increasing with IPR iteration, which intrinsically provides a way to limit the occurrence of ghost particles in the first iterations, where the residual image is still very populated.

Other approaches than a linear increase of epsilon with the outer iteration were also explored: a quadratic function ($\epsilon(i) = \epsilon(0) + (i/(N_{IPR} - 1))^2 * \Delta\epsilon$) and a square-root function ($\epsilon(i) = \epsilon(0) + \sqrt{i/(N_{IPR} - 1)} * \Delta\epsilon$). We saw that the

Fig. 14 Convergence of IPR reconstructions of case III (heavier noise) using either a fixed triangulation error of $\epsilon = 0.6$ px and $\epsilon = 0.8$ px or a linear ramp between $\epsilon = 0.4$ px and $\epsilon = 0.8$ px, increasing with the IPR iterations

square-root function led to faster convergence at noise-free conditions and at relatively low seeding densities, where the quicker progression to higher values of epsilon does not lead to critical ghost levels. At higher seeding densities and image noise though, a slower increase of epsilon better suppresses ghost particles and converges to higher values of found particles, thereby favoring the linear and the quadratic approach. Overall, the linear approach showed the most balanced results.

The process of finding the best value of ϵ for the current iterations could even be automated, further reducing the number of parameters and iterations.

4.9 About derivatives

As laid out in Sect. 3.3, this IPR implementation applies an analytic derivation of the cost function in order to determine the optimal direction and step for the repositioning of each particle in each time-step. Here we want to characterize the differences when comparing this approach to the one applied in Wieneke (2013), who used numeric derivatives by moving the particle $\Delta_p \pm 0.1$ px in all three directions of space, fitting a second-order polynomial to the three values in each direction and finally moving the particle to the minimum of this polynomial (or apply the maximum step of 0.1 px, in case no minimum is found). This ‘classic’ position optimization approach was implemented in the current framework, and two versions were used for reconstruction: one always using a $\Delta_p \pm 0.1$ px, the other using $\Delta_p \pm 0.23$ px and reductions to $\Delta_p \pm 0.1$ px at iteration 15 and to $\Delta_p \pm 0.02$ px at iteration 18. These values are comparable to the ones used for the analytic approach, with a maximum step with of 0.23 px in each direction of space corresponding to a maximum step width

of 0.4 px in 3D space. Please note that Wienekes intensity update was not adopted—the one described in Sect. 3.3.2 is applied in all cases.

Figure 16 compares results from two reconstructions using numeric derivation to an otherwise identical reconstruction using analytic derivatives. As can be seen, the differences are minor. Only at very high particle image densities ($N_I = 0.16$ ppp), the curves of found true and ghost particles differ appreciably. For this case, the reconstruction using numeric derivatives with a small step size ($\Delta_p \pm 0.1$ px) converges slowest; an increase to $\Delta_p \pm 0.23$ px accelerates convergence nearly to the level of the reconstruction using analytic derivation. At lower particle image densities, all algorithms perform basically identical. The positional accuracy of the case using a static value of $\Delta_p \pm 0.1$ px is typically lower than for the cases which reduce the step width for the final iterations. The analytic approach always retains a small accuracy benefit over both numeric approaches.

The findings above indicate that using numeric derivations is a viable alternative to analytic derivations in terms

of reconstruction performance. However, the availability of analytic derivations allows for the application of other concepts to solve the reconstruction problem, like the ‘global shake’ bundle adjustment scheme discussed in the next chapter.

5 Decalibration correction using global shake

All previously presented reconstructions were using a perfectly calibrated camera system. In real-life experiments however, decalibrations are common. These might occur as low-frequency vibrations induced, e.g., by the apparatus (e.g., a wind tunnel) or as by slight movements of the camera system by temperature differences or simply people walking by. For the latter, it is possible to conduct a volume-self-calibration prior to the evaluation of every measurement run; however, this procedure might prove cumbersome. For time-resolved series of images, the calibration is typically corrected only for the beginning or the average of the evaluation; the individual time-steps might still suffer from vibrations occurring within the run-time of the experiment (Michaelis and Wolf 2011).

Due to all these reasons, an automatic calibration correction for every evaluated time-step is desirable. As laid out in Sect. 3.6, the ‘global shake’ method, consisting of a joint solution of the cost-function derivatives of all particles and additional parameters for the camera shift, can be applied. In order to document its functionality, test cases with intentional decalibrations were created. To this end, the four cameras were subjected to a shift of Δ_c in both image space directions prior to the reconstruction. Cases with $\Delta_c = 0.5$ px and $\Delta_c = 0.1$ px were considered. The sign of Δ_c was varied to avoid canceling-out of the shift between the different cameras. In order to better show the influence of the calibration correction, the first five IPR iterations of the reconstruction followed the standard procedure as introduced in Sect. 4.2. From the sixth iteration, five iterations of global shake were performed after the normal shaking using the Steepest-Descent Method. Further reconstructions were conducted, which completely lacked the calibration correction via global shake.

Figure 17 compares exemplary results at $N_I = 0.08$ ppp of case I (no noise) in terms of found true particles, ghost particles, reprojected positional error and the decalibration of the camera system (as the average distance of reprojections of known 3D points in original and decalibrated system).

When looking at the results for the higher decalibration (Fig. 17, top), we see an initially high positional error of $\Delta_f = 0.84$ px, a low fraction of identified particles ($R_f = 0.06$) and quite high ghost levels ($R_g = 0.63$). When continuing the reconstruction without calibration correction

Fig. 15 Convergence of found true and ghost particles for case III at $N_I = 0.12$ ppp with fixed triangulation error $\varepsilon = 0.6$ px (top) and $\varepsilon = 0.8$ px (bottom). Particle statistics have been evaluated after triangulation and every inner (shake-)iteration (10 for each of the 20 outer iterations). The blue and orange curves show inner-iteration-resolved results, while black and red curves connect the end-result of each iteration

(dashed lines), Δ_f is not significantly reduced, the fraction of ghost particles rises further to $R_g = 1.40$ due to the creation of particle peak residuals by the decalibration. More true particles are found, however still only a fraction ($R_f = 0.38$ at the end of the reconstruction). If on the other side global shake is performed, starting from iteration 6 (solid curves), the decalibration is corrected after approx. 10 further iterations. At the same time, the number of found true particles rises, the number of ghost particles is reduced until it reaches 0 and the positional error shrinks to $\Delta_f = 0.02$ px. It can be seen that despite the high level of ghost particles, the global shake method was able to completely correct the decalibration.

Decalibration of an order of $\Delta_c = 0.5$ px is frequently observed by us in experimental setups that have been left at rest for a couple of hours or days. Vibrations on the other hand can be of very different scale, ranging from tens of pixel in large industrial wind-tunnels (Novara et al. 2019b) to fine-scale shifts, induced, e.g., by a fan within the camera. Instantaneous large-scale shifts are out of scope for the

global shake method—it relies on lines-of-sight that are still overlapping with the particle image. Shifts up to 1.2 px have been successfully corrected.

Larger vibrations can be handled if a time-resolved measurement is available. As the frequency of vibrations is typically lower than the recording rate, only small variations from time-step to time-step are encountered (once an initial shift is corrected). Such a situation is simulated with the second scenario ($\Delta_c = 0.1$ px), results of which are shown in Fig. 17, bottom. It can be seen that the convergence behavior is much less affected than for the higher decalibration. Initially much more true particles are identified ($R_f = 0.43$), which converge to almost the full solution within the first five time-steps. The initial ghost particle level ($R_g = 0.43$) is quickly reduced, even without calibration correction (contrary to the higher decalibration case). Without corrections, the position accuracy remains limited to approx. $\Delta_f = 0.12$ px—the same as the system decalibration, which is expected in a no-noise scenario, given the close relation of the two measures. As soon as the global shake is applied, the decalibration is corrected almost instantly and the positional error is reduced to the values seen for a perfectly calibrated case. This study demonstrates that an instantaneous correction of small vibrations using the global shake approach is feasible.

The shown reconstructions are from the noise-free case I (due to a better vividness); however, the method works the same for the noisy cases II and III.

The application of the camera correction using global shake in a real-life experiment is depicted in Fig. 18. Shown are the detected camera shifts for three cameras (out of six) for a time-series of 200 time-steps, which was recorded at a large-scale Rayleigh-Bénard Convection (RBC) experiment (see Bosbach et al. 2019). For Camera 1, an initial decalibration of $(-0.15$ px, 0.19 px) is found for the x - and y -image-coordinate, which is stable over the run-time. For Camera 2, a larger initial decalibration $(-0.20$ px, -0.64 px) is seen. Additionally, this camera exhibits a clear vibration in the x -component with an amplitude of approx. 0.03 px and a frequency of approx. 0.69 Hz. The third camera exhibits only a small initial decalibration and shows a low-amplitude vibration at a frequency of approx. 15 Hz, that is modulated by a lower frequency, which could be an effect of cross-talk with the vibration of camera 2. This experimental test case demonstrates the useful application of global shake to correct for both initial decalibrations (eliminating the need for a separate volume self-calibration step before the evaluation of each measurement run) and vibrations occurring during the actual runtime.

One thing to keep in mind is that by (repeatedly) correcting the camera shifts, it is theoretically possible that the recalibrated camera system might be shifted/rotated/

Fig. 16 Convergence of IPR reconstructions using shaking with (i) numeric derivatives and a maximum step width of $d_{\max} = 0.1$ px; (ii) numeric derivatives and $d_{\max} = 0.23$ px with two reductions of $d_{\max}$ at iterations 15 and 18; (iii) analytic derivatives with $d_{\max} = 0.4$ px. Shown are the fraction of found true and ghost particles and the positional error for particle image densities of $N_1 = 0.12$ ppp (top) and $N_1 = 0.16$ ppp (bottom) from case I (no noise)

Fig. 17 Convergence of IPR at $N_1=0.08$ ppp for case I (no noise) with initially decalibrated camera system (top: 0.5 px; bottom: 0.1 px per camera and image dimension). Dotted lines: without correction; straight lines: In addition to normal shake, global shake with decalibration correction performed after iteration 5

stretched in relation to original (the same is true for VSC). As a counter-measure, an affine transformation can be determined, which maps from the transformed coordinate system onto the old one.

6 Experimental validation

The extended IPR processing presented in this work has already been applied to a variety of experimental evaluations as part of the shake-the-box (Schanz et al. 2016) method. Therefore, we refrain from presenting a dedicated case here. Instead we refer to the following examples where the method has been applied: the study of a rotor wake (Wolf et al. 2019); homogeneous turbulence in water (Schröder et al. 2019); the flow over a surface mounted cube (Schröder et al. 2020); a large-scale turbulent boundary layer investigation (Schanz et al. 2019); a large-scale investigation of Rayleigh-Bénard Convection (Bosbach et al. 2019)—see also (Godbersen et al. 2020) for a Gallery of Fluid Motion presentation;

Examples for the use of advanced IPR in applications of multi-pulse STB (Novara et al. 2016a) are turbulent boundary layer investigations (Novara et al. 2019a), measurements of the flow on the suction side of a half wing model (Novara et al. 2019b) or investigations of a subsonic jet (Manovski et al. 2021).

7 Summary and outlook

Iterative particle reconstruction (IPR, Wieneke 2013) enabled the reconstruction of particle distributions at much higher particle image densities than previous methods. Together with the shake-the-box (STB, Schanz et al. 2016) algorithm, it initiated a paradigm shift in the field of volumetric flow measurements—replacing MART-based particle reconstructions and correlation-based 3D PIV methods by a direct determination of particle trajectories, while on the other side allowing to use particle image densities which are 10 to 30 times higher than possible for classical PTV approaches. The original method introduced by Wieneke allowed processing instantaneous images up to particle image densities of $N_1=0.05$ ppp; coupled with STB processing, this range could be extended to $N_1=0.125$ ppp when applied to time-series (Schanz et al. 2016), see also results of Case D at 0.1 ppp of the 4th Int. PIV Challenge (Kähler et al. 2016).

In this work, an extended IPR is introduced, which amends the original method in several areas. Combined, these enhancements lead to a significant jump in performance, allowing to fully reconstruct particle distributions with up to $N_1=0.14$ ppp on a single time-step. In perfect imaging conditions, nearly all particles are reconstructed, no ghost particles occur and the average positional error is below 1/100th of a pixel, even at such high particle image densities. With low levels of image noise and particle size variations, only the positional error is affected; for higher noise levels and a broad distribution of particle intensities, the detection of true particles is aggravated; however, the occurrence of ghost particles is still low. Such cases are much closer to experimental cases than a noise-free situation. In order to find all particles in such conditions, additional information from the temporal domain is required, which the STB method can provide.

Several changes in relation to the original IPR method, showing a great impact on the reconstruction capabilities, should be straight forward to integrate into existing implementations:

- In the triangulation step, the cameras are queried for peaks at a certain order. This order has been shown to have a huge influence if a single particle can be successfully identified due to the special role of the first

Fig. 18 Time-resolved camera shifts (decalibrations) for three cameras of large-scale RBC experiment (Bosbach et al. 2019), detected and corrected using the global shake method

two cameras. Replacing this single triangulation by multiple triangulations, each with a different camera order and a subsequent combination of the triangulated particle clouds, proved to significantly increase the convergence in all image conditions and particle image densities.

- Another measure, that showed an even greater impact, is to modify the intensity determination for a triangulated particle. Instead of using the average intensity found at the reprojection point on all cameras, the N_{Ci} cameras that show the highest local intensity are ignored. The intensity is determined by averaging the remaining cameras. A value of $N_{Ci} = 2$ has shown the best results for the examined four-camera case. Compared to the standard processing ($N_{Ci} = 0$), a very efficient suppression of ghost particles is found, which allows convergence even at high particle image densities.
- Using an allowed triangulation error ε that is linearly increasing with the number of iterations proved to be beneficial, especially for cases with high noise levels. Compared to several cases using a fixed value of ε lower levels of ghost particles, better convergence and a higher stability is found.

All of these measures are easy to implement and come with little to no computational effort. Instead, due to the higher convergence rate, the number of applied iterations could be reduced.

A study of the number of outer (IPR) iterations and inner (shaking and filtering) iterations showed that increasing both is beneficial when aiming for very high particle image densities, but they can be kept at reasonable numbers for densities up to $N_I = 0.1$ ppp.

Another novel feature of the presented implementation is the analytical derivation of the cost function in order to determine the direction in which a particle has to be shifted in the shake procedure. This approach showed surprisingly

low benefits compared to the original method of using numerical derivation. However, with the knowledge of analytical derivatives it is possible to formulate a global cost function that can be minimized for all particles simultaneously, using standard methods like LBFGS. At the moment, this approach turned out to not be able to compete against the single-particle based shake algorithms due to the intensive optimization possibilities of the latter; however, this might change if specialized versions of the solver are developed and applied.

Still, this ‘global shake’ approach already has a useful purpose, as shifts of the cameras in relation to the original calibration can be introduced into the cost function. This allows to correct for both initial decalibrations—caused by slight movements of the cameras prior to a measurement—and vibrations occurring during the measurement runtime.

With an effective implementation, the IPR can be quite fast. Using optimal settings for inner and outer iterations, the current version converges in noisy conditions (case III) within 5.2 s at $N_I = 0.06$ ppp and within 19.4 s at $N_I = 0.1$ ppp (97,483 particles) on an AMD Threadripper 1950× machine. Within a time-resolved STB evaluation, less iterations could be applied, reducing the computational cost significantly. The current implementation has been applied to many experimental cases, presented in other publications.

Further improvements of the IPR procedure are feasible. Some of these are:

- For single-particle-based shaking, the information of the second derivative of the cost function could be used to find an optimal step width, eliminating the need to sample the cost function on multiple locations.
- The global shake method could be adapted, such that addition and removal of particles are supported by the solver or that individual particle properties are better integrated, such as identifying the brightest cameras for the intensity update of each particle. Also, the cost func-

tion could not only include a two-dimensional shift per camera, accounting for vibrations. Instead, more complex corrections of the camera could be realized, such that instantaneous and local distortions of the calibration, induced, e.g., by density gradients or curved refractive index changes, could be captured.

- An important task is an improved automatization of the parameter settings. Currently most parameters are tuned by hand, requiring an experienced user and time, especially if conditions change within several realizations of an experiment. An automated detection of parameters like the peak detection intensity threshold, the allowed triangulation error, the maximum step size of the shake algorithm, the intensity filter and the number of inner and outer iterations would greatly simplify (and likely speed up) the evaluation.
- The filtering for ghost particles could include much more information than the particle intensity, which is the only criterion applied at the moment. Additionally, information like the intensity variations between the different cameras and the local derivative of the cost function could be used to further discern between true and ghost particles. Machine-learning-based methods, incorporating these information, are currently being tested.
- Peak detection methods could be developed that are less susceptible to image noise could be applied. The currently used approach to look for image peaks above an intensity threshold does not take the particle image shape into account. As this shape is already known via the calibrated OTF, this additional information should be harnessed. Also in this case, Machine-Learning approaches might be a promising approach.

Acknowledgements Open Access funding provided by Projekt DEAL.

Funding Open Access funding enabled and organized by Projekt DEAL. This work was partly supported by the Deutsche Forschungsgemeinschaft (DFG) through Grant No. SCHR 1165/5-1 as part of the Priority Programme on Turbulent Superstructures (DFG SPP 1881, 2nd period) and the HOMER project from the European Union's Horizon 2020 research and innovation program under Grant Agreement No. 769237.

Availability of data and material Data and material are available on request.

Compliance with ethical standards

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source,

provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- 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 (2019) Dense Lagrangian particle tracking of turbulent rayleigh Bénard convection in a cylindrical sample using shake-the-box. In: 17th European turbulence conference (ETC 2019), Torino, Italy
- Chiu WC, Rib LN (1956) The rate of dissipation of energy and the energy spectrum in a low-speed turbulent jet. *Eos, Trans Am Geophys Union* 37(1):13–26
- Cierpka C, Lütke B, Kähler CJ (2013) Higher order multi-frame particle tracking velocimetry. *Exp Fluids* 54(5):1533
- Cornic P, Leclaire B, Champagnat F, Le Besnerais G, Cheminet A, Illoul C, Losfeld G (2020) Double-frame tomographic PTV at high seeding densities. *Exp Fluids* 61(2):1–24
- Elsinga GE, Scarano F, Wieneke B, van Oudheusden BW (2006) Tomographic particle image velocimetry. *Exp Fluids* 41:933–947
- Elsinga GE, Westerweel J, Scarano F, Novara M (2011) On the velocity of ghost particles and the bias errors in tomographic-PIV. *Exp Fluids* 50(4):825–838
- Fletcher R, Powell M (1963) A rapidly convergent descent method for minimization. *Comput J* 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. In: 18th International symposium on application of laser and imaging techniques to fluid mechanics. Lisbon, Portugal
- Godbersen P, Bosbach J, Schanz D, Schröder A (2020) The beauty of turbulent convection: a particle tracking endeavor. *Gallery Fluid Motion*. <https://doi.org/10.1103/APS.DFD.2020.GFM.V0074>
- Huhn F, Schanz D, Gesemann S, Dierksheide U, van de Meerendonk R, Schröder A (2017) Large-scale volumetric flow measurement in a pure thermal plume by dense tracking of helium-filled soap bubbles. *Exp Fluids* 58:116
- Jahn T, Schanz D, Gesemann S, Schröder A (2017) 2-Pulse STB: 3D particle tracking at high particle image densities. In: 12th international symposium on PIV-PIV17. Busan, Korea
- Kähler CJ, Astarita T, Vlachos PP, Sakakibara J, Hain R, Discetti S, La Foy R, Cierpka C (2016) Main results of the 4th International PIV Challenge. *Exp Fluids* 57(6):97
- Lasinger K, Vogel C, Pock T, Schindler K (2019) 3D fluid flow estimation with integrated particle reconstruction. *Int J Comput vis* 128:1012–1027. <https://doi.org/10.1007/s11263-019-01261-6>
- Lynch K, Scarano F (2015) An efficient and accurate approach to MTE-MART for time-resolved tomographic PIV. *Exp Fluids* 56:66
- Maas HG, Grün A, Papantoniou D (1993) Particle tracking in three dimensional turbulent flows—part I: photogrammetric determination of particle coordinates. *Exp Fluids* 15:133–146
- Malik N, Dracos T, Papantoniou D (1993) Particle tracking in three dimensional turbulent flows—part II: particle tracking. *Exp Fluids* 15:279–294

- Manovski P, Novara M, Mohan NKD, 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. *Exp Therm Fluid Sci* 123:110346
- Michaelis D, Novara M, Scarano F, Wieneke B (2010) Comparison of volume reconstruction techniques at different particle densities. In: 15th international symposium on applications of laser techniques to fluid mechanics
- Michaelis D, Wolf CC (2011) Vibration compensation for tomographic PIV using single image volume self calibration. In: 9th international symposium on particle image velocimetry—PIV11. Kobe, Japan
- More JJ (1978) The Levenberg-Marquardt algorithm: implementation and theory. In: Watson GA (ed) *Numerical analysis. Lecture Notes in Mathematics*, vol 630. Springer, Berlin, Heidelberg. <https://doi.org/10.1007/BFb0067700>
- Naylor JL, Frazer BA (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, p 6
- Neeteson NJ, Bhattacharya S, Rival DE, Michaelis D, Schanz D, Schröder A (2016) Pressure-field extraction from Lagrangian flow measurements: first experiences with 4D-PTV data. *Exp Fluids* 57(6):1–18
- Nishino K, Kasagi N, Hirata M (1989) Three-dimensional particle tracking velocimetry based on automated digital image processing. *Trans ASME J Fluid Eng* 111:384–390
- Novara M, Batenburg KJ, Scarano F (2010) Motion tracking enhanced MART for tomographic PIV. *Meas Sci Technol* 21:035401
- Novara M, Scarano F (2013) A particle-tracking approach for accurate material derivative measurements with tomographic PIV. *Exp Fluids* 54:1584
- Novara M, Schanz D, Reuther N, Kähler CJ, Schröder A (2016a) 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 KP, Schröder A (2016b) Lagrangian 3D particle tracking for multi-pulse systems: performance assessment and application of shake-the-box. In: 18th International symposium on applications of laser and imaging techniques to fluid mechanics. Lisbon, Portugal
- Novara M, Schanz D, Geisler R, Gesemann S, Voss C, Schröder A (2019a) Multi-exposed recordings for 3d lagrangian particle tracking with multi-pulse shake-the-box. *Exp Fluids* 60(3):1–19
- Novara M, Schanz D, Geisler R, Gesemann S, Philipp F, Agocs J, Schröder A (2019b) Instantaneous self-calibration of a 3D imaging system in industrial facility with strong vibrations. In: Proceedings of the 13th international symposium on particle image velocimetry. No 180, pp 1133–1143
- Raffel M, Willert CE, Scarano F, Kähler CJ, Wereley ST, Kompenhans J (2018) *Particle image velocimetry: a practical guide* (Digital Content). Springer, Berlin
- 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 Symposium 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>
- Schanz D, Schröder A, Novara M, Geisler R, Agocs J, Eich F, Bross M, Kähler CJ (2019). Large-scale volumetric characterization of a turbulent boundary layer flow. In: Proceedings of the 13th international symposium on particle image velocimetry. No 182, pp 251–265
- Schneiders JFG, Scarano F (2016) Dense velocity reconstruction from tomographic PTV with material derivatives. *Exp Fluids* 57:9
- Schröder A, Geisler R, Staak K, Wieneke B, Elsinga G, Scarano F, Henning A (2011) Lagrangian and Eulerian views into a turbulent boundary layer flow using time-resolved tomographic PIV. *Exp Fluids* 50:1071–1091
- Schröder A, Schanz D, Gesemann S, Huhn F, Paz DG, Perez-Munuzuri V, Bodenschatz E (2019) Measuring the full velocity gradient and dissipation rate tensor in homogeneous turbulence using shake-the-box and flowfit. In: 17th European turbulence conference 2019. Torino, Italy
- 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 Fluids* 61(9):1–22
- Tan S, Salibindla A, Masuk AUM, Ni R (2020) Introducing OpenLPT: new method of removing ghost particles and high-concentration particle shadow tracking. *Exp Fluids* 61(2):1–16
- van Gent PL, Michaelis D, Van Oudheusden BW, Weiss P-É, De Kat R, Laskari A, Yj J, 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 (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
- Wolf CC, Schwarz C, Kaufmann K, Gardner AD, Michaelis D, Bosbach J, Schanz D, Schröder A (2019) Experimental study of secondary vortex structures in a rotor wake. *Exp Fluids* 60(11):1–16

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.