Optimizations of Optical Flow Measurement Systems

Dissertation

zur Erlangung des mathematisch-naturwissenschaftlichen Doktorgrades
„Doctor rerum naturalium“
der Georg-August-Universität Göttingen

im Promotionsprogramm für Informatik (PCS)
der Georg-August University School of Science (GAUSS)

vorgelegt von
Sebastian Gesemann

aus Minden
Göttingen, 2017

Acknowledgements

Abstract

In this cumulative thesis contributions are presented that cover the computational aspects of optical measurement systems for flow investigations in aerodynamics and fluid mechanics. Flow measurements are useful in gaining an understanding of flow dynamics and validating numerical models for flow dynamics simulations. With the goal of accurately capturing and visualizing details of unsteady and moving flow structures based on recorded digital images of small tracer particles, the contributions cover a wide range from digital image processing, camera calibration, 2D and 3D reconstruction, estimation of position, velocity and acceleration, noise reduction and spatial interpolation of scattered 3D data using physically-motivated regularizations.

Contents

III. Curriculum Vitae   **205**

1. Introduction

To this day wind tunnel experiments are performed to improve our understanding of unsteady flows and turbulence and there is no reason to expect that these kinds of experiments will be completely replaced by numerical simulations in the near future due to the computational complexity of certain flow dynamics problems. Real flow experiments are also necessary to validate numerical models that are used as simplifications to reduce the computational complexity of some fluid mechanics simulation problems.

Flow can be made visible using small tracer particles that follow the flow and are illuminated with a laser or another bright light source. Digital cameras can be used to observe the movement of the particles. This is regarded as a non-invasive method as long as the tracer particles don't affect the flow. For one possible measurement setup see Figure 1.1. It shows how five numbered high resolution cameras that are targeted at a cylindrical convection chamber which contains many small Helium-filled soap bubbles. The image was taken from publication [J] referenced in Section 2.2.

The scope of this thesis covers the computational aspects of optical flow measurement in a 3D measurement volume for the purpose of flow dynamics research. The goal was to improve upon what was previously possible in terms of accuracy and spatial resolution. The computer aided processing of digital data for this type of optical measurement involves the following steps:

- camera calibration – deriving camera model parameters based on recorded images of a calibration object and/or the measurement images themselves

- preprocessing of the measurement images (for example to remove any background-related intensity)

- a form of 3D reconstruction yielding a discretized 3D intensity volumes or a set of identified 3D particles with their tracks

- possibly a noise reduction step

- derivation of the flow field based on the result of the previous 3D reconstruction

For volumetric flow measurements aiming at high spatial resolution Tomographic Particle Image Velocimetry (*TomoPIV*, see [Elsinga2006]) methods have been most promising up until now. In TomoPIV a spatially discretized 3D volume of light intensities is reconstructed from multiple 2D images by solving a constraint linear equation system. This equation system is derived under the assumption that light intensities accumulate along lines of sight and depends heavily on the camera calibration of the optical measurement system.
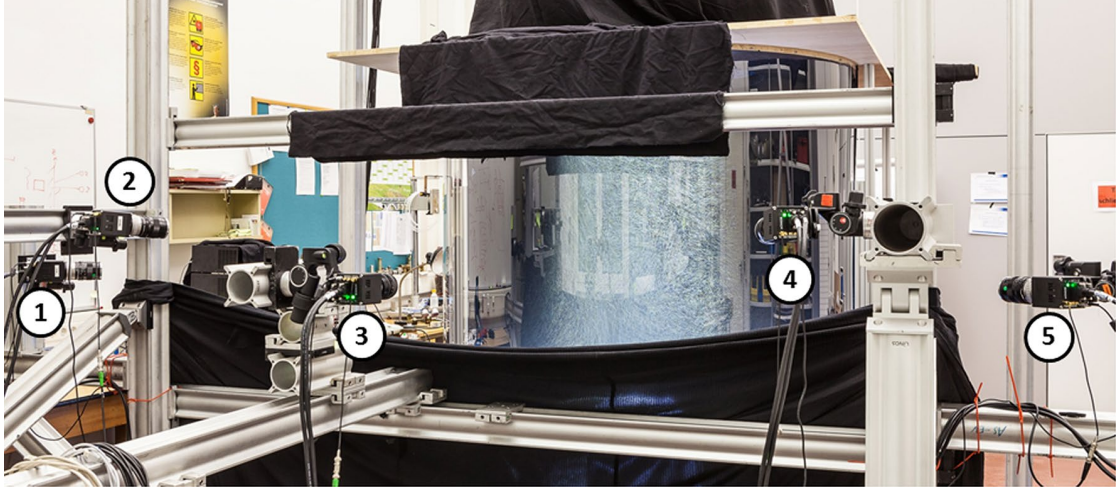
Figure 1.1.: Possible volumetric flow measurement setup as used in [J] with five numbered cameras. They are targeted at a large measurement volume filled with small particles which are illuminated from the top to observe the flow of a thermal plume.

As part of this thesis a new camera calibration procedure has been developed which accounts for a space-varying point spread function (or optical transfer function) often observed under real experimental conditions. Up until now, camera calibration for this type of flow measurement problem was limited to the geometric aspect. Geometric camera calibration is only concerned about how 2D image coordinates and 3D world coordinates are linked. However, camera lenses, light diffraction and refraction at interfaces between different media such as air, water and glass affect how tiny particles appear in the camera images. They tend to appear blurry and possibly elongated in one direction due to astigmatism. The shape of the imaged particles also depend on their location in 3D. Accounting for this type of distortion results in a projection model that is closer to reality and will be beneficial for the accuracy of the following 3D reconstruction step.

The defacto standard method for reconstructing a discrete 3D intensity signal in TomoPIV is the Simultaneous Multiplicative Algebraic Reconstruction Technique (SMART) which was first applied to the TomoPIV reconstrution problem in [Atkinson2009]. An alternative reconstruction called *S-SPG* based on the Nonmonotone Spectral Projected Gradient method (SPG, see [Birgin2000]) combined with a nonlinear substitution of the search space (S) has been developed as part of this thesis. Compared to SMART it is more robust in the presence of noise and converges faster for certain cases.

A very different approach to 3D flow measurements is based on identifying and tracking particles. Up until now, particle tracking techniques were not sophisticated enough to handle high particle densities which are otherwise necessary for a high spatial resolution and a faster convergence of statistical properties of the investigated flow. This thesis covers a novel particle tracking technique named "Shake-the-Box" (STB) for time-resolved

measurements that exploits the temporal coherence of the particles' trajectories and is able to deal with high particle densities. Compared to a tomographic reconstruction that is used in TomoPIV this method avoids the extra discretization of the measurement volume and reconstructs 3D locations and intensities of particles directly which improves the accuracy of 3D position estimates. This tracking method also does not suffer from a kind of reconstruction artefact (*ghost particles*) that are common in TomoPIV at high particle concentrations. In order to analyze flow structures given a finite amount of particle tracks a spatial interpolation of the scattered data is desired. This thesis describes a novel interpolation scheme called *FlowFit* developed for this purpose that is able to recover spatial details beyond the sampling limit by incorporating physical constraints into the interpolation method such as conservation of mass and momentum using the Navier-Stokes equation for incompressible flows. Both a velocity field and a pressure field (up to an unknown offset) is reconstructed with this tracking and fitting approach.

In Chapter 2 an overview of the different contributions is given. Chapter 3 discusses the results. The impact of the contributions is covered in Chapter 4. Chapter 5 gives a conclusion and an outlook.

1. Introduction

2. Overview of own contributions

The following sections give an overview of the author's contributions grouped into major and other own contributions.

2.1. Major own contributions

These contributions are mainly concerned about furthering the development of the algorithms and techniques for measurement data processing with respect to 3D flow measurements. The publications include verifications of the new techniques based on simulations and/or real experimental data. The different processing steps for two approaches are shown in Figure 2.1 along with an indication in square brackets of which publications have that particular processing step as their focus.

[A] Sebastian Gesemann, Daniel Schanz, Andreas Schröder, Stefania Petra, Christoph Schnörr. (2010). Recasting TomoPIV reconstruction as a constrained and L1-regularized nonlinear least squares problem. *15th Int. Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon*

[B] Daniel Schanz, Sebastian Gesemann, Andreas Schröder, Bernhard Wieneke, Matteo Novara. (2013). Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. *Meas. Sci. Technol. 24(2):024009*

[C] Daniel Schanz, Andreas Schröder, Sebastian Gesemann, Dirk Michaelis, Bernhard Wieneke. (2013). Shake The Box: A highly efficient and accurate Tomographic Particle Tracking Velocimetry (TOMO-PTV) method using prediction of particle positions. *10th Int. Symposium on Particle Image Velocimetry, Delft*

[D] Daniel Schanz, Sebastian Gesemann, Andreas Schröder. (2016). Shake-The-Box: Lagrangian particle tracking at high particle image densities. *Exp. Fluids, 57:70*

[E] Sebastian Gesemann, Florian Huhn, Daniel Schanz, Andreas Schröder. (2016) From Noisy Particle Tracks to Velocity, Acceleration and Pressure Fields using B-splines and penalties. *18th Int. Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon*

The tomographic reconstruction algorithm described in [A] is based on ideas by Stefania Petra et al [Petra2008]. They were the first to apply the Spectral Projected Gradient (SPG) method to solve the discrete 3D reconstruction problem. In this publication a
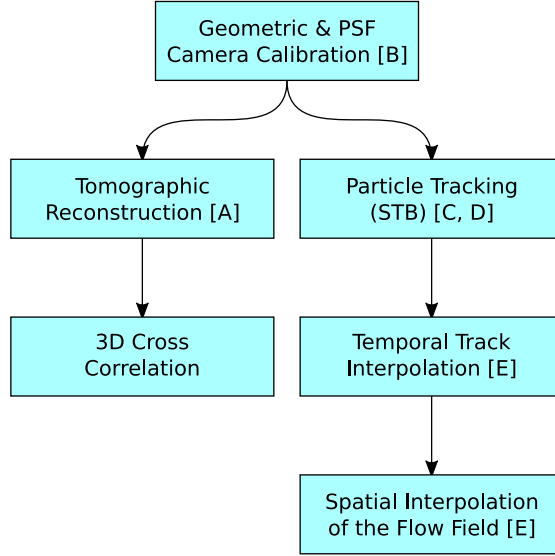
Figure 2.1.: Processing steps involved based on tomography (left path) and Lagrangian Particle Tracking (right path)

method is described (*S-SPG*) and investigated that combines the SPG with a nonlinear warping of the search space. The implementation of these methods as well as the simulation and evaluation has been done by the main author. With the help of Daniel Schanz data from a real experimental has been gathered for testing.

Publication [B] covers a new method for the calibration of the point-spread functions of cameras with the goal of improving the reconstruction quality of 3D intensity volumes. The design and implementation of the calibration procedure in software has been done by the author of this thesis. The approach is based on the self calibration method that Wieneke developed for correcting line-of-sight misalignments and thus works on the measurement images, see [Wieneke2008]. The extraction of particle shape parameters from such images within such a self-calibration is new. Daniel Schanz and Matteo Novara performed the experiment at TU Delft which resulted in data this calibration method was tested on. Daniel Schanz also evaluated the benefit of this new calibration approach.

In publication [C] and [D] a new particle tracking method is described. It is based on the iterative particle reconstruction (IPR, see [Wieneke2013]) technique developed by Wieneke for separate time steps and extended to exploit the temporal relationships between consecutively recorded measurement images in a time-resolved measurement setup. The temporal prediction of particle candidates in order to reduce successively the reconstruction burden and likelihood of ghost particles along the time series of particle images was the idea of Andreas Schröder. The identification of particles in 2D images, their triangulation, the prediction of their movement and the particle track postprocessing for denoising and computing the temporal derivatives was developed by the author of this thesis. Daniel Schanz implemented the "shaking" part of the IPR, the data structure for the particle tracks, the algorithm to extend tracks with new particle positions

and put everything together. He performed the evaluation of the procedure based on simulated and experimental data and the comparison with SMART.

Publication [E] describes a method for particle track smoothing called "TrackFit" and a new spatial interpolation scheme called "FlowFit" that takes scattered particle data (position, velocity and acceleration for one point in time) and reconstructs a flow field while exploiting physical constraints such as mass preservation and momentum preservation. It is based on an earlier simpler reconstruction developed by the author of this thesis. The flow field is reconstructed by using 3D uniform B-splines and determining its parameters by minimizing a cost function.

Andreas Schröder and Florian Huhn provided inspiration for incorporating the particles' estimated accelerations to improve the reconstruction of the flow field, specifically, Andreas Schröder suggested that the availability of acceleration data should help in reconstructing the velocity field and Florian Huhn pointed out a possible simplification of the cost function which got rid of a 3rd order spatial derivative of the velocity field from appearing in the cost function. The author of this thesis developed the model, wrote the code for evaluating the nonlinear cost function and its gradient and evaluated this reconstruction method based on simulated data from the Johns Hopkins Turbulence Database and additional experimental data provided by Daniel Schanz.

The results of these contributions are summarized in Chapter 3 and their impact is discussed in Chapter 4. In addition, appendix I of this summary covers the FlowFit method in more detail.

2.2. Other own contributions

This list of contributions primarily represents applications of the methods that were developed and described in the publications referred to in the previous sections.

[F] Hauke Ehlers, Andreas Schröder, Reinhard Geisler, Sebastian Gesemann. (2012). Combined Time-Resolved PIV and structure deformation measurements for aeroelastic investigations. *STAB Symposium*

[G] Andreas Schröder, Daniel Schanz, Reinhard Geisler, Sebastian Gesemann, Christian Willert. (2015). Near-wall turbulence characterization using 4D-PTV Shake-The-Box. *11th Int. Symposium on Particle Image Velocimetry, PIV15*

[H] Daniel Schanz, Florian Huhn, Sebastian Gesemann, Uwe Dirksheide, Remco van de Meerendonk, Peter Manovski, Andreas Schröder. (2016). Towards high-resolution 3D flow field measurements at cubic meter scales. *18th Int. Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon*

[I] Florian Huhn, Daniel Schanz, Sebastian Gesemann, Andreas Schröder. (2016). FFT integration of instantaneous 3D pressure gradient fields measured by Lagrangian particle tracking in turbulent flows. *Exp. Fluids, 57:151*

2. Overview of own contributions

[J] F. Huhn, D. Schanz, S. Gesemann, U. Dirksheide, R. van der Meerendonk, A. Schröder. (2017). Large-scale volumetric flow measurement in a pure thermal plume by dense tracking of helium-filled soap bubbles. *Exp. Fluids, 58:116*

Publication [F] describes an experiment about the deformation of a fluttering thin plate in a flow. Based on time-resolved particle image velocimetry (PIV) measurements and under the assumption of a 2D flow pressure gradients have been computed from velocity data and integrated spatially. The author of this thesis contributed the code for the spatial integration step of the pressure gradients to compute the pressure. The discrete 2D pressure field was reconstructed by solving an overdetermined linear equation system in the least squares sense which covered equations for the estimated gradients as well as equations for the pressure boundary condition.

Publications [G] and [H] and [J] represent specific applications of the methods developed and described in publications [B], [C] and [D].

The focus of publications [I] is on 3D pressure reconstruction from particle-based flow measurements. The author of this thesis contributed a FlowFit based interpolation of a curl-free pressure gradient field based on scattered acceleration data from observed tracer particles. This pressure gradient field was then spatially integrated using an FFT-based approach developed by Florian Huhn.

3. Discussion of results

Various data processing building blocks for optical flow measurements have been improved or developed as part of this thesis. Each development improved upon what was previously possible in state-of-the-art methods.

3.1. Point spread function calibration and its effect

This extended calibration method that accounts for the possibly space-varying point spread function (PSF) allows the use of a more accurate model of how particles are imaged. See Figure 3.1 for a real-life example of possible PSF-related distortions in particle measurement images due to limited depth of field.
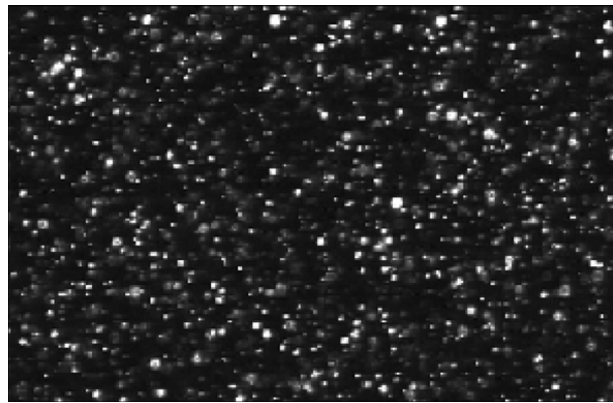
Figure 3.1.: Unprocessed camera image of a TomoPIV experiment in air at f-number 5.6. Focused as well as blurred particle images can be observed.

The developed methods for estimating these possibly space-varying point spread functions of the optical systems were based on existing self-calibrating methods that already necessitated the identification and matching of particles between different camera perspectives. The detection of particles was extended to estimate shape parameters as well as their positions. The benefit of such an extended calibration was evaluated using both synthetic as well as real experimental data.

The results of the synthetic test cases show that the choice of how discrete voxels are projected onto pixels with respect to the shape of the voxel spreading function does play an important role in the quality of reconstruction as well as the quality of the computed flow fields based on the reconstructed volumes.
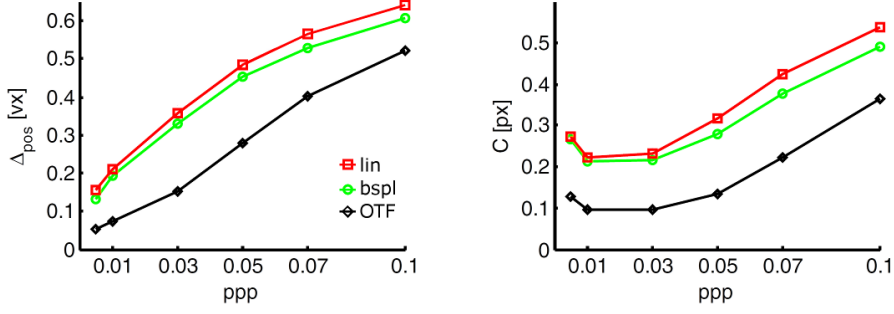
Figure 3.2.: Simulation results for the astigmatic case: error of detected particle position over different particle image densities (left), average velocity field error (right)

Figure 3.2 shows results for the case of astigmatism. Here, *lin* and *bspl* refer to fixed choices of the particle shapes based on linear interpolation and B-splines. The *OTF* variant refers to the new method using calibrated point spread functions. Its name just reflects the spectral representation (optical transfer function) of the point spread function. In this test, the estimated spatially dependent point spread function has reduced the error of detected particle positions by about 0.2 voxels and in the case of a limited depth of field (not shown here) the error is reduced by about 0.1 voxels. Comparing the computed velocity fields, the average error is only about a third in the astigmatic case and only about one half in the limited depth of field case.

For more details see publication [B].

3.2. Alternative tomographic reconstruction based on the Spectral Projected Gradient method

The SMART algorithm is the defacto standard reconstruction algorithm for solving the constrained linear equation system in TomoPIV applications. However, inconsistencies due to image noise or small calibration errors pose a problem for SMART. The algorithm is also known for its linear convergence property. This sparked the interest in looking for alternatives. The original motivation to use the SPG method for solving this reconstruction problem is due to its super linear convergence property. But in practice with typical particle images SMART performs better in terms of the computational cost than the plain SPG method.

The described S-SPG algorithm combines the spectral projected gradient method (SPG) with the multiplicative nature of the SMART algorithm. As such the S-SPG combines the robustness of the SPG method with respect to measurement noise and image preprocessing artefacts with the speed of SMART. Depending on the particle image
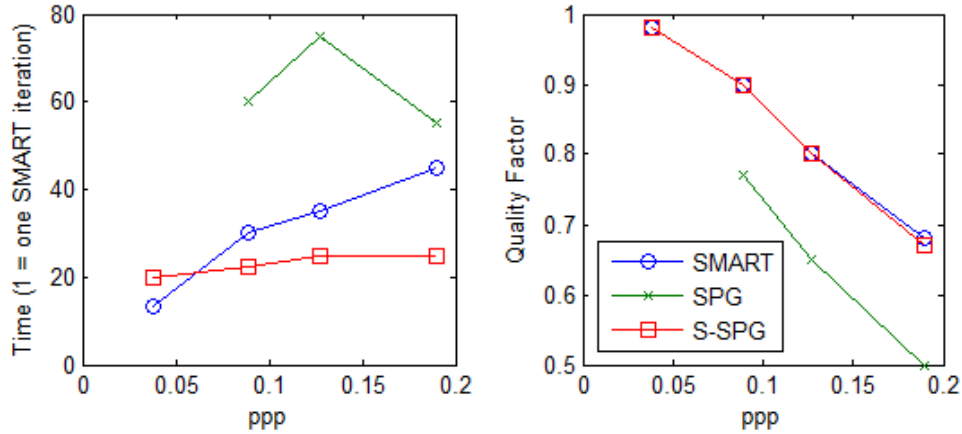
Figure 3.3.: Comparison between SMART, SPG, and S-SPG for varying particle image densities (ppp) of a simulated test case with four cameras. Left: particle image density versus the time needed to reach the optimization goal. Right: particle image density versus "quality factor".

density and the assumed point spread functions S-SPG might even outperform SMART in terms of computation time. This has been evaluated using synthetically generated image data with and without image noise. In many of the observed cases S-SPG even outperformed SMART in terms of reconstruction speed especially at higher particle image densities measured in particles per pixel (ppp). Figure 3.3 shows one of the results for a simulated setup with four cameras.

For more details see publication [A].

3.3. Shake-the-Box particle tracking and FlowFit

The Shake-the-Box (STB) particle tracking method has been demonstrated to be able to track a very high number of particles given time-resolved particle image sequences. At these particle image density only tomographic approaches have been used previously (TomoPIV). But compared to the tomographic reconstruction the STB method does not suffer from *ghost particles* and therefore allows for much higher reconstructed particle position accuracies along trajectories. Due to the advantages of the STB method in terms of speed and accuracy, the TomoPIV method does not seem like a good choice anymore for time-resolved cases, nowadays. In addition, having velocity and acceleration available at many individual particle positions allow for one- and multi-point statistics in subpixel resolution impossible for correlation-based velocity measurements.

Given the scattered particle locations and their velocities and accelerations as determined by the STB method, the FlowFit method performs a spatial interpolation of velocity and pressure reconstruction by solving a nonlinear least squares problem. The representation of the velocity and pressure field is based on 3D cubic B-splines enabling

a spatially continuous representation and analytical spatial derivatives of velocity and pressure. The associated cost function of this minimization problem includes deviations from measured and fitted quantities (velocity and acceleration) as well as a physical model (expected shape of the wave number spectrum) and physical laws (mass conservation and the Navier-Stokes equation) in form of additional costs for regularization. The simplest first generation FlowFit for interpolating velocities given scattered velocities only was shown to outperform tomographic and correlation-based methods within an international competition by a large margin. In publication [E] the simple reconstruction approach is compared with the full nonlinear optimization recovering pressure as well. It has been demonstrated via simulations that using this method the spatial resolution of the velocity field reconstruction can be improved using the acceleration data unless the real velocity field was already sufficiently sampled by the particles. For cases where the flow field is undersampled by particles, the nonlinear FlowFit is worth using to visualize small-scale flow structures.

Figure 3.4 shows the results of a synthetic test using different FlowFit variants: *div0* represents a plain spline fit only regularized by penalizing high frequency components slightly. Reconstruction *div1* builds on *div0* in that it includes penalizazion of the divergence of the velocity field which is very useful for incompressible flows. It can be seen that this recovers more details compared to *div0*. Reconstruction *div2* employs acceleration data of the particles as well. In a sense it does not only reconstruct a velocity field at one point in time but also its temporal derivative. The cost function that is minimized in *div2* builds on *div1* and makes sure that the temporal derivative of the velocity field is also mostly divergence-free. This kind of reconstruction is more complicated and results in a nonlinear least squares problem but it is able to improve the reconstruction quality and recover more structures without having to increase the particle image density. The nonlinear least squares problem is solved via the Limited-Memory version of the Broyden-Fletcher-Goldfarb-Shanno method (L-BFGS, see [Liu1989]).

For more details about the STB method see publications [C], [D]. Appendix I of this summary covers the FlowFit model and cost functions in more detail than in publication [E].
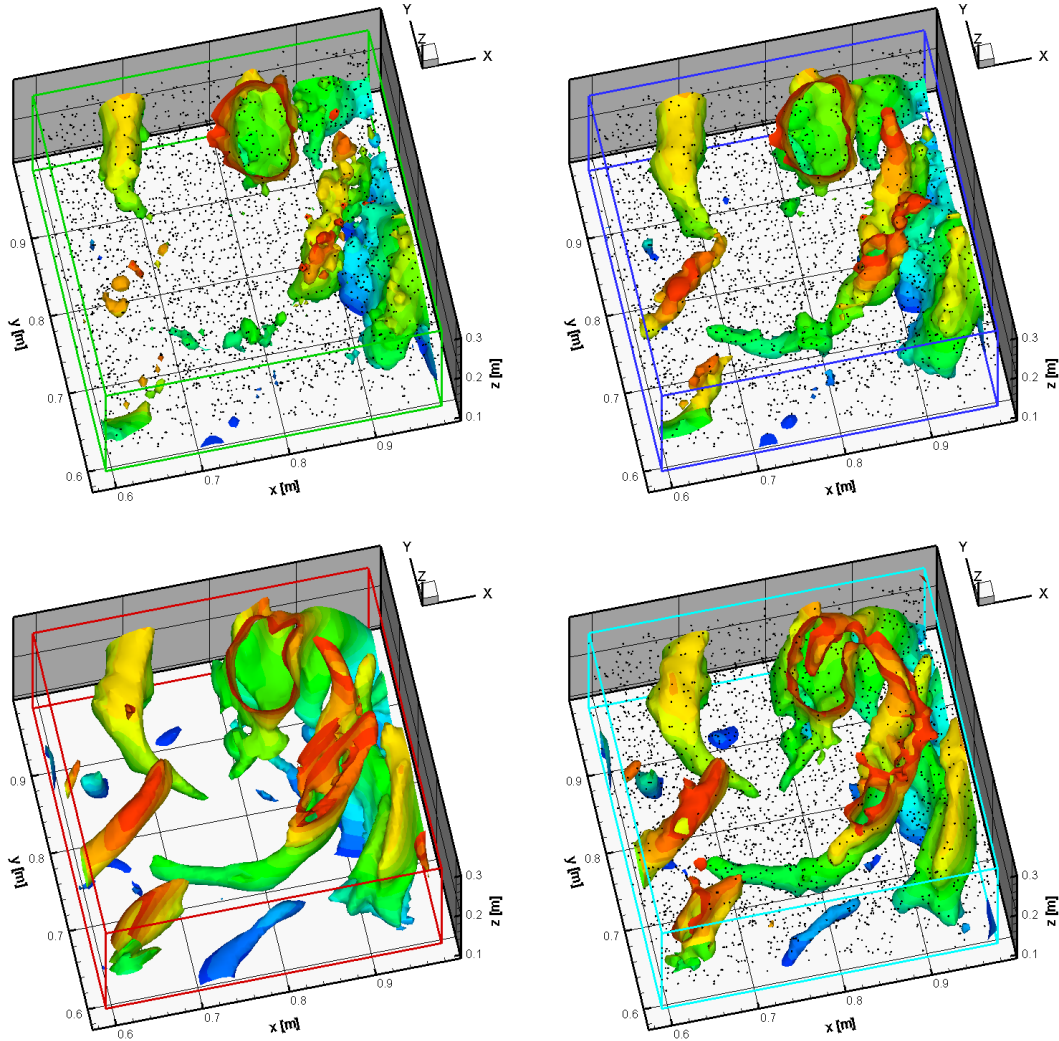
Figure 3.4.: Reconstructions of vorticity based on 3277 simulated random particles (black dots) of a subvolume of $64 \times 64 \times 32$ DNS grid points compared to the ground truth. From top left to bottom left clockwise: div0, div1, div2, and ground truth (color-coded z component)

3. Discussion of results

4. Impact of the contributions

The discussed methods in the field of optical 3D flow measurement have been successfully used on real experimental data and on published particle image data within international challenges to verify their advantages over previous state of the art methods. The following sections cover the results from two international challenges as well as results from the thermal plume experiment mentioned in the introduction chapter.

4.1. Participation in the 4th International PIV Challenge

To evaluate and compare different reconstruction techniques that take digital particle images as input and recover velocity fields, the curl of velocity and pressure two international challenges have been conducted recently.

The 4th International Particle Image Velocimetry challenge (PIV challenge) began in 2014. The aim of this challenge was to assess the recent developments in the global PIV community. This challenge was divided into multiple testcases. Test case D was a simulated time-resolved 3D PIV experiment based on a direct numerical simulation of an isotropic incompressible turbulent flow. Particle images for four simulated perspectives have been generated. Due to the synthetic nature of the generated images it was possible to compare all submissions with the known ground truth velocity field.

At the time of the submission deadline for the reconstruction results the FlowFit method for scattered interpolation done by the DLR was the first version which used quadratic instead of cubic splines and ignored acceleration data resulting in a simpler linear least squares problem to solve. Using the STB particle tracking method in combination with the simpler linear FlowFit, the DLR submission for test case D was shown to lead the competition in various respects which are summarized on the following paragraphs.

One evaluation criterium was the number of detected true and ghost particles. Ghost particles are a typical artefact of the ill-posed 3D reconstruction problem given only a few number of projections and a high concentration of particles. Only two teams from Göttingen (DLR and LaVision) were able to recover almost all true particles while keeping the number of ghost particles below 1%. The submissions based on independent reconstructions of time steps showed a much higher ghost particle count of 30% to 1000%. This can be seen in Figure 4.1.

Another evaluation criterium was how close the reconstructed velocity fields are to the ground truth velocity field. An isosurface of the Q criterion was computed based on finite differences and compared with the ground truth. By visual inspection the STB+FlowFit approach came closest to the ground truth.
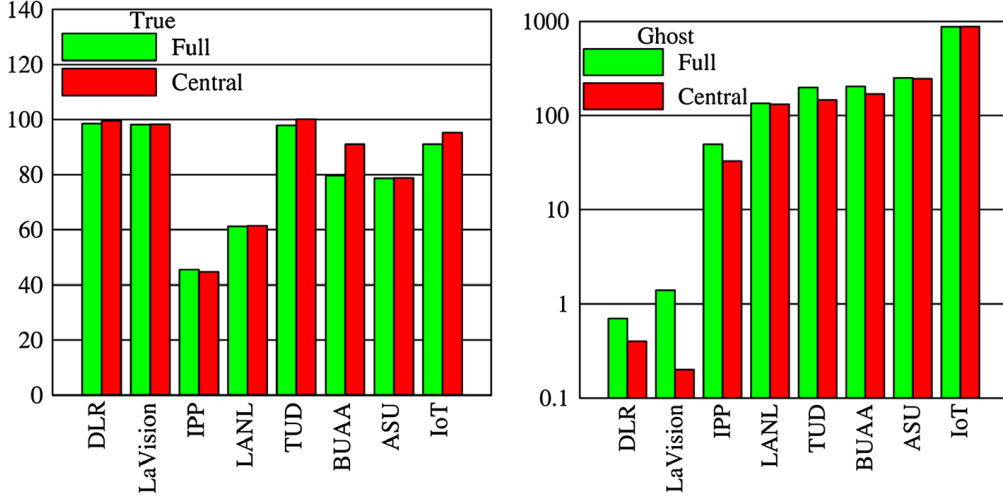
Figure 4.1.: Percentage of true (left) and ghost particles (right) for the full volume and for the central part obtained by cutting the borders (image from [Kaehler2016])

A third criterium was the temporal development of the three velocity components and the vorticity magnitude at a fixed point in the measurement volume and their comparison to the ground truth. Figure 4.2 shows the comparison graphs between different submissions for the vorticity magnitude. This comparison suggests that the STB+FlowFit method is superior to the traditional tomographic approach in terms of spatial and temporal resolution. It is the only submission that recovers the full vorticity magnitude.

A fourth criterium was the spatial spectrum of a reconstructed time step. Spectra were plotted and visually inspected. In addition, a weighted spectral power integral was computed and related to the ground truth value as a fraction (spectral energy fraction). This spectral energy was integrated over the low to mid frequencies. The STB+FlowFit approach scored best and was closely followed by team LaVision and team TU Delft. For higher frequencies the shown spectra varied much more but the DLR submission kept closest to the ground truth spectrum.

Another criterium was the total error of velocity. This was shown as scatter plots where the STB+FlowFit approach achieved the lowest error.

The results of this competition have been published in more detail in [Kaehler2016].

Since the competition the FlowFit method for spatial interpolation has been refined to use both velocity and acceleration data to fit velocity and pressure fields exploiting the Navier-Stokes equation as part of the cost function. The STB method has also been further tuned and improved.
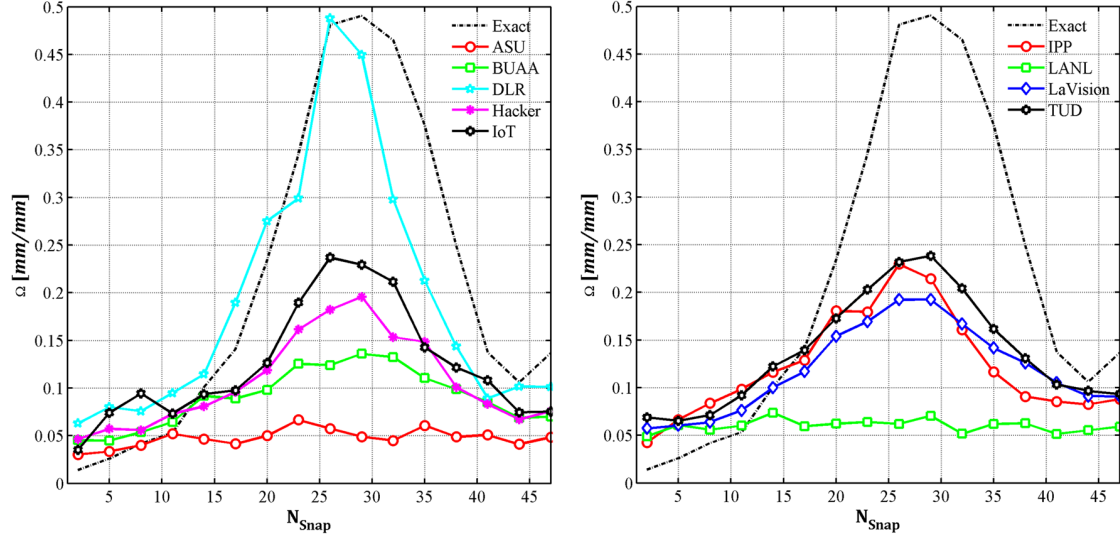
Figure 4.2.: Temporal profiles for the vorticity magnitude at a specified point in the volume (image from [Kaehler2016])

4.2. The NIOPLEX Pressure Challenge

As part of the European NIOPLEX project (Non-intrusive Optical Pressure and Loads Extraction for Aerodynamic Analysis) 10 approaches by 9 different research groups for velocity, density and pressure reconstruction based on particle images have been evaluated.

For the purpose of this challenge a zonal detatched eddy simulation (ZDES) was used. Based on these results the movement of particles have been simulated to create the synthetic measurement images with and without simulated image sensor noise.

With respect to the chosen metric the top five methods were all based on particle tracking (performed by the STB method at the DLR) and the bottom five where based on conventional PIV methods. Our method based on STB particle tracking, FlowFit interpolation of the acceleration field and an FFT-based spatial integration described in publication [I] landed on a close second place right behind a similar method called VIC+ developed by Schneiders et al from TU Delft, [Schneiders2016].

Two versions of the FlowFit have been used during this challenge. The first version was used to interpolate the acceleration field with penalization of the acceleration field's rotation resulting in a linear least squares problem. The result was then used in an FFT-based spatial integration to recover pressure which was described in [I]. The second version of the FlowFit directly reconstructed pressure and thus included the spatial integration as part of the equation system. Assumptions about how velocity relates to pressure were made in a certain region of the reconstruction volume and about how velocity relates to both pressure and density throughout the volume. We decided not to impose the zero-divergence penalizations of the velocity field as part of the FlowFit cost functions due to doubts regarding their impact in this compressible flow case. The VIC+

method which achieved the best score was very similar to our approach but assumed zero-divergence of the velocity field. For more details see publication [vanGent2017].

4.3. Example result from the thermal plume experiment

The experimental setup shown in the introduction and described in more detail in publication [J] referenced in section 2.2 is a remarkable example of what the newly developed tools allow us to do. Handling such a large measurement volume of about 0.6 m^3 at such a spatial resolution and accuracy is very hard to achieve using conventional tomographic techniques. In comparison, the new particle tracking approach has much lower time and memory requirements for 3D reconstruction. The applied techniques led to estimates on the locations, velocities and accelerations of about 275.000 tracked particles per time step. Tracking this high amount of particles is untypical of previous particle tracking approaches and allows recovering the flow structures in more detail. Combined with the regularized nonlinear spatial interpolation method developed and described in [E] small flow structures beyond the sampling limit can be recovered to some extent between known particle locations. The regularization takes physical laws and constraints into account such as conservation of mass and momentum.

Figure 4.3 shows one of the results previously published in [J]. A subset of the tracked particles along with isosurfaces of the Q-criterion is shown. Determining and visualizing the Q-criterion necessitates a spatially continuous reconstruction of the flow field. This visualization allows identifying moving vortex structures with a high spatial and temporal resolution.

Figure 4.3.: Isosurfaces of Q-criterion $(5/s^{-2})$ color-coded by streamwise velocity and selected particle tracks (11 time steps, slice of 50 mm in depth) with color-coded stream-wise acceleration (image taken from contribution [J])

4. Impact of the contributions

5. Conclusions and Outlook

The developments in optical flow measurement systems described and discussed in this cumulative thesis present a significant advancement in terms of accuracy and spatial resolution compared to what was previously possible. The results have been independently confirmed via international challenges in which multiple advanced research groups participated. Furthermore, recently the STB and FlowFit technique achieved two awards (STAB Forschungspreis 2016 and the DLR Wissenschaftspreis 2017) based on the evaluation and assessment from independent juries.

The techniques involve particle tracking of time-resolved measurements followed by a temporal smoothing of the particle tracks and spatial interpolations for multiple time steps independently. One possible approach to improve on this would fuse numerical simulations with measurements into a global Bayesian state estimation problem. Instead of reconstructing each time step independently one could try to take a flow field estimate from a previous time step as prior for reconstructing the next time step. Unfortunately, this is made difficult due to the high number of state dimensions (3D velocity and pressure field). But techniques like these are already applied for weather forecasting. Further research is necessary in order to evaluate and test similar techniques for the purpose of flow field measurements.

Another area of further research is uncertainty estimation. After smoothing the particle tracks, estimates for the uncertainty of particle locations, velocities and accelerations are already available based on a physical model for the particles' motion and model parameters derived from particle data statistics. The same approach could be taken for uncertainty estimation after the spatial interpolation. The first step would be to estimate a wavenumber spectrum of the velocity field based on two-point correlations that can be computed using pairs of particles at varying distances. With this information, the choice of the regularization parameter pen_{hf} in the cost function of FlowFit could be optimized and uncertainty of the reconstructed velocity field could be estimated.

5. Conclusions and Outlook

Bibliography

[A] Sebastian Gesemann, Daniel Schanz, Andreas Schröder, Stefania Petra, Christoph Schnörr. (2010). Recasting TomoPIV reconstruction as a constrained and L1-regularized nonlinear least squares problem. *15th Int. Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon*

[B] Daniel Schanz, Sebastian Gesemann, Andreas Schröder, Bernhard Wieneke, Matteo Novara. (2013). Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. *Meas. Sci. Technol. 24:024009*

[C] Daniel Schanz, Andreas Schröder, Sebastian Gesemann, Dirk Michaelis, Bernhard Wieneke. (2013). Shake The Box: A highly efficient and accurate Tomographic Particle Tracking Velocimetry (TOMO-PTV) method using prediction of particle positions. *10th Int. Symposium on Particle Image Velocimetry, Delft*

[D] Daniel Schanz, Sebastian Gesemann, Andreas Schröder. (2016). Shake-The-Box: Lagrangian particle tracking at high particle image densities. *Exp. Fluids, 57:70*

[E] Sebastian Gesemann, Florian Huhn, Daniel Schanz, Andreas Schröder. (2016) From Noisy Particle Tracks to Velocity, Acceleration and Pressure Fields using B-splines and penalties. *18th Int. Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon*

[F] Hauke Ehlers, Andreas Schröder, Reinhard Geisler, Sebastian Gesemann. (2012). Combined Time-Resolved PIV and structure deformation measurements for aeroelastic investigations. *STAB Symposium*

[G] Andreas Schröder, Daniel Schanz, Reinhard Geisler, Sebastian Gesemann, Christian Willert. (2015). Near-wall turbulence characterization using 4D-PTV Shake-The-Box. *11th Int. Symposium on Particle Image Velocimetry, PIV15*

[H] Daniel Schanz, Florian Huhn, Sebastian Gesemann, Uwe Dirksheide, Remco van de Meerendonk, Peter Manovski, Andreas Schröder. (2016). Towards high-resolution 3D flow field measurements at cubic meter scales. *18th Int. Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon*

[I] Florian Huhn, Daniel Schanz, Sebastian Gesemann, Andreas Schröder. (2016). FFT integration of instantaneous 3D pressure gradient fields measured by Lagrangian particle tracking in turbulent flows. *Exp. Fluids, 57:151*

Bibliography

[J] F. Huhn, D. Schanz, S. Gesemann, U. Dirksheide, R. van der Meerendonk, A. Schröder. (2017). Large-scale volumetric flow measurement in a pure thermal plume by dense tracking of helium-filled soap bubbles. *Exp. Fluids, 58:116*

[Franka88] Leopoldo P. Franca, Thomas J. R. Hughes. (1988). Two classes of mixed finite element methods. *Comput. Methods Appl. Mech. Engrg., 69(1):89–129*

[Liu1989] Dong C. Liu, Jorge Nocedal. (1989). On the limited memory BFGS method for large scale optimization. *Mathematical Programming, 45(1):503–528*

[Birgin2000] E. G. Birgin, J. M. Martinez, M. Raydan. (2000). Nonmonotone spectral projected gradient methods on convex sets. *SIAM Journal on Optimization, 10(4):1196-1211*

[Elsinga2006] G. E. Elsinga, F. Scarano, B. Wieneke, B.W. van Oudheusden. (2006). Tomographic Particle Image Velocimetry. *Exp. Fluids, 41:933–947*

[Wieneke2008] Bernhard Wieneke. (2008). Volume self-calibration for 3D particle image velocimetry. *Exp. Fluids, 45(4):549-556*

[Petra2008] Stefania Petra, Constantin Popa, Christoph Schnörr. (2008). Enhancing Sparsity by Constraining Strategies: Constrained SIRT versus Spectral Projected Gradient Methods. *7th Workshop on Modelling of Environmental and Life Sciences Problems*

[Atkinson2009] Callum Atkinson, Julio Soria. (2009). An efficient simultaneous reconstruction technique for tomographic particle image Velocimetry. *Exp. Fluids, 47:563-578*

[Wieneke2013] Bernhard Wieneke. (2013). Iterative reconstruction of volumetric particle distribution. *Meas. Sci. Technol., 24(2):024008*

[Kaehler2016] Christian J. Kähler, Tommaso Astarita, Pavlov P. Vlachos, Jun Sakakibara, Rainer Hain, Stefano Discetti, Roderick La Foy, Christian Cierpka. (2016) Main results of the 4th International PIV Challenge. *Exp. Fluids, 57(6):97ff*

[Schneiders2016] Jan F. G. Schneiders, Fulvio Scarano. (2016). Dense velocity reconstruction from tomographic PTV with material derivatives. *Exp. Fluids, 57(9):139ff*

[vanGent2017] P. L. van Gent, D. Michaelis, B.W. van Oudheusden, P. É. Weiss, R. de Kat, A. Laskari, Y. J. Jeon, L. David, D. Schanz, F. Huhn, S. Gesemann, M. Novara, C. McPhaden, N. J. Neeteson, D. E. Rival, J. F. G. Schneiders, F. F. J. Schrijer. (2017) Comparative assessment of pressure field reconstructions from particle image velocimetry measurements and Lagrangian particle tracking. *Exp. Fluids, 58(4):33ff*

Appendices

I. FlowFit model and cost functions

The problem that arises after Lagrangian Particle Tracking (LPT), which the FlowFit method is intended to solve, is recovering a spatially highly-resolved flow field (velocities, accelerations and pressure up to an unknown offset) based on scattered data that is given by the particles' trajectories. With FlowFit time snapshots are reconstructed independently. For one time snapshot the locations, velocities and accelerations of the particles are assumed as given including absolute error estimates of velocity (σ_u) and acceleration (σ_a) within some measurement volume. This volume is discretized using a uniform Cartesian grid where each grid point is associated with four B-spline weighting coefficients, three for the three velocity components and one for the pressure field. The velocity field as well as the pressure is therefore represented within the cuboid measurement volume as a weighted sum of cubic 3D B-splines. Figure I.1 shows an example of such a grid in two dimensions and the volume it defines.
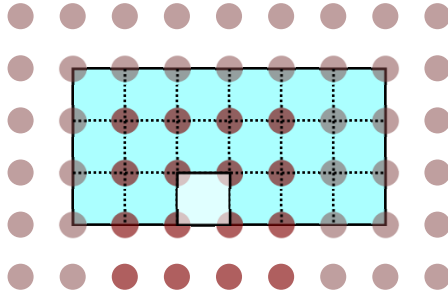
Figure I.1.: Example of a 2D grid with the volume it defines. The third cell of the lower row is highlighted with the surrounding 4×4 grid points which affect this cell with their weighting coefficients due to the overlapping cubic B-spline functions.

Given this model, velocity $\vec{u}$ and pressure $\boldsymbol{p}$ including their spatial derivatives up to the second order can be expressed as a linear combination of the respective weighting components of up to $4 \times 4 \times 4$ grid points in 3D. This is due to the fact that the cubic B-spline function has a support of $4 \cdot h$ in each dimension where h is the distance between neighbouring grid points. Evaluation of these functions at the exact grid point coordinates is simpler because it involves only the respective weighting coefficients of the surrounding $3 \times 3 \times 3$ grid points. The contribution of the weighting coefficients of grid points that are farther away vanishes. Figure I.2 shows the cardinal cubic B-spline β_3 for one dimension and a grid point spacing of $h = 1$. The base functions are all shifted versions of each other covering the whole volume.
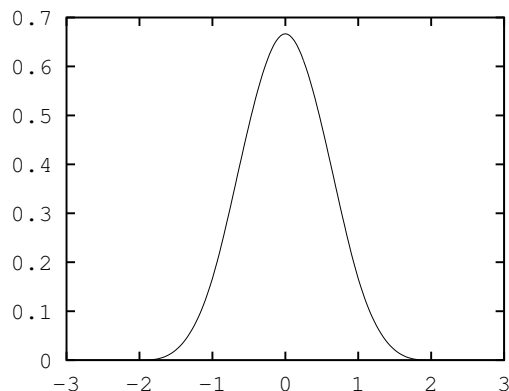
Figure I.2.: Cubic B-spline β_3 for grid point spacing $h = 1$

The function $\beta_3 : \mathbb{R} \mapsto \mathbb{R}$ can be defined as follows:

$$
\beta_3(t) = \begin{cases} 0 & \text{for } 2 \leq \|t\| \\ \frac{1}{6} \left(2 - \|t\|\right)^3 & \text{for } 1 \leq \|t\| < 2 \\ \frac{4}{6} - \|t\|^2 + \frac{1}{2}\|t\|^3 & \text{for } \|t\| < 1 \end{cases} \tag{I.1}
$$

In higher dimensions such as 3D the dimensions can be treated separately successively reducing it dimension by dimension to one in a way similar to trilinear interpolation. This is equivalent to the use of the 3D base function B given as follows:

$$
B : \mathbb{R}^3 \longmapsto \mathbb{R} \;,\; B\left((x_1, x_2, x_3)^T\right) = \prod_{i=1}^{3} \beta_3(x_i) \tag{I.2}
$$

The unknown function parameters—the B-spline weighting coefficients—are recovered by minimizing a cost function that references the measured particle locations, velocities and accelerations. In addition to the measurement data, one or more regularizations are used to both render the reconstruction problem uniquely solvable as well as improve the reconstructions' quality with respect to the spatial resolution.

Suppose $c_d^{i,j,k} \in \mathbb{R}$ refers to the d-th component of a 4D B-spline vector that is associated with a grid point indexed by i, j, k along the three spatial axes. Components $d = 1, 2, 3$ are used to represent a velocity field and component $d = 4$ is used to represent a pressure field. Each grid point has an associated location $L^{i,j,k} \in \mathbb{R}^3$ in meters with a distance between neighboring grid point of h in meters. Then, the velocity field $\vec{u}_c$ and the pressure field p_v defined as a functions on the 3D coordinates $x \in \mathbb{R}^3$ within the reconstruction volume and depending on the weighting coefficients c can be written as the following superposition of weighted cubic 3D B-splines:

$$\vec{u}_c(x) = \sum_{i,j,k} B\left(\frac{x - L^{i,j,k}}{h}\right) \cdot \left(c_1^{i,j,k}, c_2^{i,j,k}, c_3^{i,j,k}\right)^T \tag{I.3}$$

$$p_c(x) = \sum_{i,j,k} B\left(\frac{x - L^{i,j,k}}{h}\right) \cdot c_4^{i,j,k} \tag{I.4}$$

In order to use the measured particles' velocities and accelerations as part of the cost function we need to be able to express velocities and accelerations in terms of the unknown B-spline weighting coefficients. Incorporating measured velocities is simple. Velocity at a particular point in the measurement space is directly available via equations (I.3) and (I.4). With the help of the momentum equation of the Navier-Stokes equations, acceleration can be expressed in terms of pressure and velocity. FlowFit has been developed for incompressible flow cases with constant density. Under this assumption and no other forces being present, the momentum equation simplifies to the following equation. For simplicity, p is referred to as pressure but is strictly speaking pressure divided by the density. Here, ν refers to the kinematic viscosity:

$$\vec{a} = \frac{\partial \vec{u}}{\partial t} + (\vec{u} \cdot \nabla)\,\vec{u} = -\nabla p + \nu \Delta \vec{u} \tag{I.5}$$

Incompressibility can be used as part of the cost function to improve the spatial resolution of the reconstruction. The obvious condition to exploit is (I.6):

$$\nabla \cdot \vec{u} = 0 \tag{I.6}$$

The momentum equation (I.5) allows expressing the temporal derivative of velocity in terms of pressure and velocity:

$$\frac{\partial \vec{u}}{\partial t} = -\nabla p + \nu \Delta \vec{u} - (\vec{u} \cdot \nabla)\,\vec{u} \tag{I.7}$$

Since equation (I.6) is satisfied not only for a single point in time but globally in time, the divergence of the temporal derivative of velocity has to be zero as well everywhere in the measurement volume. With the help of equation (I.7), this constraint can be expressed as follows:

$$\Delta p + \nabla \cdot (\vec{u} \cdot \nabla)\,\vec{u} = 0 \tag{I.8}$$

Equation (I.8) presents one possible approach for recovering pressure based on a velocity field. However, in FlowFit it is used as part of the cost function with the intention of recovering both p and $\vec{u}$ simultaneously in a physically consistent way and at a possibly higher spatial resolution.

Equations (I.6) and (I.8) are used to penalize the divergence of velocity and its temporal derivative at every grid point in the reconstruction volume. Since this would not necessarily enforce zero divergence in the space between the grid points, penalization of the gradient of divergence of velocity is applied as well. This has been observed to

improve the reconstruction and to attenuate the divergence uniformly instead of just at the grid points. A similar technique is known as *grad-div stabilization* which was originally proposed in [Franka88] to improve the conservation of mass in finite element methods. However, the penalization of the gradient is not done for the constraint (I.8) since it would require third order spatial derivatives at the grid points (which are not defined for cubic B-splines) and would result in a much more complicated optimization problem to solve.

FlowFit is usually applied with a grid point to particle ratio of around 10:1—meaning 10 grid points for every particle. Such a high ratio is used to retain all the information present in the particle data. This avoids unintentional spatial smoothing and keeps enough degrees of freedom to include other regularizations with the intention of recovering some small scale structures beyond the sampling limit. The problem of computing the weights based on the measured data without further regularizations, however, is an underdetermined linear equation systen with no unique solution. A popular approach to deal with this is an appropriate Tikhonov regularization. In the case of FlowFit this is done by a light penalization of small scale structures in the cost function via a spatial high pass filter so that ambiguities are resolved in favour of more dominant large scale structures. This is not uncommon for splines and will minimize their curvature (*natural spline*). Within FlowFit, this regularization is deemed appropriate because that large scale structures are expected to be more dominant in real velocity and pressure fields than small scale structures.

The light penalization of small scale structures is implemented with a spatial highpass filter that works directly on the B-spline weighting coefficients c. This highpass filter is implemented as difference between the original and a spatially smoothed version of the weighting coefficients. The component of the spatially smoothed 3D weight array is determined using a separable 3D convolution with the convolution kernel of $[0.25, 0.5, 0.25]$ in each dimension. For the smoothed weighting coefficients the used notation in the remainder of this section is $\widehat{c}$. The following equation details the relationship between the original and spatially smoothed coefficients:

$$\widehat{c}_d^{\,i,j,k} = \sum_{\delta_i=-1}^{1} \sum_{\delta_j=-1}^{1} \sum_{\delta_k=-1}^{1} \frac{\left(2 - |\delta_i|\right)\left(2 - |\delta_j|\right)\left(2 - |\delta_k|\right)}{4^3} \, c_d^{\,i+\delta_i, j+\delta_j, k+\delta_k} \tag{I.9}$$

Applying the above equation requires special treatment for the grid points that are outside of the reconstruction volume because they lack some neighbors. In this case, linear extrapolation of c is used. The combination of linear extrapolation and this kind of spatial smoothing is equivalent to not smoothing at all in that particular direction. So, the smoothing is effectively reduced to two, one or zero dimensions for grid points on one of the six sides, on one of the twelve edges or on one of eight corners of the cuboid, respectively.

The full and nonlinear cost function F for the *div2* FlowFit variant is the sum of different components. For the purpose of defining F and its components some notation is introduced first: l_{p_i}, u_{p_i} and a_{p_i} refer to the measured location, velocity and acceleration of the i-th of N_p observed particles with $1 \leq i \leq N_p$, respectively. Let Ω be the set

of *all* 3D grid point indices (i, j, k), $\partial\Omega$ be the subset of grid point indices outside of the reconstruction volume and $\Gamma = \Omega \setminus \partial\Omega$ be the subset of grid point indices that are contained in the reconstruction volume. Estimates for the error of measured velocity and acceleration are given as σ_u and σ_a, respectively. The divergence of velocity is referred to as d_0 and its temporal derivative is referred to as d_1:

$$d_0 := \nabla \cdot \vec{u} \tag{I.10}$$

$$d_1 := \Delta p + \nabla \cdot (\vec{u} \cdot \nabla)\, \vec{u} \tag{I.11}$$

Finally, there are the parameters pen_{hf} and pen_{div} which control the contribution of small scale structures and divergences to the cost function. With these notations and definitions the cost function F in terms of the unknown B-spline weighting coefficients c can be defined as follows:

$$F(c) = V(c) + A(c) + H_u(c) + H_p^{\partial\Omega}(c) + D_0(c) + D_1(c) \tag{I.12}$$

$$V(c) = \sum_{i=1}^{N_p} \left(u_{p_i} - u_c(l_{p_i})\right)^2 \tag{I.13}$$

$$A(c) = \sum_{i=1}^{N_p} \left(\frac{\sigma_a}{\sigma_u} \left(a_{p_i} - a_c(l_{p_i})\right)\right)^2 \tag{I.14}$$

$$H_u(c) = \sum_{(i,j,k)\in\Omega} \sum_{d=1}^{3} \left(pen_{hf} \left(c_d^{i,j,k} - \hat{c}_d^{i,j,k}\right)\right)^2 \tag{I.15}$$

$$H_p^{\partial\Omega}(c) = \sum_{(i,j,k)\in\partial\Omega} \left(\frac{\sigma_a}{\sigma_u} \frac{1}{h}\, pen_{hf} \left(c_4^{i,j,k} - \hat{c}_4^{i,j,k}\right)\right)^2 \tag{I.16}$$

$$D_0(c) = \sum_{(i,j,k)\in\Gamma} \left(h\, pen_{div}\, d_0(L^{i,j,k})\right)^2 + \\ \sum_{(i,j,k)\in\Gamma} \left(\frac{1}{4}h^2\, pen_{div}\, \left\|\nabla d_0(L^{i,j,k})\right\|_2\right)^2 \tag{I.17}$$

$$D_1(c) = \sum_{(i,j,k)\in\Gamma} \left(\frac{\sigma_a}{\sigma_u} \frac{h}{10}\, pen_{div}\, d_1\left(L^{i,j,k}\right)\right) \tag{I.18}$$

Here, V (I.13) and A (I.14) refer to the sum of squared deviations between the measurements and the fitted values at particle locations for velocity and acceleration, respectively. H_u (I.15) and H_p (I.16) are small scale structure penalizations for velocity and pressure. The small scale structure penalization for pressure is only applied at the border. A penalization over the whole volume is not necessary since pressure is already

determined by $\vec{u}$ via the cost component D_1 and the small scale penalization for $\vec{u}$ is already applied via H_u (I.15) on the whole volume. Both, D_0 and D_1 are derived from the incompressibility constraint which penalize nonzero divergence of the velocity field (D_0) or its temporal derivative (D_1). D_1 is based on equation (I.8) which links the divergence of the temporal derivative of velocity to the pressure and velocity variables via the Navier-Stokes momentum equation.

Choosing a factor of $\frac{1}{10}$ in the definition of D_1 (I.18) was done to be conservative about this penalization since it involves high order spatial derivatives that may not be well-represented given this model. The factor $\frac{1}{4}$ in D_0 (I.17) for the gradient of the divergence leads to a rather uniform penalization of divergence instead of just at the grid points locations. The usual choice of the weighting parameter pen_{div} is 1 and for pen_{hf} is a lower value such as 0.05 to avoid any unwanted spatial smoothing.

Due to the nonlinearity of D_1 in $\vec{u}$ and therefore its nonlinearity in c this least squares problem is a nonlinear one which requires the appropriate solver. In FlowFit, the Limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS, see [Liu1989]) method is used with a starting point of all zero. In the interest of fast convergence and due to the fact that large scale structures are expected to dominate the the solution FlowFit can be instructed to reconstruct a lower resolution grid first (half the resolution in each dimension) which is then interpolated to the full grid as the starting point for a second round of minimizing the cost function.

The reconstruction variant *div1* ignores the D_1 component of the cost function which turns this optimization problem into a linear least squares problem. As such, other solvers such as the conjugated gradient method (CG) for least squares are viable. The elimination of D_1 also removes the strong the connection between the velocity and pressure weighting coefficients so they become independent with respect to the cost function in case of zero viscosity. For this reason the regularization $H_p^{\partial\Omega}$ is replaced by H_p^{Ω} which includes all the inner grid points. The variant *div0* eliminates both D_0 and D_1 from the cost function and thus does not enforce zero divergence at all. The comparison between these variants for reconstructing a velocity field can be seen in Figure 3.4. It is clear that with the more elaborate cost function (*div2*) FlowFit is able to recover more details.

II. Publications

The five major and five other contributions mentioned in chapter 2 are following in order on the following pages.

II. Publications

Recasting Tomo-PIV reconstruction as constrained and L1-regularized nonlinear least squares problem

Sebastian Gesemann[1], Daniel Schanz[1], Andreas Schröder[1], Stefania Petra[2], Christoph Schnörr[2]

1: Dept. of Experimental Fluids, Institute of Aerodynamics and Flow Technology, German Aerospace Center (DLR)
{sebastian.gesemann, daniel.schanz, andreas.schroeder}@dlr.de
2: Dept. of Mathematics and Computer Science, University of Heidelberg, Germany
{petra,schnoerr}@math.uni-heidelberg.de

Abstract. Tomo-PIV is a volumetric flow measurement system. It is able to obtain instantaneous 3D-3C velocity fields using multiple cameras that observe the same volume of two subsequently illuminated particle distributions from different angles. Part of the measurement technique is the 3D reconstruction of the particle volume from 2D images. This work presents a promising alternative approach to the popular reconstruction algorithms MART and SMART. Simulations and experimental data of a turbulent free jet at Ma = 0.7 are used in order to assess and demonstrate the performance of the newly developed reconstruction algorithm based on constrained least squares strategies and L1-regularization.

1. Introduction

In Tomo-PIV the reconstruction problem is typically modeled as linear equation system where the volume is discretized (Elsinga et al, 2005). This system relates each unknown "voxel" intensity to a small subset of pixels under the assumption that light intensities accumulate along lines of sight and, typically, that the voxels' intensities are not negative. The relationship between voxels and pixels is completely determined by the system's matrix which encodes geometric as well as spreading function information about the measurement configuration and optical systems involved. In an ideal case this matrix is correctly determined by means of calibration. This linear model is appropriate for small particles since virtually no occlusion occurs. The size of the imaged particles is mainly due to diffraction and refraction effects of the optical system.

Unfortunately, these equation systems are severely underdetermined due to the low number of projections. Typically, only three or four cameras are used to record the illuminated particle volume from different perspectives. However, the nonnegativity constraint for voxel intensities seems to be one of the main reasons why reconstructions work reasonably well at least for low to moderate particle densities. The sheer size of these equation systems poses a problem in terms of computation time and working memory requirements. Iterative solvers are often applied with a low fixed number of iterations instead of iterating until a certain convergence criterion is satisfied. The performance of reconstruction algorithms during the first iterations is therefore an interesting study.

Recently, Petra et al (2008) drew attention to a generic optimization algorithm with interesting properties (superlinear convergence for a certain problem set) for the reconstruction problem called nonmonotone Spectral Projected Gradient Method – "SPG" (Birgin et al, 2000). The reconstruction problem is formulated as a linear least squares problem which accounts for errors in the pixel intensity measurements and constrains voxel intensities to be nonnegative. The least squares approach and superlinear convergence are attractive features of the algorithm as they could increase robustness and reduce reconstruction time.

In this work we investigated and compared the performance of SMART (Mishra et al, 1999) as one of the state-of-the-art reconstruction algorithms for Tomo-PIV (Atkinson, Soria, 2009) and the nonmonotone SPG algorithm through various simulations. We were able to improve the performance of reconstruction via SPG with a nonlinear substitution that warps the search space.

The algorithm that combines this substitution and the SPG solver is referred to as S-SPG throughout the rest of this paper. In addition, we tested the use of the L1-regularization as part of the cost function. The L1-regularization plays an important role in recovering sparse signals from undersampled data (Compressed Sensing paradigm, see Petra et al 2009). The combination of S-SPG with L1-regularization is referred to as L1S-SPG.

Finally, SMART and S-SPG are used to reconstruct 40 x 40 x 6 mm^3 volumes from a real free jet experiment with a magnification factor of approximately 30 voxels per millimeter in each dimension. The images show blurred as well as very sharp particles (aliasing) due to the trade-off between large depth of field and high light intensity. This typically calls for a slight lowpass filter as part of the image preprocessing but it gave us the chance to compare the performance between SMART and the SPG-based solvers under these difficult conditions.

2. Reconstruction algorithms

Given an M times N weighting matrix A with coefficients $a_{i,j}$ and a measurement vector b encoding the recorded pixel intensities we seek to compute the discrete volumetric intensity distribution x – typically a regular grid of voxels in the measurement volume – so that Ax approximates b, x is reasonably sparse and nonnegative.

The algorithms MART and SMART inherently enforce this nonnegativity constraint due to their multiplicative nature and the restriction of nonnegative weights in A. For reference, these are the SMART equations we used:

$$r_i^{(k)} = \frac{b_i}{\sum\limits_{j=1}^{N} a_{i,j} x_j^{(k-1)}}$$

$$x_j^{(k)} = x_j^{(k-1)} \prod_{i,a_{i,j}>0} \left(r_i^{(k)} \right)^{a_{i,j}\mu}$$

Here, the vector $r^{(k)}$ refers to the ratios between measured pixel intensities and current projections in the k-th iteration. The update of the volume is done by simultaneously applying a multiplicative corrective term which, for a choice of μ with $\mu\sum_i a_{i,j} = 1$, is a weighted geometric average of the previously computed intensity ratios. Smaller values for μ slow down convergence and higher values tend to produce oscillating sequences of intensity distributions which is why we chose it to satisfy the aforementioned equation.

The goal of the Spectral Projected Gradient method is to find a minimum of an objective function. The obvious choice for an objective function in this case is the sum of squared pixel intensity errors:

$$f(x) = \frac{1}{2}\|Ax - b\|_2^2$$

$$\nabla f(x) = A^T (Ax - b)$$

To enforce the nonnegativity constraint we can use a projection P as part of the algorithm that replaces negative coefficients with zero. The algorithm basically computes gradients of the objective function and derives a descent direction d according to the current gradient, a scale factor of $1/\sigma^{(k)}$ where $\sigma^{(k)} I_N$ is supposed to approximate the objective function's Hessian matrix at the current location, and the projection P:

$$d^{(k)} = P\left(x^{(k-1)} - \frac{1}{\sigma^{(k-1)}} \nabla f\left(x^{(k-1)} \right) \right) - x^{(k-1)}$$

$$x^{(k)} = x^{(k-1)} + \alpha^{(k)} d^{(k)}$$

$$\sigma^{(k)} = \max\left\{ \sigma_{\min}, 0.1 \cdot \sigma^{(k-1)}, \frac{<d^{(k)}, \nabla f\left(x^{(k)}\right) - \nabla f\left(x^{(k-1)}\right)>}{<d^{(k)}, \alpha d^{(k)}>} \right\}$$

For $\alpha^{(k)}$ with $0 < \alpha^{(k)} \le 1$ the next $x^{(k)}$ is guaranteed to have nonnegative coefficients due to the projection *P*. In most iterations $\alpha^{(k)} = 1$ is used – even tolerating a nonmonotone development of the objective function's cost values. However, global convergence can be guaranteed if $\alpha^{(k)}$ is occasionally set to lower values so that the new cost stays below the maximum cost value seen during the last *T* iterations. In our case we used *T = 4*. For the next value of σ the difference between the new gradient and old gradient is taken into account as well as the last step αd. Here, $<\square>$ refers to the scalar product. With a lower bound of $0 < \sigma_{\min}$ the step length is restricted and the algorithm can even navigate nonconvex regions of the objective function. In addition, we prevented new values for σ to go below one tenth of the previous value to reduce the chance of an unnecessarily large descent step *d* during the next iteration. For more details on the nonmonotone Spectral Projected Gradient Method see Birgin et al, 2000.

3. Evaluation criteria

Due to the low number of projections the quality of a reconstruction is not completely determined by the residual alone. Solutions to the reconstruction problem may differ greatly in quality while producing the same images when applied to the weighting matrix *A*. For simulations with known volumetric intensity distribution and correct matrix *A* – correct in the sense that synthetically generated images correspond to *Ax* except for optionally added noise – we not only can check the development of the residual of intermediate solutions but also compare the current iteration's reconstruction with the original volume. To measure the similarity between two intensity distributions we simply used the normalized cross correlation. This cross correlation is sometimes known as "quality factor". The use of the normalized cross correlation seems reasonable as it is expected to be indicative of the correlation peak heights during correlation-based displacement estimation.

Since all algorithms discussed here can be written in terms of matrix vector products – even SMART via *M* additional logarithms and *N* exponentiations per iteration – and these products are dominating the computation time, a fair comparison of computation time can be done by counting the number of matrix vector products each algorithm uses internally. A SMART iteration always uses two of these "costly operations" per iteration – one for the projection and one for the update step. The SPG algorithm requires slightly more of these operations at average due to the occasional step size reduction and recomputation of the residual with another step size.

4. Warping the search space as preconditioning for SPG

After initial tests we noticed the superiority of SMART over SPG when directly applied on the constrained linear least squares problem in two aspects: SMART managed to reduce the residual more quickly during the first iterations in many cases and the intermediate solutions $x^{(k)}$ produced by SMART had higher quality factors compared to the SPG reconstructions at *similar* residual levels. This suggests that in these instances there was no unique solution to find and that the natures of the respective algorithms affected the search paths differently. These results are included in section 5.

An attempt to combine the good features of both algorithms led to a nonlinear substitution of the search space. Instead of applying the SPG directly on the linear least squares problem, we substitute x for z^2 (componentwise, $\forall j, 1 \le j \le N : x_j =: z_j^2$) and let the SPG find the optimal z. This substitution affects the gradients in a way that higher voxel intensities change more quickly than lower voxel intensities just like it is the case with multiplicative reconstruction approaches. However, the objective function it not convex quadratic in z anymore. Turning a constrained linear least squares problem into a constrained nonlinear least squares problem seems counterproductive at first, but the SPG algorithm is still applicable and shows a better convergence behavior as can be seen in the next section. The substitution does not affect the complexity of the implementation in any significant way.

5. Simulations

We simulated discrete particle volumes of 256 x 256 x 128 voxels with varying particle densities and image noise levels. Four virtual cameras observed the volume under parallel projection with viewing directions $(\pm 1; \ \pm 1; \ 2)^T$ and a 1:1 voxel-pixel size ratio. The particle positions were determined by a pseudo random number generator using uniform distributions in the volume including nonzero sub-voxel positions. For each generated particle up to 4^3 voxels have been altered by sampling and adding a three-dimensional cubic B-Spline to the discrete volume. The projection of each voxel with nonzero intensity is done similarly. The three-dimensional voxel coordinate is mapped to a pixel coordinate and a two-dimensional cubic B-Spline is sampled, thus, affecting up to 4^2 pixels per camera. This results in a projection matrix A where each of the 2^{23} columns contains up to 4^3 nonzero entries which always sum up exactly to the number of cameras due to the partition of unity property of B-Splines. We chose the SMART parameter $\mu = 1/4$ accordingly.

As an initial guess for all algorithms we computed $x^{(0)}$ to be

$$x_j^{(0)} := s\left(\min_c \left\{b^T A_{|,j}^{(c)}\right\} + \varepsilon\right),$$

where the scale factor s minimizes $\left\| Ax^{(0)} - b \right\|_2^2$,

$$\text{and} \quad A = \begin{bmatrix} A^{(1)} \\ \vdots \\ A^{(4)} \end{bmatrix}.$$

$A^{(1)}$, $A^{(2)}$,…, $A^{(4)}$ are the respective submatrices for all four cameras and ε is a small positive number to prevent initializing a voxel intensity with zero.

5.1 Simulation of various particle densities

Four cases have been simulated: 3 000, 7 000, 10 000 and 15 000 random particles have been rendered into a volume and projected to four image planes. Approximately 79 000 pixels are potentially affected by particles in each camera. This corresponds to particle per pixel (ppp) values of approximately 0.038, 0.089, 0.127, and 0.190, see Fig. 1 for a combination of all four cases in one image of the first virtual camera. The tests have been repeated with different initialization parameters for the random number generator to compute representative residual and quality factor curves against the number of costly operations.
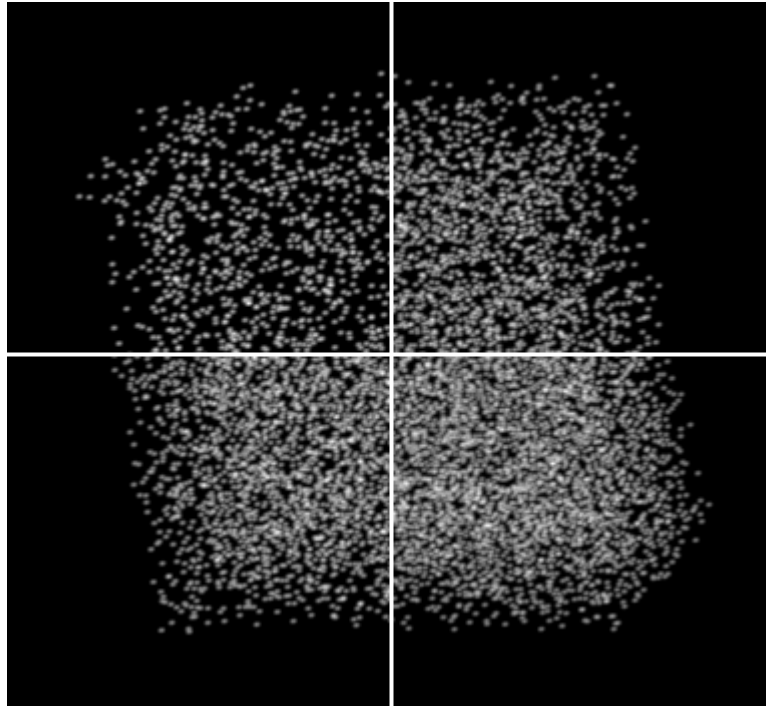
Fig. 1 Simulated first camera showing ppp = 0.038 (top left), ppp = 0.089 (top right),
ppp = 0.127 (bottom left), ppp = 0.190 (bottom right)

Fig. 2 through Fig. 5 show the average results of the simulations. In Fig. 2 and Fig. 3 we can clearly see that SMART tends to outperform both of the alternative reconstruction approaches during the first couple of iterations. After about 10 costly operations in the first case with few particles, S-SPG and SMART share the same slope of the residual curve. SPG applied on the original least squares problem (red curve) has a much weaker residual curve slope. The difference to S-SPG is apparent. The residual level of 10^2 is reached by SMART after 10 costly operations, by S-SPG after 16 costly operations while SPG reaches this level after 65 costly operations. But also the quality factors of the reconstructed intermediate volumes at this residual level differ greatly. For a residual of 10^2 SPG reconstructs an intensity distribution with a quality factor below 0.7 while the quality factors of the reconstructed volumes for both of the other algorithms are above 0.9. SMART and S-SPG inherently picked solutions that are closer to the original than the reconstructions of SPG at the same residual level.
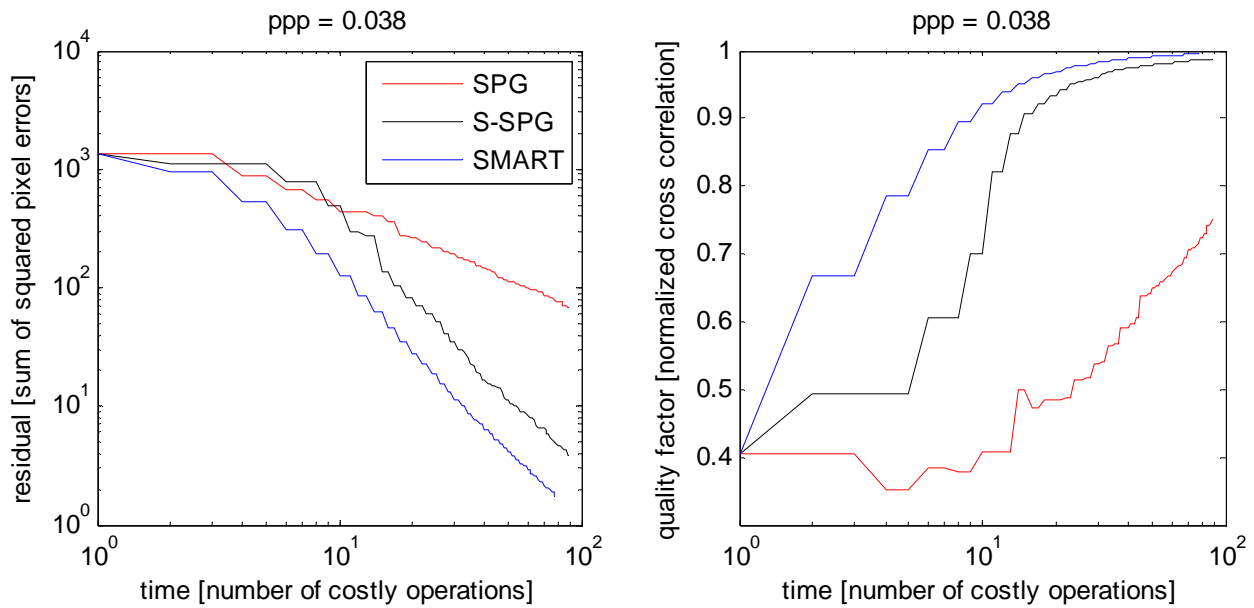
**Fig. 2 Results for the first case with 3 000 particles and 40 iterations,
residual (left) and quality factor (right)**
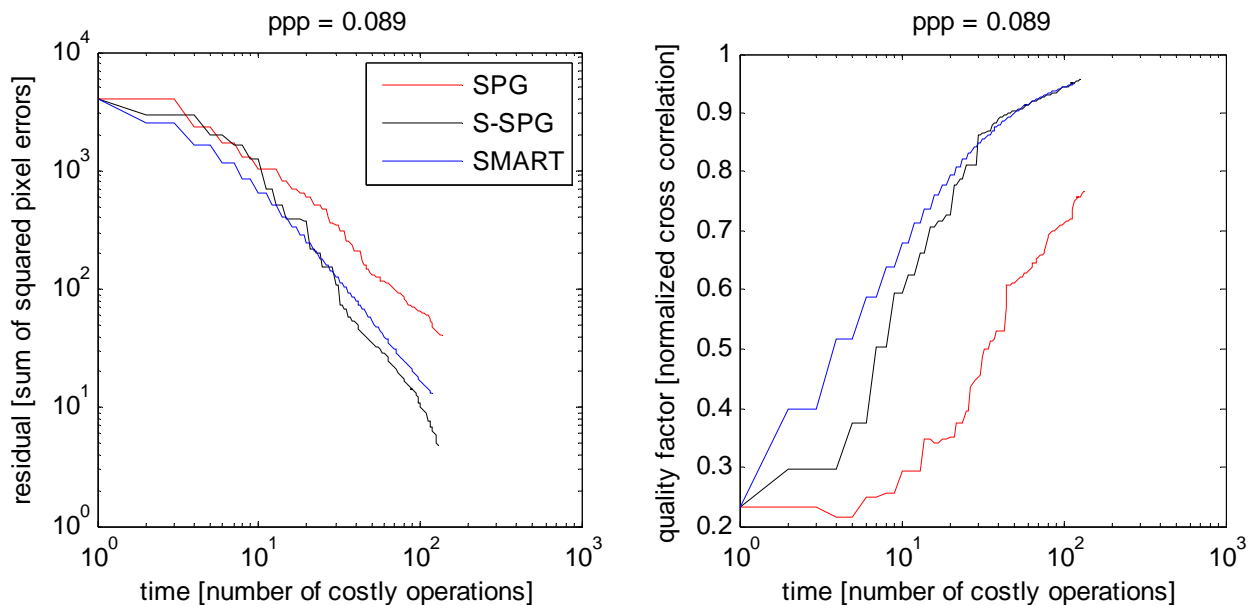
**Fig. 3 Results for the second case with 7 000 particles and 60 iterations,
residual (left) and quality factor (right)**

In the second case with 7000 particles (ppp = 0.106) S-SPG starts slow but can catch up with SMART after about 20 costly operations. After 30 costly operations (15 SMART iterations), S-SPG manages to outperform SMART by approximately 50% in terms of residual minimization. The quality factors for volumes derived by SMART and S-SPG are comparable after that time. SPG is still behind in terms of the residual as well as the quality factor.
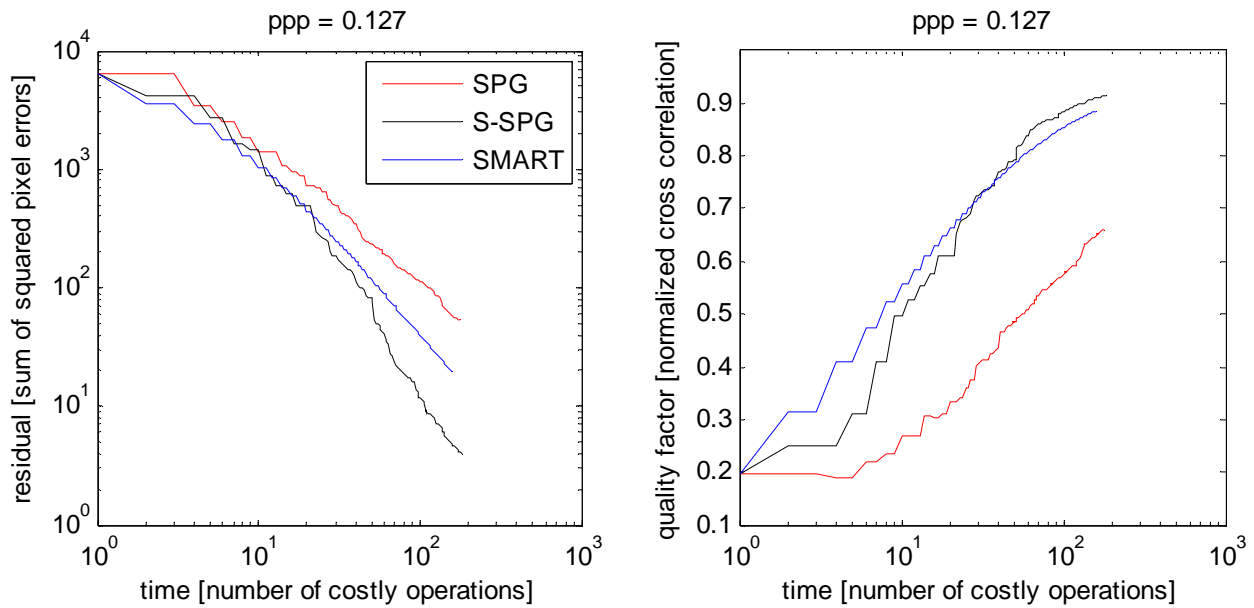
**Fig. 4 Results for the third case with 10 000 particles and 80 iterations,
residual (left) and quality factor (right)**
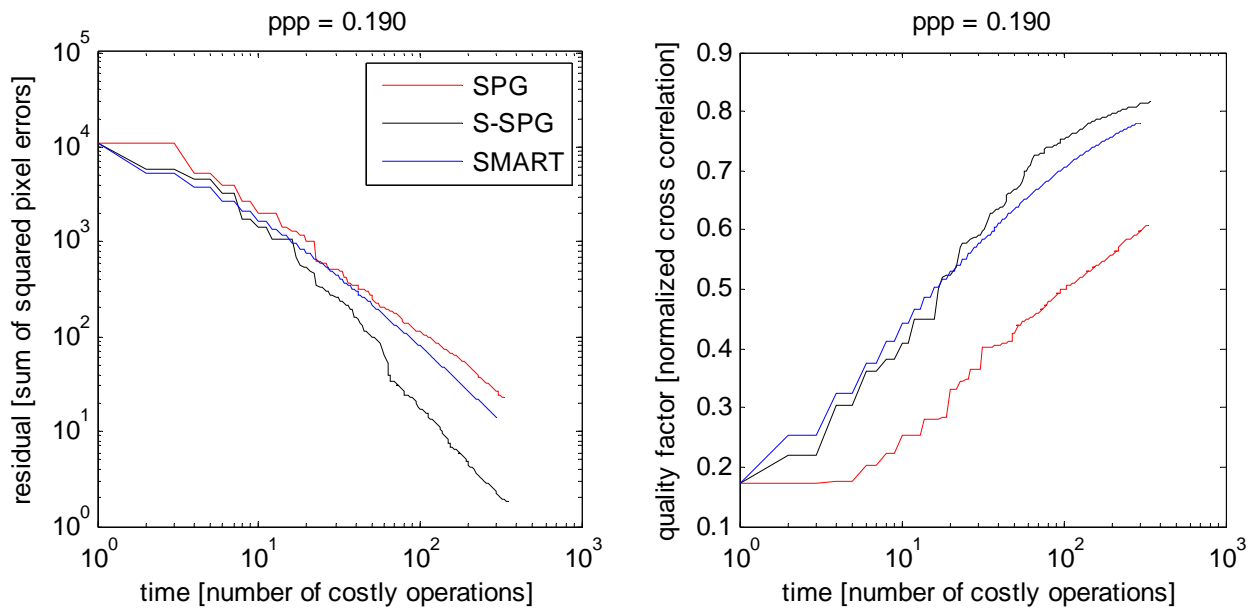
**Fig. 5 Results for the fourth case with 15 000 particles and 150 iterations,
residual (left) and qualify factor (right)**

With increased particle density S-SPG extends its lead in terms of residual minimization and quality factors of reconstructed volumes. Even though SPG (applied on the original least squares problem) catches up in terms of residual, it computes volumes of lower quality. At such high particle densities there probably is no unique solution for the minimization problem. Nevertheless SMART and S-SPG manage to reconstruct volumes that are remarkably close to the original.

5.2 Simulation of noisy images

An additional L1-regularization of the objective function f did not seem to help much in the simulated cases with noise-free images. But we noticed a slight quality improvement of the reconstructed volumes when it was applied on problem instances with noisy images. L1-regularization introduces a new parameter λ which controls the effect of the sum of voxel intensities on the cost value of the objective function. In the following formulas 1_N refers to a column vector containing N entries of all ones.

$$f(x) = \frac{1}{2}\|Ax - b\|_2^2 + \lambda 1_N^T x$$
$$\nabla f(x) = A^T(Ax - b) + \lambda 1_N$$

A simulation with 15 000 particles in a 256 x 256 x 128 volume has been repeated six times with different initializations of the pseudo random number generator. Fig. 6 shows one of the images for the first camera. Each projected particle contributes exactly an intensity of one to the sum of pixel intensities for one camera. The generated noise was Gaussian with a standard deviation of 0.05. Pixels with negative intensity have been set to zero prior to reconstruction.
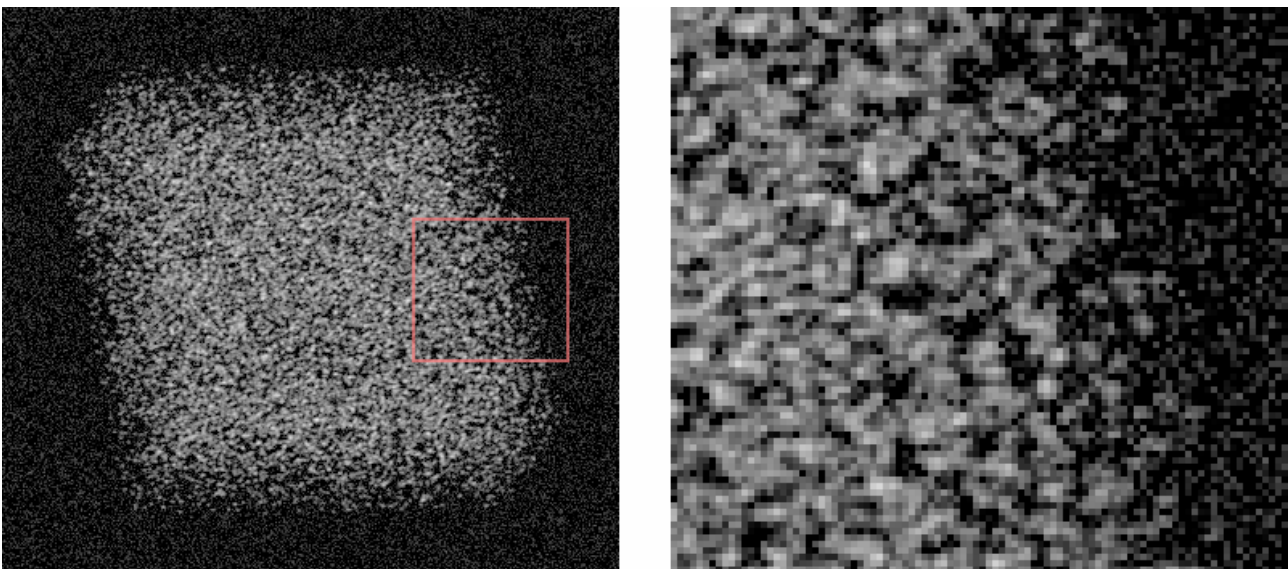
Fig. 6 Simulation with 15 000 particles (ppp = 0.190),

images with Gaussian noise (standard deviation of 0.05)

Fig. 7 shows again the development of the residuals and quality factors over time for three algorithms. This time, SPG has been replaced with L1S-SPG. Due to the penalization of high voxel intensity sums via L1-regularization L1S-SPG allows the residual to be higher if the L1-norm can be reduced to a certain extent. This can be seen on the left side of the figure. S-SPG reconstructed volumes after 50 costly operations have a lower residual than corresponding volumes reconstructed by L1S-SPG. But we can see on the right side that this L1-regularization actually improves the quality factor of the reconstructed volumes slightly. The regularization parameter λ was chosen to be 0.1.
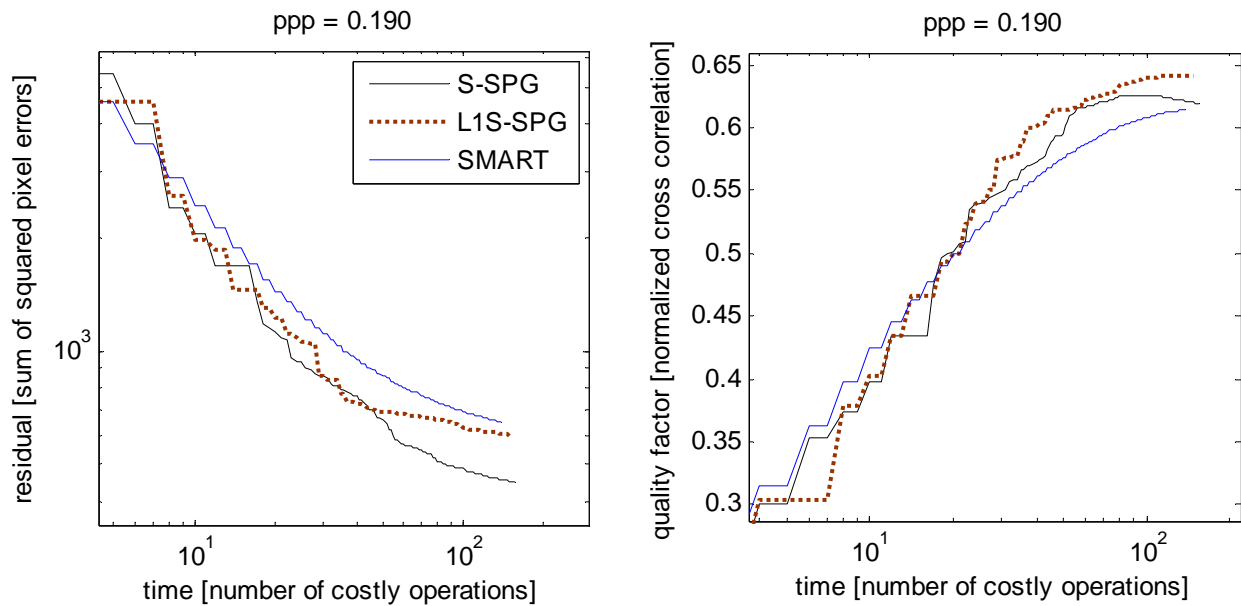
**Fig. 7 Results for a simulation with noisy images and 70 iterations,
residual (left) and quality factor (right)**

6. Application on real experimental data

To verify the suitability of the new algorithm in practice we prepared a Tomo-PIV experiment measuring a free turbulent jet at Ma = 0.7. Four high resolution *PCO 4000* cameras equipped with f = 100 mm *Zeiss* lenses arranged in a pyramidal viewing set-up observing an investigation area of 120 x 50 x 5 mm^3 that is located in a region between x/d = 18 to 22 downstream a round nozzle of 15 mm diameter. As tracer particles 1µm diameter DEHS droplets were used which were illuminated by two combined *BigSky CFR200* double cavity lasers with 400 mJ pulse energy each in an extended and collimated laser beam, which was back-reflected in itself using a mirror. The illuminated volume was confined using two 50 x 5 mm² knife-edges, see Fig. 8. Several typical imaging problems especially when using Tomo-PIV in air flows have been modeled experimentally. Due to the change of the local light scattering angle a relatively strong intensity variation of the particle images can be observed along the field of views. Several aperture numbers between $f_\#$ = 2.8 and 8 have been tested in order to change the dynamic range of particle image intensities and densities. At the same time this $f_\#$ range produces a wide variation of particle image diameters and introduces background noise due to limited depth-of-focus. Reconstructions with self-calibration will be compared in order to assess the robustness of the individual algorithms.

Fig. 9 shows one of the images recorded by the first camera where parts not affected by the discretized volume have been masked black. We can clearly see the shape of the volume and some slight perspective distortion. The volume that is selected here is 40 x 40 x 6 mm^3 which is about a millimeter thicker than the light sheet. In the close-up on the right side we can see some very sharp particles (aliasing) as well as blurred particle images.
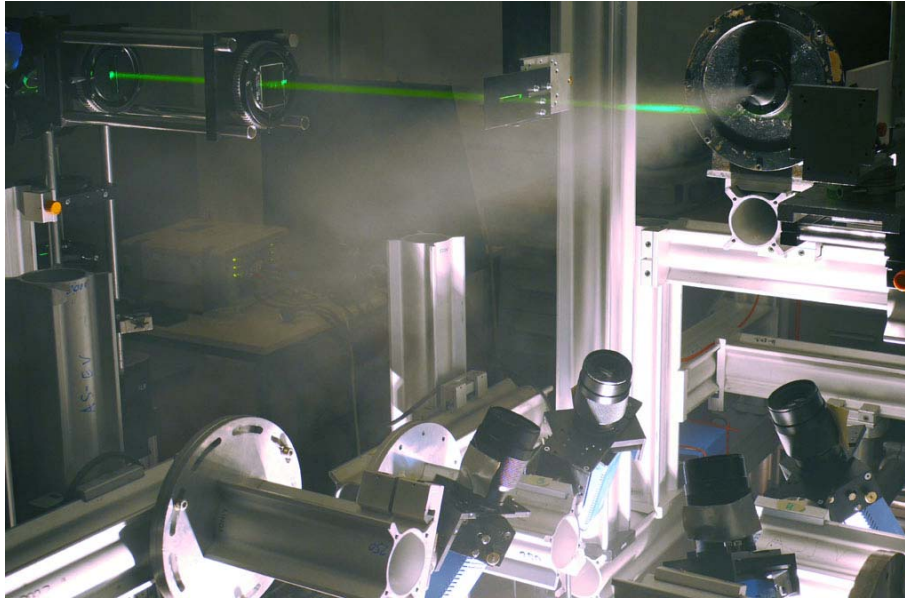
Fig. 8 Experimental setup: light sheet (top), nozzle (top right),
four cameras in a pyramidal setup (bottom right)
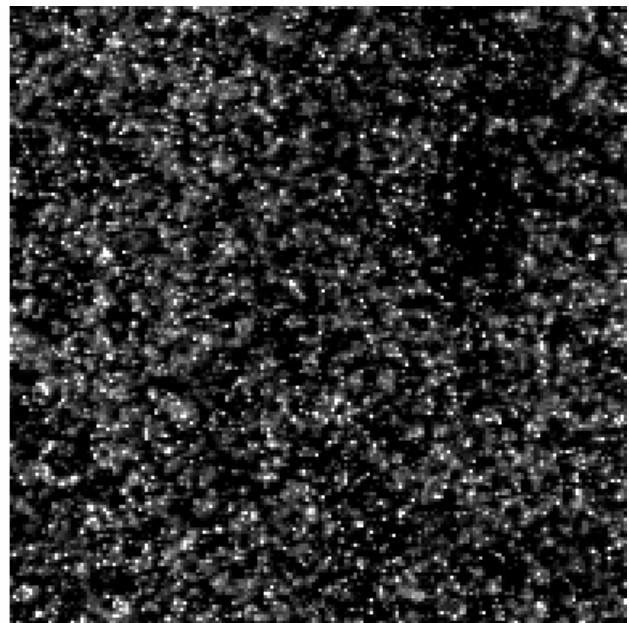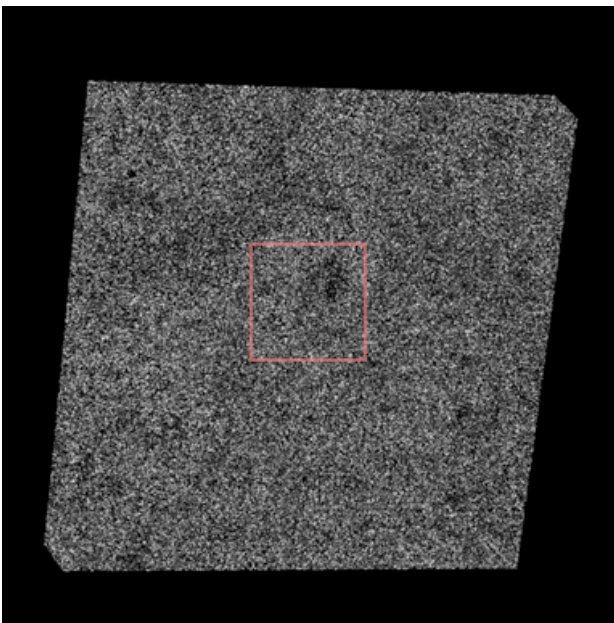
Fig. 9 Masked image of the first camera, overview (left) and close-up (right)

With no further image processing such as smoothing SMART tends to remove such sharp particles since the black surroundings of these isolated bright pixels lead to multiplicative corrective terms close to zero and this in turn sets many voxel intensities near the respective lines of sight irrevocably to zero. The least squares based approach on the other hand computes an intensity distribution which, when projected back to the image plane according to the projection matrix A, approximates the original image within the limits of A. This can be seen in Fig. 10. While the projection of the reconstructed volume using SMART is slightly less bright and missing some particles, the projection of the reconstructed volume using S-SPG preserves the overall brightness and particles.
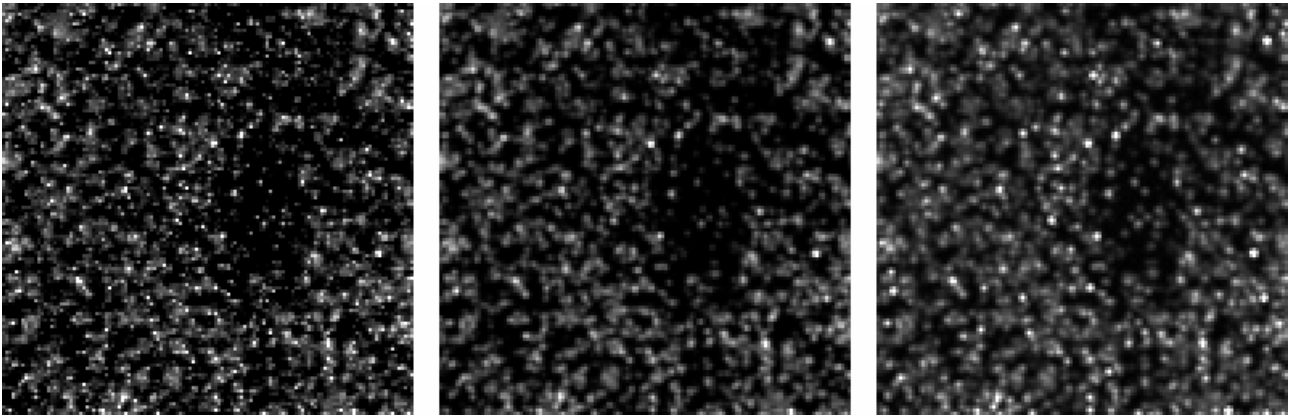
**Fig. 10 Original recording (left), projection of SMART reconstruction (center),
projection of S-SPG reconstruction (right), in both cases after 20 iterations**
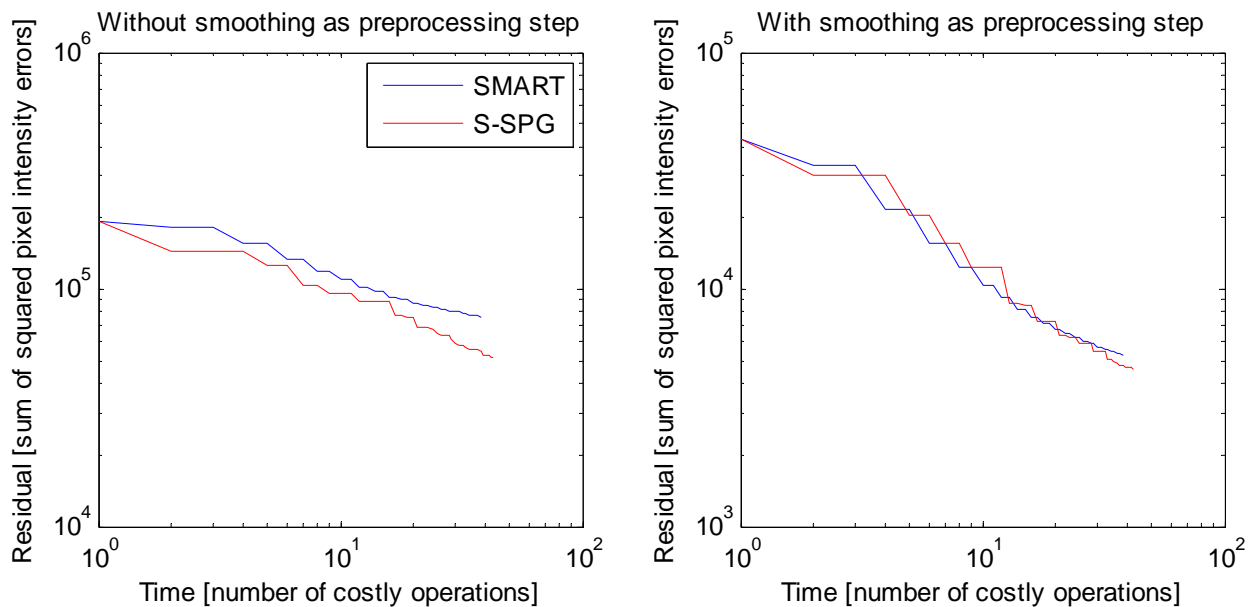
**Fig. 11 Residual curves (average over two frames) for
both preprocessing cases and algorithms**

This effect can also explain the development of the residuals over time, see Fig. 11. In the first case (no smoothing as preprocessing step, left) the S-SPG approach is clearly superior in terms of residual minimization. This can be attributed to the "aggressive" behavior of SMART when it comes to computing and applying the correction terms which seems to be counterproductive in cases where the system of equations is inconsistent.

In the preprocessing case this is less pronounced possibly due to more nonzero pixel intensities surrounding particles. The performances of SMART and S-SPG are similar. There is hardly a noticeable difference between the projections of both reconstructed volumes in this case.

7. Conclusions

Tomographic reconstruction is a vital part of the Tomo-PIV flow measurement technique. High quality reconstructions are just as desirable as fast reconstructions. In this work we presented alternatives to MART and SMART with promising results. Our L1S-SPG implementation requires about twice as much working memory compared to SMART-based solvers but gives better results with respect to our evaluation criteria after a few iterations in cases of high particle density or noisy images. For low particle densities the SMART-based algorithms perform better during early iterations which suggests that a combination of both, a SMART variant for initial iterations followed by L1S-SPG iterations, can be beneficial.

S-SPG and L1S-SPG are also applicable when optical transfer functions have been calibrated to optimize the weights of the equation system (see Schanz et al, 2010). The use of altered weights affects the condition number of the linear equation system and in cases of blurry imaging of particles the condition number is expected to increase. This typically slows down convergence of iterative solvers. In preliminary tests with these kinds of equation systems the SPG-based solvers managed to reduce the residual more quickly than SMART which can be attributed to SPG exploiting information about the objective function's curvature.

We would also like to point out that the generality of the optimization algorithms allows the inclusion of more sophisticated regularizations as well as new degrees of freedom of a more realistic imaging model. For example, additional parameters could be introduced to compensate for unaccounted inter-camera intensity differences or sub-pixel shifts due to vibrations of the optical system.

References

Elsinga G, Scarano F, Wieneke B, van Oudheusen B W (2006) Tomographic Particle Image Velocimetry. Exp Fluids 41:933–947

Mishra D, Muralidhar K, Munshi P (1999) A Robust MART Algorithm For Tomographic Applications. Numerical Heat Transfer, Part B: Fundamentals, 35:4, 485-506

Atkinson C, Soria J (2009) An efficient simultaneous reconstruction technique for tomographic particle image Velocimetry. Exp. Fluids, 47:563-578

Birgin E G, Martinez J M, Raydan M (2000) Nonmonotone spectral projected gradient methods on convex sets. SIAM J. Optimiz. 10:1196-1211

Petra S, Popa C, Schnörr C (2008) Enhancing Sparsity by Constraining Strategies: Constrained SIRT versus Spectral Projected Gradient Methods. In: 7[th] Workshop on Modelling of Environmental and Life Sciences Problems (WMM 08), Ed Acad Romane, Bucharest, Romania

Petra S, Schröder A, Schnörr C. (2009) 3D Tomography from Few Projections in Experimental Fluid Dynamics. In: Imaging Measurement Techniques in Flow Analysis Notes on Numerical Fluid Mechanics and Multidisciplinary Design, Vol. 106. Springer-Verlag Berlin Heidelberg, pages 63-72, ISBN 978-3-642-01105-4

Petra S, Schnörr C (2009) TomoPIV meets compressed sensing. University of Heidelberg, Germany, to appear in Pure Mathematics and Applications: Special issue on Discrete Tomography

Schanz D, Gesemann S, Schröder A, Wieneke B, Michaelis D (2010) Tomographic reconstruction with non-uniform optical transfer functions. Proceedings of 15[th] Int Symp on Applications of Laser Techniques to Fluid Mechanics, paper #1709 (submitted)

Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction

Daniel Schanz[1], **Sebastian Gesemann**[1], **Andreas Schröder**[1],
Bernhard Wieneke[2] **and Matteo Novara**[3]

[1] Department of Experimental Methods, German Aerospace Center (DLR), Institute of Aerodynamics and Flow Technology, Germany
[2] LaVision GmbH, Göttingen, Germany
[3] Aerospace Engineering Department, Delft University of Technology, 2629, The Netherlands

E-mail: daniel.schanz@dlr.de

Abstract

A new approach to the weighting function, which describes particle imaging in tomographic reconstruction, is introduced. Instead of assuming a spatially homogeneous mapping function of voxels to the images, a variable optical transfer function (OTF) is applied. By this method, the negative effects of optical distortions on the reconstruction can be reduced considerably. The effects of these improvements in reconstruction quality on the methods of tomographic particle imaging velocimetry, as well as 3D particle tracking are investigated. A method to calibrate the OTF to experimental circumstances is proposed as an additional step to the volume self-calibration. It is shown that this kind of calibration is able to capture the predominant particle imaging both for simulated as well as experimental data. The most common distortions of particle imaging are blurring due to a small depth of field and astigmatism due to imaging optics. The effects of both of these distortions on reconstruction and correlation quality are investigated via simulated data. In both cases, a strong influence on relevant parameters can be seen. Reconstructions using a spatially varying OTF, calibrated to the imaging conditions, show a significant improvement in reconstruction quality and the accuracy of the particle peak position, as well as in the accuracy of the gained displacement vector field when using two time steps. Evaluation of experimental data by PTV methods shows a reduction in ghost particle intensity and improvements in peak position accuracy. A computationally efficient method of applying the OTF to tomographic reconstruction is introduced.

Keywords: tomographic PIV, tomographic reconstruction, MART, PTV, particle imaging, weighting function

(Some figures may appear in colour only in the online journal)

1. Introduction

The reconstruction of the particle distribution within a volume of a flow is a key step in the tomographic particle imaging velocimetry (TOMO-PIV) technique (Elsinga *et al* 2006a). Different approaches have been in discussion over the last years, mainly of algebraic nature (ART, SART, MART)

(Herman and Lent 1976). Atkinson and Soria reported the successful use of a simultaneous MART (SMART) algorithm on tomographic reconstruction data, as well as a reduction of reconstruction time by using a multiplicative line-of-sight (MLOS) initialization of the volume (Atkinson and Soria 2009).
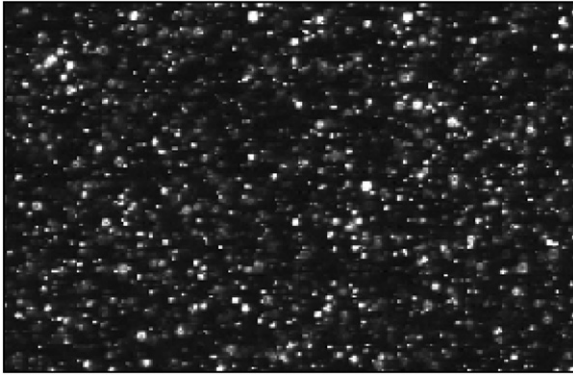
Figure 1. Unprocessed camera image of a TOMO-PIV experiment in air at $f_\# = 5.6$. Focused as well as blurred particle imaging can be observed.

Figure 2. Corner areas of a single image from a TOMO-PIV experiment in a water tunnel. Left: lower-left corner, right: upper-right corner. Astigmatic distortions in different directions can be seen.

Reconstruction algorithms have to assume a mapping function of voxels onto the pixels of the cameras. This mapping function can be, for instance, a linear interpolation between neighboring voxels on the line-of-sight of the current pixel (Bosbach *et al* 2009) or a Gaussian function of a certain width (Elsinga *et al* 2006a). In all recent works, this mapping function was chosen to be uniform for the whole volume and for all cameras. Therefore, if one or more cameras see different shapes of the same particle due to the circumstances of the experimental setup, reconstruction will suffer: as the shape of the voxel region, representing a particle, results from an overlay (multiplication) of the different particle images, the reconstructed particle might be deformed, enlarged, reduced in intensity or displaced. The probability of the occurrence of ghost particles will also be increased due to the scatter of particle image energy into the volume (Worth and Nickels 2010).

With respect to homogeneous particle imaging, the applications of TOMO-PIV in water and in air pose different problems. When using TOMO-PIV in air, the seeding particles have to be very small (typically DEHS droplets with a diameter of around 1 μm) in order to be able to follow the flow. Because of this, the intensity of the scattered light is several orders of magnitude lower than in the water case. This necessitates the use of low f-numbers (typically $f_\# = 4$–11), as well as moving the cameras close to the interrogation volume. Both measures lead to a reduction of depth of field and therefore limit the usable thickness of the volume.

Figure 1 shows an unprocessed image of a TOMO-PIV experiment on a free turbulent jet in air, where the cameras were set to $f_\# = 5.6$. Even though the volume was a thick sheet of only 5 mm in depth, it can clearly be seen that only particles in a well-defined area of focus are imaged sharply. Particles out of this plane appear either as Gaussian blobs of different sizes or as donut-shaped images.

In water, depth of field is a smaller problem, as particles with a relatively large diameter (10–100 μm) can be used without introducing too much particle slip. The amount of light scattered by such particles allows the usage of high f-numbers ($f_\# = 16$–32). Still, when examining deep volumes, depth of field can become a problem. TOMO-PIV in water has to deal with another obstacle in uniform particle imaging:
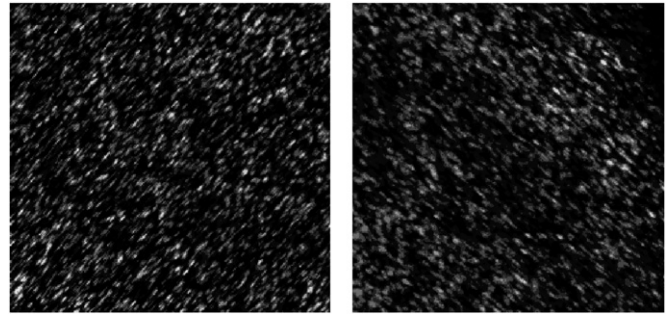
the necessity of imaging the volume through an interface (air/glass/water) with a change of refractive index.

The cameras have to be placed outside the water reservoir used, typically observing the interrogation volume through glass plates. Often it is not possible to position all cameras in such a way that they are level with the glass plate. Instead, some cameras have to look at an angle to achieve a certain viewing position. This situation will cause a certain amount of astigmatism (caused by different focal points for sagittal and transverse rays), especially in the corners of the image. As a consequence, the particle image is elongated in one direction (the direction of the sagittal plane), leading to an elliptical shape. The effect of astigmatism can be so pronounced that it can be used to carry out volumetric flow measurements using only one camera (Cierpka *et al* 2010).

The direction of elongation can change over the image. Figure 2 displays an example from a TOMO-PIV experiment carried out in the water tunnel at TU Delft (Schröder *et al* 2011). Both images show different corners of the same recording of a single camera, and the spatial inhomogeneity of astigmatic distortions can clearly be seen. Other cameras will show different image shapes of the same particles. When using a uniform imaging function for the volume reconstruction process, the reconstructed particles will result from a superposition of the different distortions.

These considerations led to the idea that a spatially varying mapping function is needed, which describes the voxel–pixel-relationship in dependence on the voxel position. Its inclusion in the reconstruction process is illustrated in section 3.

2. SMART reconstruction

Atkinson and Soria applied a SMART algorithm to TOMO-PIV reconstruction and found that by using an MLOS initialization of the voxel-space, higher reconstruction speeds—while maintaining the same quality compared to a standard MART approach—can be reached (Atkinson and Soria 2009).

SMART calculates the update of a voxel as a product of the ratio of the recorded intensity to the projected intensity for all pixels influenced by the voxel. All relevant pixels of
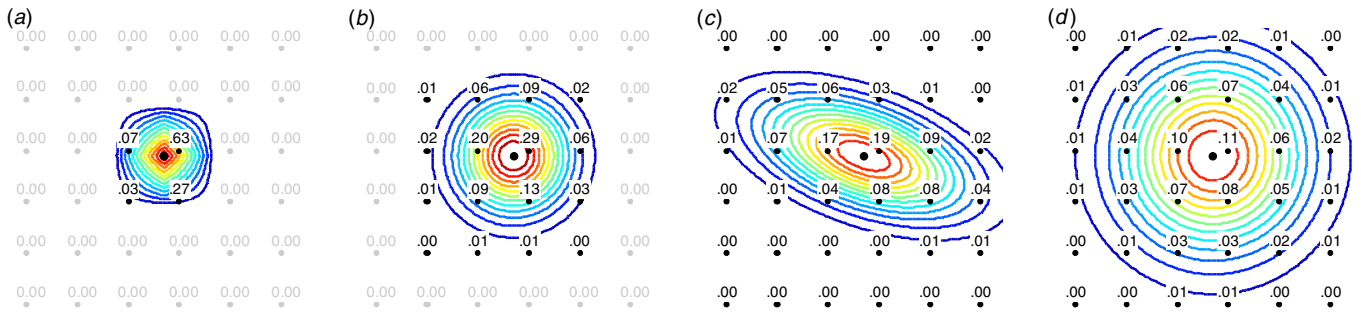
Figure 3. Examples of weighting functions, sampled on a 6×6 pixel-grid around the projection point. The shape of the weighting function is shown as a contour plot, overlaid by the sampling points (pixels) and their corresponding weight W_{ij} (normalized such that $\sum W_{ij} = 1$). (*a*) Bilinear interpolation of the neighboring pixels; (*b*) Gaussian function realized as cubic B-spline of width 3; (*c*) OTF as in equation (2) with $a = 0.8$, $b = 0.7$, $c = 2.5$; (*d*) OTF as in equation (2) with $a = 0.8$, $b = 0.0$, $c = 0.7$.

all cameras are included simultaneously. The update of the intensity I_j of voxel j in iteration $k + 1$ is calculated as follows:

$$I_j^{k+1} = I_j^k \cdot \prod_i^{N_i} \left[\left(\frac{P_i}{\sum_n W_{in} I_n^k} \right)^{\mu W_{ij}} \right]^{1/S_W}. \qquad (1)$$

Here, N_i is the number of pixels observing the current voxel (sum over all cameras), P_i the recorded intensity, μ a relaxation parameter ($\mu \leqslant 1$), W_{ij} the weight (influence) of pixel i on voxel j and S_W the sum of weights of all N_i pixels. The sum in the denominator represents the projection of all involved voxels on pixel i.

The volume is initialized using the MLOS approach: every voxel is given an initial intensity, calculated by the product of the intensities of the pixels observing the current voxel. By this method, voxels that do not see any intensity are zeroed already during initialization and need not be processed anymore (Worth and Nickels 2008).

For all calculations in this paper, a SMART algorithm with MLOS initialization was used. Different weighting functions, as described in the next section, were applied during the projection- and the update-steps.

3. From weighting functions to optical transfer functions

The SMART algorithm uses a projection of the voxels onto the pixels of the images (as opposed to the MART algorithm, which follows the reverse direction): for every voxel, the exact point of projection on all cameras is calculated in sub-pixel accuracy, according to the calibration. The pixels in the vicinity of this projection point, whose intensity is influenced by the voxel, are those with non-negative weights W_{ij} in equation (1). The spreading of the pixels influenced by a voxel is given by the *weighting function* used in the reconstruction. In the present reconstruction algorithms, this weighting function was taken to be uniform in space and the same for all cameras.

In this investigation, three different types of weighting functions are used: first, a *bilinear interpolation* between neighboring pixels (resulting in a total of 2×2 pixels per camera)—consecutively referred to as the 'lin' case (see figure 3(*a*)). Second, an approximation of a Gaussian function, implemented as two-dimensional *cubic B-spline* of width 3 (resulting in a total of 4×4 pixels per camera), consecutively

called the 'B-spline' case (figure 3(*b*)). Both cases take the weighting function to be constant for all voxels and cameras, representing the current standard approach.

In order to be able to describe varying imaging conditions for particular regions of the volume and the different cameras, a parameterized weighting function is introduced:

$$W(x, y) = p \cdot \exp\left(-\frac{1}{2} \begin{pmatrix} x \\ y \end{pmatrix}^T \begin{pmatrix} a & b \\ b & c \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \right). \qquad (2)$$

Here, x and y denote the current pixel coordinates, a, b and c describe the shape of the image, and p the peak height. By variation of the parameters, two-dimensional Gaussian and elliptical functions of arbitrary shape and size can be described. The shape given by the parameters is sampled on a pixel-grid of selectable dimension (see figures 3(*c*) and (*d*) for examples using a grid of 6×6 pixels) around the projection point of the voxel. The obtained pixel weights are used as the weighting coefficients W_{ij} in the SMART reconstruction. In the following, this parameterized weighting function will be called the *optical transfer function (OTF)* approach.

The determination of the parameters a, b, c and p from the original camera images is done via a calibration procedure, which is described in the next section.

4. OTF calibration

In order to be able to evaluate a spatially varying transfer function for the different cameras of a given experimental setup, a procedure has to be introduced which calibrates the OTF parameters to the imaging conditions. This OTF calibration was realized as an add-on step to the *volume self-calibration (VSC)*, which is nowadays a standard operation during the evaluation of TOMO-PIV data (Wieneke 2008). The scheme of this process will be explained in the following.

The goal of the VSC is to correct small errors in the geometrical calibration. These errors in the magnitude of few pixels are present in nearly every experimental setup. TOMO-PIV experiments are very sensitive to such decalibrations, as the reconstruction step depends on intersecting lines-of-sight of all cameras for a given particle.

VSC uses experimental particle recordings to correct the small errors from the geometrical calibration. Figure 4 illustrates the concept for a single particle: the 2D position
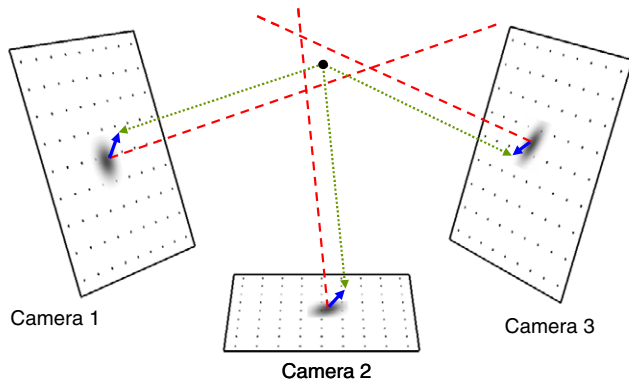
Figure 4. Scheme of the volume self-calibration. Dashed: original projection; black dot: assumed correct particle position; dotted: corrected backprojection; solid: disparity vector.

Figure 5. Basic discretization model.

Figure 6. Implementation with the OTF-accounted voxel pulse spreading function.

of the particle image is identified in each camera by a peak-finder algorithm. From the detected particle images, the particle position in space is triangulated, following the lines-of-sight as given by the geometrical calibration (dashed lines in figure 4). Typically the lines-of-sight of the different cameras will not intersect at one point of space due to decalibrations—therefore the point in space that is closest to all lines-of-sight is assumed to be the correct particle position (black dot). The position of this point is backprojected onto the camera images (dotted lines); the difference of the projected point from the original position of the particle image gives a disparity vector. Averaging the disparity over many particles gives an accurate measure for the amount of decalibration for each camera. To account for spatially varying decalibrations (introduced by lens distortions), the interrogation volume is split into several subvolumes and an averaged disparity vector is calculated for each of these.

The VSC process yields two pieces of information, which are valuable for the calibration of the OTF to a given experiment: the approximated 3D coordinates of a large number of particles and the coordinates of the particle images on the various cameras. It is therefore possible to pick particles that are known to be placed in a certain region of space and identify the shape of their images on every camera using a two-dimensional Gaussian peak-fitter, yielding the parameters a, b, c and p given in equation (2). By averaging the found peak parameters over many particles originating from the same region of space, the prevalent particle imaging in the given subvolume can be identified for all cameras.

It should be noted in this context that the choice of the weighting function has a direct influence on the size of the reconstructed particles. The wider the weighting function for a given reconstruction, the more pixels are influenced by every given voxel—leading to a reduction in the size of the reconstructed particles. The described fitting of the particle shape yields parameters that assume the particles to be singular points in space—therefore the reconstruction would normally lead to only one non-zero voxel per reconstructed particle. For the use of TOMO-PIV this effect is counterproductive (though it is physically correct, as seeding particles normally are smaller than 1 voxel), because the sub-pixel information is lost and peak locking occurs. To deal with this issue, the found
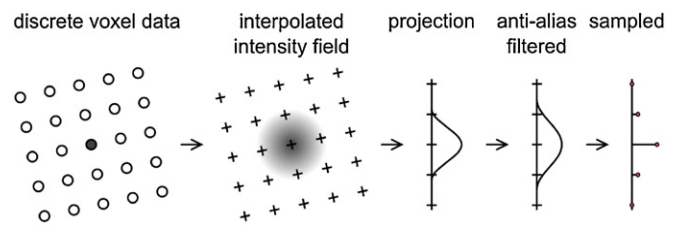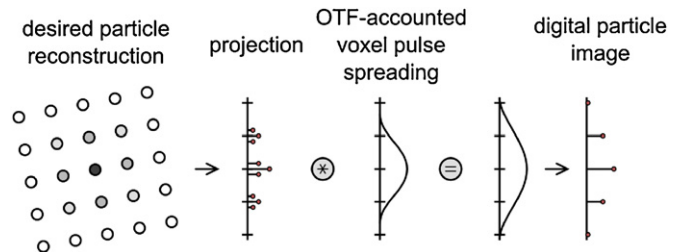
weighting functions are modified in a way that a typical particle diameter of 2 to 3 voxels is achieved—avoiding problems of peak locking. To further explain the details of this step, a look at the mathematical projection method used in our SMART implementation is necessary.

The basis for our projection model is shown in figure 5. Conceptually, a discrete scalar field of light intensities is interpolated to form a continuous scalar field. In digital signal processing terms, this step is equivalent to applying a lowpass filter that removes mirrored frequencies outside the range of wave numbers that are representable in a space discrete signal. This continuous scalar field is then projected onto an image plane by integrating the light intensity along the projection direction resulting in a continuous scalar field of reduced dimension. To satisfy the conditions of the sampling theorem for avoiding aliasing artifacts, we need to make sure that the signal is appropriately bandlimited before we sample it. This is done by applying an anti-aliasing filter to the projection.

The last curve named 'anti-alias filtered' in this example can be thought of as an image domain *voxel pulse spreading function* for a certain subvolume. The curve is the sharpest possible spreading function that avoids imaging artifacts due to the volume domain interpolator and aliasing artifacts due to the image domain anti-aliasing filter. For a reasonably small subvolume, the described mapping can be implemented without volume domain interpolation as long as the interpolator's spreading effect is accounted for in the image domain as part of the voxel pulse spreading function.

To accommodate to real optical systems with possibly non-uniform OTFs, the voxel pulse spreading function can be derived from the detected shape of real particle images the way they are recorded by the optical system. Figure 6 illustrates the relationship between the desired discrete particle reconstruction (particles with a diameter of 2 to 3 voxels), the OTF-accounted voxel pulse spreading function and a digital particle image. A smooth particle reconstruction is desirable because it will likely support the representation of the particle
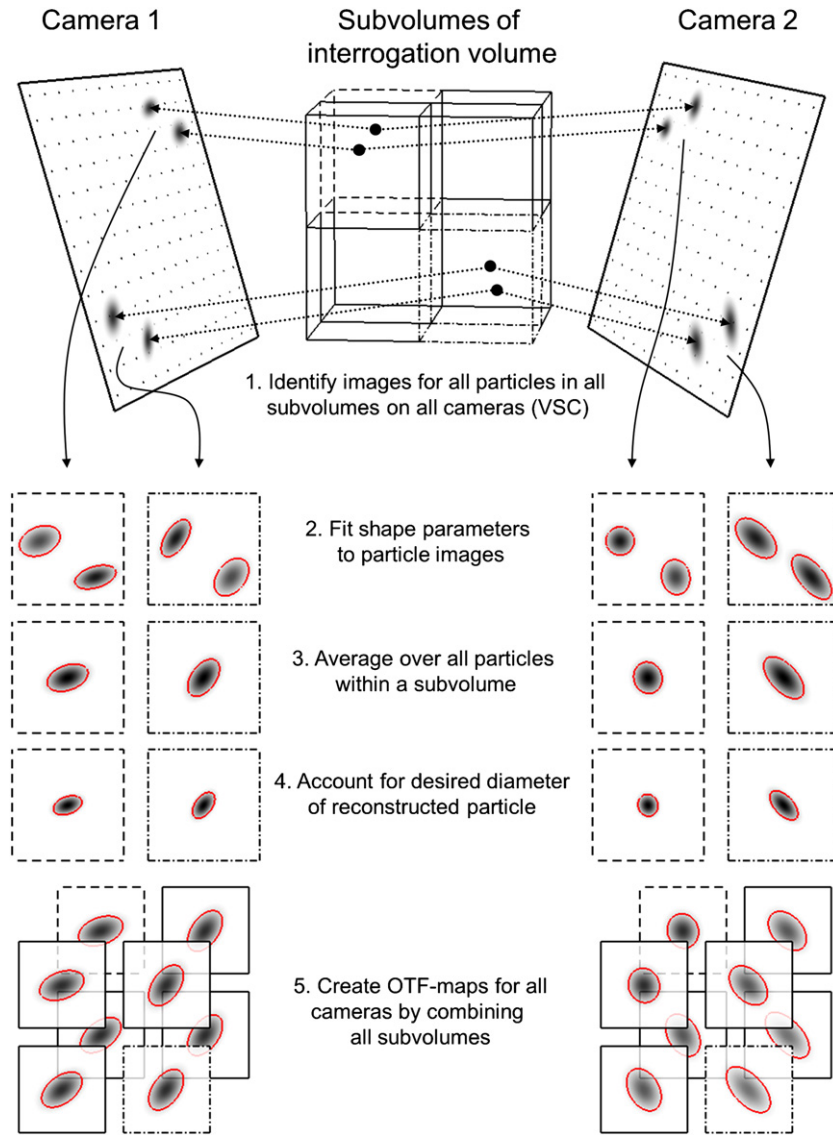
Figure 7. Scheme of the OTF calibration. The interrogation volume—here imaged by two cameras—is divided into subvolumes (here 2 × 2 × 2, exemplary subvolumes with two particles each are highlighted by dashed and dashed-dotted lines, respectively). The following steps are carried out: (1) during the process of the VSC, the position of particles in space, as well as the position of their images on all cameras is known. The particles are sorted into different subvolumes. (2) For all found particles the shape parameters of the particle image are determined by a two-dimensional Gaussian peak-fitter, independently for every camera. (3) The shape parameters for all particles within a certain subvolume are averaged, giving the OTF parameter for a point reconstruction. (4) To account for the desired diameter of the reconstructed particle (2–3 voxels) the OTF-accounted voxel pulse spreading function is determined. (5) Organizing the found shape parameters according to the subvolume-division yields spatial maps for the OTF calibration for every camera.

positions more accurately. To finally compute the voxel pulse spreading function we need to take the detected particle shape in the image domain, compensate for the spreading effect of multiple non-zero voxels of the desired reconstruction and make sure that the resulting voxel pulse spreading function will not be any narrower than the sharpest possible one from the basic discretization model to avoid imaging/aliasing artifacts.

The result of applying the described steps to all particles is the knowledge of the OTF-accounted voxel pulse spreading function (in short, OTF) for all subvolumes and cameras. It is now possible to create OTF maps for the whole volume on a coarse grid of e.g. 5 × 5 × 3 subvolumes. An example for such an OTF grid is given in figure 10 for the jet experiment

described in section 4.1. The scheme of the OTF calibration is summarized in figure 7.

Due to the linear nature of the OTF parameterization it is possible to determine the coefficients for every point of the interrogation volume by using trilinear interpolation of the sampling points of the neighboring OTF-cells.

When performing a SMART reconstruction, the weighting coefficients W_{ij} and W_{in} used in equation (1) are calculated in a window around the projection point of each voxel according to equation (2) with the coefficients gained from the OTF calibration. The width of the OTF window can be chosen in dependence on the particle image size of the experiment. Values ranging from 4 × 4 pixels to 16 × 16 pixels have been used. The larger the size of the OTF window, the
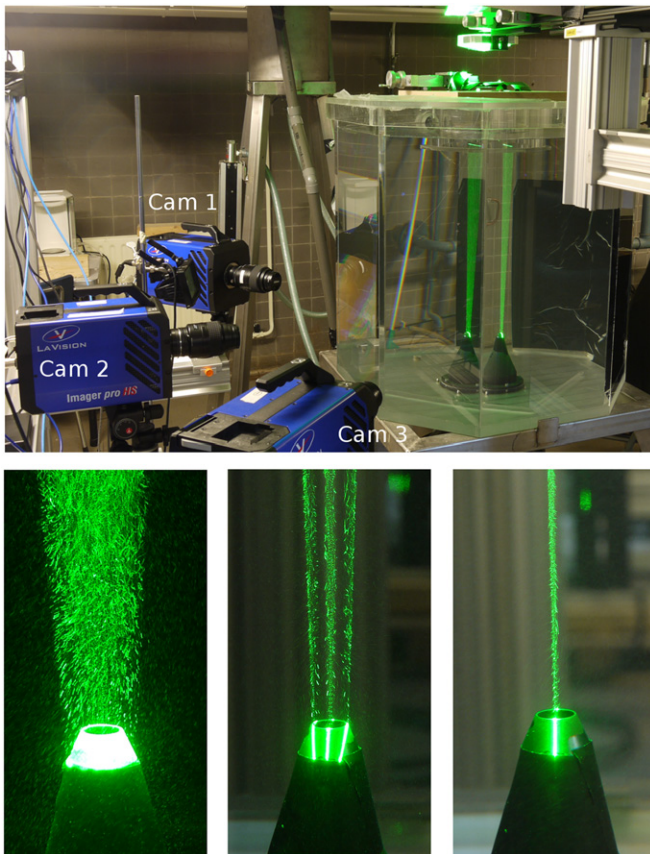
Figure 8. Top: set-up of the TOMO-PIV experiment at the water-jet facility at TU Delft. Bottom row: illumination using the full volume (left), the three-slit mask (middle) and the one-slit mask (right).

larger the computational cost is, since more pixels have to be considered for every voxel (see section 7 for more thoughts on that topic).

Even more than the VSC, the OTF calibration relies on finding clear, distinct particle images on the camera recordings. It is therefore very beneficial for a successful OTF calibration to process images with a low seeding density. Above a certain threshold of particle density, when particles start to overlap, a reliable OTF-determination will not be possible. This limit is of course dependent on the particle image diameter. For most cases, seeding densities between ppp = 0.005 and ppp = 0.01 (ppp: 'particles-per-pixel') have been found to give a reasonable balance of availability and distinctness of particle images.

4.1. Calibration of experimental data

In order to judge the applicability of the described OTF calibration to experimental data, the method was applied on data gained from a TOMO-PIV experiment on a free jet in a water tank. The experiment was carried out in the jet facility of TU Delft, the Netherlands. Figure 8 illustrates the experimental setup, which consisted of a water jet created by a nozzle with a diameter of $d = 10$ mm and an exit-velocity of $v_S = 0.5$ m s^{-1}. The jet was illuminated from above using a light volume of cylindrical shape. The illumination was chosen to diverge with growing distance to the nozzle in order to account

for the spreading of the jet flow and vortices with time. As light source, a Quantronix Darwin-Duo Nd-YLF high-speed laser with a repetition rate of 1 kHz and a light output of 2×25 mJ was used. The water was seeded using polyamide particles with an average diameter of 56 μm, the measured seeding density was approximately ppp = 0.047. The illuminated jet was imaged by three LaVision Imager pro HS high-speed cameras, observing an interrogation volume extending from $y = 13$ mm to $y = 30.6$ mm above the nozzle. The illuminated diameter was approx. 26 mm (at the top) to 23 mm (at the bottom). The water tank had an octagonal shape to allow all cameras to image the volume perpendicular to the air/glass/water interfaces. The cameras were equipped with 105 mm Nikon lenses. As the illumination was cylindrical, no Scheimpflug-adapters were needed. A cropped resolution of 672 $\times$ 512 pixels was used. In the context of this paper the focus is not on the investigation of flow parameters, but on particle imaging—therefore a small volume with small recording sizes was chosen to allow faster data processing. A more detailed description of a similar experiment can be found in Violato and Scarano (2011).

In order to generate various particle-imaging conditions, different f-numbers (ranging from $f_\# = 4$ to $f_\# = 32$) were used. The aim was to quantify the negative effects of having difficult imaging conditions (blur because of limited depth of field) compared to a case with optimal imaging conditions.

Reducing the amount of seeding within the water tank was not possible in-between the different measurements, as a full cleaning would be required. VSC and OTF calibration require a low seeding density, therefore a way to temporarily reduce the number of imaged particles was found: a plate with three thin slits was placed in the light path, leaving only three illuminated lightsheets within the volume (see figure 8). On images of this case, a VSC with OTF calibration, spanning three z-planes, was performed. In order to be able to compare the OTF maps found to a real recording, another set of images was taken with only the central lightsheet illuminated (see figure 8). Therefore, three recordings were done within a short time for all considered cases: first the full volume (on which the final reconstruction is performed); second with the three-slit mask installed (for calibration purposes) and third with the one-slit mask installed (for comparison with the OTF calibration). Examples of the single-slit recordings using $f_\# = 5.6$ are shown in figure 9. It is clearly visible that even in a single plane the effects of limited depth of field are very pronounced. For cameras 1 and 3, looking at the lightsheet at an angle of approx. 45°, only the central regions of the particle field are in focus, while the outer regions become more and more blurred. In camera 1 effects of astigmatism are visible, even though the central line-of-sight was perpendicular to the glass plates.

Using the recordings with three illuminated planes (located approx. at $z = -6$, 0, 6 mm) an OTF calibration was carried out. 500 images recorded at $f_\# = 5.6$ were used to obtain the results shown in figure 10. The interrogation volume was split into 9 $\times$ 3 $\times$ 3 subvolumes, shown in figure 10 as three OTF-z-planes for every camera. Looking at the central z-plane ($z = 0$) it is possible to compare the results
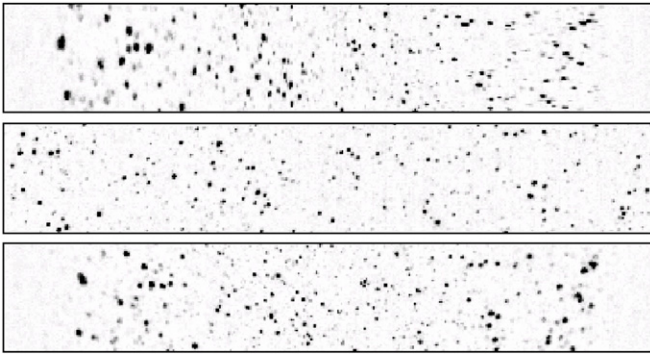
Figure 9. Color-inverted excerpts of recordings for all three cameras at $f_\# = 5.6$ for a single lightsheet in the central plane of the volume. Cameras 1 (top) and 3 (bottom) show sharp imaging in the middle of the plane, while the outer regions are blurred due to depth of field. Camera 1 additionally shows astigmatisms. Camera 2 (center) is uniformly sharp, as it is positioned perpendicular to the illuminated plane.

to the single-slit images shown in figure 9. For camera 1, the effects of blurring and astigmatism are captured—with particle imaging elongated in the *y*-direction on the left of the image, sharp imaging in the center and elongation in the *x*-direction to the right. For camera 2, the found OTFs are nearly uniform and sharp, as seen in the recorded images. Camera 3 shows blurred particle imaging on the left and right borders, but no astigmatism—also consistent with the single-slit recordings. Looking at the other *z*-planes the trends continue: for cameras 1 and 3, the region of sharp imaging shifts left/right, as is to be expected, while the blurred particle imaging increases. The focus of camera 2 seems to have been adjusted between $z = 0$ mm and $z = 6$ mm, as these two planes both show similar sharp imaging, while the $z = -6$ mm plane is uniformly blurred.

In conclusion it is obvious that the calibration of the OTF worked for the relatively extreme example shown here. Cases

with less pronounced differences in particle imaging should be easier to calibrate.

5. Application to simulated data

In order to quantify the effects of different particle imaging three test scenarios were created, simulating the conditions of (a) perfect imaging without any broadening of the particle images, (b) astigmatic distortion and (c) Gaussian blurring due to limited depth of field.

Using a particle image generator of the Davis 7.4 software (LaVision), particle distributions within a source-voxel-space of $400 \times 400 \times 160$ voxels were created. The number of particles was varied so that ppp-values ranging from ppp = 0.005 to ppp = 0.1 were realized. In all cases, the particle diameter was 3 voxels. For a second image (time step), the initial particle distribution was displaced by a vector field, describing a sinusoidal wave in the *y*-direction with a wavelength of 128 voxels. The direction and amplitude (5 pixels maximum) of this wave change in the *z*-direction. For both time steps, the particles are projected onto virtual camera images (camera configuration: 30° left, right, up and down relative to the *x*–*y*-plane) under different imaging conditions.

The test images created in all scenarios were used for reconstructions of the original volume by the three SMART variants, using the bilinear interpolation, the B-spline and the calibrated OTF as weighting functions. For comparison, a MART reconstruction was added, which was conducted by a commercial software package. MART reaches convergence after fewer iterations than SMART does. For the current reconstructions, MART used 8 iterations, while the SMART types used 25 iterations. As a single SMART iteration is around 1.5–3 times faster (depending on the weighting function used) than one MART iteration, total reconstruction times were comparable.
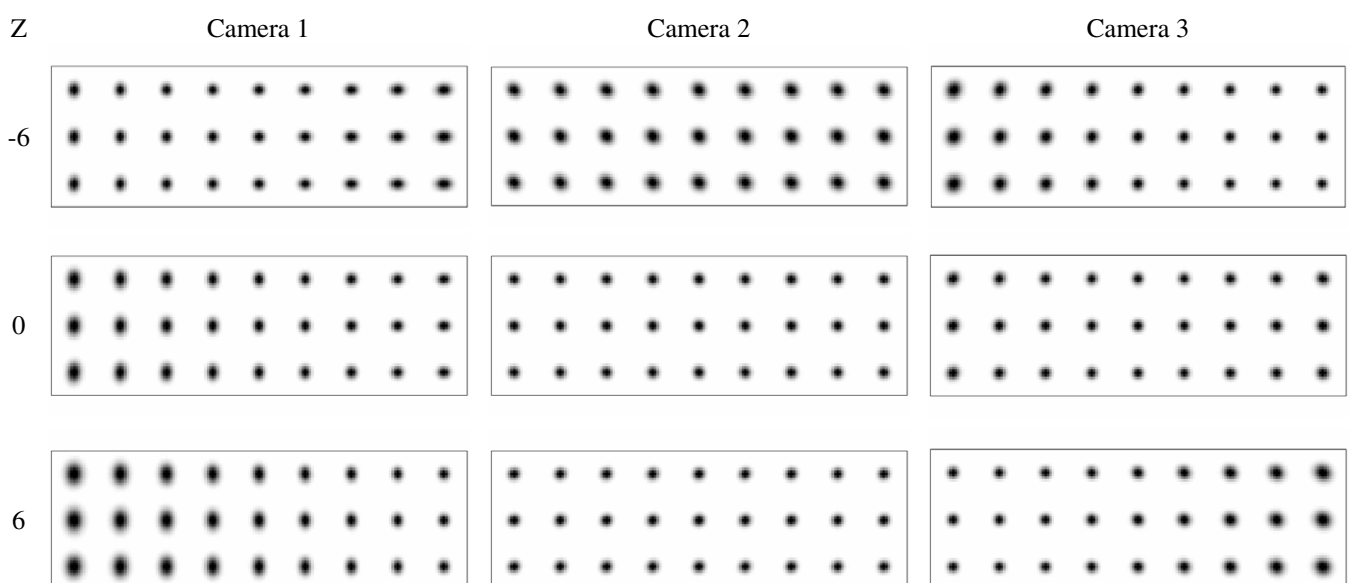
Figure 10. Averaged transfer functions for all cameras at $f_\# = 5.6$ in three *z*-layers of the interrogation volume of the jet experiment, magnified by a factor of 2.0. *Z* given in mm.

Three different criteria were used to assess the quality of the calculations: the measure for the reconstruction quality is the so-called quality factor q,

$$q = \frac{\sum_j I_{j,\text{rec}} I_{j,\text{orig}}}{\sqrt{\sum_j I_{j,\text{rec}}^2 \sum_j I_{j,\text{orig}}^2}}, \tag{3}$$

describing the correlation coefficient between the original volume and a reconstructed one. If $q = 1$, the two volumes are identical; a value of $q = 0.75$ is regarded as sufficient for a successful correlation (Elsinga *et al* 2006b).

While the quality factor is a good measure to assess reconstruction quality, for PIV purposes other factors may be more important. For instance, if the reconstructed particles are much wider than the original ones, q will suffer. For the accuracy of the correlation, however, this might be of no importance at all. Reconstructing wider particles might even yield smoother (albeit not more correct) results, as the correlation peaks get wider, making the occurrence of multiple correlation peaks less likely in the case of large gradients within the correlation volume. In order to have a quality measurement independent of the particle size, the correlation quality based value C is introduced: both the reconstructed volumes, as well as the original volumes, are correlated using the Davis 7.4 volume correlation (two passes with interrogation volumes of 32^3 voxels, 75% overlap). The resultant vector fields are subtracted from each other; the average vector length of this vector field gives the value C:

$$C = \frac{1}{N_V} \sum_{i=1}^{N_v} |\overrightarrow{v}_{R,i} - \overrightarrow{v}_{O,i}|, \tag{4}$$

where N_V is the total number of vectors, $\overrightarrow{v}_{R,i}$ is the ith vector gained from the correlation of the reconstructed voxel-spaces, $\overrightarrow{v}_{O,i}$ is the ith vector gained from the correlation of the source-voxel-spaces.

Another parameter of interest when performing particle reconstructions is the accuracy with which the particle position is represented by the voxel values during the reconstruction. In particular, when trying to perform particle tracking in 3D, the error introduced by the reconstruction process is of huge interest. After the reconstructions, a particle finder was used within the volume to identify particle positions and their diameter with sub-pixel accuracy. These results were compared to the original particle positions. For every original particle, a reconstructed particle was searched within a distance of 1 voxel in each direction. If such a particle was found, its difference in position to the original particle was computed. If no such particle was found, it counted as a 'no match'. Averaging all displacements gives the value Δ_{pos}. In this process the average particle diameter was also evaluated.

The three test scenarios and the results for the different quality criteria are described in the following.

5.1. Perfect imaging conditions

In this case, the particles were projected onto the camera images using parallel projection—no blurring in any direction was introduced. This results in very sharp images, which gives an ideal case for reconstruction. Figure 11(*a*) shows parts of

(a) *(b)*

Original   OTF
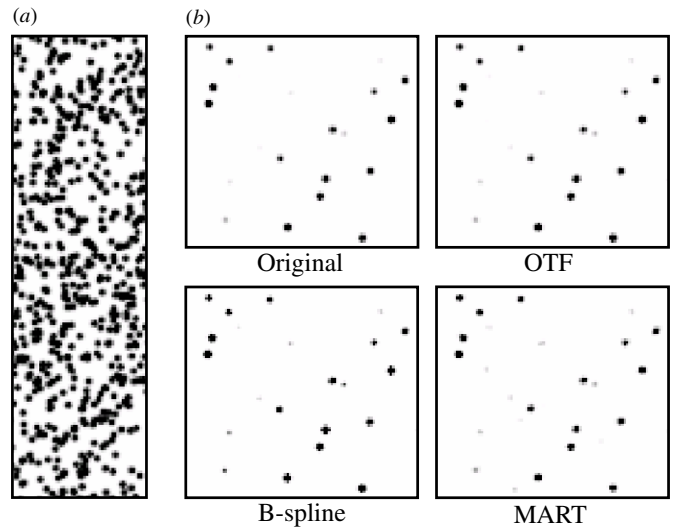
B-spline   MART

Figure 11. (*a*) Detail of camera 1 for ppp = 0.05 under perfect imaging conditions. (*b*) Section of plane $z = 123$ of the original and reconstructed volumes (SMART with OTF and B-spline weighting functions, as well as MART).

the image of camera 1 for ppp = 0.05. The images from the other cameras look comparable.

For such a scenario, a very narrow weighting function is the optimum. The OTF, described by equation (2), consisted of a small Gaussian peak with width $\sigma = 0.2$ voxels. Figure 11(*b*) shows small two-dimensional sections of the different reconstruction results. As a reference, the image from the original volume is given. It can be seen that the particles reconstructed by the OTF case are a little wider than those reconstructed by the B-spline case. This result can be explained by the fact that the B-spline describes a wider weighting function than the chosen OTF; therefore fewer voxels are needed to represent a certain particle image. The lin case (not shown) looks very comparable to the OTF case. The particles reconstructed by MART are even wider than those reconstructed by the lin case, indicating a very narrow weighting function. All reconstructions show results very close to the original volume.

Figure 12(*a*) shows the quality factor q reached after 25 iterations (MART: 8 iterations) as a function of ppp. In all cases, q stays well above 0.9, indicating a very good reconstruction.

The OTF and the lin cases are very close, showing that a narrow weighting function can be described well by the method presented in section 3. The B-spline case tends to perform a bit below the other approaches, as its weighting function is less suited to the imaging conditions. When approaching ppp = 0.1, the effects of particle overlap begin to become obvious. The lin and OTF cases seem to have the most problems with these effects, as q starts to fall off more quickly.

Due to the perfect imaging conditions, the quality factor remains much higher than commonly assumed for a certain ppp. A value of ppp = 0.05 is considered to be a threshold, above which the correlation results become unreliable. In this scenario, this threshold will be much higher.

The results for the correlation quality C are presented in figure 12(*b*). The lin and OTF cases perform best for ppp up to
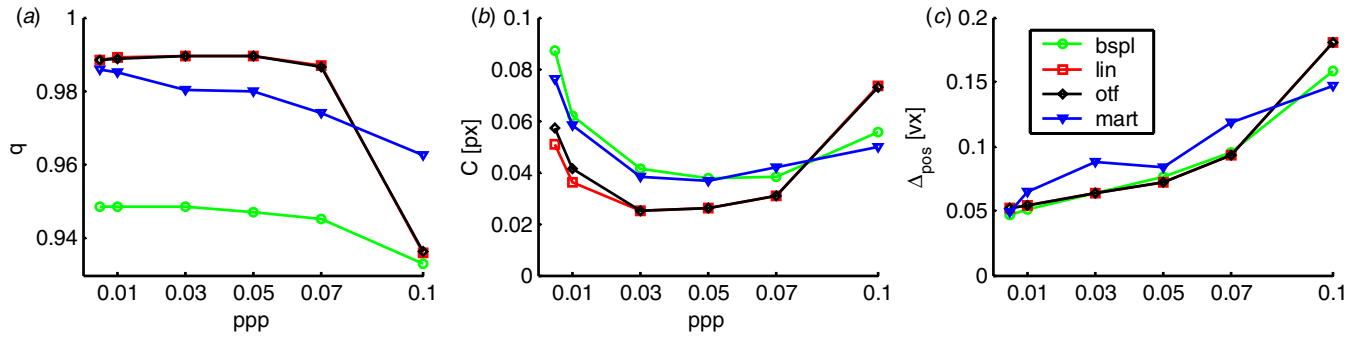
Figure 12. Perfect imaging conditions: (*a*) quality factor after 25 iterations (MART: 8); (*b*) correlation quality value; (*c*) average difference of detected peak position to original peak.

0.07. The MART and the B-spline cases are very comparable. For low values of ppp, the particle density is too low to allow a good correlation for the fixed interrogation window size, leading to errors that also show up in C, as small differences to the original volume have a larger influence in these cases. The best results can be seen between ppp = 0.03 and ppp = 0.05, where an average vector difference of around 0.025 px is reached in the lin and OTF cases.

Figure 12(*c*) shows the results of the peak position accuracy. The lower the particle density, the better the peak position is reconstructed. All methods give comparable results, and the minimum average particle displacement is around 0.05 voxels.

The averaged particle diameters for the different calculations are as follows: B-spline: 2.3 px; lin: 2.8 px; OTF: 2.8 px; MART: 2.9 px. As already indicated, the B-spline case yields the smallest particles, while MART produces the largest and is closest to the original particle size of 3 voxels.

5.2. Astigmatisms

The simulation of astigmatic distorted particle imaging was performed in a way that all cameras see astigmatism, whose direction and magnitude change with depth. In the $z = 1$ plane, axis '*a*' of the elliptical particle image is elongated by 3 pixels in both directions and axis '*b*' is untouched. In the $z = 160$ plane, the imaging parameters are reversed: now axis '*b*' is elongated. In the planes in-between, a linear transition of these extremes takes place. In the middle of the volume, the particles are imaged circularly. The angle α of the ellipse is chosen to be $\alpha = [45°, -45°, 0°, 90°]$ for the four cameras, so that every camera sees a different image of a given particle.

Figure 13(*a*) shows parts of the images of camera 3 (looking at the volume from the top) for that case. The image shows particle images with a strong elliptical shape at an angle of 45°. These are images of particles deep in the volume, near $z = 160$. When looking down, images of particles from planes further in front appear, that are either imaged more or less circular or show an elliptical shape elongated along the other principal axis. The images of the other cameras are comparable, except for showing astigmatism in other directions.

From these images, reconstructions with the different methods were calculated. The OTF case was calibrated using
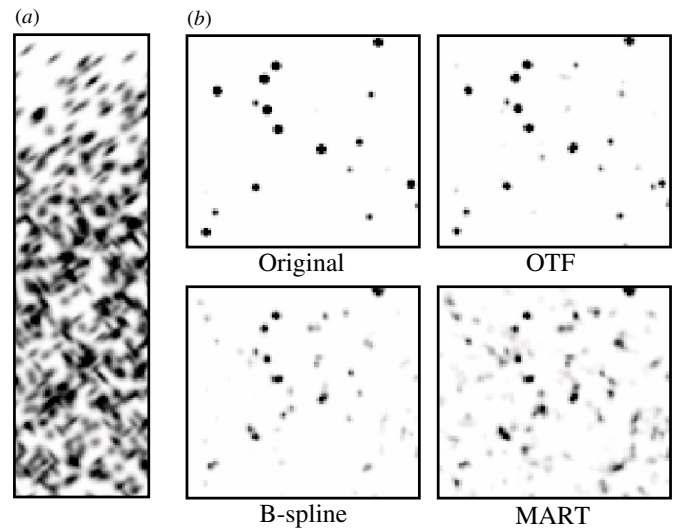
Figure 13. (*a*) Upper-left corner of camera 3 for ppp = 0.05 with simulated astigmatism. (*b*) Section of plane $z = 5$ of the original and the reconstructed volume (SMART with OTF and bilinear weighting functions, as well as MART).
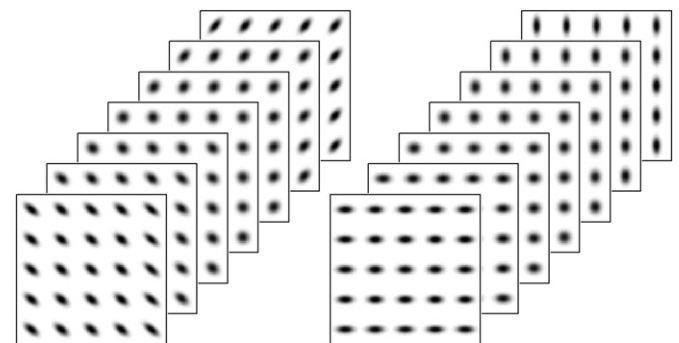
Figure 14. Results of the OTF calibration for camera 1 (left) and camera 4 (right) for the astigmatism test case.

a grid of $5 \times 5 \times 7$ subvolumes. Images with low seeding density (ppp = 0.005) were used. Results for two cameras of this calibration are shown in figure 14. The calibration was able to capture the imaging situation very well.

Figure 13(*b*) shows small two-dimensional sections of the different reconstruction results near the front of the volume for ppp = 0.05. It can clearly be seen that the OTF case yields the best representation of the original volume. The particle shapes
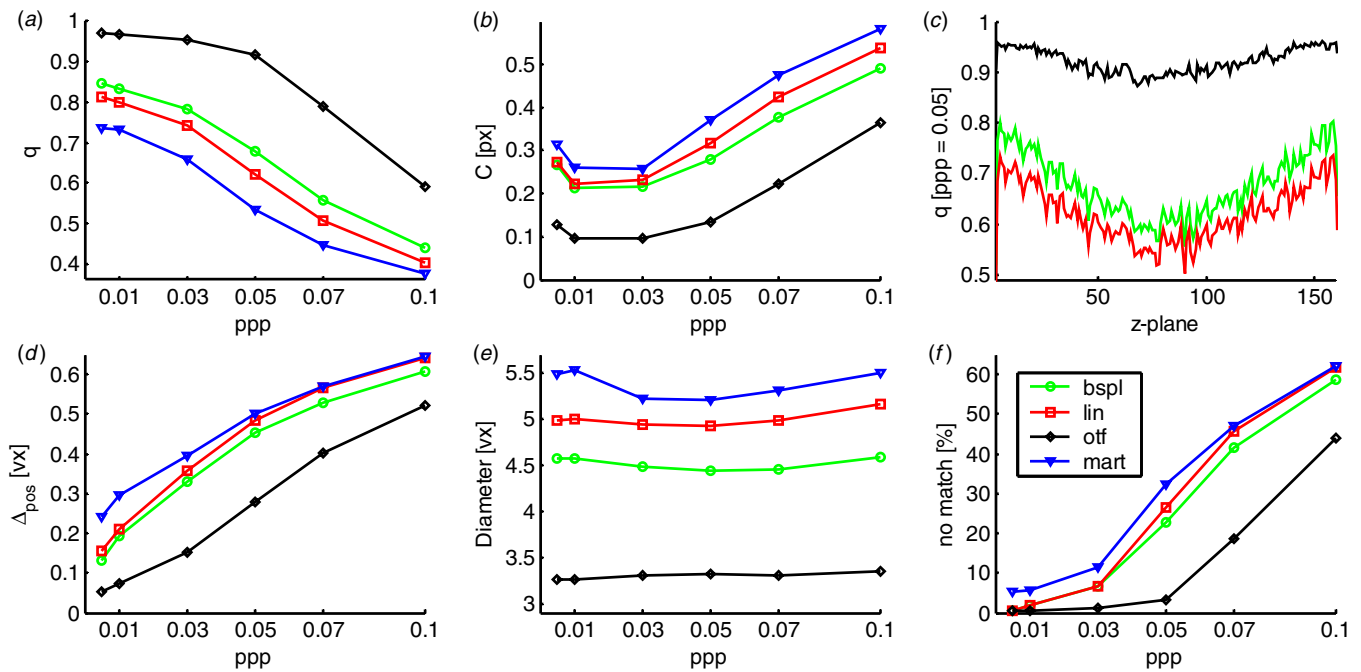
Figure 15. Astigmatic imaging conditions. (*a*) Quality factor after 25 iterations (MART: 8) over ppp; (*b*) quality factor development over iterations for ppp = 0.05; (*c*) correlation quality value; (*d*) quality factor profile over all *z*-planes.

are better defined and sharper compared to the B-spline and the MART cases. While the OTF case shows little ghost intensity, the B-spline case shows signs of ghost particles, as distinct peaks start to appear where no particle should be. The MART reconstruction shows even more ghost intensity, though the particle peak intensity seems better defined, compared to the B-spline case.

More specific results are shown in figure 15. The quality factor reached after 25 SMART iterations (figure 15(*a*)) supports the findings of figure 13(*b*): the OTF case is clearly better than all other cases with uniform weighting functions. For low ppp, *q*-values well above 0.9 are reached; up to ppp = 0.07 the quality factor remains above 0.8. Over ppp = 0.05, particle overlap becomes increasingly problematic, so that the quality values of all cases start to fall off more quickly. The best uniform transfer function is the B-spline case. This can be understood, as the average particle image diameter is wide; therefore the narrow transfer functions of the lin and, even more so, the chosen MART cases are not well suited.

The situation with respect to the correlation quality *C* (figure 15(*b*)) is comparable to *q*: the OTF case shows a distinct advantage over all other methods. For lower values of ppp—up to ppp = 0.05—the OTF case reaches average vector difference displacements that are about half as great as the best other case.

When looking at the profile of the quality factor along the *z*-axis (figure 15(*c*)) it becomes obvious that the center parts of the volume are the hardest to reconstruct precisely. In these regions, the particles are imaged nearly circular, but with a relatively wide diameter. The cases with uniform weighting functions (being comparatively narrow) will reconstruct particles with a large diameter, which leads to a reduction of *q*. For the OTF case this effect is much less pronounced. At the outer limits of the volume, where elliptical imaging occurs,

the narrow second principal axis of the images—overlapped at different angles—helps to limit the size of the reconstructed particles.

The results of the peak-finder calculations are presented in figures 15(*d*) to (*f*). The accuracy of the peak position is much improved by the use of an OTF: for ppp = 0.005, an accuracy of 0.05 voxels is reached, which is comparable to the values of the perfect imaging test case in section 5.1. The B-spline case shows an average error of 0.13 voxels at this ppp; the other algorithms are worse. Up to ppp = 0.03, the OTF case exhibits errors that are less than half of the ones of the other cases. Over ppp = 0.05, the advantage diminishes, because other effects, such as particle overlap, begin to dominate.

The percentage of non-matched particles (original particles for which no corresponding reconstructed particle could be found, figure 15(*f*)) shows profound advantages for the OTF cases. For ppp = 0.05, the OTF reconstruction was able to reconstruct nearly 97% of all particles so that they could be found by the peak detection. The B-spline case had a success rate of only 78%, and MART of 68%.

The average detected particle diameter (figure 15(*e*)) shows that MART again reconstructs the largest particles, at around 5.3 voxels. The lin case (∼5.0 voxels) and the B-spline case (∼4.6 voxels) are lower, but the OTF case (∼3.3 voxels) is much closer to the original value of 3 voxels.

All the results of this test case show the profound advantages of using a spatially varying OTF when dealing with astigmatism. All examined parameters are significantly improved by the use of such a weighting function.

5.3. Gaussian blurring (limited depth of field)

In order to simulate the effects of limited depth of field, the test images were created with a linear increase of blurred
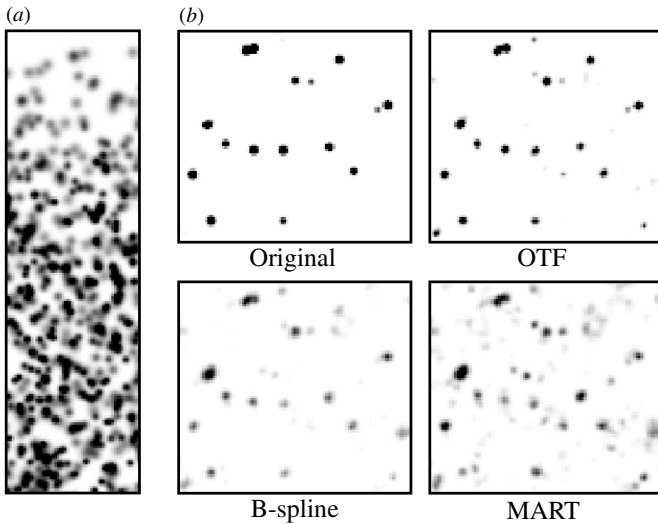
Figure 16. (*a*) Upper-left corner of camera 1 for ppp = 0.05 with limited depth of field. (*b*) Section of plane *z* = 150 of the original and the reconstructed volumes.

imaging in *z*: at *z* = 1 the particle images are widened by 0.3 pixels in every direction, while at *z* = 160, the spread is 2.0 pixels. This leads to camera images shown in figure 16(*a*): at the top, blurred particle images, originating from the deeper *z*-planes, can be seen. Further down, these are superposed by sharply imaged particles from the higher *z*-planes. In this case, the cameras are focused on the same plane; thus every single particle is imaged equally by all cameras.

Figure 16(*b*) shows results from the different reconstructions of these images for ppp = 0.05. A plane near the lower end of the volume is shown, where particle imaging

is blurred. As in section 5.2, the OTF reconstruction shows the best reproduction of the original volume. The particles are well defined and ghost intensity seems to be low, though a little higher compared to the astigmatism scenario. The B-spline and the MART cases show much higher ghost intensity. The particle shape is widened, and the peaks are less distinct.

The results for the quality factor *q*, shown in figure 17(*a*), illustrate these findings: for the OTF case, a much higher reconstruction quality can be achieved. For ppp = 0.05, a quality factor of 0.94 is seen, while the best non-OTF method reaches 0.80. Again, the B-spline case gives better results, compared to the lin or MART case, as the average particle image is relatively wide.

The *q*-profile given in figure 17(*c*) illustrates the dependence of the reconstruction quality on the particle image diameter: the lin and B-spline cases show a steep falloff in *q*, almost linearly following the increase in blurred particle imaging. In the OTF case, this effect is much less pronounced.

The important parameter of correlation quality (figure 17(*b*)) also shows a benefit when using the OTF method. With *C* = 0.058 px in the OTF versus *C* = 0.089 px in the B-spline case (for ppp = 0.03), this advantage is less pronounced than in the astigmatism calculation, but still relevant up to ppp = 0.07.

The accuracy of the particle peak position, displayed in figure 17(*d*), shows an interesting aspect: although the quality factor and the correlation quality are clearly better in the OTF case, the particles are not found with better accuracy. All methods perform pretty much the same in this regard. This can be explained by the fact that particle imaging may be blurred, but is still spherical and identical for all cameras. The reconstructed particles will differ in size, but their centre is well represented nevertheless.
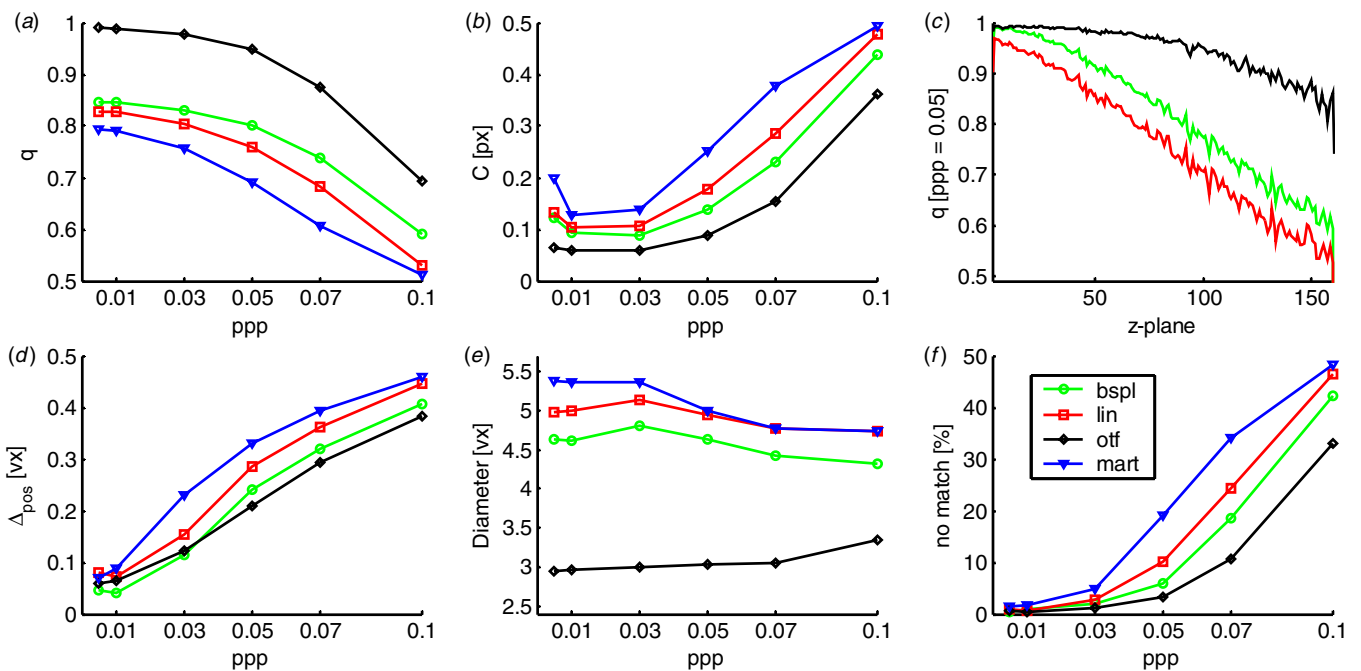
Figure 17. Limited depth of field. (*a*) Quality factor after 25 iterations (MART: 8) over ppp; (*b*) correlation quality value over ppp; (*c*) quality factor profile over all *z*-planes; (*d*) average difference of detected peak position from original peak; (*e*) average detected particle diameter; (*f*) percentage of original particles with no reconstructed partner.
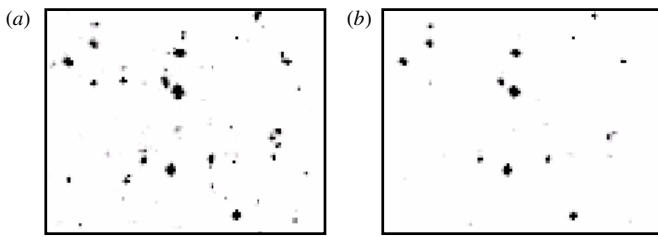
Figure 18. Details of plane 128 of reconstructions from the jet experiment (lenses at $f_\# = 5.6$) using (*a*) bilinear interpolation; (*b*) calibrated OTF (see figure 10).

Still, the OTF approach has an advantage for particle identification: more particles are correctly found (see figure 17(*f*)). The large diameter—and therefore reduced core intensity—of the particles reconstructed by non-OTF methods seems to prevent reliable particle detection in the low *z*-planes. The oversized particle reconstruction of the non-OTF methods can be seen in figure 17(*e*): the OTF case reconstructs particles with average diameters close to the original value, while the other algorithms show significantly higher numbers.

Altogether it can be noted that the prepared test cases showed significantly better results when using the OTF method. The reconstructed volumes were chosen comparably small in order to allow efficient data processing. An example for a larger volume is presented in section 7, showing that the improvements introduced by the OTF translate very well to larger volumes.

6. Application to experimental data

It is much harder to show quality differences of reconstructions for experimental data, as the ground truth is not known. It is often difficult to judge which is a better reconstruction/vector-field without a reference. One method would be to compare results originating from calculations performed with a reduced camera set to those performed with all available cameras (Novara and Scarano 2011). The experiment on a free turbulent jet in water, presented in section 4.1, used a three-camera setup—which makes the use of a camera-subset not suitable. Therefore a PTV approach was chosen to compare the different calculations.

The interrogation volume of $30 \times 25 \times 30$ mm^3 was discretized into $534 \times 445 \times 543$ voxels. Reconstructions were performed using the bilinear interpolation and the OTF, calibrated according to the considerations in section 4.1. The volume was initialized by performing one iteration of MLOS, followed by ten iterations of SMART. An example of the reconstruction results for a plane located approx. 7.8 mm behind the central plane is given in figure 18. The OTF case seems to produce sharper, more defined particles. Overall fewer peaks are visible for the OTF case.

A 3D particle finder was used to identify particle positions within the voxel space with sub-pixel accuracy (comparable to the method used on the synthetic data). 50 successive time steps were reconstructed and treated by the particle finder. The resulting list of particles for these time steps was fed into the 'time series 3D-PTV' algorithm of LaVision Davis 8.0.4, where continuous particle tracks were detected. Figure 19
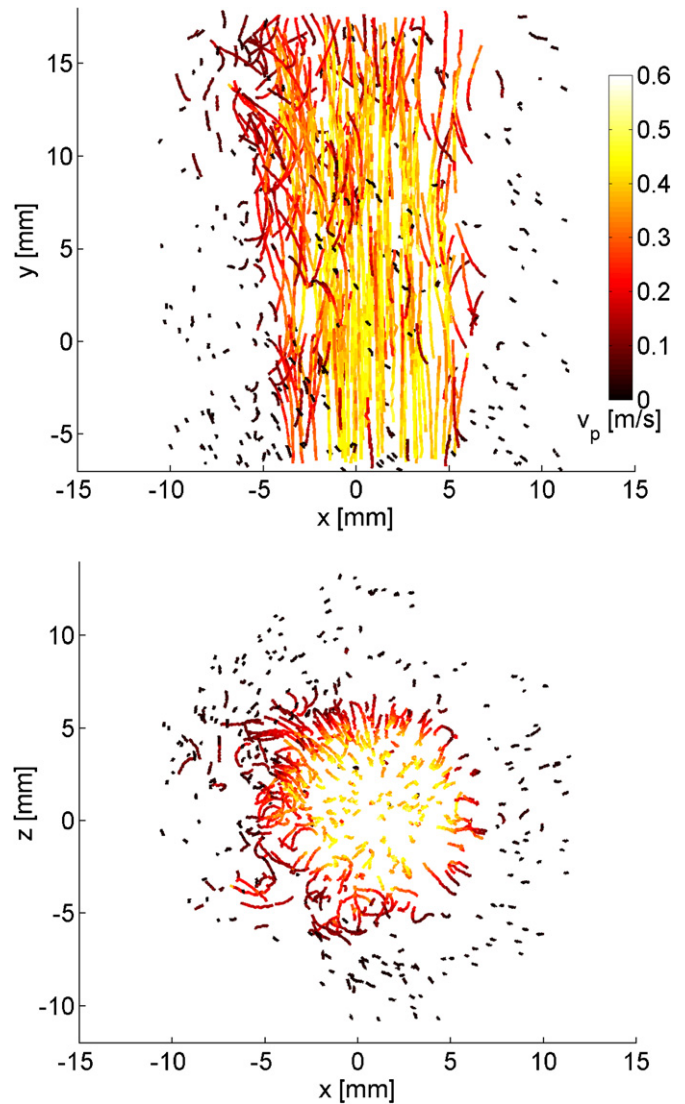
Figure 19. Particle tracks found by 3D-PTV for the jet experiment. Color indicates particle velocity. Top: side view; bottom: top view.

shows visualizations of all particles that were found in more than 16 consecutive volumes in the OTF reconstructions. Fast particles in the center of the jet can be seen, as well as slowly moving particles in the outer regions of the jet. Some of these are being sucked into the central regions and are accelerated in the process.

The particle finder identified considerably more particles for the lin case (approx. 48 000 per volume) compared to the OTF case (approx. 26 000 per volume). The track-finder algorithm found around 164 000 tracks with lengths of 2 or 3 for the lin case, while only around 64 000 of these were found in the OTF case. Such a short track length indicates either a wrong pairing or that the found particle was a ghost particle and its track ended as soon as the original particles had moved far enough to resolve the ambiguity. Therefore the smaller number of very short tracks indicates a reduced ghost level for the OTF case. Long tracks (>12 consecutive time steps) are found slightly more often in the OTF case (1567 versus 1473).

In order to judge the accuracy of the particle position the following method was applied: for every found track above

Table 1. Comparison of PTV accuracy for reconstructions performed with bilinear interpolation ('Lin') and calibrated OTF. Given are the averaged rms-values for the three components of velocity (u, v, w; in m s^{-1}) and position (x, y, z; in mm) in relation to a second-order polynomial fit to the according values for every found track.

	u_{rms}	v_{rms}	w_{rms}	x_{rms}	y_{rms}	z_{rms}
Lin	0.0238	0.0251	0.0292	0.0203	0.0234	0.0242
OTF	0.0205	0.0208	0.0269	0.0168	0.0187	0.0225
Gain	13.9%	17.1%	7.9%	17.2%	20.1%	7.1%

a length of 12, a polynomial of second order is fitted to every component of velocity (u, v, w) and position (x, y, z) of the track. The rms-values of all components relative to the fitted polynomial are evaluated and averaged over all tracks. This method assumes that the particles should move on smooth trajectories and should not be rapidly accelerated. Small accelerations and trajectory changes are accounted for by the polynomial fit. The results of this analysis are presented in table 1. For the three components of velocity a gain of accuracy between 8% and 17% is reached by using the OTF method. The position accuracy is improved between 7% and 20%. The component in the streamwise direction (y) is always the one that benefits the most, while the z-component is least affected. This finding can be explained by the in-line camera setup: reconstruction artifacts (elongation of the reconstructed particles along camera lines-of-sight) are, in this setup, present only in the x- and—most pronounced—in the z-directions. The achievable accuracy is therefore least compromised in the y-direction. The higher absolute rms-values in y can be explained by the orientation of the jet along the y-axis, leading to higher accelerations of the particles along this axis.

To summarize the findings of the evaluation of the jet experiment it can be said that the OTF method yielded fewer ghost particles and improved the accuracy of the peak detection by 7% to 20%, depending on the component.

7. Reconstruction speed considerations

Investigations using tomographic PIV have been and still are sensitive with regard to the calculation time necessary for the operations of reconstruction and correlation. It is therefore of practical interest not to introduce additional time consuming operations. This section intends to address the potential impacts on computation time by applying the OTF approach to tomographic reconstructions.

When assuming a constant number of non-zero voxels the computation time for the projection- and the voxel-update-steps is heavily dependent on the number of camera-image pixels influenced by each voxel. Therefore it is evident that the move from a bilinear interpolation of the nearest neighbors (2×2 pixel interrogation size) to an OTF with a size of e.g. 8×8 pixels comes at a huge computational cost. It is however possible to significantly reduce the additional cost for both the projection and the voxel update.

In the case of the projection step it has been found effective to not apply the OTF parameters to every single voxel, but to cluster voxels into small subsets, where the OTF
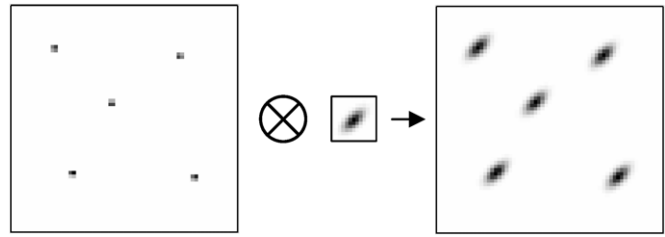
Figure 20. Application of the OTF by convolving an 8×8 pixel kernel (size can be freely chosen) with a temporary projection image (here 64×64 pixels).

can be assumed to be constant: a predefined subvolume (e.g. $15 \times 15 \times 15$ voxels) is taken from the full voxel-space. All voxels within this small subvolume are then projected onto a small temporary image, whose size is chosen according to the maximum spread of the projections of the involved voxels. The key step is that the projection onto this temporary image is done using *bilinear interpolation*, being the smallest weighting function that allows sub-pixel accuracy. After the intensities of all voxels have been projected, the OTF is applied simultaneously for all these voxels by using a two-dimensional convolution of the temporary image with a kernel, representing the OTF in the interrogated region of space. The process is illustrated in figure 20. After this step, the resulting image is simply added pixel-by-pixel to the appropriate area of the main projected image. The computational cost of the added convolution step is much lower compared to applying the full-size OTF for all voxels in the subvolume. The introduced differences due to the assumption of constant OTF within the subvolume have been found to be insignificant. By varying the kernel size the grid-size of the OTF can be freely chosen, without any consequences for the computation time.

The convolution method introduced in the last paragraph allows a speed-up of the OTF application to the projection step of the SMART algorithm. For the update step an easier practice was found: during the analysis of the different synthetic test cases presented in section 5, it was noticed that the application of the OTF *only during the projection step* yields nearly the same results as applying the OTF in both steps. It is sufficient to use a bilinear interpolation during the voxel update—while using the OTF for the projection—to fully benefit from the OTF approach. This finding leads to a very significant reduction in computation time needed for the voxel update. Combined with the application of the OTF by convolution a very efficient inclusion of the OTF in the reconstruction process was realized.

In order to document the impact on computation time for the different types of reconstruction, a variation of the astigmatism test case, presented in section 5.2, was used. All basic parameters were left untouched, and only the dimensions of the source-voxel-space were enhanced to $1000 \times 1000 \times 400$ voxels; the resulting test images looked comparable to figure 13(*a*) and had dimensions of 1200×1200 pixels. For this test a single particle density of ppp = 0.03 was used. Reconstructions were performed using (*a*) the bilinear interpolation for both steps of projection and voxel update; (*b*) the cubic B-spline for both steps; (*c*) the calibrated OTF (width: 8×8 pixels) for both steps (referred
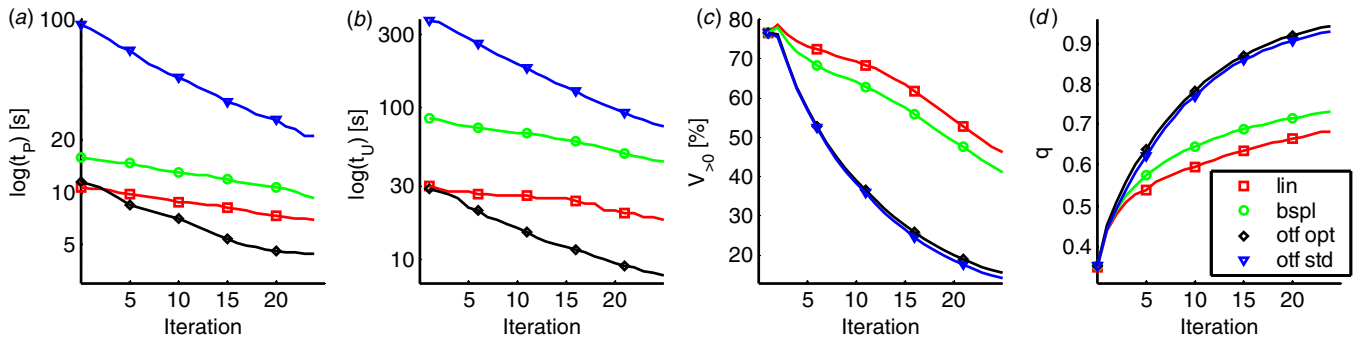
Figure 21. Results from reconstructions of a synthetic test case using astigmatic imaging conditions (see section 5.2) with enhanced volume (1000 × 1000 × 400 voxels). Shown are (on logarithmic time scales) (*a*) the computation time per iteration for the projection t_P and (*b*) for the voxel update t_U, as well as (*c*) the number of non-zero voxels $V_{>0}$ and (*d*) the quality factor q. Reconstructions were performed using the following weighting functions: (i) bilinear interpolation, (ii) cubic B-spline and (iii) calibrated OTF. For the OTF a standard method and a method optimized for computation speed were compared.

to as 'OTF standard'); and (*d*) the calibrated OTF (width: 8 × 8 pixels) with the convolution method for the projection and the bilinear interpolation for the voxel update (referred to as 'OTF optimized'). 25 iterations of SMART followed an MLOS initialization. The calculations were performed using 40 cores on a Dell PowerEdge 815 server. The time needed for the projection (t_P) and the voxel update (t_U) was measured for every iteration, as well as the quality factor q and the percentage of non-zero voxels $V_{>0}$.

Results from these investigations are shown in figure 21. The cases using bilinear interpolation and cubic B-spline show a slow decline of t_P and t_U with the iteration number, as more and more voxels are zeroed during the reconstruction—decreasing the number of voxels that need to be processed. The voxel update comes at a higher computational cost compared to a projection, because of the power operation (x^y) needed in equation (1) (for the projection the intensities are added). After the MLOS initialization $V_{>0}(0) = 76.4\%$ of the voxels still have intensities higher than zero. The bilinear interpolation needs $t_P(1) = 10.7$ s for the first projection and $t_U(1) = 30.0$ s for the first update. After 25 iterations, $V_{>0}(25)$ is down to 46.0% and the calculation times are reduced to $t_P(25) = 6.9$ s and $t_U(25) = 18.2$ s. In the B-spline case calculation times are higher, as 4 × 4 pixels have to be considered per camera. This effect is more pronounced in the voxel update ($t_U(1) = 84.1$ s), as the power operation takes a dominant role in calculation time. The B-spline case converges to a solution with fewer non-zero voxels ($V_{>0}(25) = 41.1\%$), as the weighting function is a slightly better fit to the imaging conditions, compared to the bilinear interpolation. This is also measurable in the quality factor, where $q(25) = 0.728$ for the B-spline, while the bilinear interpolation only reaches $q(25) = 0.678$.

When examining the two OTF cases it is evident that these cases show different convergence behavior, compared to the cases with uniform weighting function. The standard OTF case starts with very high computation times per iteration for both the projection ($t_P(1) = 93.1$ s) and the voxel update ($t_U(1) = 371.9$ s). The wider weighting function (8 × 8 pixels) comes at a high computational cost. However, the system converges much faster and many voxels are zeroed rapidly compared to the lin and B-spline cases. In the last

iteration a value of $V_{>0}(25) = 14.3\%$ is reached. The reduced number of voxels needed to be processed leads to a rapid reduction of computation time per iteration ($t_P(25) = 20.8$ s; $t_U(25) = 72.7$ s). Still, the total computation time for the standard OTF case (5631 s) is much higher compared to the bilinear interpolation (826 s) and the B-spline (1894 s). The quality factor improves significantly over the other two cases: $q(25) = 0.928$ is reached.

Looking at the results for the optimized OTF case it can be seen that the time needed for the first projection ($t_P(1) = 11.4$ s) is slightly higher than for the bilinear interpolation. This is to be expected, as the weighting function used is the same, but some extra calculation time is needed for the convolutions of the temporary images. The time needed for the first voxel update ($t_U(1) = 29.8$ s) is identical to the bilinear interpolation, as expected. The evolution of the number of non-zero voxels ($V_{>0}(25) = 16.1\%$) and the quality factor ($q(25) = 0.943$) is nearly identical to the standard OTF case. Therefore the computation time per iteration decreases rapidly also in this case ($t_P(25) = 4.4$ s; $t_U(25) = 7.8$ s). The total calculation time amounts to 551 s, which is approx. 33% faster compared to the bilinear interpolation.

The correlation quality factor C, introduced in section 5, confirms the trends seen from q and $V_{>0}$ ($C_{\text{LIN}} = 0.565$; $C_{\text{BSPL}} = 0.511$; $C_{\text{OTF STD}} = 0.265$; $C_{\text{OTF OPT}} = 0.244$). As with q, the quality achieved with the time-optimized OTF implementation is even slightly higher compared to the standard implementation.

Summing up this section, it can be said that a method was found to apply the OTF in a way that does not compromise on quality, but proves to be very cost effective. In the (rather extreme) synthetic case presented here it was possible to achieve a significant reduction in computation time, not only compared to the standard OTF method, but even compared to the fastest uniform weighting function, because the number of voxels that need to be processed is reduced much more efficiently (due to less ghost energy scattered in the volume). Comparable behavior was found in all synthetic test cases. It remains to be seen if such computational gains can also be realized in experimental data, but an inclusion of the OTF method without computational penalty seems to be easily achievable with the proposed methods.

8. Summary and outlook

A spatially varying OTF for tomographic reconstruction has been introduced to account for changes in particle imaging over space. The concept has been implemented by the use of a parameterized weighting function in the SMART algorithm.

A method for calibrating the parameters of particle imaging on a coarse volumetric grid has been introduced. This OTF calibration was realized as an enhancement to the VSC, which is already a standard step of every TOMO-PIV analysis. It was shown that the proposed calibration method works for simulated as well as experimental data.

Test cases, simulating the effects of astigmatism and Gaussian blurring, have been created. Reconstructions of these cases using the OTF, as well as standard approaches, were conducted and compared to a test case with ideal imaging conditions. It was found that by using a weighting function that is spatially calibrated to the imaging conditions, the negative effects of distorted imaging on reconstruction quality, correlation quality and accuracy of particle detection can be greatly reduced. For high particle densities (ppp > 0.07) the effects of particle overlap, which are amplified by large particle images, begin to reduce the gains of the OTF approach.

Evaluation of experimental images by PTV algorithms showed a reduction of the number and intensity of ghost particles when using a calibrated OTF. Improvements of the accuracy of the peak position determination could be seen in the order of 7–20%, depending on the spatial component.

A straightforward implementation of the OTF into existing algorithms might cause severe impact on computational cost (depending on the chosen grid-size of the OTF), as a lot of camera pixels have to be processed per voxel. A method was proposed to reduce calculation times, while maintaining all benefits introduced with the OTF. Reconstructions of synthetic cases using this method showed that it was possible to realize reconstruction speeds that are even faster than those achieved using the smallest uniform weighting function (a bilinear interpolation of the nearest neighbors).

As uniform particle imaging for all cameras over the whole volume is not easy to accomplish, the OTF method has the potential to improve experimental results in nearly all tomographic PIV setups. Combination with other recent approaches to improve reconstruction quality, for instance *motion tracking enhancements* (Novara *et al* 2010), seems feasible.

A further refinement of the proposed method would be to establish a system that allows the use of arbitrary weighting functions. Not all optical distortions can be described by Gaussian functions (caustically distorted or donut (ring)-like particle imaging comes to mind). With the current approach these kinds of distortions cannot be accounted for.

An application of the OTF method to other measurement techniques is conceivable: in this work it has been shown that the use of an OTF calibration has significant benefits for the peak accuracy in PTV applications. It should be a key point in the recently introduced method of *iterative reconstruction of volumetric particle distributions* (Wieneke 2013), which relies heavily on the knowledge of the correct weighting function— as particles are no longer reconstructed as clusters of voxels, but as point particles.

Kähler *et al* showed that the achievable resolution of planar PIV recordings can be enhanced by using PTV methods or single pixel correlation, when dealing with averaged results (Kähler *et al* 2012). In this scenario the uncertainty can be pushed below 1/100th of a pixel under certain conditions. For a reliable PTV result a precise determination of the center of mass of a particle image is required. In the case of distorted particle imaging the accuracy of this measurement can be compromised. The knowledge of the prevalent particle imaging, gained by performing an OTF calibration, could be used in the peak fitting process to further increase the achievable resolution.

References

Atkinson C and Soria J 2009 An efficient simultaneous reconstruction technique for tomographic particle image velocimetry *Exp. Fluids* **47** 563–78

Bosbach J, Kühn M and Wagner C 2009 Large scale particle image velocimetry with helium filled soap bubbles *Exp. Fluids* **46** 539–47

Cierpka C, Segura R, Hain R and Kähler C J 2010 A simple single camera 3C3D velocity measurement technique without errors due to depth of correlation and spatial averaging for microfluidics *Meas. Sci. Technol.* **21** 045401

Elsinga G, Scarano F, Wieneke B and van Oudenheusen B W 2006a Tomographic particle image velocimetry *Exp. Fluids* **41** 933–47

Elsinga G, van Oudenheusen B W and Scarano F 2006b Experimental assessment of tomographic-PIV accuracy *13th Int. Symp. on Applications of Laser Techniques to Fluid Mechanics (Lisbon)*

Herman G T and Lent A 1976 Iterative reconstruction algorithms *Comput. Biol. Med.* **6** 273–94

Kähler C J, Scharnowski S and Cierpka C 2012 On the uncertainty of digital PIV and PTV near walls velocimetry *Exp. Fluids* **52** 1641–56

Novara M, Batenburg K J and Scarano F 2010 Motion tracking-enhanced MART for tomographic PIV *Meas. Sci. Technol.* **21** 035401

Novara M and Scarano F 2011 Performances of motion tracking enhanced Tomo-PIV on turbulent shear flows *Exp. Fluids* **52** 1027–41

Schröder A, Geisler R, Staack K, Elsinga G E, Scarano F, Wieneke B, Henning A, Poelma C and Westerweel J 2011 Eulerian and Lagrangian views of a turbulent boundary layer flow using time resolved tomographic PIV *Exp. Fluids* **50** 1071–91

Violato D and Scarano F 2011 Three-dimensional evolution of flow structures in transitional circular and chevron jets *Phys. Fluids* **23** 124104

Wieneke B 2008 Volume self-calibration for 3D particle image velocimetry *Exp. Fluids* **45** 549–56

Wieneke B 2013 Iterative reconstruction of volumetric particle distribution *Meas. Sci. Technol.* **24** 024008

Worth N A and Nickels T B 2008 Acceleration of Tomo-PIV by estimating the initial volume intensity distribution *Exp. Fluids* **45** 847–56

Worth N A and Nickels T B 2010 A tomographic PIV resolution study based on homogeneous isotropic turbulence DNS data *Exp. Fluids* **49** 637–56

'Shake The Box': A highly efficient and accurate Tomographic Particle Tracking Velocimetry (TOMO-PTV) method using prediction of particle positions

**Daniel Schanz[1], Andreas Schröder[1], Sebastian Gesemann[1],
Dirk Michaelis[2], Bernhard Wieneke[2]**

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

[2] LaVision GmbH, Göttingen, Germany

ABSTRACT

A novel approach to the evaluation of time resolved particle-based tomographic data is introduced. By seizing the time information contained in such datasets, a very fast and accurate tracking of nearly all particles within the measurement domain is achieved at seeding densities comparable to (and probably above) the thresholds for tomographic PIV. The method relies on predicting the position of already tracked particles and refining the found position by an image matching scheme ('shaking' all particles within the measurement 'box' until they fit the images: 'Shake The Box' - STB). New particles entering the measurement domain are identified using triangulation on the residual images.

Application of the method on a high-resolution time-resolved experimental dataset showed a reliable tracking of the vast majority of available particles for long time-series with many particles being tracked for their whole length of stay within the measurement domain. The image matching process ensures highly accurate particle positioning. Comparing the results to tomographic PIV evaluations by interpolating vector volumes from the discrete particles shows a high conformity of the results. The availability of discrete track information additionally allows for Lagrangian evaluations not possible with PIV data, as well as easy temporal smoothing and a reliable determination of derivations.

The processing time of a not fully optimized version of STB proved to be a factor of 3 to 4 faster compared to the fastest methods available for TOMO-PIV.

1 INTRODUCTION AND MOTIVATION

Since its introduction by Elsinga in 2005, Tomographic PIV (TOMO-PIV) [1, 2] has been rapidly accepted as a reliable and accurate mean of 3D-flow measurements. Applications range from highly resolved measurements in air [3] to time-resolved measurements in water [4] and in air [5, 6]. Like nearly all three-dimensional measurement techniques, TOMO-PIV has to deduct the position in space of the used particle tracers from two-dimensional camera images. The use of an iterative approach to this reconstruction, using algorithms like MART or SMART [7, 8] that reconstruct particles as intensity peaks in a voxel space, allows for much higher seeding densities compared to other approaches, such as three-dimensional particle tracking, based on particle triangulation [9]. Using 3d-correlation methods after the reconstructions process ensures a robust deduction of velocity information from the data, reducing negative effects of ghost particle as long as their intensity is below the real particles' intensity.

However, some drawbacks are associated with the technique: Ghost particles will always have influence on the vector result, especially when using high seeding densities. Furthermore, results gathered from cross-correlation represent averages over interrogation volumes and therefore smooth out velocity gradients and fine flow structures. This effect might be overcome by the use of adaptive weighting in the correlation process [10] Another downside is the large amount of computational time needed for the data processing, as well as large amounts of data that need to be (at least temporarily) saved to hard disk. When dealing with time-resolved data, it is difficult to use information gained from other time-steps in the processing of the current one. Methods like 'Motion Tracking Enhancement' [11] do that, but at a high computational cost.

These considerations show that it would be desirable to move from the representation of particles as intensity clusters in a huge voxel-space to direct knowledge of particle positions in space. Tracking such particles in time enables precise velocity determination, without the need of a spatial average. Lagrangian measurements would easily be possible. The computation time could probably be reduced, as the amount of data to be processed is dramatically reduced compared to a voxel space. Three-dimensional Particle Tracking Velocimetry (3D PTV) [9] does exactly that by triangulating particles in each time-step and then trying to find matching particles in the different time-steps. However, the triangulation process is limited by seeding density, allowing only about an order of magnitude fewer particles compared to TOMO-PIV.

The method of 'Iterative reconstruction of Volumetric Particle Distribution' (IPR), recently introduced by Wieneke [12] overcomes the problem of limited particle density: in contrast to conventional triangulation methods, an iterative approach of particle placement is applied, which allows to process particle numbers that are comparable to typical TOMO-PIV experiments. The working principle is to compute a distribution of discrete particle positions by iteratively adding particles, refining their position by moving ('shaking') the particle around in small steps, until an optimum is found in the particle projection relation to the original images (an image matching approach). Using this method, highly populated particle distributions can be reconstructed on a particle basis. Wieneke created voxel spaces from the gained particle distribution and showed via 3D-correlation, that the results of IPR are comparable to those of TOMO-PIV. Still, the obtained particle distributions exhibit the problem of ghost particles, possibly interfering with tracking processes. Due to the iterative nature, the processing time of IPR showed to be comparable to a tomographic reconstruction.

The method introduced in this paper combines the IPR method with an effective way of seizing the time-information in time-resolved PIV measurements. By this a method is created, which allows a very fast processing of highly seeded three-dimensional data, while capturing the movement of the vast majority of real particles and creating virtually no ghost particles. The key step is to produce a prediction of the particle distribution in the currently processed step, using extrapolation of existing particle tracks. This predicted particle distribution is used as an initialization to the IPR process and allows a severe reduction of iterations and therefore processing time. Willneff [13] also used the prediction of particle positions in space as a mean of improving particle tracking results, albeit only to close occurring gaps in conventionally created tracks using triangulation.

The general working principle of the method is given in paragraph 2, more detailed explanations are given in paragraph 4, where the application to experimental data is illustrated.

2 THE 'SHAKE THE BOX' METHOD

Conventional methods of evaluating highly seeded three-dimensional particle-based measurements rely on an individual treatment of every single snapshot of the particle distribution:

Applying TOMO-PIV, a tomographic reconstruction of every time-step is performed, with a subsequent correlation of two consecutive voxel spaces. The IPR-method computes the particle distribution from scratch for every snapshot, requiring many iterations until converging to the solution. When dealing with non-time-resolved data, typically obtained by low-repetition rate double-frame cameras, such an approach seems reasonable, as only two frames with closely related solutions to the reconstruction problem are available.

As soon though as the data at hand is sufficiently time-resolved, the approaches based on strictly singular image processing neglect the possibility of utilizing already processed data to extract information on the currently processed step. Therefore, the evaluation of a series from time-resolved experiments proves to be a lengthy process, typically taking weeks or months to process on a modern computer cluster.

The method presented here seizes the time-information by building predictions of the particle distribution and effectively refining this initial distribution by image matching, as described by Wieneke [12]. The method was termed 'Shake The Box' (STB) due to the procedure of producing an educated guess of the particles within the 'box' (measurement domain) and then 'shaking' the particles around, until an optimal distribution is reached. STB aims to be efficient in terms of calculation time, memory requirements and hard disk space, as well as being highly precise in respect to the investigated data and producing widely usable data output.

Moving from a huge voxel-space to discrete particle positions during the reconstruction phase has several advantages: The amount of main memory needed by the voxel-space can become quite significant; saving the data to hard disk is time consuming and can use vast amounts of space for long time-series. When correlating two voxel spaces both have to be loaded into the main memory again, requiring read time as well as double the amount of memory. When looking at the experimental data, presented later in this paper, we see that the used voxel space consists of around 2000x2000x400 voxels, equating to 12 GB of data. Saving the reconstructed volumes to disk for later use takes up around 1 GB per volume (reduction due to data compression) – for the 3000 subsequent images of one run, a total amount of 3 TB is needed.

In comparison, the representation of particles via positions takes up very little space: Typically 6 to 10 values (coordinates, velocities, intensity and other parameters) are associated to a particle per time-step. This means that typical numbers of particles (50.000 to 100.000, depending on the camera system and volume size) can be stored in less than 5 MB of RAM. Writing and reading from hard disk is very fast and the total amount of data for a time-resolved run is normally a few GB.

When particle positions are known for single snapshots it is a relatively easy step to find matching partners in successive time-steps (as the volume is very sparse compared to the images) – therefore tracking the particles in space is possible. The knowledge of particle tracks, spanning over multiple time-steps, allows further processing of the data: Fitting the positions in space with suitable functions (polynomials or splines) for a certain amount of time-steps allows accounting for noise introduced during the process of particle position identification. Derivations of the velocity (e.g. accelerations) can be computed with better accuracy from such fitted data.

Another feature of reliably tracked particles plays a key role in the STB-method: It is possible to extrapolate the particle position with quite good accuracy for the next, unprocessed, time-step. Fitting polynomials to the last few time-steps of the known particles paths yields a good approximation of the particle distribution in the current step. As the particles do not move steadily, there will still be errors in the particle placement, but these are mostly small (typically less than a voxel). Such a particle distribution, which is already very close to the real distribution, is a very good starting point for the IPR-method. Using this image-matching scheme it is possible to refine the particle placement until the error falls below a desired threshold. The computational cost of this process is much lower compared to the effort of computing the particle distribution from scratch, which is normally done for every snapshot. Furthermore, the process of finding matching particles in the new distribution can be omitted, as the partners are directly known.

In every step, a certain amount of particles will leave and a (normally similar) number of particles will enter the interrogation area. Particles leaving the AOI can just be removed from the tracking process and their tracks will end at the volume border. Newly emerging particles will have to be identified and eventually integrated into the tracking process. As the number of new particles is small compared to the total amount of particles (which are already tracked and their positions predicted for the current step), it is relatively easy to determine the 3D coordinate of new particles. A normal triangulation process is able to identify such particles, as their density is low enough to allow the triangulation to work reliable (this is not the case for the original seeding density, which is far above the threshold for non-iterative triangulation methods). Newly found particles that reoccur in a certain number of steps can be treated as identified particle tracks and be integrated into the prediction process.

In order to work efficiently, the algorithm requires a certain fraction of particles to be tracked; otherwise the triangulation process will yield too many particles, resulting in many ghost particles, difficult identification of new tracks and increased computation time. It is therefore very beneficial if some kind of initialization is done, providing enough track information for the first images of a run to allow an adequate prediction of the particle distribution of the next time-step.

A suitable track-initialization can be attained in several ways: On one hand, IPR could be used to process the first steps of a time-series. These time-steps would receive a full treatment with the algorithm, assuring a reconstruction close to the maximum quality obtainable by the method. It is sufficient to treat five to ten time-steps, in which a particle tracking algorithm would detect connected particle trajectories. In most cases, the examined flow is turbulent enough to avoid ghost particles moving along the flow for more than two or three time-steps, so that it can be assumed that mostly real particles are tracked, as soon as the track-length exceeds e.g. four time-steps.

Another method to obtain a track-initialization is to use voxel spaces, produced by tomographic reconstructions within the TOMO-PIV method. After several iterations of the reconstruction algorithm the intensity distribution (roughly) resembles particle shapes within the volume. A three-dimensional particle identification method (e.g. a Gaussian peak fitter) can be used to identify particle candidates within the voxel space. These positions can then again be treated by a tracking algorithm, identifying connected particle tracks and effectively removing ghost particles. Such methods were already used to obtain Lagrangian statistics from the flow and to judge reconstruction quality [6, 14]. The application of STB to experimental data, presented in paragraph 4, used an initialization relying on particle identification in voxel spaces – mainly because that data was already available. Using IPR to create an initialization should be equally successful.
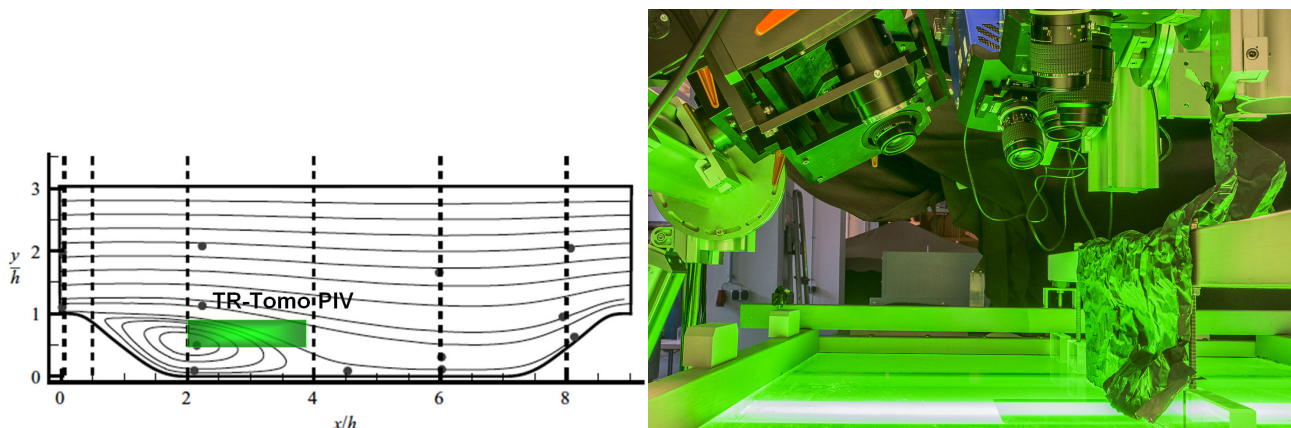
Figure 1 (Left): Schematic of the channel with hill height h = 50 mm configuration (distance between hills Lx = 9h = 450 mm), submerged in the water tunnel facility with a channel height of Ly = 3.035h = 151.75 mm and a width of 900 mm (based on test case ERCOFTAC 81, for further information see [15])
(Right): Six high speed cameras (four cameras in Scheimpflug arrangement) imaging a light volume within the water flow below.

Using the particle tracks for time-steps 1-n gained from the initialization it is possible to predict the particle distribution for time n+1 and from there on iterate the depicted algorithm, extending the known particle track from step to step, optimizing the new positions by shaking (image matching) and successively adding newly found particle tracks. The STB-method works its way through a time-series, always creating the information needed to effectively process the next time-step by refining the result of the currently processed step.

A more detailed description of the methods' application to an experimental dataset - introduced below - is presented in paragraph 4.

3 EXPERIMENTAL SETUP

An experiment, conducted within the scope of the EU-FP7 project AFDAR ('Advanced Flow Diagnostics for Aeronautical Research'), will be used to demonstrate the applicability of the STB-method to real experimental data.

The experiment took place in the water tunnel facility at Technical University of Munich. The flow behind a series of identical longitudinal hills ('periodic hills', ERCOFTAC test-case 81 [15]) was investigated using a high-speed tomographic PIV system. Six Imager pro HS 4M (PCO Dimax) cameras were used to observe a measurement volume of 80x80x20 mm^3, located 2 h (100 mm) downstream of the seventh hill (the experiment uses of a total of ten consecutive hills). Wall-normal height spans from 25 to 45 mm in order to capture the shear layer (see Fig. 1). The water was seeded using ~30 μm polyamide particles. Illumination was realized using a Quantronix Darwin Duo continuous laser, provided by UNIBWM. The laser beam was widened by two successive telescope optics using cylindrical lenses, resulting in an oval light profile. The profile was cut in rectangular shape by a passe-partout that was fixed at the side wall of the channel. This volume light sheet passes through the interrogation volume and is back-reflected into itself using an end-mirror located directly outside of the opposite wall of the tunnel [5]. A second passe-partout is installed there. This setup enables all cameras to be in forward scattering and thus gather a maximum of light. In order to assure sufficient contrast for the imaged particles, a sheet of black adhesive foil was installed below the illuminated area.

Due to space restrictions four cameras were placed in line, whereas two of them observed the measurement volume in an off-axis arrangement. The four outermost cameras had to be equipped with Scheimpflug-adapters due to their viewing angle relative to the measurement volume. Five cameras used 105 mm Nikon Micro Nikkor lenses, while one camera used a 100 mm Zeiss Distagon Macro lens. The strong scattering of light by the particles allowed closing the apertures to $F_\#$ = 22, minimizing particle blurring effects due to limited depth of field or astigmatisms. An average resolution of approx. 21.5 pixels per mm was achieved.

Calibration was done using a 3D-calibration-plate, providing two planes of calibration markers, thus requiring no wall-normal movement of the plate. A carrier was constructed, securing the plate firmly between the two neighbouring hills (see Fig. 2). The carrier was fixed in spanwise direction using strong magnets at the side walls of the tunnel. Small, inevitable errors of the calibration were corrected by applying the method of volume-self-calibration (VSC) [16] to the particle images. Back-projection errors of around 1 pixel were found and corrected to values below 0.02 pixels by this method. Alongside the VSC, a calibration of the optical transfer function (OTF) [14] was performed, gathering the averaged particle imaging form different areas of the measurement domain on the camera images. As indicated by Wieneke, the use of a calibrated OTF is very beneficial to the accuracy of particle placement using IPR.
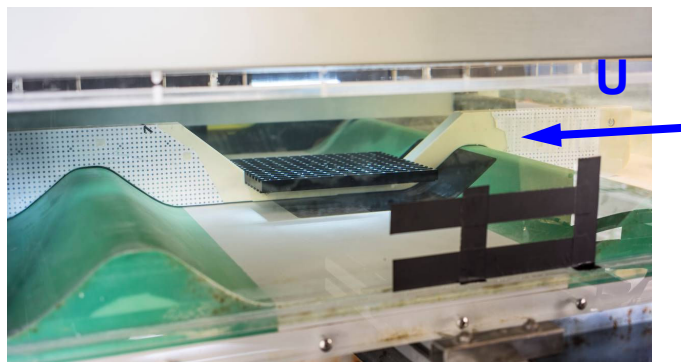
Figure 2: Two-plane calibration target positioned in the middle of the test section on a carrier with negative hills and distance holders to the side-walls. Upper plane located at y = 35 mm.
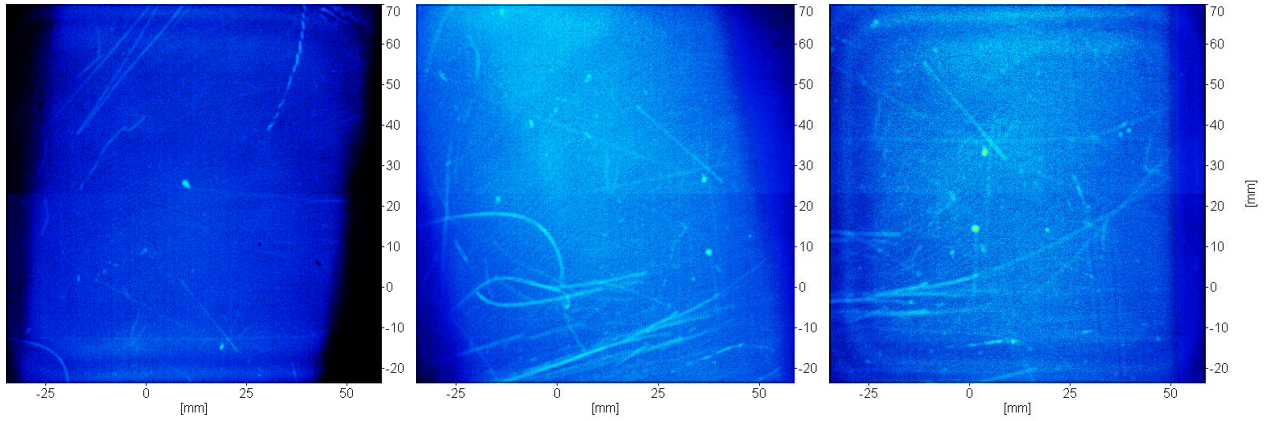
Figure 3: Minimum images, computed over a run of 3.000 images for three of the six used cameras. Scratches on the plexiglass surface are clearly visualized, bubbles sticking to the plexiglass surface can be spotted (see right figure).

The cameras observed the interrogation volume from the top of the channel through a plexiglass plate. Due to the experimental setup being in use for several years, this plate showed small scratches, which could potentially interfere with the successful particle reconstruction in certain regions of space. Additionally small bubbles were sometimes produced by the flow mechanism, which passed above the interrogation volume from time to time or even sticked to the plexiglass plate through which the cameras were observing the measurement area.

To visualize the impact of these viewing obstacles on the camera images, Fig. 3 shows the minimum images (minimal intensity of all images over a run of 3000 images). A multitude of scratches of different sizes can be seen, as well as two stationary bubbles. In normal particle images the scratches are visible as regions of unsharp imaging or as sources of particle image displacements, bubbles are regions of totally obstructed particle imaging.

Two flow speeds, corresponding to Re = 8.000 and Re = 33.000, were measured at a repetition rate of 500 Hz and 1000 Hz, respectively. Results shown in this paper were obtained from a run at Re = 8.000. Due to the low fluid velocity, the flow is well resolved temporally: On average, the particles move approx. 6.0 voxel between successive frames for Re = 8.000 and 8.6 voxel for Re = 33.000.

In addition to particle tracking by application of STB, conventional TOMO-PIV evaluations were carried out. Tomographic reconstruction from the camera images (each with a resolution of 2016 x 2016 pixel) was performed using a SMART algorithm with MLOS initialization - yielding voxel spaces of 1951 (flow direction) x 2035 (spanwise) x 405 (wall normal) voxels. Using 3D-correlation, vector volumes with 163x170x33 vectors could be realized, the final size of an interrogation volume (48^3 voxel) is approx. 2x2x2 mm^3. Instantaneous vector volumes for both cases show highly three-dimensional flow (see Fig. 4).
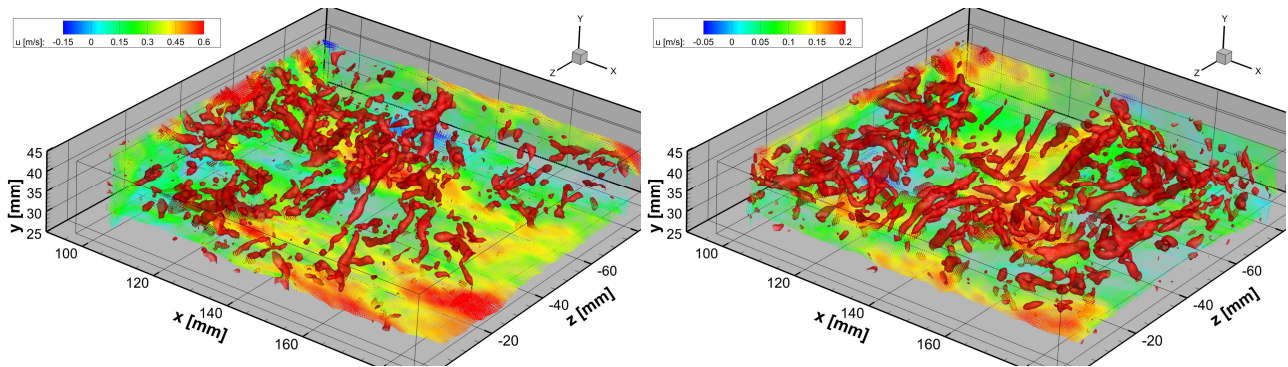
Figure 4: Instantaneous TOMO-PIV result for Re = 33.000 (left) and Re = 8.000 (right). Isosurfaces of 3D swirl strength λ2 (different thresholds for the two plots). Additional vector slices, color-coded by streamwise velocity (u). Flow in positive x-direction. Please note that the coordinate system was turned and shifted, compared to other images below.

4 APPLICATION TO EXPERIMENTAL DATA

The data gained during the experiment described in the last paragraph poses several difficulties to an evaluation method: The scratched surface and occurring bubbles lead to areas within the measurement volume, where not all cameras have recorded information on the particles present. The algorithms used in tomographic reconstruction tend to not reconstruct any particle information at all in such regions, as they compute the voxel intensity as a product of the intensity seen by all cameras. The IPR method on the other hand could successfully determine the particles within such regions, as particle triangulation is done both with a full set of cameras, but also with certain cameras missing. If a particle is obstructed only in one camera, the reduced triangulation process will still pick it up. STB could track a particle, even if it is not visible in more than one camera. The position prediction should be sufficiently accurate, such that the image matching process should be able to position the particle based on the information of the unobstructed cameras. It is however possible, that wrong peaks in cameras viewing, e.g. a bubble image, draw the particle away from its correct position, especially if it is obstructed for multiple time-steps.

The particle images on the cameras are of very different intensities, as additionally to the used ~30 μm polyamide particles, seeding residuals of earlier measurements, as well as dust/dirt-particles were present, all with different light scattering properties.

The STB-method was applied to a run at Re = 8.000, recorded at a frequency of 500 Hz. In order to produce a track initialization for the experimental data, it was chosen to use existing tomographic reconstructions of the first five time-steps of this run. The particle distribution was reconstructed using a SMART algorithm with MLOS initialization with inclusion of the calibrated OTF into the reconstruction process by an appropriate parameterization of the weighting functions [14]. To ensure optimal quality 15 iterations of SMART were performed. Reconstruction times were approximately 30 minutes per volume on a 48-core cluster.

Particles within these volumes were identified using a 3D particle peak detection algorithm of Lavision Davis 8.1.2. A threshold of 1.000 counts (of a maximum of 64.000 counts) was applied, above which it was assumed that a particle was found. A Gaussian 5x5x5 fit was used to determine the sub-voxel accurate position of the particle. It was necessary to choose such a low threshold, as the very inhomogeneous particle image intensity is reflected in the reconstructed volume, leading to real particles having a low intensity. It was clear, that a lot of ghost particles would be identified for every single step. Around 750.000 particle candidates were found within each volume, which is considerably more than the actual number of particles: Counting the particles on the camera images with a peak fitting method shows that only around 75.000 particle images are present on the active image area of the camera with the lowest particle count (differences in cameras due to different viewing- and Scheimpflug-angles). Considering the used image area of around 2.5 Megapixels per camera, the effective seeding density is around 0.03 – 0.04 ppp (particles per pixel) for the different cameras. The fact that 90 percent of the identified particles were ghost particles illustrates how prominent the ambiguity problem is in tomographic measurements, even when a high quality reconstruction is applied. However, as the examined flow is highly turbulent, occurring ghost particles decorrelate within few frames, therefore it is possible to separate real particles from ghost particles by searching for continuous tracks. A particle tracking algorithm was applied to the found particle distributions, searching for tracks with at least four steps within the five processed snapshots. By additionally filtering for sufficiently smooth tracks, the number of identified tracks was around 66.000 – a number close to the maximum number of expected particles (75.000) and a good starting point for the STB method.

The found tracks for t = 1 – 5 were used to construct a prediction for the particle distribution at t = 6. To this end, a polynomial of order n is fitted to at least n+1 previous particle positions and extrapolated to the next time-step. In the beginning, it was chosen to use a polynomial of order 1 (linear) to fit the four previous time-steps. The particles intensity is set to the average intensity of the last four steps of the same particle. It seems obvious to use higher-order fitting on more predecessors as soon as the processed track exceeds a certain length. However, for this first examination of the method, the linear fit on 4 preceding particle positions is kept during the whole process.

The extrapolation to t = 6 yields new particle positions that are back-projected to virtual camera images. When subtracting these virtual images I_{proj} from the original recordings I_{orig}, small deviations of the predicted particle positions from the real position become obvious in the residual image I_{res}. The IPR algorithm then tries to minimize the residual R by moving the particle in small steps in space ('shaking' the particle). R is computed as the difference of the particle-augmented residual to the projection of the current particle position: $R[x',y',z'] = I_{res+p}-I_{Part[x',y',z']}$, where $I_{res+p} = I_{res} + I_{part}$ (I_{part}: projection of the particle being processed using its initial coordinates $[x,y,z]$; $I_{Part[x',y',z']}$: projection of the particle being processed using its modified coordinates $[x',y',z']$). R is always computed as the sum of all pixels in a certain neighborhood (typically around the size of the particle image diameter, as given by the OTF-calibration) of the particle projection point of all cameras. Please see [12] for more details.
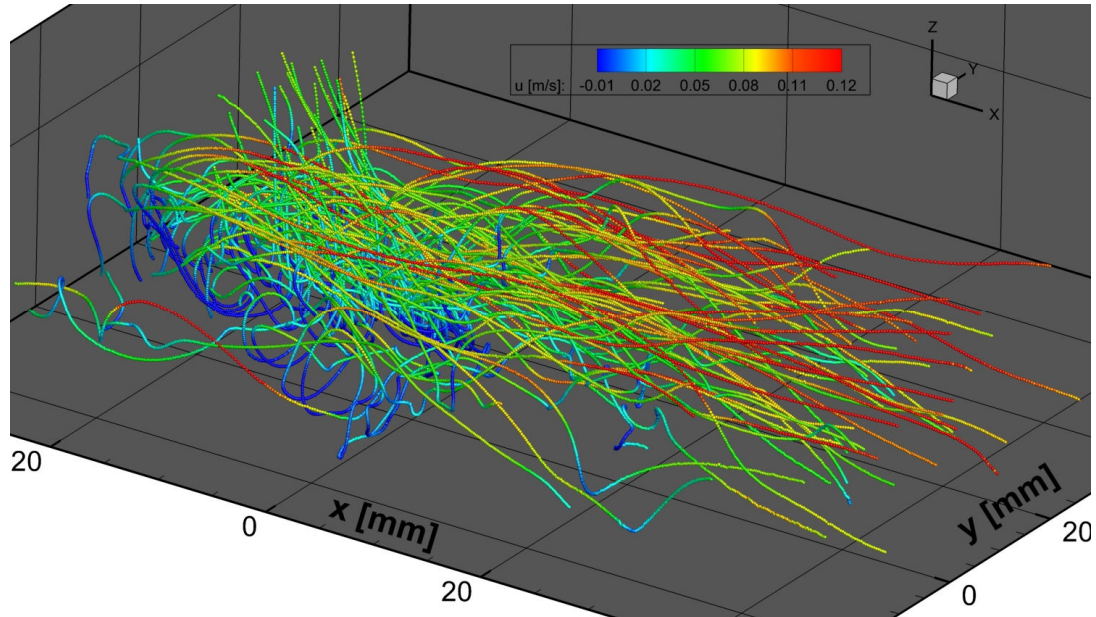
Figure 5: Particle Tracks with lengths between 596 and 600 steps. These particles were detected by the track initialization and have been tracked over the whole run of 600 images. 74 tracks are shown. Flow in positive x-direction

In this case, the algorithm is applied in the following way: The particle is moved in x-direction to the positions $x' = p_x-0.2, p_x-0.1, p_x, p_x+0.1, p_x+0.2$; residual R is calculated for all positions. A polynomial of order 2 is fitted to the five values of R and the new x-position of the particle is set to the minimum of the fitted function. The procedure is repeated for y- and z- direction. Therefore, corrections of up to 0.2 voxels in each direction of space are possible per iteration. The shaking-process is iteratively applied until I_{res} cannot be minimized further.

Alongside the shaking of the particles to their correct positions, the particle intensity i_p is iteratively updated to best fit the intensities found in the original recordings for the current step: $i_p' = i_p \cdot$ sqrt($\Sigma_k I_{res+p}/ \Sigma_k I_{part}$), where k signifies all pixel on all cameras influenced by the current particle (again following Wieneke [12]).

New particles – those who are entering the measurement domain, as well as previously untracked ones – can be identified via triangulation on the residual images I_{res}. On I_{res} peaks are visible that either belong to new particles or to already captured particles whose intensity does not completely match the recorded intensity of all cameras (this is often the case, as the light scattering behavior is complex and cannot completely be calibrated – especially when dealing with particles of different sizes). However, only new particles will be successfully triangulated: particles that are already tracked, are given an intensity that represents the average of all cameras. Therefore, when there is a residual peak belonging to such a particle in a number of cameras, the other cameras must have zero (actually negative) residual from the particle, and it will not be triangulated. An allowed triangulation error of $\varepsilon = 1.5$ pixel was chosen. This value allows for small deviations caused by overlapping particles, but does not lead to an excessive amount of particle candidates. The intensity is initialized with the minimum of the OTF-weighted intensity of the particle image on all cameras.

As already described by Wieneke [12], the triangulation process can be repeated with certain cameras left out. By this, effects of particle overlap can be reduced. In the experimental data assessed here the additional advantage is that the effects of scratches on the surface and air bubble can be reduced.

Especially in the first iteration the triangulation process identifies quite large numbers of particle candidates: On average, the first execution of the triangulation after the prediction finds around 4.000 - 8.000 particles, depending on image preprocessing. The combined triangulations with one camera missing can add another 10.000 to 20.000 candidates. Only such particles that are not within a distance of 1 voxel of an existing particle are accepted. As still a large amount of ghost particles is present in the newly triangulated particles, these are written to a separate particle list and are not yet considered in the prediction process. Comparable to the creation of the track initialization, new particles are only accepted if four consecutive occurrences of a particle are identified by the tracking algorithm. To help with the identification of matching partners, an estimated velocity is computed from neighboring tracked particles. A search radius of 7 voxels is applied to the point given by the estimator. After this process, around 500-1.000 new particle tracks are identified for each step. The rest of the newly triangulated particles is discarded as ghost particles.

The IPR-method as described in [12] always reconstructs all particle positions without the use of an initialization. Therefore, the number of iterations needs to be high: 8 triangulation processes (n_1), as well as 8 triangulation processes with reduced camera number (n_2) are carried out, each of these followed by 6 iterations of particle-shaking. Due to the prediction process in the STB-method, the number of iterations can be significantly reduced: for the results presented later in this paragraph, values of $n_1 = 2$ and $n_2 = 1$ are used, each followed by 4 iterations of particle-shaking.
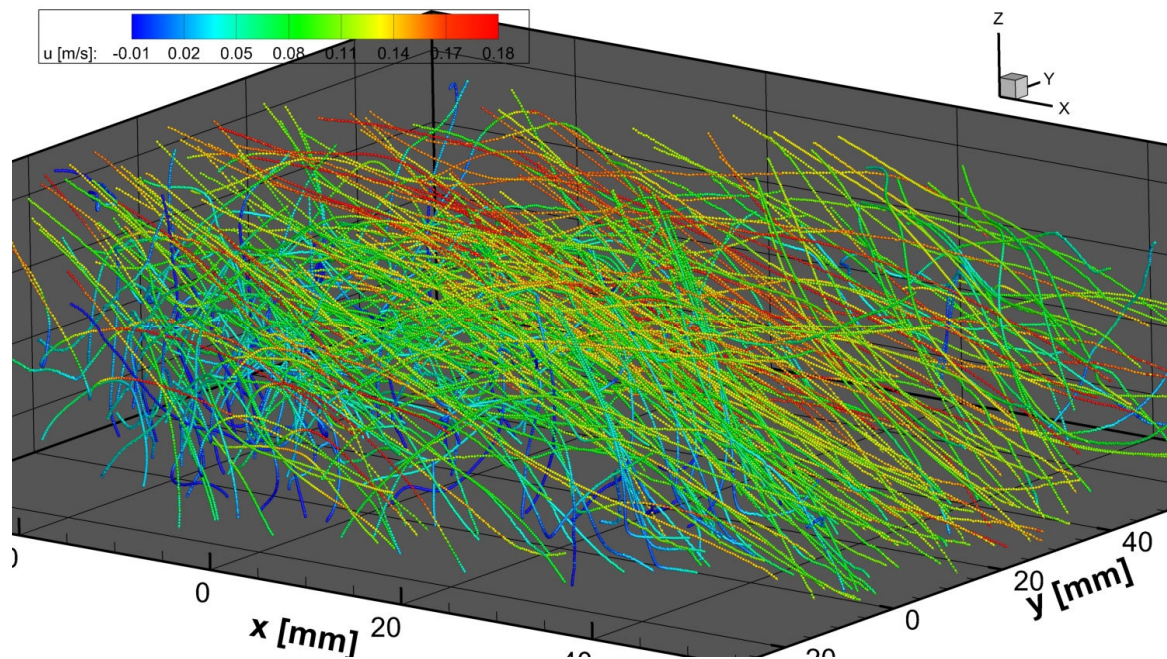
Figure 6: Particle tracks with length 200. 338 Tracks out of a total of 450.000 are shown. Flow in pos. x-direction.

If the shake process puts a newly triangulated particle within a radius of one voxel of a tracked particle, it is deleted to avoid effects of multiple particles representing one real particle. If, on the other hand, two tracked particles come that close to each other, they are kept in the particle distribution - when looking at whole tracks, the likelihood of particles passing close to each other is much larger as when looking at a single time-step. Particles are deleted from the current time-step, if their intensity falls below a certain threshold. For newly triangulated particles this threshold was chosen to be 10 percent of the average particle intensity i_{avg}. For tracked particles a threshold of 2 % of i_{avg} was applied. As the intensity of the real particles varies significantly (due to different particle sizes), the low threshold was needed for tracked particles in order not to delete real particles. A particle is completely deleted from the tracking process either if it leaves the measurement domain or if the 2%-threshold is undercut for two successive time-steps. In this case it has to be assumed that the particle was drawn to a wrong position at a certain point and the position prediction is not accurate enough anymore. If the particle still exists (which it should) it will be picked up by the triangulation process and a new track will be formed for the same particle. On average 300-550 particles leave the AOI per time-step and 250-500 tracks are ended because of too little particle intensity. After around 80 time-steps equilibrium of newly detected and ended tracks is reached and the total number of tracked particles does not change significantly anymore. This equilibrium number of tracked particles (N_P) depends on a number of factors, but is especially sensitive to the preprocessing applied to the camera images. For a preprocessing typically done for TOMO-PIV evaluation (sliding minimum subtracted, normalization with local average, Gaussian smoothing and sharpening - called 'PP1' in the following) $N_P \approx 55.000$ was reached. A more conservative preprocessing (Subtract minimum image, Gaussian smoothing and sharpening, subtract a constant of 15 counts – called 'PP2' in the following) resulted in $N_P \approx 71.000$, which is already close to the maximum number of expected particles. It seems that PP1 erased real particle information from the images, even though some of the peaks retained in PP2 are of a magnitude comparable to remaining noise.

Reconstruction times per time-step were around 2 minutes for PP1 and around 3 minutes for PP2 (due to higher number of peaks for the triangulation and higher total number of particles) when using a 48 core cluster. The triangulation process can be optimized further, therefore some performance gains should still be possible. To put these numbers into perspective: The currently fastest methods of TOMO-PIV evaluation take around 9 minutes per time-step on the same cluster (Davis 8.1.2 FAST MART method using 3 iterations for reconstruction (around 2.5 minutes) and Davis 8.1.2 3D Direct Correlation with volume rescaling (around 6.5 minutes for correlation volumes of 48^3 voxels with 75 % overlap)). Using standard methods, such as 6 iterations of SMART with MLOS initialization and conventional FFT-based 3D Correlation total processing time is around 40 minutes per time-step (10 minutes SMART, 30 minutes correlation).

Two reconstruction runs were performed, using PP1 (600 time steps) and PP2 (150 time steps). Tracks in a wide range of lengths were identified. Fig. 7 shows the distribution of track lengths for both runs. The 600-step run identified a total number of around 450.000 tracks, with a peak track-length of approximately 31 time-steps. Tracks around this length are typically tracks of low-intensity particles, which tend to be lost easier due to overlap with higher intensity particles or due to random noise peaks. The track of such particles gets picked up by the track finding algorithm some time-steps later, therefore multiple track entries of shorter length exist for one particle.

a)

b)
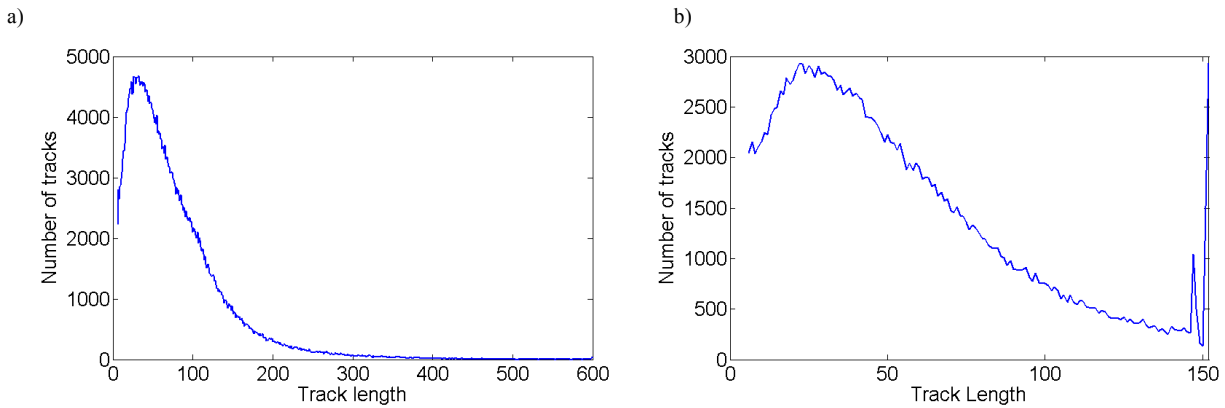
Figure 7: Distribution of track lengths for STB-runs over 600 steps using PP1 (a) and 150 tracks using PP2 (b).

The 150-step run identified 210.000 tracks with a peak in track lengths around 25 steps. This shows that the run using less aggressive image preprocessing was able to identify more particles, but these were on average tracked for a shorter time, as the noise level is already near to the intensity of these additional particles (their tracks will consist of more sub-tracks). As a future development of the algorithm, a method to identify separate track fragments and connect them - as proposed by Willneff [13] - is feasible.

Alongside low intensity particles with multiple track fragments, a large amount of particles is reliably tracked through the interrogation volume. Looking at the 600-step run a total number of 117.000 tracks of length 100 or higher was found. 74 tracks were recorded with lengths between 596 and 600 time steps (see Fig. 5). These are particles that were first identified in the track initialization process and never lost during the whole run. The vast majority of the particles resides in the measurement volume for a shorter time than 600 images (corresponding to 1.2 seconds of measurement time). Particles mostly enter the measurement domain either from the upper or the upstream border. The mean flow is directed downwards, therefore particles travel only an average distance of 3 to 5 cm before leaving the domain again. Fig. 1 illustrates the flow topology. As the main flow velocity is around 0.2 m/s, most particles pass the interrogation volume in 100 to 200 time-steps. Particles with very long track lengths mostly reside in low speed streaks, where particle can be observed much longer. Fig. 5 illustrates that behavior: most of these very long tracks start in a low-speed region that was present in the lower upstream corner of the measurement domain at the beginning of the run. The particles spiral around the low-speed structures, while slowly moving downstream. Most of them are drawn up at some point, where they are accelerated by higher-speed streaks. Other tracks can be seen that start nearly at the top of the volume but are quickly drawn into the low-speed region and reside there for a large number of time-steps. Only for particles, whose tracks started to the right (looking in stream-wise direction) of the volume, very long tracks were found. The left side of the volume was dominated by a high-speed streak, therefore all particles found there left the volume before the end of the evaluation run.

Fig. 6 shows tracks with a length of 200 time-steps, 338 of which were found within the dataset. Different kinds of particle movements are visible: many particles are following the main flow direction - entering the volume at the top or upstream side and leaving at the bottom or the downstream border. These are typically particles with relatively high velocities, following the downward streaks. Other particles reside in low speed regions, much like the particles in Fig. 5. They mostly leave the volume at the bottom border, being transported into regions near the wall, probably with even lower velocities. Particles with track length 200 were found all over the volume, as is expected. The tracks shown in Fig. 5 and Fig. 6 were visualized using the raw data coming from the reconstructions – no smoothing or fitting in time was applied.

Fig. 8 illustrates the total amount of particles tracked within one time-step, showing step 150 of the run using PP2. Viewing from the top, around 71.000 particles are displayed with color coded stream-wise velocity. Flow structures can be identified, with a strong streak of fast fluid at the left side of the volume (looking stream-wise), while the right side is dominated by a low speed region. This flow topology was already apparent when looking at Fig. 5, showing that all very long tracks originate from the right side of the volume, where particle speeds are slow.
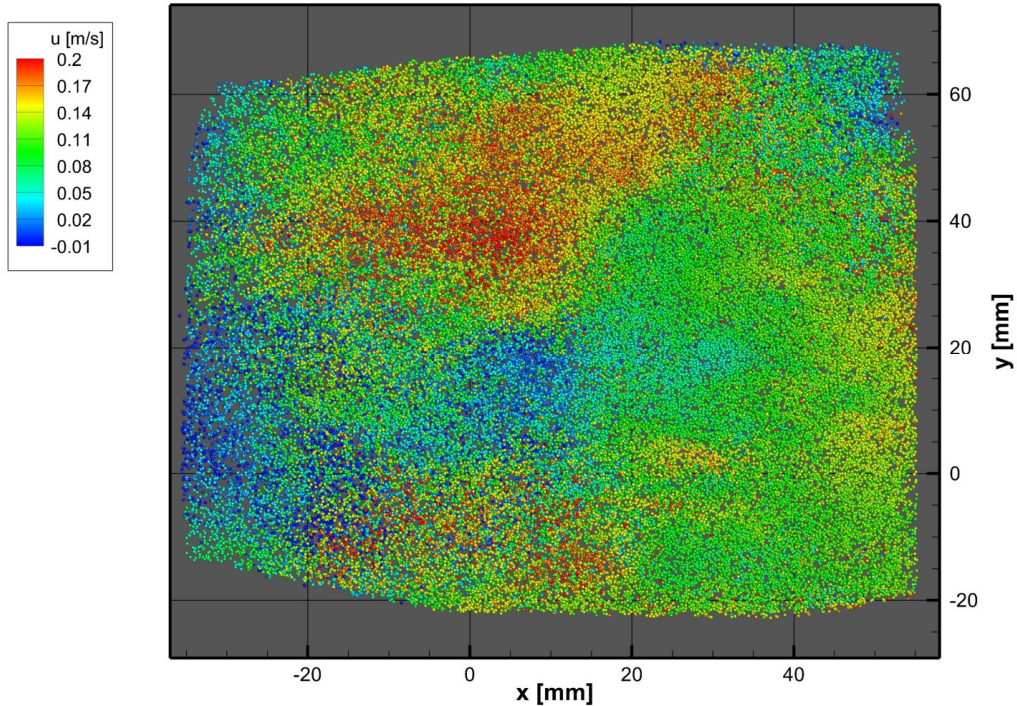
Figure 8: 71.000 Particles identified for time-step 150 of an STB-run of length 150, using images preprocessed with PP2. Color coded streamwise velocity. Flow in positive x-direction.

In order to compare the quality of the reconstructed particle distributions, vector volumes were created, using interpolation between adjacent particles. These vector volumes can be compared to similar ones, gained by classical TOMO-PIV (reconstruction and 3D Correlation). Vectors were computed by averaging the velocity values of all particles found within a distance of 48 voxels to the given point in space. The contribution of a particle to the velocity average is weighted with a Gaussian profile in dependency on the distance to the vector point. The vectors were calculated on a three-dimensional with 12 voxels spacing, conforming to a cross-correlation with 48^3 interrogation cubes and 75 % overlap. Comparisons were done to TOMO-PIV data, originating from MLOS-SMART reconstructions with 15 iterations and two different cross-correlation schemes (Davis 8.1.2 FFT-based 3D cross correlation with outlier detection, as well as 3D Direct Correlation). Both runs, using PP1 and PP2 preprocessed images, were included. Fig 8 shows vector volume results of time-step 150 for the four mentioned evaluations. After 145 independent time-steps it can be assumed that the STB method operates completely detached from the initialization and the results are representative for the general quality to be expected.

Looking at the vector volumes it can be said that the results are very much alike. The general flow topology is gathered equally by all methods; also fine-scale structures are mostly reproduced similarly. The TOMO-PIV evaluation using FFT cross-correlation (Fig. 9b) shows a few visible signs of remaining outliers. This method is the only one that does not use a Gaussian weighted correlation approach and is therefore the most susceptible to inhomogeneous seeding densities. Both 3D Direct Correlation and the vector-determination method used on the STB-data, use interrogation windows that are actually wider than 48^3 voxels, albeit with very low weights for particles farther away. This method decreases the probability of outlier creation, as well as increases the accuracy when particle are close to the vector point. The Direct Correlation additionally prevents outliers from occurring due to multiple steps of resampling the volume.

Both correlation results show a slightly higher dynamic range compared to the STB-results. This finding is a bit surprising, because it was expected that the correlations will smooth out gradients more than a particle-based approach. If this observation is due to the process of vector-deduction from the particle data will be examined in future steps.

The PP2 STB-result seems to be slightly noisier compared to the PP1-STB and the Direct Correlation result. It might be that the less aggressive image preprocessing left enough noise on the images to influence the results. On the other hand it is possible that smaller structures are captured due to the increased number of particles.

In summary, it can be said that both runs of STB produced results that compare well to TOMO-PIV data when interpolating vector volumes from the particle data. Additionally, the discrete track information is available, which allows ways of data analysis not easily obtained for PIV data (Lagrangian statistics, easy temporal smoothing, derivations) and gives a better capture of gradients.
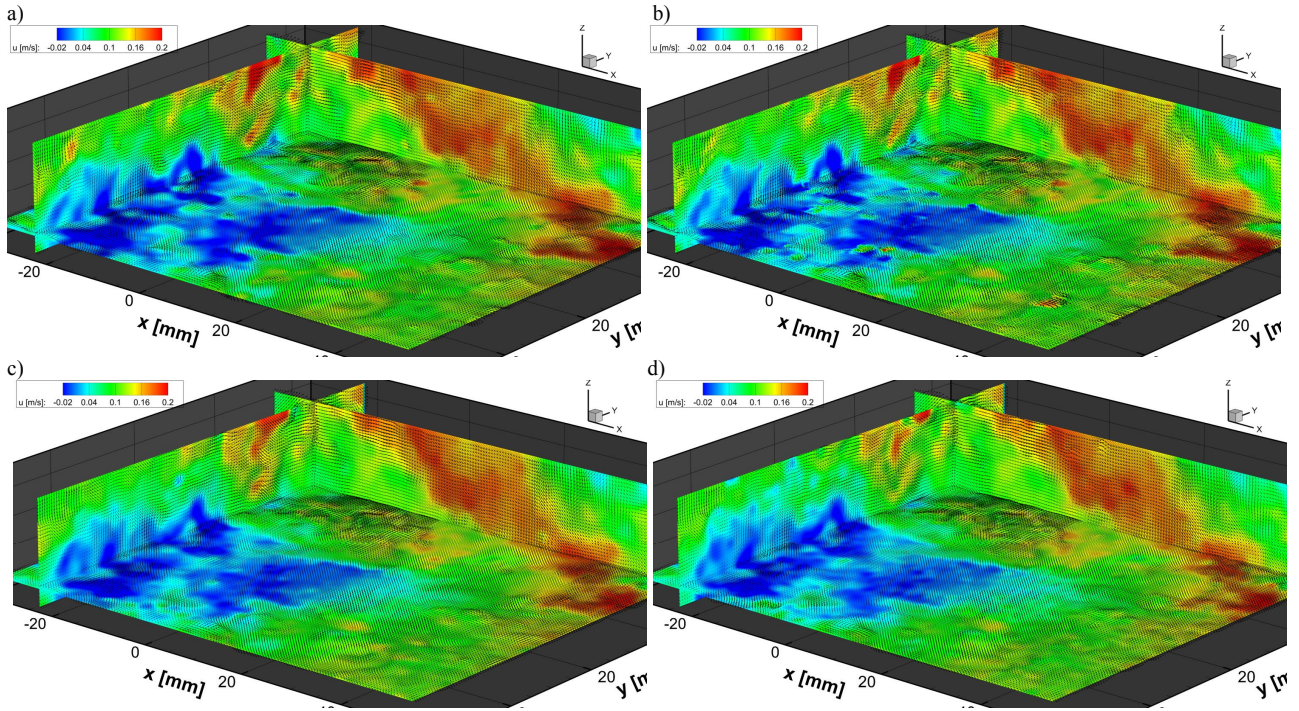
Figure 9: Comparison of vector results for step 150 out of a total of 3.000 of an experimental run with Re = 8.000.
a) SMART reconstruction, Davis Direct Correlation (Multigrid $128^3 \rightarrow 96^3 \rightarrow 64^3 \rightarrow 48^3$ interrogation volumes, 75 % overlap);
b) SMART reconstruction, Davis FFT 3D Correlation (Multigrid $128^3 \rightarrow 96^3 \rightarrow 64^3 \rightarrow 48^3$ interrogation volumes, 75 % overlap);
c) STB TOMO-PTV, images using PP1 preprocessing (54.000 particles in this step), vectors interpolated from discrete particles, velocity averaged for particles within a distance of 48 voxels, Gaussian weighting with distance
d) STB TOMO-PTV, images using PP2 preprocessing (71.000 particles in this step), vectors interpolated from discrete particles, velocity averaged for particles within a distance of 48 voxels, Gaussian weighting with distance

5 SUMMARY AND CONCLUSIONS

A new 3D PTV method of evaluating time resolved tomographic data is presented. The method allows very fast processing of data with (at least) the same seeding densities compared to TOMO-PIV and tracking the motion of the vast majority of particles imaged by the camera system. The tracking of ghost particles is effectively avoided by the creation of long particle tracks.

The method relies on the prediction of the particle distribution for the currently processed step by extrapolating the path of the particles already tracked. As long as the bulk of the available particles are tracked, this prediction yields a particle distribution that is very close to the real one. The method if 'Iterative reconstruction of Volumetric Particle Distribution' (IPR) [12] is used to refine the predicted particle position by moving the particle in small steps around the volume ('shaking' the particle), until the reprojection errors are minimized. The number of iterations needed is small due to the good prediction of the particle distribution. Particles newly entering the measurement domain are identified using triangulation on the residual images, which only show a low seeding density due to the images of tracked particles being reduced or completely erased by subtracting the particles' image. Newly triangulated particles are added to the list of tracked particles if the particle appears on at least four consecutive time-steps on a reasonable trajectory.

A particle is removed from the tracking process if it reaches the volume border or if it's intensity falls below a threshold for two consecutive time-steps (in this case it is assumed the track was lost due to adverse imaging conditions and it will most likely be picked up again by the triangulation process for new particles).

The algorithm was casually termed 'Shake The Box' (STB) because of the process of predicting all particles in the measurement 'box' and 'shaking' them into place by the IPR method.

In order to give the algorithm a good starting point it is beneficial to create a track initialization for the first few time-steps of a time series. This initialization can either be realized by applying IPR with a full set of iterations to ensure maximum quality or by performing tomographic reconstructions and identifying reconstructed particles in the voxel space. In both cases, a tracking algorithm finds connected particles in the considered time-steps, effectively getting rid of ghost particles. With a good initialization, the STB-method rapidly converges to equilibrium of newly found and lost

particles. The number of identified particles is close to the total number of particles within the measurement domain, if image preprocessing is applied that does not eradicate images of real particles.

Due to the reliable prediction of particle positions, the method performs very fast. The processing time of the used (not completely optimized) version of STB was a factor of 3 to 4 shorter compared to the currently fastest TOMO-PIV evaluation methods. When comparing to standard methods, the factor is around 14 to 20.

The STB-method was applied to an experimental dataset on the flow behind a series of periodic hills at Re = 8.000, recorded at a frequency of 500 Hz. On average, the particles move around 6 voxels between frames. A track initialization was created using existing tomographic reconstructions of the first five images. Two runs of STB were performed using different preprocessing settings for the camera images (length: 150 and 600 frames). It was shown that both runs converge to a stable number of tracked particles (55.000 for a rather strong preprocessing, 71.000 for a more conservative one). The maximum number of expected particles is around 75.000 (as shown by a peak finder applied to the camera images). Particle intensities of the investigated data vary significantly, thus a stronger preprocessing leads to a loss of real particles. The remaining particles are tracked reliable, however. The occurrence of multiple particle tracks for one particle is decreased, compared to the run using mild preprocessing. Many particles are correctly tracked for their whole length of stay in the measurement domain. The 600-step run showed 117.000 tracks with length above 100 steps, some particles were tracked for all 600 images.

A comparison to TOMO-PIV evaluations was done by creating vector volumes from the discrete particle distributions. These were compared to similar volumes calculated by state-of-the-art TOMO-PIV evaluation algorithms. Only very minor differences were found, showing that the quality of the data gathered by the STB-method is comparable to TOMO-PIV, even when giving up the greatest advantage – the knowledge of real particle trajectories – for space averaged vector representation.

The whole process of applying STB to experimental data showed a very reliable and stable behavior of the method. The results were less depending on the different parameters of the algorithm than on the input (namely the preprocessing used on the images). Considering the problematic quality of the images (scratches on the plexiglass interface, bubbles on the surface, varying particle intensities) this is a good sign for the easy applicability of the method on various datasets, as long as the time-resolution is sufficient.

6 OUTLOOK

The features of the STB-method that until now were captured only qualitatively using experimental data will be quantified using a synthetic dataset with known ground truth. STB exhibits some more advantages over other evaluation schemes not seizing the time information, which will be investigated in the future. The method should be able to deal with images with higher seeding densities, compared to TOMO-PIV or single-frame IPR. The knowledge of particle tracks enables an efficient reduction of the parameter space and reduces the 'effective' seeding density. Using synthetic data it will be assessed how STB compares to TOMO-PIV and IPR at high seeding densities. As indicated by Wieneke [12], the accuracy of particle placement using IPR (image matching) should be better compared to TOMO-PIV - at least up to seeding densities of 0.1 ppp. It will be assessed if this still holds true for the STB-method and if further improvements can be made when using high seeding densities. A method to reconnect track fragments originating from one particle, like proposed by Willneff [13], should be able to further increase the number of completely tracked particles.

The approach of the method itself can be varied: It is conceivable to not only start the track finding process at the beginning of the run, but also at multiple points within the time-series. From there, the method could run forwards and *backwards*, until the evaluation processes meet at some point. Most tracks should find a partner in the two processes, but there could also be tracks that were missed (or lost) when coming from one direction, which could be recovered by calculating from the other direction.

The application of the STB-method (or even IPR) to conventional two-frame PIV-data is doubtful, as no prediction is available and particle matching is difficult due to many ghost particles. A dual-volume TOMO-PIV system, as presented by Schröder et. al. ('Dual-Volume and Four-Pulse Tomo PIV using polarized laser light', Contribution 148 for PIV 13), however lends itself to be treated by some kind of tracking procedure. In such a system, the particles are traced over four time-steps, therefore accelerations could be computed from particle tracks. Furthermore, the ghost particles between successive frames are uncorrelated due to different camera systems, allowing much easier particle tracking.

ACKNOWLEDGEMENTS

REFERENCES

[1] Elsinga G, Scarano F, Wieneke B and van Oudheusden BW "Tomographic Particle Image Velocimetry" Exp. Fluids **41**:933–947 (2005)

[2] Scarano F "Tomographic PIV: principles and practice" Meas. Sci. Tech. **24** 012001 (2013)

[3] Schanz D, Schröder A, Heine B, Dierksheide U "Flow structure identification in a high-resolution tomographic PIV data set of the flow behind a backward facing step" 16th Int Symp on Applications of Laser Techniques to Fluid Mechanics, Lisbon, (2012)

[4] Violato D, Moore P, Scarano F "Lagrangian and Eulerian pressure field evaluation of rod-airfoil flow from time-resolved tomographic PIV" Exp. Fluids **50** 1057-1070 (2011)

[5] Schröder A, Geisler R, Elsinga G E, Scarano F, Dierksheide U "Investigation of a turbulent spot and a tripped turbulent boundary layer flow using time-resolved tomographic PIV" Exp. Fluids **44** 305-316 (2008)

[6] Schröder A, Geisler R, Staak K, Wieneke B, Elsinga G, Scarano F and Henning A "Lagrangian and Eulerian views into a turbulent boundary layer flow using time-resolved tomographic PIV" Exp. Fluids, **50** 1071-1091 (2011)

[7] Herman GT, Lent A "Iterative reconstruction algorithms" Comput Biol Med **6** 273–294(1976)

[8] Atkinson C Soria J "An efficient simultaneous reconstruction technique for tomographic particle image Velocimetry" Exp. Fluids **47** 563-578 (2009)

[9] Maas H G, Gruen A, Papantoniou D, "Particle tracking velocimetry in three-dimensional flows" Exp. Fluids **15** 2 133-146 (1993)

[10] Novara M, Ianiro A, Scarano F "Adaptive interrogation for 3D-PIV" Meas. Sci. Tech. **24** 024012 (2013)

[11] Novara M, Batenburg KJ and Scarano F "Motion tracking-enhanced MART for tomographic PIV" Meas. Sci. Technol. **21** 035401 (2010)

[12] Wieneke B "Iterative reconstruction of volumetric particle distribution" Meas. Sci. Technol. **24** 024008 (2013)

[13] Willneff J "A Spatio-Temporal Matching Algorithm for 3D Particle Tracking Velocimetry" Swiss Federal Institute of Technology Zurich Diss. ETH No. 15276 (2003)

[14] Schanz D, Gesemann S, Schröder A, Wieneke B, Novara M "Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction" Meas. Sci. Technol. **24** 024009 (2013)

[15] Rapp C, Manhart M "Flow over periodic hills: an experimental study" Exp.Fluids **51** 247–269, DOI 10.1007/s00348-011-1045-y (2011)

[16] Wieneke B. "Volume self-calibration for Stereo PIV and Tomographic PIV" Exp. Fluids **45** 549-556 (2007)

CrossMark

RESEARCH ARTICLE

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

Daniel Schanz[1] · Sebastian Gesemann[1] · Andreas Schröder[1]

Abstract A Lagrangian tracking method is introduced, which uses a prediction of the particle distribution for the subsequent time-step as a mean to seize the temporal domain. Errors introduced by the prediction process are corrected by an image matching technique ('shaking' the particle in space), followed by an iterative triangulation of particles newly entering the measurement domain. The scheme was termed 'Shake-The-Box' and previously characterized as '4D-PTV' due to the strong interaction with the temporal dimension. Trajectories of tracer particles are identified at high spatial accuracy due to a nearly complete suppression of ghost particles; a temporal filtering scheme further improves on accuracy and allows for the extraction of local velocity and acceleration as derivatives of a continuous function. Exploiting the temporal information enables the processing of densely seeded flows (beyond 0.1 particles per pixel, ppp), which were previously reserved for tomographic PIV evaluations. While TOMO-PIV uses statistical means to evaluate the flow (building an 'anonymous' voxel space with subsequent spatial averaging of the velocity information using correlation), the Shake-The-Box approach is able to identify and track individual particles at numbers of tens or even hundreds of thousands per time-step. The method is outlined in detail, followed by descriptions of applications to synthetic and experimental data. The synthetic data evaluation reveals that STB is able to capture virtually all true particles, while effectively suppressing the formation of ghost particles. For the examined four-camera set-up particle image densities N_I up to 0.125 ppp could be processed. For noise-free images, the attained accuracy is very high. The addition of synthetic noise reduces usable particle image density ($N_I \leq 0.075$ ppp for highly noisy images) and accuracy (still being significantly higher compared to tomographic reconstruction). The solutions remain virtually free of ghost particles. Processing an experimental data set on a transitional jet in water demonstrates the benefits of advanced Lagrangian evaluation in describing flow details—both on small scales (by the individual tracks) and on larger structures (using an interpolation onto an Eulerian grid). Comparisons to standard TOMO-PIV processing for synthetic and experimental evaluations show distinct benefits in local accuracy, completeness of the solution, ghost particle occurrence, spatial resolution, temporal coherence and computational effort.

1 Introduction and motivation

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

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

✉ Daniel Schanz
daniel.schanz@dlr.de

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

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

1.1 Tomographic PIV

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

1.2 TOMO-PIV and temporal information

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

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

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

1.3 Limitations of TOMO-PIV

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

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

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

1.4 Lagrangian particle tracking

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

1.5 Alternative particle detection methods

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

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

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

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

1.6 Shake-The-Box (STB)

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

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

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

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

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

1.7 Features of STB

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

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

1.8 Outline of the manuscript

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

2 The Shake-The-Box method

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

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

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

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

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

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

2.1 Initialization phase

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

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

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

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

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

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

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

2.2 Convergence phase

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

2.2.1 Predicting next position of a tracked particle

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

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

2.2.2 Position refinement ('shaking')

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

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

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

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

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

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

Evaluating R at the positions x_1', x_2' and x_3' yields three residual values, which are fitted using a polynomial of second order. The extremum $x_{R,ex}$ of this parabola is determined, and the new particle position $x_{I,new}$ is set to:

$$x_{I,new} = \begin{cases} x_I - \delta_S & \text{if } |x_{R,ex}| > \delta_S \text{ and } R[x_1', y, z, I_P] > R[x_3', y, z, I_P] \\ x_{R,min} & \text{if } x_I - \delta_S < x_{R,min} < x_I + \delta_S \\ x_I + \delta_S & \text{if } |x_{R,ex}| > \delta_S \text{ and } R[x_1', y, z, I_P] < R[x_3', y, z, I_P] \end{cases} \quad (2)$$

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

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

2.2.3 Intensity correction

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

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

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

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

2.2.4 Deleting particles

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

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

2.2.5 Identifying new particle candidates

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

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

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

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

2.2.6 Adding and deleting tracks

After the complete processing of time-step t_n, a mixture of particle candidates and tracked particles is available. The particle candidates may still contain a significant amount of ghost particles, while the tracked particles should represent true particles. Using the particle candidates of time-steps t_{n-3} to t_n, additional tracks of length 4 are searched. One difference to the initialization phase is the origin of the predictor for the track searching. Here, no predictor field from cross-correlation is available, but a predictor can be constructed from neighbouring tracked particles. If at least two tracked particles are found within a neighbourhood (e.g. three times the average particle distance), the predictor is calculated as a Gaussian-weighted average of the velocities of these particles. If not enough neighbouring particles are found, a general, larger search radius (e.g. corresponding to the largest expected particle shift) is applied to the position of the particle.

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

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

2.3 Converged phase

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

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

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

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

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

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

2.4 Outlier removal

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

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

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

2.5 Multi-pass processing

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

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

3 Synthetic data validation

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

3.1 Creation of synthetic tracks

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

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

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

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

3.2 Application of STB to synthetic images

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

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

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

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

Fig. 2 Details of virtual camera image for ppp = 0.01, ppp = 0.05 and ppp = 0.125 (from *left* to *right*)
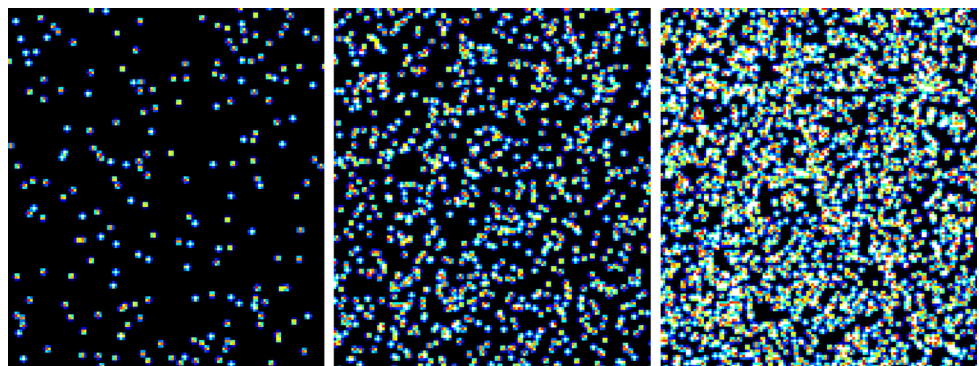
Table 1 STB parameters applied to synthetic image data from time-step 5 onwards

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

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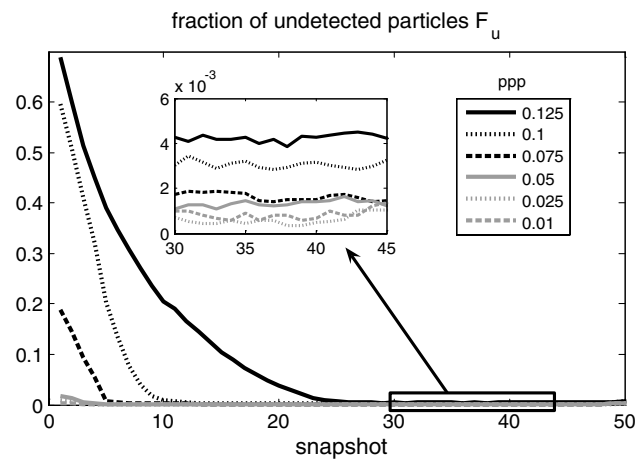

Fig. 4 Number of *ghost particles* within the system of *tracked* particles relative to the number of true particles

Fig. 3 Number of *undetected particles* relative to the true particle number
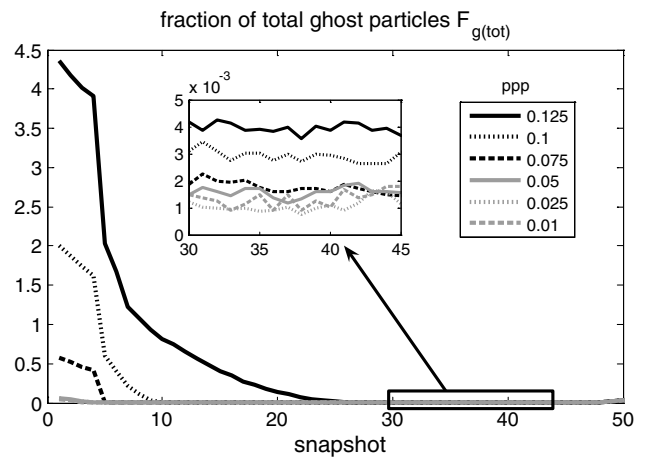
Fig. 5 *Total* number of reconstructed *ghost particles* (tracked and untracked) relative to the number of true particles

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

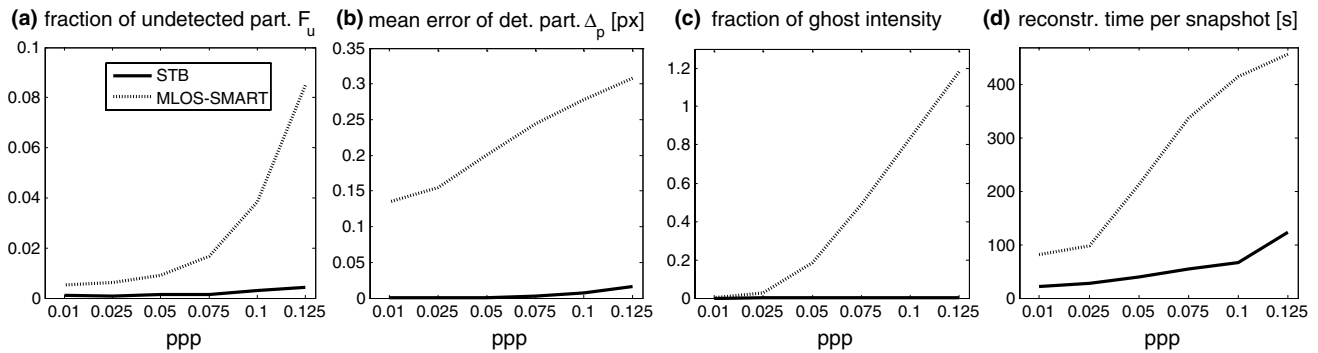### 3.3 Comparison to tomographic reconstruction

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

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

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

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

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

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

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

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

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

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

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

3.4 Influence of image noise

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

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

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

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

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

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

Fig. 8 Detail view from camera image for $N_I = 0.01$ ppp for different levels of artificially added noise: $\sigma = 0.03 \cdot I_{p,avg}$, $\sigma = 0.1 \cdot I_{p,avg}$, $\sigma = 0.2 \cdot I_{p,avg}$ (from *left* to *right*). σ is the variance of the normal distribution used for the random noise generation; $I_{p,avg}$ denotes the average peak intensity of a particle image
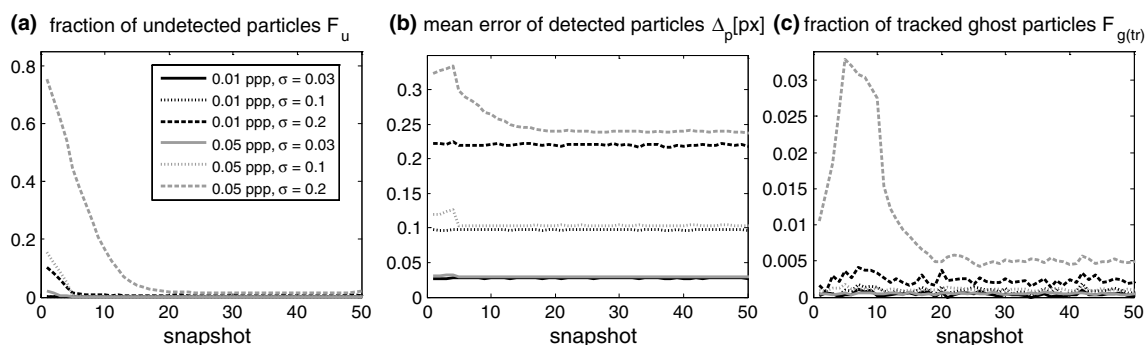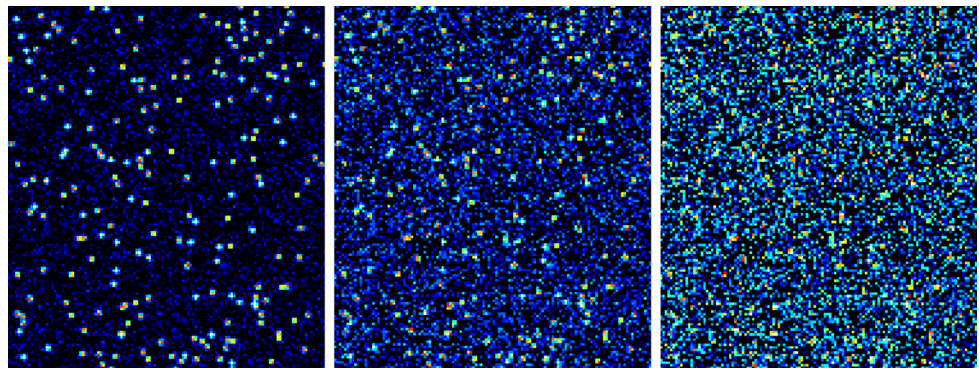
Fig. 9 Temporal development of STB runs for 0.01 and 0.05 ppp at different image noise (parameterized by σ in units of $I_{p,avg}$). **a** Fraction of non-detected particles; **b** particle position accuracy; and **c** fraction of ghost particles within the tracked particles

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

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

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

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

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

In order to improve the positional accuracy and to extract velocity and acceleration information from the tracks, temporal filtering is applied. To this end, a third-order B-spline
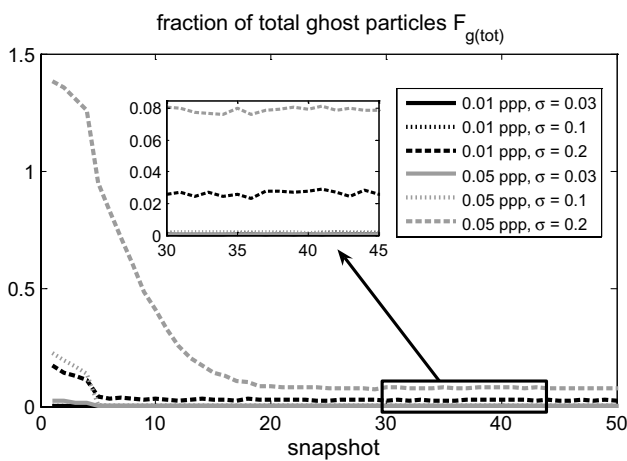
Fig. 10 Total (tracked and untracked) number of ghost particles relative to the true particle number for different noise levels (σ in units of $I_{p,avg}$) and particle image densities
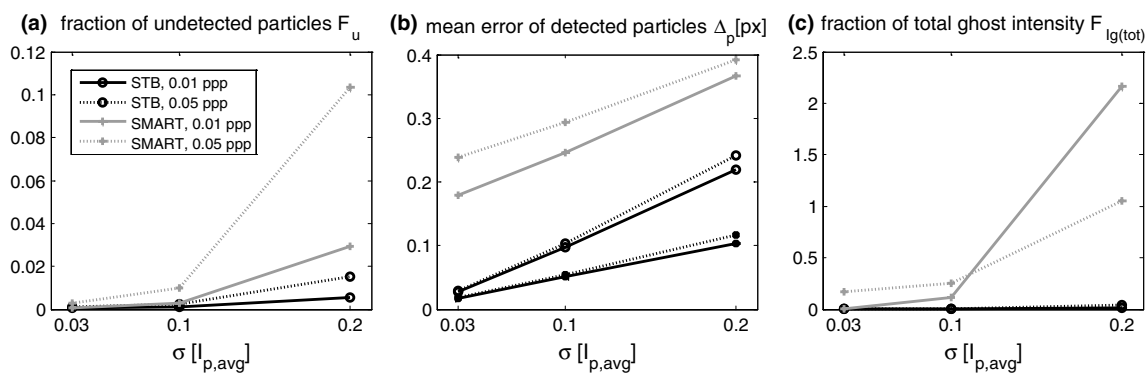
Fig. 11 Comparison of results gained by tomographic reconstruction with subsequent particle peak identification and tracking results by STB for varied amounts of image noise. **a** Fraction of non-detected particles; **b** mean error of detected particles; for STB, two curves are supplied: mean positional error of the raw data (*black curves*, *empty circle*) and mean positional error after applying a temporal fit (*black curves*, *filled circle*); and **c** fraction of intensity sum over all ghost particles in relation to the intensity sum of all true particles. Values averaged over images 40–44 of the time series analysed in Fig. 9

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

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

Table 3 Arithmetic mean (Δ) and root mean square (rms) of magnitude of position error (in px), velocity error (in px/time-step) and acceleration error (in px/time-step^2) for STB and SMART—reconstructions at 0.05 ppp for different noise levels

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

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

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

The temporal fit to the particle trajectory additionally yields values for Lagrangian acceleration (material derivative). The accuracy of the gained results was assessed similarly to velocity by computing the root mean square of the acceleration error (rms_a). The different noise levels show values of rms_a between 0.021 and 0.039 px/time-step^2. Following Adrians definition of DVR (1997), a *dynamic acceleration range* (DAR) is calculated as the ratio of rms_a and the maximum value of acceleration (~0.9 px/time-step^2, taken from the ground truth tracks). This yields values for DAR between 43 and 23 for the different noise cases (see Table 3).

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

4 Application to experimental data

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

Figure 12a shows the experimental set-up, which consists of a water jet created by a nozzle with a diameter of $d = 10$ mm and an exit velocity of $v_s = 0.43$ m/s. The jet
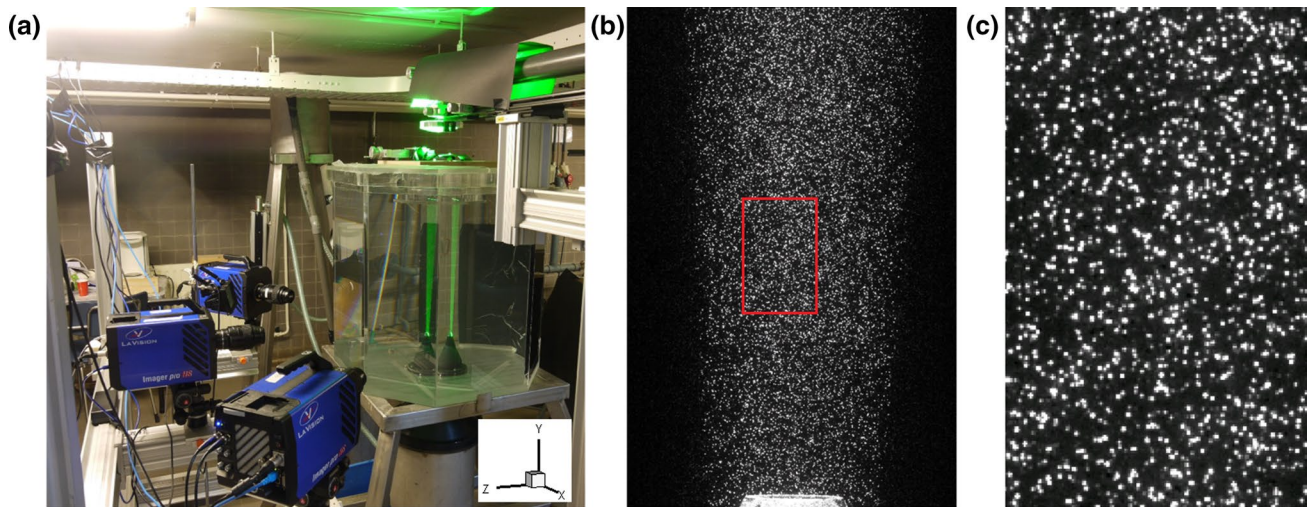
Fig. 12 **a** Set-up of the TOMO-PIV experiment at the water jet facility at TU Delft, the Netherlands. Three high-speed cameras image a water jet, using circular laser illumination; **b** full image of one exemplary camera; and **c** camera image detail

is illuminated from above using a light volume of cylindrical shape. As light source, a Quantronix Darwin-Duo Nd-YLF high-speed laser with a repetition rate of 1 kHz and a light output of 2·25 mJ is used. The water is seeded using polyamide particles with an average diameter of 56 μm. Imaging is realized by three LaVision Imager pro HS high-speed cameras, observing an interrogation volume extending from $y = 1.4$–54 mm above the nozzle. The illuminated diameter is approximately 28 mm (at the top) to 22 mm (at the bottom). The octagonal shape of the water tank allows all cameras to image the volume perpendicular to the air/glass/water interfaces. The cameras are equipped with 105 mm Nikon lenses, using $f_{\#} = 22$. A cropped resolution of 672×1024 pixels is used. Figure 12b, c shows particle images, as seen by one camera.

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

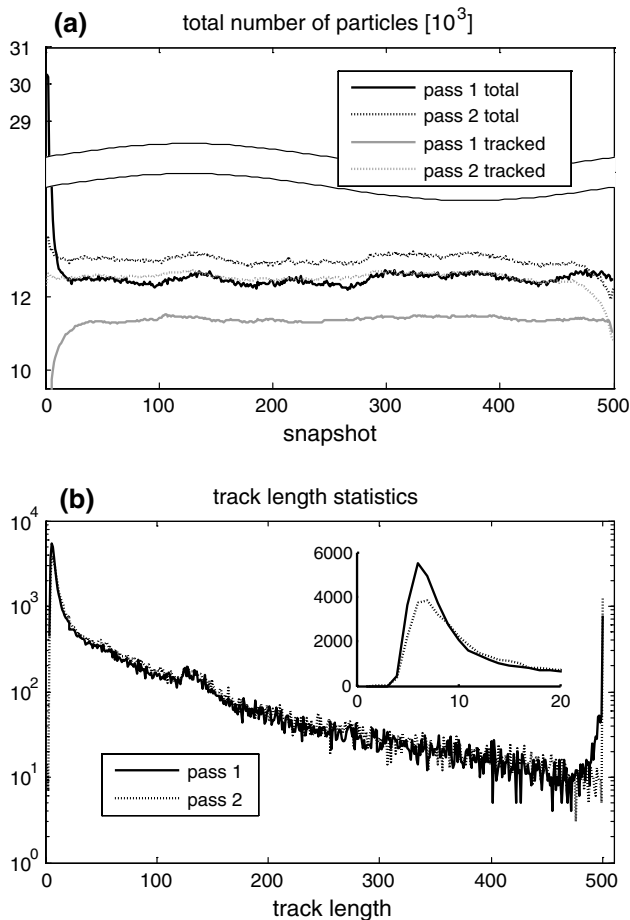### 4.1 Algorithm parameters and tracking evaluation

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

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

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

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

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

(a) total number of particles [10³]

pass 1 total
pass 2 total
pass 1 tracked
pass 2 tracked

snapshot

(b) track length statistics

pass 1
pass 2

track length

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

$f_{Pt} = 1.0$ $f_{Pt} = 2.0$

camera image

reprojected image

residual image

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

4.1.1 Remarks on residual images

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

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

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

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

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

4.1.2 Track length statistics

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

4.2 Track filtering

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

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

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

the fast particles coming from the nozzle are surrounded by a field of low-velocity particles that describe a slow circular motion around the jet. The detail plot in Fig. 15b shows how the particles are excited to curling trajectories by passing vortex rings that originate from Kelvin–Helmholtz (KH) instabilities at the shear layer (Violato et al. 2011). When using window methods like correlation, such details are lost. The shear layer is thin near the nozzle, but quickly widens with increasing distance. Figure 16 gives a better impression of the steepness of the velocity drop-off, by displaying radial bin-averaged (ensemble-averaged) profiles of streamwise velocity v at different heights. To create this plot, the velocity of all particles from the 500 time-steps located in a height interval of ±2 mm of the given y-value (e.g. for $y = 1.0$ d, all particles within $y = 8$–12 mm) were averaged in bins, according to their distance to the jet axis. The bins have similar volume; therefore, the spacing is closer in the outer regions of the jet. 100 bins are distributed over $r = 0$–15 mm, leading to particle numbers of around 5000–10,000 per bin in each of the six heights. Very close to the nozzle ($y = 0.2$ d), a steep velocity falloff
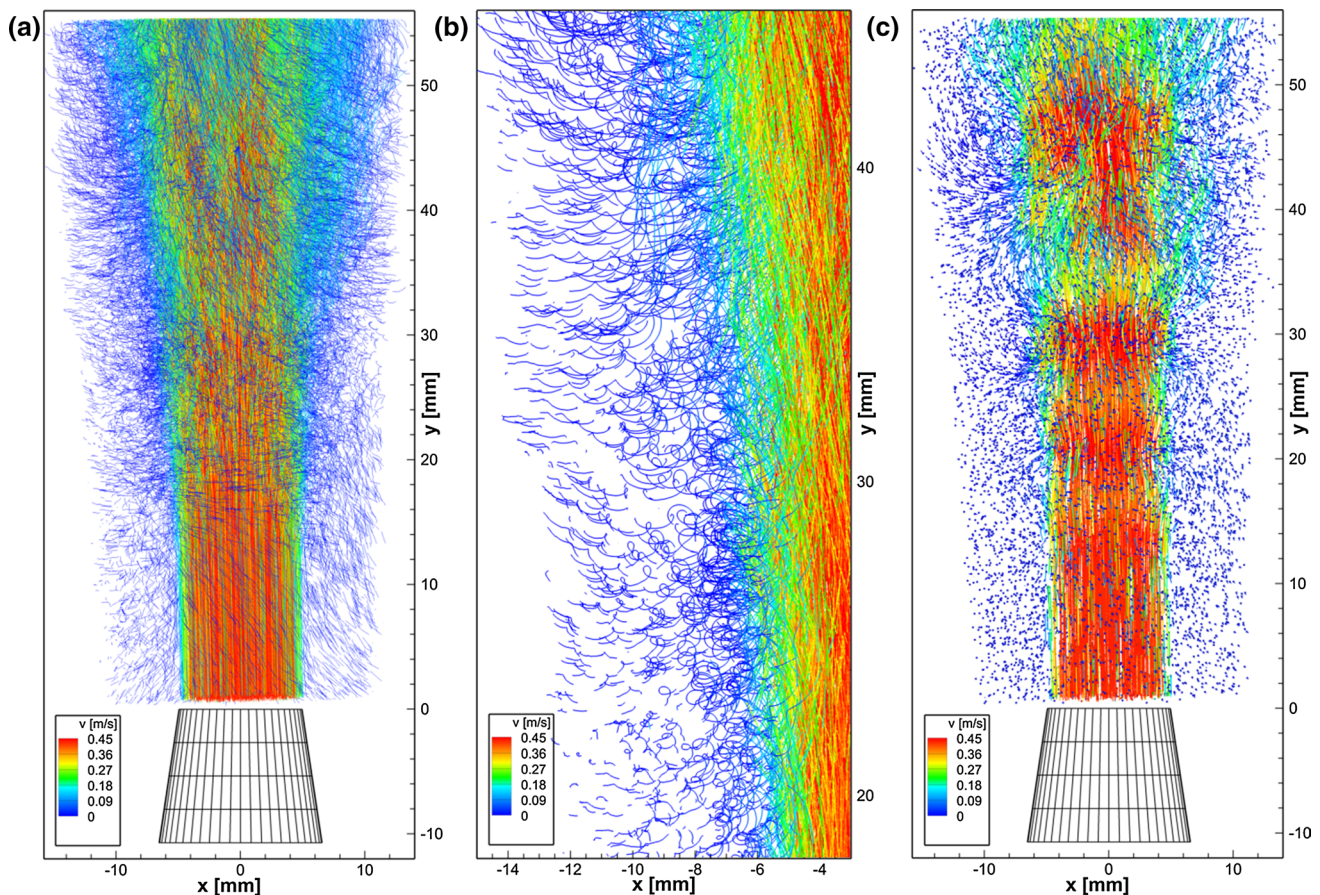
Fig. 15 a Particle tracks of water jet, reconstructed by STB. Overlay of 100 time-steps, *colour coded* by streamwise velocity; b detail of (a), showing the entrainment of particles; and c tracked particles for single time-step (t_n, given by *dots*) with a tail of 15 time-steps (reaching back to t_{n-14})
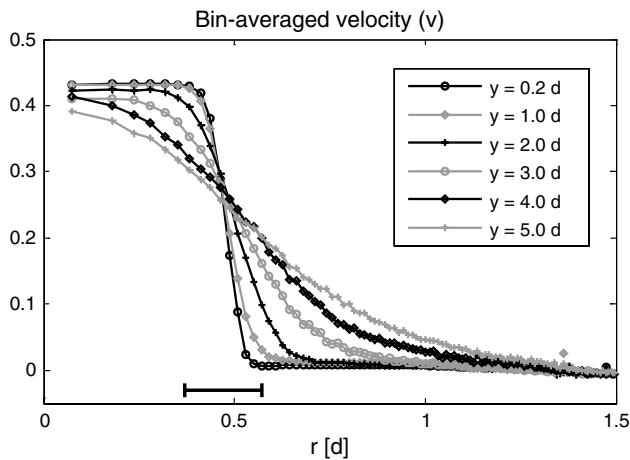
Fig. 16 Radial profiles of streamwise velocity (*v*) at different heights of the jet. The size of a 36-pixel wide correlation window is given as a reference

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

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

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

4.3 Interpolation to Eulerian grid ('FlowFit')

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

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

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

Evaluating the jet case, the measurement domain was divided into a cell system of quadratic B-splines, such that on average every eighth cell contains a particle (0.125 ppc, 'particles per cell'). The pitch of the resulting cell system is 0.62 mm; each cell represents a volume of approximately. 0.25 mm³; each particle represents approximately 2.0 mm³. For reference, a 36³ px correlation volume applied for TOMO-PIV processing comprises approximately 8.4 mm³. The solution of the resulting equation system is sampled

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

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

4.4 Lagrangian accelerations

Applying a temporal fit to the particle tracks not only yields Lagrangian velocities, but also accelerations. These values are of special interest, as it has proven to be difficult to extract accurate acceleration data from spatially smoothed velocity fields gained from TOMO-PIV. Shake-The-Box allows for the accurate extraction of Lagrangian accelerations, while remaining spatially well resolved. Figure 17b displays the acceleration in streamwise direction for the same time-steps as given in Fig. 15c, visualizing how the particles are accelerated when drawn into one of the large ring vortices and decelerated when they are ejected again. The same effect is visible for the larger vortices that occur after the KH breakdown.

One application of acceleration (or material derivative) data is the extraction of pressure distribution, which is actively been worked on recently. While some methods directly work on discrete particle tracks (Neeteson and Rival 2015), most of the developed methods require acceleration data on an Eulerian grid (Violato et al. 2011; Novara et al. 2013; Huhn et al. 2015). The FlowFit interpolation scheme can be used to create Eulerian gridded data from the discrete acceleration values. The same basic
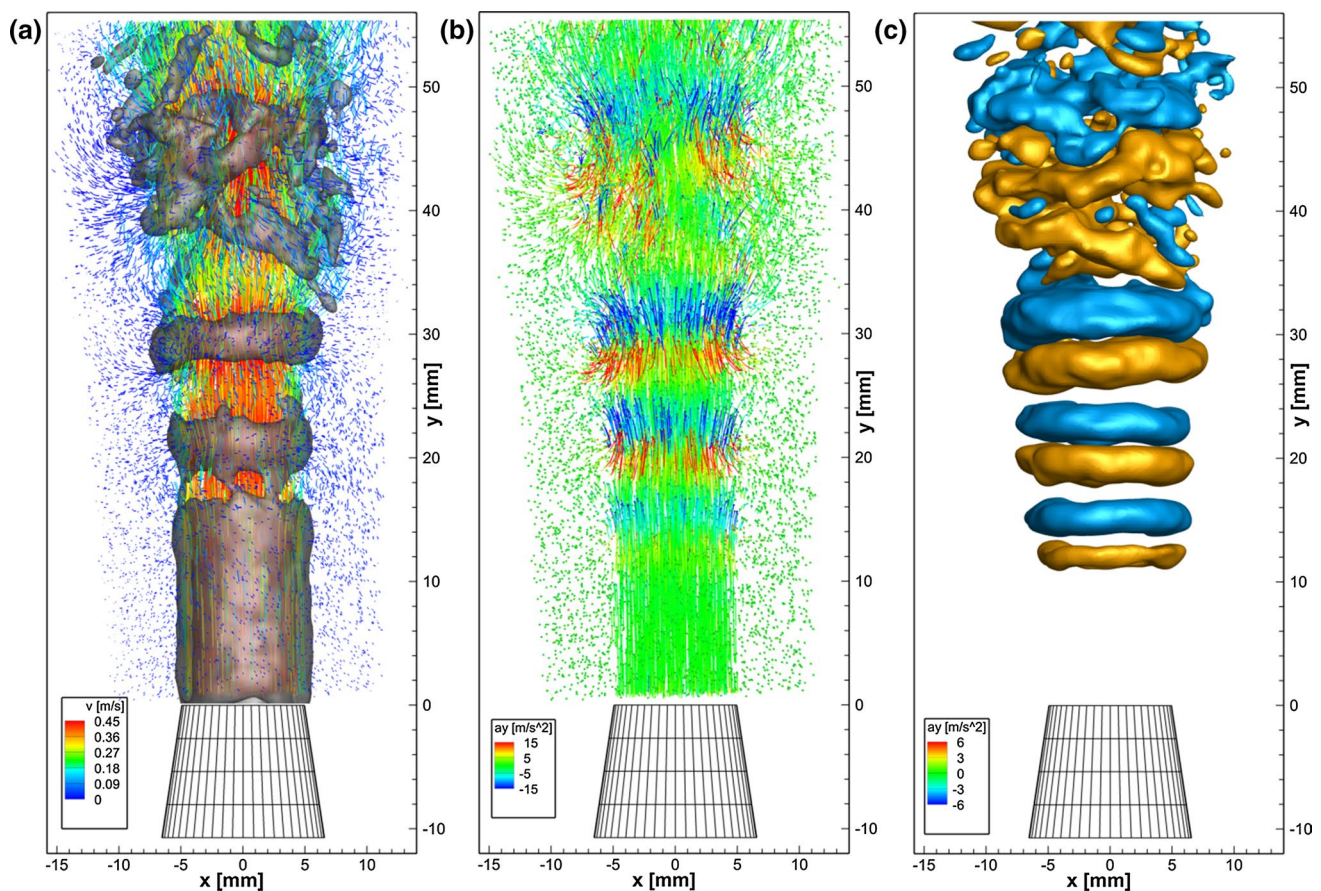
Fig. 17 **a** Isosurfaces of vorticity ($\omega = 175/s$) as calculated by *FlowFit* at single time-step (t_n) with superimposed particle tracks, extending 5 time-steps back in forth in time; **b** tracked particles for single time-step (t_n, given by spheres) with a tail of ten time-steps, *colour coded* by streamwise acceleration; and **c** isosurfaces of streamwise acceleration (ay $= 4$ m/s^2 and ay $= -4$ m/s^2) at single time-step t_n as calculated by *FlowFit* with penalization of rotation

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

4.5 Comparison to TOMO-PIV

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

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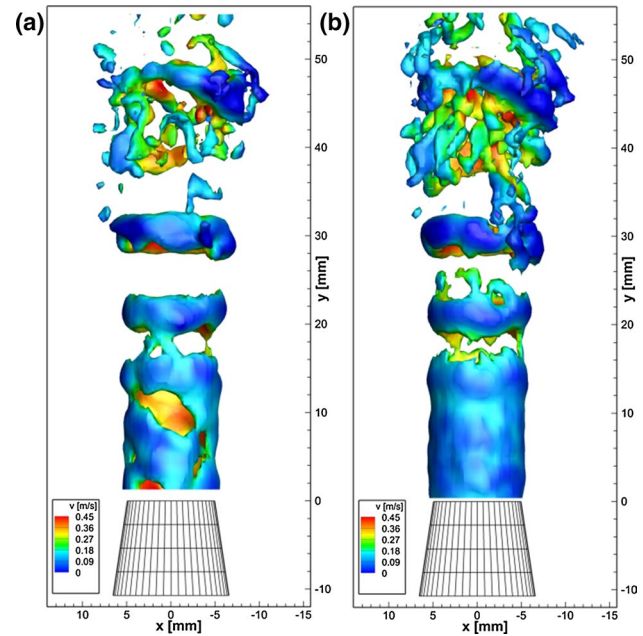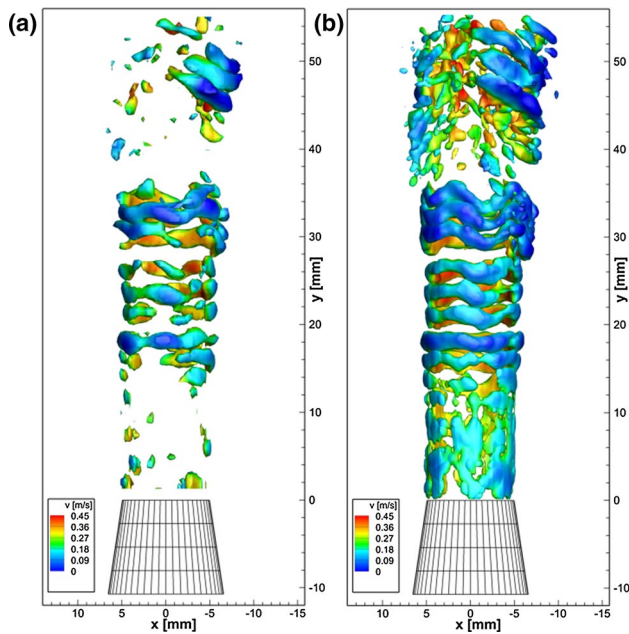

Fig. 18 a Isosurfaces of vorticity ($\omega = 175/s$), colour coded by streamwise velocity, gained by TOMO-PIV processing; **b** result gained from STB and grid interpolation using FlowFit at the same spatial resolution

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

4.6 Computational effort

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

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

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

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

5 Conclusions

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

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

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

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

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

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

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

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

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

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

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

References

Adrian RJ (1997) Dynamic ranges of velocity and spatial resolution of particle image velocimetry. Meas Sci Technol 8:1393–1398

Atkinson C, Soria J (2009) An efficient simultaneous reconstruction technique for tomographic particle image velocimetry. Exp Fluids 47:563–578

Atkinson C, Coudert S, Foucaut JM, Stanislas M, Soria J (2011) The accuracy of tomographic particle image velocimetry for measurements of a turbulent boundary layer. Exp Fluids 50:1031–1056

Bastiaans RJM, van der Plas GAJ, Kieft RN (2002) The performance of a new PTV algorithm applied in super-resolution PIV. Exp Fluids 32:345–356

Ben-Salah R, Alata O, Tremblais B, Thomas L, David L (2015) Particle volume reconstruction based on a marked point process and application to Tomo-PIV. EUSIPCO 2015. Nice, France, 31 Aug–3 June 2015

Bourgoin M, Ouellette NT, Xu H, Berg J, Bodenschatz E (2006) The role of pair dispersion in turbulent flow. Science 311(5762):835–838

Cierpka C, Lütke B, Kähler CJ (2013) Higher order multi-frame particle tracking velocimetry. Exp Fluids 54:1533

Cornic P, Champagnat F, Plyer A, LeClaire B, Cheminet A, Le Besnerais G (2014) Tomo-PTV with sparse tomographic reconstruction and optical flow. In: 17th international symposium on applications of laser techniques to fluid mechanics. Lisbon, Portugal, 07–10 July 2014

Cornic P, Champagnat F, Cheminet A, LeClaire B, Le Besnerais G (2015) Fast and efficient particle reconstruction on a 3D grid using sparsity. Exp Fluids 56:62

Dalziel SB (1992) Decay of rotating turbulence: some particle tracking experiments. Appl Sci Res 49:217–244

Discetti S, Astarita T (2011) A fast multi-resolution approach to tomographic PIV. Exp Fluids 52:765–777

Discetti S, Astarita T (2012) Fast 3D PIV with direct sparse cross-correlations. Exp Fluids 53:1437–1451

Discetti S, Natale A, Astarita T (2013) Spatial filtering improved tomographic PIV. Exp Fluids 54:1505

Discetti S, Agüera N, Cafiero G, Astarita T (2015) Ensemble 3D-PTV for high resolution turbulent statistics. In: 11th international symposium on PIV–PIV15. Santa Barbara, USA, 14–16 Sept 2015

Earl A, Cochard S, Thomas L, Tremblais B, David L (2015) Implementation of vibration correction schemes to the evaluation of a turbulent flow in an open channel by tomographic particle image velocimetry. Meas Sci Technol 26:015303

Eilers PHC, Marx BD (1996) Flexible smoothing with B-splines and penalties. Stat Sci 11(2):89–121

Elsinga GE, Tokgoz S (2014) Ghost hunting—an assessment of ghost particle detection and removal methods for tomographic-PIV. Meas Sci Technol 25:084004

Elsinga GE, van Oudheusden BW, Scarano F (2006b) Experimental assessment of tomographic-PIV accuracy. In: 13th international symposium on applications of laser techniques to fluid mechanics. Lisbon, Portugal, 26–29 June 2006

Elsinga GE, Scarano F, Wieneke B, van Oudheusden BW (2006a) Tomographic particle image velocimetry. Exp Fluids 41:933–947

Elsinga GE, Westerweel J, Scarano F, Novara M (2011) On the velocity of ghost particles and the bias errors in Tomographic-PIV. Exp Fluids 50:825–838

Gesemann S (2015) From particle tracks to velocity and acceleration fields using B-splines and penalties. arXiv 1510.09034

Ghaemi S, Scarano F (2011) Counter-hairpin vortices in the turbulent wake of a sharp trailing edge. J Fluid Mech 689:317–356

Hain R, Kähler CJ, Michaelis D (2008) Tomographic and time resolved PIV measurements on a finite cylinder mounted on a flat plate. Exp Fluids 45:715–724

Henningsson P, Michaelis D, Nakata T, Schanz D, Geisler R, Schröder A, Bomphrey RJ (2015) The complex aerodynamic footprint of desert locusts revealed by large-volume tomographic particle image velocimetry. J R Soc Interface 12:108

Herman GT, Lent A (1976) Iterative reconstruction algorithms. Comput Biol Med 6:273–294

Huhn F, Schanz D, Gesemannn S, Schröder A (2015) Pressure fields from high-resolution time-resolved particle tracking velocimetry in 3D turbulent flows. In: Proceedings of NIM2015 Workshop, Poitiers, France

Kähler CJ, Scharnowski S, Cierpka C (2012a) On the resolution limit of digital particle image velocimetry. Exp Fluids 52:1629–1639

Kähler CJ, Scharnowski S, Cierpka C (2012b) On the uncertainty of digital PIV and PTV near walls. Exp Fluids 52:1641–1656

Kasagi N, Nishino K (1990) Probing turbulence with three-dimensional particle tracking velocimetry. Exp Therm Fluid Sci 4:601–612

La Porta A, Voth GA, Crawford AM, Alexander J, Bodenschatz E (2001) Fluid particle accelerations in fully developed turbulence. Nature 409:1017–1019

Lüthi B, Tsinober A, Kinzelbach W (2005) Lagrangian measurement of vorticity dynamics in turbulent flow. J Fluid Mech 528:87–118

Lynch K, Scarano F (2013) A high-order time-accurate interrogation method for time-resolved PIV. Meas Sci Technol 24:035305

Lynch K, Scarano F (2014) Experimental determination of tomographic PIV accuracy by a 12-camera system. Meas Sci Technol 25:084003

Lynch K, Scarano F (2015) An efficient and accurate approach to MTE-MART for time-resolved tomographic PIV. Exp Fluids 56:66

Maas HG, Grün A, Papantoniou D (1993) Particle tracking in three dimensional turbulent flows—part I: photogrammetric determination of particle coordinates. Exp Fluids 15:133–146

Malik N, Dracos T, Papantoniou D (1993) Particle tracking in three dimensional turbulent flows—part II: particle tracking. Exp Fluids 15:279–294

Michaelis D, Wolf CC (2011) Vibration compensation for tomographic PIV using single image volume self calibration. In: 9th international symposium on particle image velocimetry—PIV11. Kobe, Japan, 21–23 July 2011

Neeteson NJ, Rival D (2015) Pressure-field extraction on unstructured flow data using a Voronoi tessellation-based networking algorithm: a proof-of-principle study. Exp Fluids 56:44

Neeteson NJ, Bhattacharya S, Rival DE, Michaelis D, Schanz D, Schröder A (2015) Pressure-field extraction from lagrangian flow measurements. In: 11th international symposium on PIV–PIV15. Santa Barbara, USA, 14–16 Sept 2015

Nishino K, Kasagi N, Hirata M (1989) Three-dimensional particle tracking velocimetry based on automated digital image processing. Trans ASME J Fluid Eng 111:384–390

Nobach H, Bodenschatz E (2009) Limitations of accuracy in PIV due to individual variations of particle image intensities. Exp Fluids 24:045302

Novara M, Scarano F (2013) A particle-tracking approach for accurate material derivative measurements with tomographic PIV. Exp Fluids 54:1584

Novara M, Batenburg KJ, Scarano F (2010) Motion tracking-enhanced MART for tomographic PIV. Meas Sci Technol 21:035401

Novara M, Ianiro A, Scarano F (2013) Adaptive interrogation for 3D-PIV. Meas Sci Tech 24:024012

Novara M, Schanz D, Kähler CJ, Schröder A (2015) Shake-The-Box for multi-pulse tomographic systems: towards high seeding density particle tracking in high speed flows. In: 11th international symposium on PIV–IV15. Santa Barbara, USA, 14–16 Sept 2015

Ouellette NT, Xu H, Bodenschatz E (2006) A quantitative study of three-dimensional Lagrangian particle tracking algorithms. Exp Fluids 40(2):301–313

Savitzky A, Golay MJE (1964) Smoothing and differentiation of data by simplified least squares procedures. Anal Chem 36(8):1627–1639. doi:10.1021/ac60214a047

Scarano F (2013) Tomographic PIV: principles and practice. Meas Sci Technol 24:012001

Scarano F, Poelma C (2009) Three-dimensional vorticity patterns of cylinder wakes. Exp Fluids 47:69–83

Schanz D, Schröder A, Heine B, Dierksheide U (2012) Flow structure identification in a high-resolution tomographic PIV data set of the flow behind a backward facing step. In: 16th international symposium on applications of laser techniques to fluid mechanics, Lisbon

Schanz D, Gesemann S, Schröder A, Wieneke B, Novara M (2013a) Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. Meas Sci Technol 24:024009

Schanz D, Schröder A, Gesemann S, Michaelis D, Wieneke B (2013b) Shake-the-Box: a highly efficient and accurate tomographic particle tracking velocimetry (TOMO-PTV) method using prediction of particle position. In: 10th international symposium on particle image velocimetry—PIV13. Delft, The Netherlands, 1–3 July 2013

Schanz D, Schröder A, Gesemann S (2014) Shake-the-Box—a 4D PTV algorithm: accurate and ghostless reconstruction of Lagrangian tracks in densely seeded flows. In: 17th international symposium on applications of laser techniques to fluid mechanics. Lisbon, Portugal, 07–10 July 2014

Schneiders JFG, Azijli I, Scarano F, Dwight RP (2015) Pouring time into space. In: 11th international symposium on particle image velocimetry—PIV15. Santa Barbara, California, 14–16 Sept 2015

Schröder A, Geisler R, Elsinga GE, Scarano F, Dierksheide U (2008) Investigation of a turbulent spot and a tripped turbulent boundary layer flow using time-resolved tomographic PIV. Exp Fluids 44:305–316

Schröder A, Geisler R, Staak K, Wieneke B, Elsinga G, Scarano F, Henning A (2011) Lagrangian and Eulerian views into a turbulent boundary layer flow using time-resolved tomographic PIV. Exp Fluids 50:1071–1091

Schröder A, Schanz D, Geisler R, Willert C, Michaelis D (2013) Dual-volume and four-pulse tomo PIV using polarized laser light. In: 10th international symposium on particle image velocimetry—PIV13. Delft, The Netherlands, 1–3 July 2013

Schröder A, Schanz D, Michaelis D, Cierpka C, Scharnovski S, Kähler CJ (2015a) Advances of PIV and 4D-PTV "Shake-The-Box" for turbulent flow analysis—the flow over periodic hills. Flow Turb Comb 95(2–3):193–209

Schröder A, Schanz D, Geisler R, Gesemann S, Willert C (2015b) Near-wall turbulence characterization using 4D-PTV Shake-The-Box. In: 11th international symposium on particle image velocimetry—PIV15. Santa Barbara, California, 14–16 Sept 2015

Sciacchitano A, Wieneke B, Scarano F (2013) PIV uncertainty quantification by image matching. Meas Sci Technol 24:045302

Violato D, Scarano F (2011) Three-dimensional evolution of flow structures in transitional circular and chevron jets. Phys Fluids 23:124104

Violato D, Moore P, Scarano F (2011) Lagrangian and Eulerian pressure field evaluation of rod-airfoil flow from time-resolved tomographic PIV. Exp Fluids 50:1057–1070

Violato D, Ianiro A, Cardone G, Scarano F (2012) Three-dimensional vortex dynamics and convective heat transfer in circular and chevron impinging jets. Int J Heat Fluid Flow 37:22–36

Wereley ST, Meinhart CD (2001) Second-order accurate particle image velocimetry. Exp Fluids 31:258–268

Wieneke B (2007) Volume self-calibration for stereo PIV and tomographic PIV. Exp Fluids 45:549–556

Wieneke B (2013) Iterative reconstruction of volumetric particle distribution. Meas Sci Technol 24:024008

Wieneke B (2015) PIV uncertainty quantification from correlation statistics. Meas Sci Technol 26:074002

Wiener N (1949) Extrapolation, interpolation, and smoothing of stationary time series, vol 2. MIT Press, Cambridge

Willneff J (2003) A spatio-temporal matching algorithm for 3D particle tracking velocimetry. Swiss Federal Institute of Technology Zurich Diss. ETH No. 15276

Xu H (2008) Tracking Lagrangian trajectories in position-velocity space. Meas Sci Technol 19:075105

Xu H, Bourgoin M, Ouellette NT, Bodenschatz E (2006) High order lagrangian velocity statistics in turbulence. Phys Rev Lett 96:024503

From Noisy Particle Tracks to Velocity, Acceleration and Pressure Fields using B-splines and Penalties

Sebastian Gesemann[*], Florian Huhn, Daniel Schanz, Andreas Schröder

German Aerospace Center (DLR), Institute of Aerodynamics and Flow Technology, Department of Experimental Methods, Germany
* Correspondent author: sebastian.gesemann@dlr.de

ABSTRACT

Recent advances in particle tracking methods for flow measurements ([Scha14], [Scha15]) allow reconstruction of particle distributions with high particle densities that could previously only be handled with tomographic reconstruction methods. High particle densities enable reconstructions of flow fields with high spatial resolution. The STB method for tracking particles has been proven to provide more accurate estimates of particle positions without reconstructing "ghost particles" while at the same time saving time and memory compared to tomographic reconstructions. In this work, B-spline based methods for noise reduction and reconstruction of velocity, acceleration and pressure fields based on the particle track data are described and evaluated using synthetic data. The goal of these methods is to reconstruct accurate and high resolution velocity, acceleration and pressure fields. This is achieved by using the particle data directly and exploiting known physical properties such as a freedom of divergence of velocity and the Navier-Stokes equation for incompressible and uniform-density flows.

1. Introduction

Determining 3D velocity fields from flow experiments is difficult especially if a high spatial resolution is desired. For optical particle-based measurement techniques, a high spatial resolution implies a high density of tracer particles that need to follow the flow and have to be observed using multiple cameras.

Instead of identifying and matching particles separately between different views, which is nontrivial at high particle densities, a tomographic approach has been successfully applied in the past to the particle distribution reconstruction problem. This method is called *TomoPIV* for Tomographic Particle Image Velocimetry. In TomoPIV the measurement volume is discretized and reconstructed as the solution to a large but sparse and constrained linear equation system resulting in a discrete volume of light intensities. One way of deriving flow fields from these volumes is to apply a cross correlation between two subvolumes of reconstructions from

neighboring points of time to detect the average flow velocity within such a subvolume. This is a robust method to compute flow velocities but results in a spatially lowpass filtered representation of the velocity field depending on the size of the subvolumes used for cross-correlation, the window size.

With recent advances in particle tracking techniques ([Scha14], [Scha15]) the density of particles that can be reconstructed directly without discretizing the volume is approaching densities that are typically used in TomoPIV measurements. Reconstructing particle locations directly instead of a discretized intensity volume has several benefits over a tomographic reconstruction: It typically requires a only small fraction of CPU and RAM resources compared to what is needed for TomoPIV to solve the large constrained equation system. Also, such a direct particle reconstruction method avoids an additional layer of spatial discretization which can be expected to improve the accuracy of the estimated particle locations. Given a sequence of time-resolved measurement images of a flow with tracer particles, particle tracks can be reconstructed with these new techniques.

In this work, we will describe a method for noise reduction of particle tracks that is inspired by Wiener/Kalman filtering [Kal80] as well as a method for a spatial interpolation of velocity, acceleration and pressure fields from the scattered particle data which we developed with the goal of preserving much of the information present in the particle tracks and avoiding any unwanted spatial lowpass filtering effect such as the one that is inherent in correlation-based methods. The reconstruction methods are assessed using synthetic data from a direct numeric simulation (DNS).

Section 2 gives an overview of the reconstruction methods. Section 3 explains how particle track interpolation and Wiener/Kalman filter based noise reduction can be combined to derive good estimates of particle velocities and accelerations. In section 4 the spatial interpolation of scattered particle data is described. Section 5 presents and discusses results of the reconstruction methods applied on synthetic data. Section 6 closes with a conclusion.

2. Overview of the reconstruction method

Our approach to compute velocity, acceleration and pressure fields based on noisy particle tracks via a cost function minimizer can be split into two parts:

1. *TrackFit* takes the noisy particle location data of a particle track and computes a B-spline curve for the track. This step includes noise reduction similar to Wiener/Kalman filtering

and allows computing particle location along with its first and second order derivatives at any point in time within the time interval of the observed particle track.

2. *FlowFit* takes particle locations along with their velocities and/or accelerations for any particular point in time and computes 3D B-spline curves for velocity, acceleration and/or pressure by minimizing a cost function that optionally accounts for constraints such as freedom of divergence for incompressible flows. The result is again a continuous function depending on a finite number of degrees of freedom.

These steps share many similarities conceptually. Both make use of uniform cubic B-splines to represent the result as a continuous function and both employ a similar form of noise reduction via penalization like it was introduced in [Eil96] for the one-dimensional case. The difference is *TrackFit* deals with data that is already equidistantly sampled and one-dimensional (time) while *FlowFit* deals with scattered data in the three dimensions.

3. TrackFit for temporal interpolation

In the interest of optimal noise reduction we performed spectral analyses of the particle tracks of different experiments. We assume that the particle location measurements can be represented as a sum of the real particle locations and white measurement noise that does not correlate with the real particle locations. Based on these spectral analyses we derived a simple physical model of how particles move: We assume the third order derivative of a particle's real location with respect to time (jolt) to be white noise as well which we refer to as *dynamic jolt noise*. In terms of the Kalman filter theory [Kal60], this noise is what excites the dynamic system and represents what is unpredictable. This particular assumption about how particles move corresponds to a $1/f^3$ shape of the real particle location's amplitude spectrum where f refers to the frequency, see figure 1. This model is a simplification of what we observed from the spectral analysis. It matches our observations reasonably well in the frequency range where the signal-to-noise ratio is close to 0 dB. It tends to overestimate the signal-to-noise ratio for other frequencies with a much higher or lower signal-to-noise ratio but this is not an issue because these differences have little effect on the magnitude response of the optimal Wiener filter. This simple model has two parameters: the level of particle position measurement noise and the level of dynamic jolt noise. The frequency f_{cross} at which both curves cross and the signal-to-noise ratio is exactly 0 dB only depends on the ratio of these two noise levels. This frequency and the level of particle position measurement noise can be obtained from a spectral analysis, specifically from the location of the kink and the level of the flat section of the amplitude spectrum.
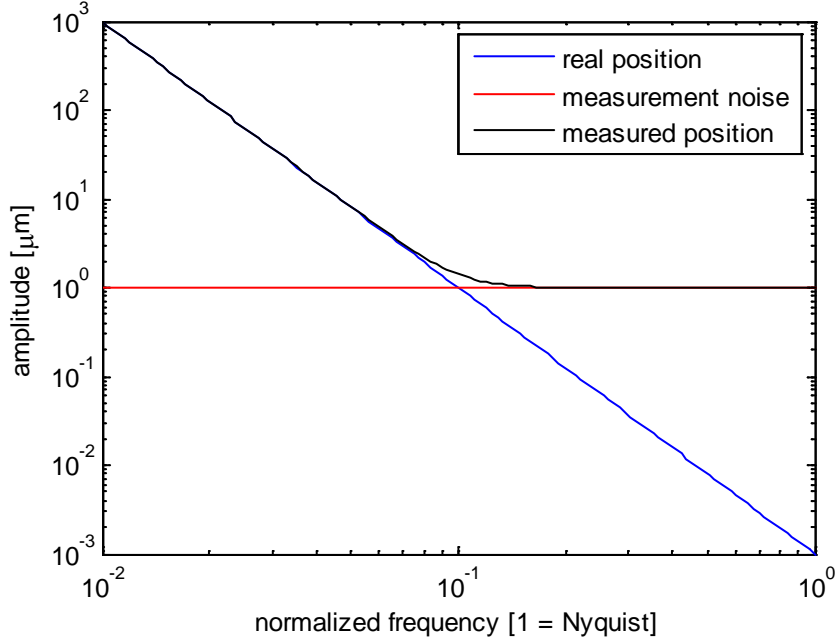
Figure 1: Example for the spectral model of measured particle position data with a normalized crossover frequency 0.1 and a position measurement error of 1 micrometer.

Assuming that the real particle track of length n can be represented as a cubic B-spline curve p where $\vec{c_i} \in \mathbb{R}^3$ for $i = 0, 1, \dots n + 1$ are weighting coefficients for the cubic base splines β_3

$$p_c(t) = \sum_{i=0}^{n+1} \vec{c_i} \cdot \beta_3(t - i)$$

$$\beta_3(t) = \begin{cases} \frac{4}{6} - |t|^2 + \frac{1}{2}|t|^3 & \text{for } t < 1 \\ \frac{1}{6}(2 - |t|)^3 & \text{for } 1 \leq t < 2 \\ 0 & \text{else} \end{cases}$$

and assuming further that the measurement noise as well as the dynamic jolt noise is Gaussian, the Maximum-Likelihood solution for the particle track given n measured positions $\vec{y_i} \in \mathbb{R}^3$ for $i = 1, \dots n$ can be computed by solving an overdetermined linear equation system in a least squares sense. The corresponding cost function F is

$$F(c) = \sum_{i=1}^{n} |p_c(i) - \vec{y_i}|^2 + \sum_{i=1}^{n-1} \left| \lambda \cdot \dddot{p_c}\left(i + \frac{1}{2}\right) \right|^2$$

Here, $\dddot{p_c}$ refers to the third order derivative of p_c. It is effectively a kind of Tikhonov regularization tailored for our physical model. This is also known as a P-spline (penalized

spline) [Eil96]. Both, $\ddot{p}_c$ and p_c can be written as weighted sum of the unknowns $\vec{c_i}$ for any particular point in time between 1 and n with some exceptions. The parameter λ relates the standard deviation of position measurement errors to the standard deviation of the dynamic jolt noise and affects the strength of the smoothing. The following choice for λ

$$\lambda = \frac{1}{(\pi \cdot f_{cross})^3}$$

will lead to an appropriate lowpass filtering effect that approximates the optimal Wiener filter for our assumptions quite well. The advantages of this procedure compared to ordinary Wiener filtering in the temporal or spectral domain is that it results in a continuous representation of the particle track that can be sampled and derived at any point in time and does not require any special treatment at the boundaries of the particle track.

With the help of the Moore-Penrose pseudoinverse of the matrix of the above mentioned linear equation system and estimates for the physical model parameters that can be derived from the magnitude spectrum it is also possible to estimate the errors of the fitted particle positions, velocities and accelerations. These error estimates can later be used for appropriate weighting of velocity and acceleration deviations during the spatial interpolation.

4. FlowFit for spatial interpolation

Our FlowFit development started out as a simple regression for interpolating vector fields using uniform 3D B-splines with optional penalization of divergence as part of the cost function. It was born out of necessity. We needed a way to interpolate scattered velocity data in order to participate in the 2014 PIV challenge [Käh16] with a particle tracking algorithm. We later extended FlowFit to allow penalization of rotation instead of divergence for interpolating particle accelerations in the cases where acceleration is dominated by the pressure gradient. Both cases end up being sparse linear least squares problems.

In an attempt to further improve the quality of the reconstruction for incompressible and constant-density flows we developed a nonlinear version of this interpolation problem that exploits additional physical constraints and recovers pressure at the same time. We refer to this as the second generation FlowFit.

4.1. First generation FlowFit

In this simple linear version we represent a vector field $\vec{v}$ for a cuboid volume as a weighted sum of 3D base splines. The 3D base spline $\overline{\beta_3}$ is the result of a 3D convolution of the 1D base spline β_3:

$$\overline{\beta_3}: \mathbb{R}^3 \longmapsto \mathbb{R}$$

$$\overline{\beta_3}(x) = \prod_{i=1}^{3} \beta_3(x_i)$$

$$\vec{v_c}: \mathbb{R}^3 \longmapsto \mathbb{R}^3$$

$$\vec{v_c}(x) = \sum_{i=-1}^{I} \sum_{j=-1}^{J} \sum_{k=-1}^{K} \overrightarrow{c_{i,j,k}} \cdot \overline{\beta_3}\left(x - \overrightarrow{p_{i,j,k}}\right)$$

In the formula for the vector field $\vec{v_c}$ above $\overrightarrow{c_{i,j,k}}$ are the weights that need to be determined and $\overrightarrow{p_{i,j,k}}$ represents the 3D coordinate of point of a uniform Cartesian grid with point spacing h and lower corner $\vec{o}$:

$$\overrightarrow{p_{i,j,k}} \equiv \vec{o} + h \cdot (i\ j\ k)^T$$

Due to the overlap of the cubic base splines the function in $\vec{v}$ the vector field can only be considered to be defined for the cuboid volume with lower and upper corners $\overrightarrow{p_{0,0,0}}$ and $\overrightarrow{p_{I,J,K}}$, respectively. For every point in this space that does not necessarily coincide with a grid point $\overrightarrow{p_{i,j,k}}$ the value of the vector field at this position including its spatial derivatives up to the 2nd order in each dimension can be expressed exactly as a linear combination of the parameters $\overrightarrow{c_{i,j,k}}$. Each detected particle within that volume will amount to three equations linking the unknown coefficients $\vec{c}$ to measured velocity or acceleration depending on whether a velocity field or acceleration field should be reconstructed.

To avoid any unwanted spatial smoothing we choose the internal grid resolution h so that the ratio between the number of detected particles and the product $I \cdot J \cdot K$ is typically in the range of [0.05, 0.33]. As a consequence the number of unknowns would be 3 to 20 times higher than the number of equations. Such an equation system is obviously singular and therefore has no unique solution. But we know in advance that most of the energy is typically concentrated in the lower wavenumbers.

With a certain kind of Tikhonov regularization we exploit this knowledge. This way, we select a solution that minimizes the L2 norm of a highpass-filtered version of the vector field that approximately represents the curvature of the field. This regularization can also act as another form of noise reduction.

The highpass filter is computed by a difference between the coefficients $\vec{c}$ and a spatially lowpass-filtered version of $\vec{c}$. This lowpass filter is separable and can therefore be computed by separate filtering in X-, Y- and Z- direction with the filter kernel [0.25, 0.5, 0.25]. At the boundaries this filtering is only applied in two or one dimension(s) for the faces of the cuboid or its edges, respectively.

It is possible to add further physically motivated regularizations such as the penalization of divergence (for a velocity field of an incompressible flow) or curl (for an acceleration field that is dominated by the pressure gradient). Divergence and curl can also be represented as a linear combination of the unknown coefficients $\vec{c}$ at any point in this volume. We apply this kind of regularization at all the *inner* grid points $\vec{p}$.

We solve the resulting optimization problems with the help of an iterative algorithm: CGLS, the Conjugate Gradient method for Least Squares problems [Pai82].

4.2. Second generation FlowFit

Reconstructing the velocity field and the acceleration field separately based on measured velocities and accelerations – like described in the previous section – is computationally cheap. In both cases this results in a weighted linear least squares problem that can be easily solved using an iterative algorithm for such problems. However, such an independent reconstruction of velocity and acceleration is not able to account for all constraints we know should be satisfied in certain cases. In case of an incompressible flow, we not only expect the divergence of the velocity field to vanish, but also the divergence of the temporal derivative of the velocity field:

$$\nabla \vec{u} = 0 \quad \Longrightarrow \quad \nabla \frac{\partial \vec{u}}{\partial t} = 0 \tag{1, 2}$$

With velocity ($\vec{u}$) and material acceleration ($\vec{a}$) available and using the following equivalence

$$\vec{a} \equiv \frac{D\vec{u}}{Dt} \equiv \frac{\partial \vec{u}}{\partial t} + \vec{u} \cdot \nabla \vec{u}$$

we can express the second divergence constraint (2) as equation (3)

$$\nabla \cdot (\vec{a} - \vec{u} \cdot \nabla \vec{u}) = 0 \qquad (3)$$

In theory, combining both minimization problems into a single one allows us to account for the second divergence constraint (2) via the cost function. Instead of penalizing the curl of the acceleration field we express acceleration in terms of pressure and velocity assuming uniform density. This removes three acceleration variables and adds a single kinetic pressure variable for each grid point (pressure divided by density). Acceleration is derived from pressure and velocity according to the Navier-Stokes equation for incompressible flows with uniform density (ignoring other forces):

$$\vec{a} = -\nabla \bar{p} + \nu \Delta \vec{u} \qquad (4)$$

From here on, we refer to $\bar{p}$ as pressure for simplicity. Substituting the right hand side of equation (4) for acceleration $\vec{a}$ in equation (3) and exploiting the fact that the velocity field we are interested in is divergence-free simplifies the second divergence constraint (2) to

$$\Delta \bar{p} + \nabla \cdot (\vec{u} \cdot \nabla \vec{u}) = 0 \qquad (5)$$

Penalizing the left side of equation (5) via the cost function for every inner grid point will lead to a nonlinear weighted least squares problem (nonlinear in the variables for $\vec{u}$) that requires a different solver such as the nonlinear CG method or the L-BFGS method [Noc80]. From the four degrees of freedom per grid point, three velocity components and one pressure component, effectively two remain due to the two divergence constraints (1) and (5). Since the number of effective degrees of freedom did not increase compared to the first generation FlowFit for divergence-free velocity fields at the same grid resolution we expect to see an improvement of the reconstruction quality now that we are able to use 6 constraints for each detected particle (velocity and acceleration) instead of just 3 (velocity).

The expected gain of this joint optimization of velocity and pressure will likely also depend on the amount of information that is present in the measured particle velocities and accelerations in terms of the dynamic ranges. For example, in case the acceleration data is very noisy we would not expect the acceleration data to help much in improving the velocity field reconstruction. To avoid an amplification of noise, it is important to properly weight velocity and acceleration deviations in the cost function according to their estimated measurement accuracy.

From this point on, the different kinds of reconstruction problems that are possible within this framework are referred to as *div0*, *div1*, *div2* and *pot*. They differ in their choice of variables, input data and cost function. Possible inputs are $\vec{u}$ and $\vec{a}$, the measured velocities and accelerations of the particles. Possible variables are B-spline base function weights for $\vec{u}$ (velocity) and $\bar{p}$ (pressure). All variants are regularized by penalizing high frequency components to some extent (referred to as HF) via the cost function. The cost functions include a weighted sum of errors between measured and reconstructed data (U, A for velocity and acceleration respectively) and possibly other errors such as nonzero divergences of velocity (D1) or errors with respect to equation (5) (D2). See Table 1 for a summary of the kinds of reconstruction problems we tested in this work. Table 2 specifies the weights used for various errors in the cost function in terms of fitting parameters *penhf*, *pendiv* and an estimate r for the ratio between acceleration measurement error and velocity measurement error.

Variant	Input (via particles)	Grid variables for	Cost Function
div0	u	u	HF + U
div1	u	u	HF + U + D1
div2	u, a	$u, \bar{p}$	HF + U + A + D1 + D2
pot	a	$\bar{p}$	HF + A

Table 1: Different reconstruction variants and their properties

Error quantity	Weighting factor in cost function
Velocity	1.0
Acceleration	$1/r$
Highpass-filtered Velocity	$penhf$
Highpass-filtered Pressure	$penhf \cdot h \cdot 1/r$
Divergence of velocity (D1)	$pendiv \cdot h$
Divergence of velocity derivative (D2)	$pendiv \cdot h \cdot 0.1 \cdot 1/r$

Table 2: Weights of errors in cost function in terms of fit parameters *penhf*, *pendiv* and an estimate *r* for the ratio between acceleration measurement error and velocity measurement error r

The two remaining degrees of freedom per grid point is a property shared between variant *div1* and *div2*. But instead of three constraints per particle for the velocity components in *div1*, the variant *div2* makes use of six constraints per particle including the acceleration components. In

this respect, the reconstruction variant *div2* is similar to the VIC+ method developed by Schneiders et al [Schn15].

For solving the linear least squares problems we created out own C++ implementation of the CGLS method [Pai82]. For solving the nonlinear least squares problems *div2* we chose a freely available C implementation [libLBFGS] of the L-BFGS method [Noc80] simply because we were not able to find an implementation of the nonlinear CG algorithm that suited our needs.

Evaluating the cost function and its gradient in the nonlinear case is more difficult to implement compared to the linear optimization problems. We have developed a custom automatic differentiation library in C++ for this special purpose.

5. Evaluation of reconstruction methods using synthetic data

The use of synthetic data and simulated particles allows us to compare the reconstructions to the known ground truth. For this we chose "Forced Isotropic Turbulence" of the John Hopkins Turbulence Database [Li08]. The data provided in this database covers velocity and pressure for a cube of 1024^3 grid points with circular boundary conditions for various points in time. The flow is incompressible and the density is uniform. Using sinc-based interpolation we derived acceleration from the provided velocity and pressure data using the Navier-Stokes equation and the given kinematic viscosity for a $64 \times 64 \times 32$ subvolume. Such an interpolator was necessary to stay consistent with the DNS simulation with respect to freedom of divergence of velocity.

For different simulated particle densities, see table 3, we each computed four reconstructions named *div0*, *div1*, *div2* and *pot* in accordance to Table 1 based. Figures 2, 3 and 4 show vorticity isosurfaces of velocity fields reconstructions with increasing particle density from figure to figure. In each of the shown particle density cases we can clearly see that the reconstruction quality improves if more and more physically motivated regularizations are applied. At a particle count of 3277 the second generation FlowFit (div2) already comes visually close to the ground truth whereas the simpler linear reconstruction that ignores the divergence constraints entirely is not able to recover many structures. The performance of the first generation FlowFit (*div1*) which accounts for the divergence of the velocity field but does not yet make use of measured accelerations tends to be between that of *div0* and *div2*. But for higher particle densities (not shown) the reconstructions using *div1* and *div2* look very similar probably because the number of particles was large enough so that any undersampling problem could be mitigated by the by both kinds of regularizations.

Particle count	3277	6554	13107	26214
Particle spacing in terms of the Kolmogorov length scale	7.3117	5.8003	4.6061	3.6558

Table 3: Connection between particle count and particle spacing in terms of the Kolmogorov length scale
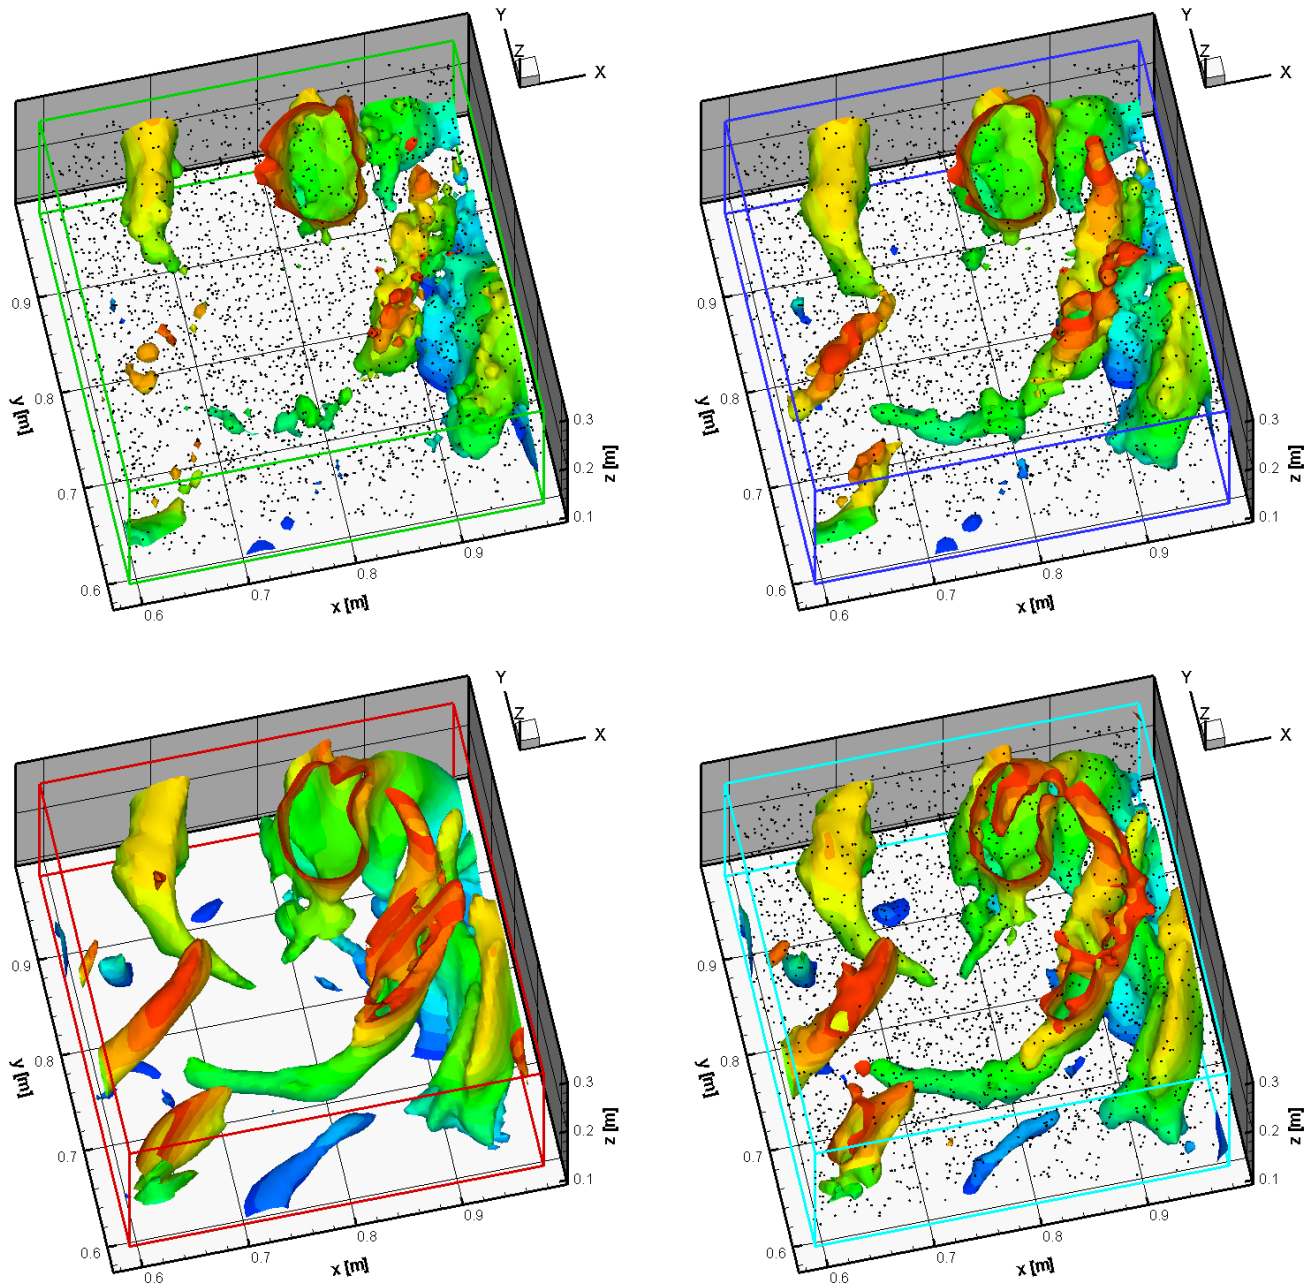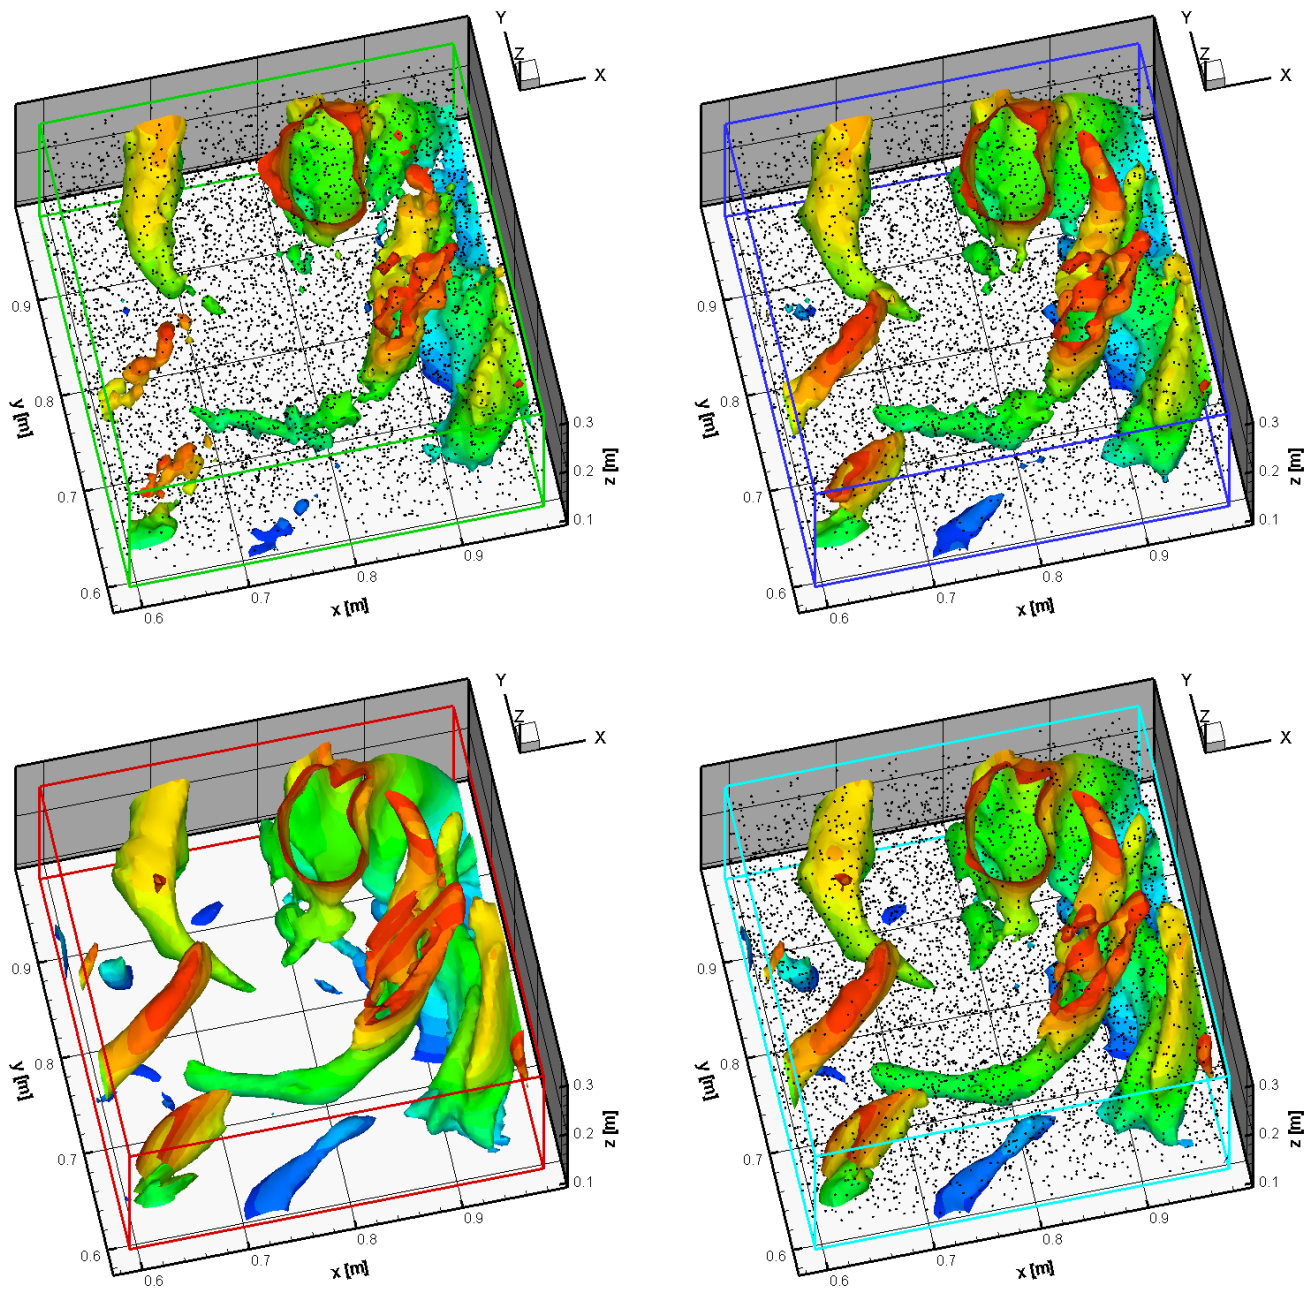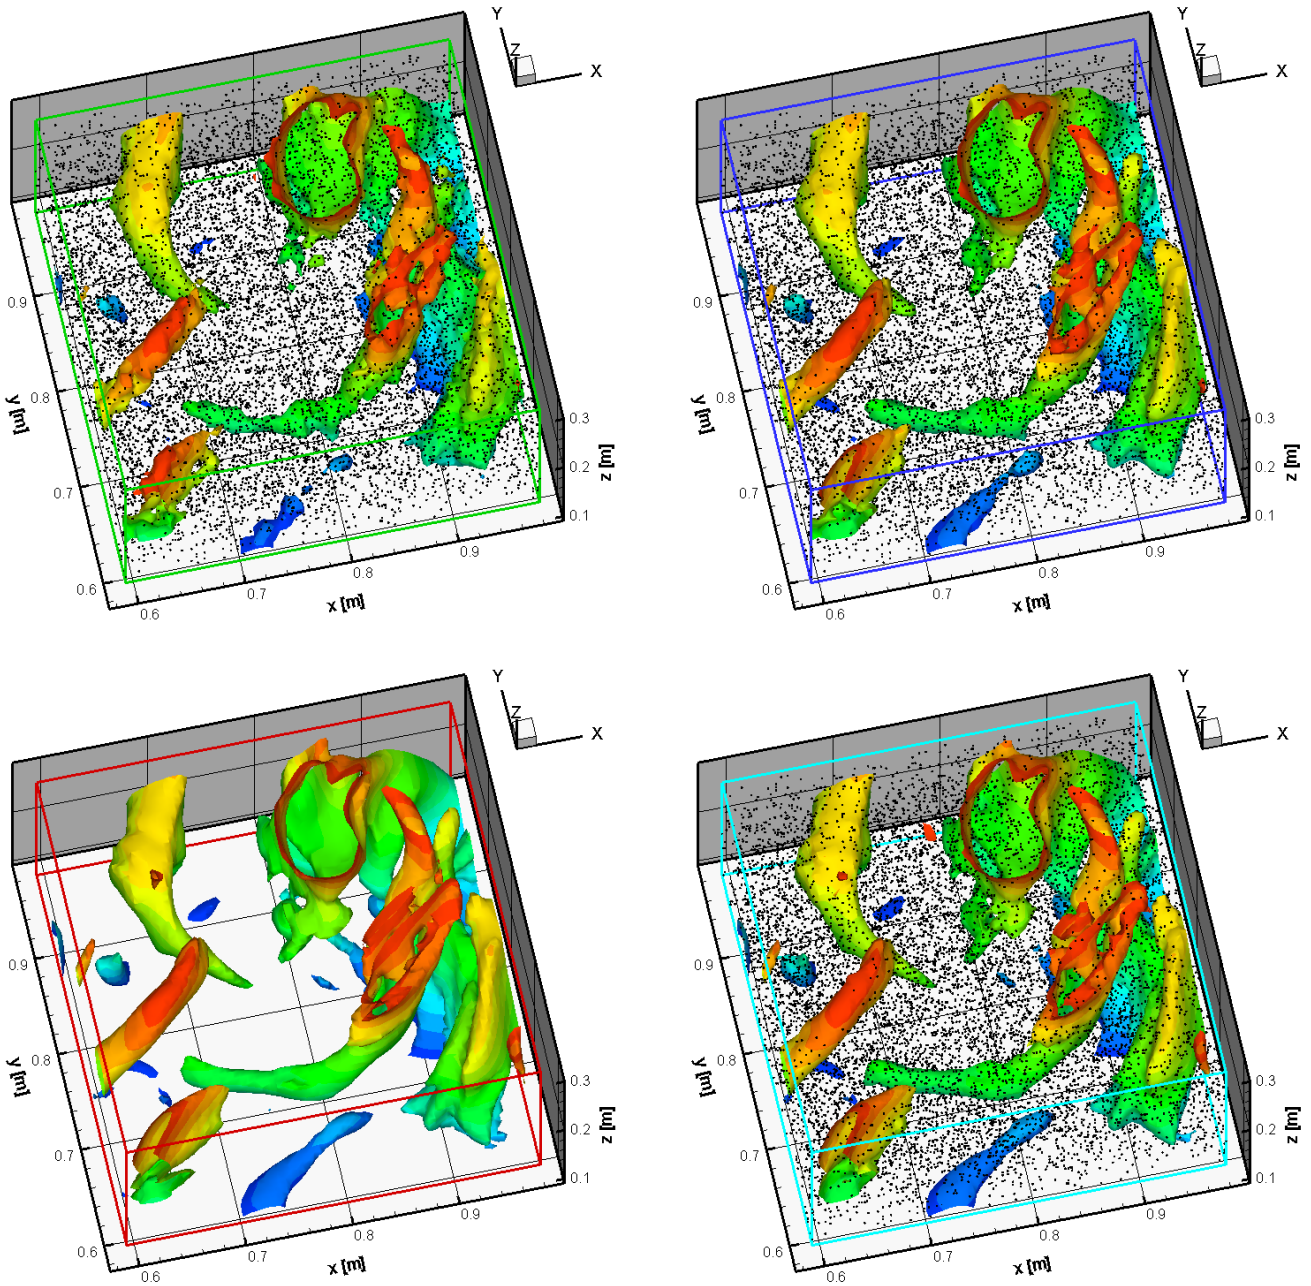

Figure 2: Reconstructions of vorticity based on **3277** simulated particles (black dots) compared to the ground truth. From top left to bottom left clockwise: div0, div1, div2, ground truth (color-coded z)

Figure 3: Reconstructions of vorticity based on **6554** simulated particles (black dots) compared to the ground truth. From top left to bottom left clockwise: div0, div1, div2, ground truth (color-coded z)

Figure 4: Reconstructions of vorticity based on **13107** simulated particles (black dots) compared to the ground truth. From top left to bottom left clockwise: div0, div1, div2, ground truth (color-coded z)

Another way to visualize the behavior of the reconstruction methods is to compute and plot the signal-to-noise ratio as a function of the wavenumber. We have done this analysis for multiple particle densities where in each case five particle distributions have been generated. The averaged signal-to-noise ratio curves are shown in figure 5.
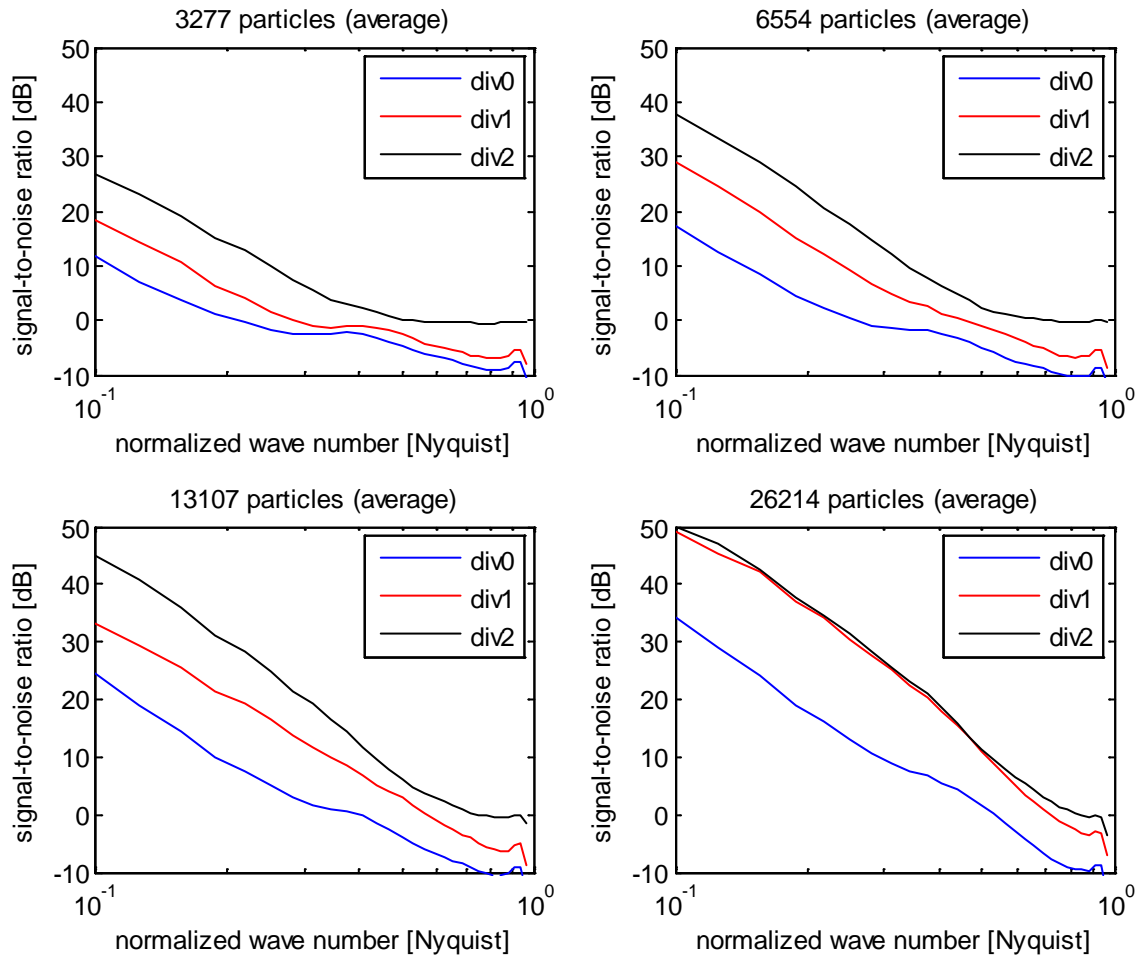
Figure 5: Signal-to-noise ratio with respect to velocity for different wavenumbers, particle densities and reconstruction variants

Penalizing the divergence of velocity locally in time (*div1*) improves the signal-to-noise ratio compared to the simpler *div0* reconstruction in all of the shown cases. Using the nonlinear optimization (*div2*) on the same particle data further improves the signal-to-noise ratio for all but the last case. The particle count in the last case seems to be large enough to avoid undersampling problems even if only the divergence of velocity is penalized locally in time (*div1*).

6. Summary and Conclusions

The presented B-spline based reconstruction algorithms yield a continuous representation in time and space for velocities, acceleration and pressure. This allows temporal and spatial supersampling. Compared to *div0*, the continuous fields computed by the *div1* variant are consistent with physical equations imposing conservation of mass ($\nabla \cdot \vec{u} = 0$) locally in time. In the synthetic test cases discussed here this method has been shown to provide better

reconstruction to a simple B-spline interpolation (*div0*). The *div2* variant accounts in addition for a vanishing temporal derivative of the divergence as well as the conservation of momentum (Navier-Stokes equation) which tends to further improve the signal-to-noise ratio. This effect can be seen in section 5 and confirms the results by Schneiders et al [Schn15].

To the best of our knowledge, this is the first time viscosity has been considered in the reconstruction of velocity and pressure based on scattered velocity and accelerations data. We have noticed that for the synthetic test case "Forced Isotropic Turbulence" ignoring the viscosity term leads to noticeably worse reconstructions.

These flow reconstruction methods complement the recent advances in particle tracking techniques on the path to high resolution 3D measurements.

References

[Eil96]   Paul H. C. Eilers, Brian D. Marx (1996) Flexible smoothing with B-splines and penalties. Statistical Science 11(2): 89-121

[Kal60]   Rudolf E. Kálmán (1960) A New Approach to Linear Filtering and Prediction Problems. Transaction of the ASME, Journal of Basic Engineering 82(1), pages 35-45.

[Käh16]   Christian J. Kähler, Tommaso Astarita, Pavlos P. Vlachos, Jun Sakakibara, Rainer Hain, Stefano Discetti, Roderick La Foy, Christian Cierpka (2016) Main results of the 4th International PIV Challenge. Exp. in Fluids, DOI:10.1007/s00348-016-2173-1

[Li08]   Y. Li, E. Perlman, M. Wan, Y. Yang, R. Burns, C. Meneveau, R. Burns, S. Chen, A. Szalay & G. Eyink (2008) A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence. Journal of Turbulence 9, No. 31

[Noc80]   Jorge Nocedal (1980) Updating Quasi-Newton Matrices with Limited Storage. Mathematics of Computation 35, pp. 773-782

[Pai82]   Christopher C. Paige, Michael A. Saunders (1982) LSQR: An Algorithm for Sparse Linear Equations and Sparse Least Squares. ACM Trans. Math. DOI: 10.1145/355984.355989

[Scha14]   Daniel Schanz, Andreas Schröder, Sebastian Gesemann (2014) Shake The Box – a 4D PTV algorithm: Accurate and ghostless reconstruction of Lagrangian tracks in densely seeded flows. 17th Int. Symp. Appl. Laser Tech. Fluid Mech., Lisbon

[Scha15]   Daniel Schanz, Sebastian Gesemann, Andreas Schröder (2015) Shake-The-Box: Accurate Lagrangian particle tracking at high particle densities. Exp. in Fluids

[Schn15]   Jan F.G. Schneiders, Ilias Azijli, Fulvio Scarano, Richard P. Dwight (2015) Pouring time into space. 11th Int. Symp. on Particle Image Velocimetry -- PIV15, Santa Barbara

[libLBFGS] http://www.chokkan.org/software/liblbfgs/ (retrieved 2015) C port of original Fortran L-BFGS implementation released under MIT license.

Combined time-resolved PIV and structure deformation measurements for aero-elastic investigations

Hauke Ehlers, Andreas Schröder,
Reinhard Geisler, and Sebastian Gesemann

German Aerospace Center (DLR), Institute of Aerodynamics and Flow Technology,
Bunsenstr. 10, 37073 Göttingen, Germany
`{hauke.ehlers,andreas.schroeder,`
`reinhard.geisler,sebastian.gesemann}@dlr.de`
`www.DLR.de`

Abstract. The aero-elastic behaviour of a thin plate forced to aero-elastic flutter has been investigated by applying two different optical measuring techniques. The aim of the experiment is to solve the collar-triangle of forces. In addition to the description of the wind tunnel experiment and the measuring techniques a method of calculating instantaneous pressure fields and aerodynamic forces from PIV data is shown and demonstrated for a selected test case.

Keywords: aero-elastic, flutter, particle image velocimetry, PIV, structure deformation, IPCT, instantaneous, time resolved

1 Introduction

The interaction of structural (elastic and inertial) as well as aerodynamic forces causes complex dynamic aeroelasticity problems. Part of this is the flutter phenomena which is potentially destructive and needs to be avoided in most cases. Great effort is made experimentally and numerically in considering flutter characteristics to avoid structural failure. This includes the safety of aeronautical light weight structures but also buildings and bridges.

Within the AFDAR project (Advanced Flow Diagnostics for Aeronautical Research) flutter of a thin plate in a low-speed wind tunnel has been investigated by using optical measuring techniques. The advantage is to measure contactless, non-destructive, time-resolved and with high accuracy. The acquired time-resolved data enables solving the collar-triangle of aerodynamic A, inertial I and elastic E forces (general aeroelastic stability problem: $A + I + E = 0$) and is also useful for computational code validation aiming in coupling CFD and CSM-codes in subsonic flows.

Recent efforts have derived pressure fields from planar velocity field data measured by Particle Image Velocimetry (PIV). The so called planar pressure imaging (PPI) is based on the momentum conservation principle and has been investigated by e.g. Bauer and Köngeter [2] and Oudheusden et al. [3] for steady

flows under approximately two-dimensional (2D) flow conditions. By measuring time resolved velocity data instantaneous pressure fields have been estimated on a square cylinder by Kurtulus et al. [4] and de Kat et al. [1].

High precision image based surface reconstruction methods are developed at DLR Göttingen. [examples]. Konrath et al. [7] investigated the simultaneous measurement of forces, model position, wing deformation and flow velocity fields for plunging wing models. An approach to determine inertial forces for a harmonically oscillating stiff wing models by using an optical measuring set-up is given by Ehlers et al. [5].

The present paper shows a generic flow-structure-interaction experiment which has been performed in a low speed wind tunnel at DLR Göttingen using a thin rectangular plate forced to aeroelastic flutter. Two optical measuring systems have been installed to synchronously measure the flow field around (high-speed PIV) and the deformation of the plate (high-speed IPCT). The PIV set-up enables for calculating time series of planar velocity vector fields. The estimation of the unsteady aerodynamic forces on the plate surfaces rely on the assumption of a 2D flow hence the thin plate is assumed to oscillate mainly in a 2D bending mode. By solving the incompressible momentum equation end spatially integrating the pressure gradients [1] for each time step and extrapolating the related pressure fields close to the plate surface along the span of the plate the time resolved aerodynamic force distribution results. The high-speed IPCT system is used to measure the 3D unsteady deformation of the thin plate hence this set-up is capable to detect higher modes of the bending oscillation. The inertia forces are calculated from the acceleration of the plate surface which is achieved by differentiating the deformation with respect to the time and by using the mass distribution of the plate [5]. The elastic forces are calculated from the measured deflection of the surface, the plate stiffness and by using an analytical thin flat plate deformation model. As the IPCT measuring set-up provides the 3D distribution of the two different structural forces, the 3D distribution of the aerodynamic forces can be calculated, too.

The focus of the present paper is in describing the experimental set-up in detail and giving some first results of the planar pressure fields...

2   Experimental Set-up and Measuring Procedures

The experiment has been performed in a circular Göttingen-Type low speed wind tunnel at DLR Göttingen with a cross section of 1.0×0.7 m. The turbulence level for flow velocities below 20 m s^{-1} is less than 0.5%. A flat plate supporting base with an elliptical leading edge has been installed within the wind tunnel test section. A thin plate model was fixed to one side downstream to the supporting base. The other edged of the model were freely suspended. The support ensures a laminar approaching flow overflowing the upper side of the thin plate only (Fig. 1). The model has a thickness of 0.05 mm and is made of hardened carbon steel. The plate dimensions are $(b \times l)$ 150×75 mm^2. By means of a FEM model the first three eigenmodes could be detected at $f = 75; 116; 221$ Hz. The model

has been specially prepared for the optical measuring techniques. Therefore the surface was coated with a random white dot pattern on a black background for the surface reconstruction measuring technique. In the central plane of the model a narrow band of high-gloss black dye was applied to optimize the PIV measurement close to the surface. Two glass plates on the lateral sides of the thin plate inhibit a 3D-flow exchange at the plate side edges. The influence of the glass end plates to the 3D characteristics of the flow is assumed to be low. The applied high-speed PIV system consists of a diode-pumped Nd:YAG Laser from Lee. Inc. (LDP-200MQG) providing 20 mJ per pulse at 5 kHz using both cavities simultaneously. The laser was guided through a set of light sheet forming lenses and a combination of mirrors enabling an illumination of a mist of 1 μm DEHS tracer particles in a central plane on the upper and lower side of the oscillating thin plate. For this purpose the light sheet was split in two by the mirror arrangement (see Fig. 2). The digital imaging system consists of a pco.dimax CMOS camera (12 bit) which has been operated at 5 kHz framing rate and a reduced resolution of 1344×688 Pixel (Pixel size: 11 μm). The camera sensor was aligned parallel to the illuminated measuring plane. A 85 mm Nikon lens at $f/\# = 2.0$ was mounted on the camera imaging a field of view of 66×130 cm^2. The recordings are processed with a 2D cross-correlation algorithm with a final interrogation window size of 18×18 pixels and an overlap of 66%. The image magnification factor was 10.22 pixel/mm which leads to a spatial resolution of 0.587×0.587 mm^2 for the resulting vector field. The multi-grid interrogation method was applied starting with an initial window size of 96×96 pixels. Assuming a precision of the determined displacements of better than 0.1 pixel, the uncertainty of the velocity vectors is less than 1% of U_∞. An automatic mask generation tool was constructed which is able to detect the line of reflection on the model surface by digital image processing (erosion/dilatation principles).

The high-speed IPCT set-up consists of two Photron SA1.1 CMOS cameras (12 bit) mounted on lens adapters allowing an adjustment of the Scheimpflug angle for oblique viewing directions. Two 60 mm Zeiss lenses have been used at $f/\# = 5.6$ to get the desired depth of field. The cameras have been operated at 2.5 kHz framing rate and a reduced resolution of 1024×544 Pixel (Pixel size: 20 μm). Two pulsed LED illuminators developed at DLR Göttingen have been used as light sources to illuminate the random pattern of small white paint dots on the upper side of the thin plate surface. The pulse length was 12 μs at 2.5 kHz. The illuminated surface has been imaged by both cameras in stereo viewing (see Fig. **??**). The size of the angle between the IPCT cameras was about 67°. A two-plane calibration target has been imaged within the measurement volume for calculating a function for all possible lines-of-sight of the two IPCT cameras. This function enables a local triangulation of the found dot correspondences from both camera viewings which have been estimated by using an initial mapping of the surface dot pattern image and a successive iterative 2D cross-correlation scheme (similar to a PIV evaluation). Hence the surface can be reconstructed. Before the dynamic measurements the static rest position of the model has been recorded as a reference position for later deformation calculations.

Both measuring systems have been operated simultaneously. For this purpose a controlling system was used consisting of a pulse generator as master clock (output frequency: 2.5 kHz) and two sequencers (see Fig. 3). *Sequencer 1* was triggered directly by the pulse generator and did provide the laser with a continuous output signal of 5 kHz. A second output signal of 2.5 kHz was used as input for the second sequencer. A manual switch was applied to interconnect *sequencer 2* which was programmed to control the exposure of the PIV camera (*pco*) synchronous to the laser and the IPCT measuring system (*Photron* cameras and *LED*s). The IPCT images have been recorded between two successive PIV recordings.
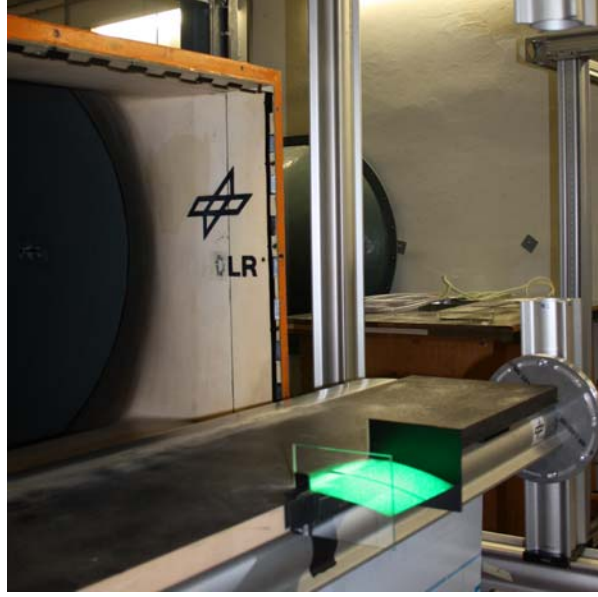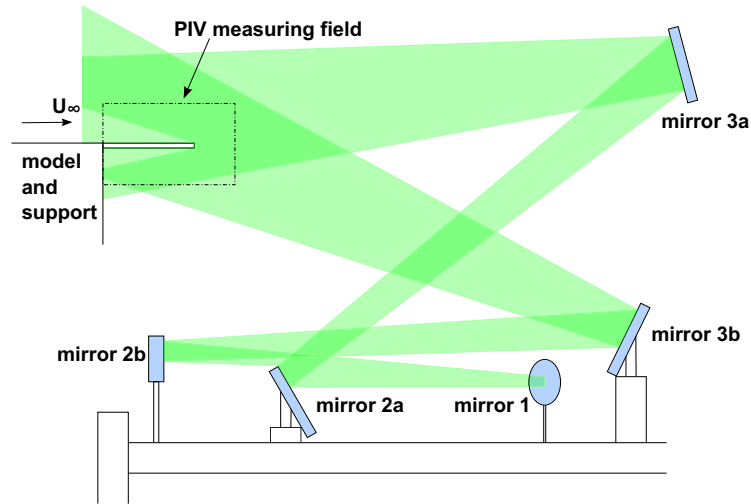

Fig. 1. Thin plate model (green illuminated) and glass end plates mounted downstream to the supporting plate; wind tunnel nozzle is seen in the background.

3   Pressure Calculation

The in-plane pressure gradients are obtained by the incompressible momentum equation. Written in conservative differential form for a fixed fluid element the equation is given as

$$\nabla p = -\rho \left[\frac{\partial \boldsymbol{u}}{\partial t} + \left(\boldsymbol{u} \cdot \nabla \right) \boldsymbol{u} - \nu \nabla^2 \boldsymbol{u} \right] \; . \tag{1}$$

Fig. 2. Part of the optical set-up for PIV measurement; light sheet is reflected into the measuring plane by *mirror 1*; *mirror 2a* only reflects the lower part of the light sheet to illuminate the upper part of measuring plane; *mirror 2b* is installed to illuminate the lower part of the measuring plane.
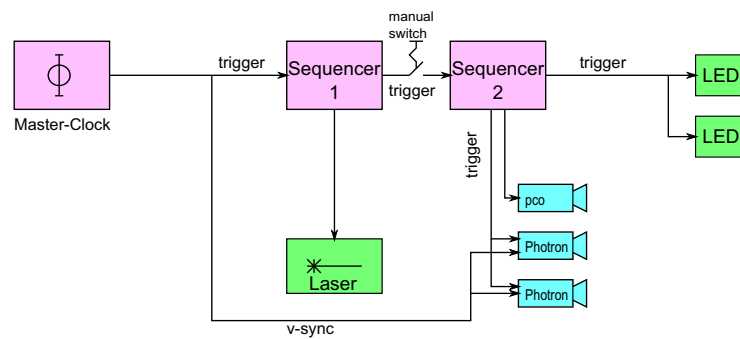
Fig. 3. Schematic of the measuring control system; the PIV system (*Laser* and *pco* camera) is operated at 5 kHz, the IPCT system (*LED*s and *Photron* cameras) is operated at 2.5 kHz.

In Eq. 1 ∇p is the pressure gradient, ρ the fluid density, $\boldsymbol{u}$ the velocity vector and ν the kinetic viscosity. If the flow field under investigation is nearly two-dimensional Eq. 1 can be reduced. In Cartesian coordinates we obtain the 2D form:

$$\frac{\partial \boldsymbol{p}}{\partial x} = -\rho \left[\frac{\partial \boldsymbol{u}}{\partial t} + \boldsymbol{u}\frac{\partial \boldsymbol{u}}{\partial x} + \boldsymbol{v}\frac{\partial \boldsymbol{u}}{\partial y} - \nu \left(\frac{\partial^2 \boldsymbol{u}}{\partial x^2} + \frac{\partial^2 \boldsymbol{u}}{\partial y^2} \right) \right] \,, \tag{2}$$

$$\frac{\partial \boldsymbol{p}}{\partial y} = -\rho \left[\frac{\partial \boldsymbol{v}}{\partial t} + \boldsymbol{u}\frac{\partial \boldsymbol{v}}{\partial x} + \boldsymbol{v}\frac{\partial \boldsymbol{v}}{\partial y} - \nu \left(\frac{\partial^2 \boldsymbol{v}}{\partial x^2} + \frac{\partial^2 \boldsymbol{v}}{\partial y^2} \right) \right] \,. \tag{3}$$

The terms on the right hand side of Eq. 2 and 3 can all be obtained from time resolved PIV measurements. For the assumption of 2D flow characteristics the 2D velocity field measured by a mono PIV set-up do suffice to determine the in-plane pressure gradients. For the derivatives with respect to the time a least squares differential operator is applied [6]. Spatial derivatives are approximated by central difference operators of 4th order.

The pressure filed is obtained by spatially integrating the pressure gradient field and using Dirichlet boundary conditions at the inflow boundary. Therefore an algorithm has been developed to solve the overdetermined system of equations ($\mathbf{M} \cdot p = \nabla p$) fast an precise. A 4th order approximation of the differentials is applied.

4 Results

For further descriptions a test case has been selected were the thin steel plate (150×75 mm^2; $t = 0,05$ mm) was forced to aero-elastic flutter in a steady-state oscillating mode. The free stream velocity was $U_\infty = 5.5$ m s^{-1}. Fig. 4 shows the measured velocity vector field for one instant. The model support is marked by a light grey box, the thin plate is shown as a black line. The boundary layer is clearly obvious. For this instant the flow tries to follow the curvature of the plate, separates near the first third of the plate length and a vortex is building up. The interaction of growing and convecting vortices leads to variations in the pressure field and stimulates the thin plate to oscillation. If the oscillation is assumed to be real periodic the deformation frequency of the model is about $1/T \approx 40$ Hz. In Fig. 5 the high resolution instantaneous pressure field calculated from PIV data by Eq. 2 and 3 is given at six different time instants of a period of motion. Colour-coded is the relative dynamic pressure. The profile of the thin plate and the model support are blanked. The time series illustrate the evolution and movement of low and high pressure regions.

5 Conclusion

A generic aero-elastic wind tunnel experiment has been carried out by means of two different optical measuring techniques. The time resolved data set enables for solving the collar-triangle of forces. The first step of calculating instantaneous pressure fields from 2D PIV data is shown. The next step will be to compare these results to the 3D approach of surface deformation measurements.
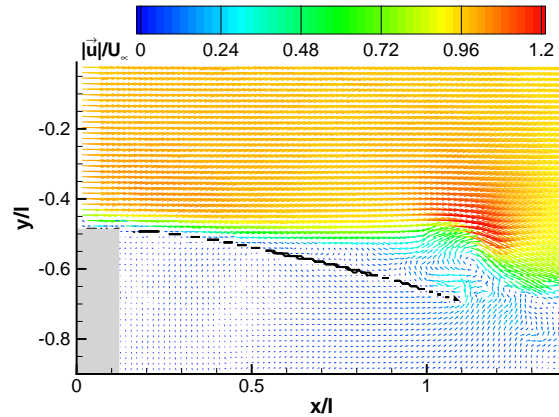
Fig. 4. PIV velocity vector field. Every second vector is plotted

Acknowledgments.

References

1. de Kat, R., van Oudheusden, B. W., Scarano, F.: Instantaneous planar pressure field determination around a square-section cylinder based on time-resolved stereo-PIV. In: 14th Int Symp on Applications of Laser Techniques to Fluid Mechanics. Lisbon (2008)
2. Bauer, T., Köngeter, J.: PIV with high temporal resolution for the determination of local pressure reduction from coherent turbulent phenomena. In: 3rd Int. Workshop on PIV, pp.671–676. Santa Barbara (1999)
3. Oudheusden, van B.W., Souverein, L.J.: Evaluation of the pressure field from PIV in a shock wave boundary layer interaction. In: 7th Int. Symposium on Particle Image Velocimetry. Rome, Italy (2007)
4. Kurtulus, D.F., Scarano, F., David, L.: Unsteady aerodynamic forces estimation on a square cylinder by TR-PIV. Experiments in Fluids. 42, 185–196 (2007)
5. Ehlers, H., Konrath, R., Agocs, J., Radespiel, R., Wokoeck, R.: Ermittlung der Massenkräfte periodisch bewegter Tragflügelmodelle unter Anwendung von optischer Messtechnik. In: Deutscher Luft- und Raumfahrtkongress 2001, Paper DLRK 2011-241269. Bremen, Germany (2011)
6. Raffel, M., Willert, C., Wereley, S., Kompenhans, J.: Particle Image Velocimetry: A Practical Guide. Springer-Verlag, Berlin, Germany (2007)
7. Konrath, R., Schlager, B., Kirmse, T., Kompenhans, J., Möller, T.J., Wokoeck, R., Emge, M., Radespiel, R.: Simultaneous Measurements of Unsteady Aerodynamic Loads, Flow Velocity Fields, Position and Wing Deformations of MAVs in Plunging Motion. In: Dillmann, A., Heller, G., Klaas, M., Kreplin, H.-P., Nitsche, W., Schröder, W. (eds.) New Results in Numerical and Experimental Fluid Mechanics VII. NNFM, vol. 112, pp. 331–338. Springer, Heidelberg (2010)
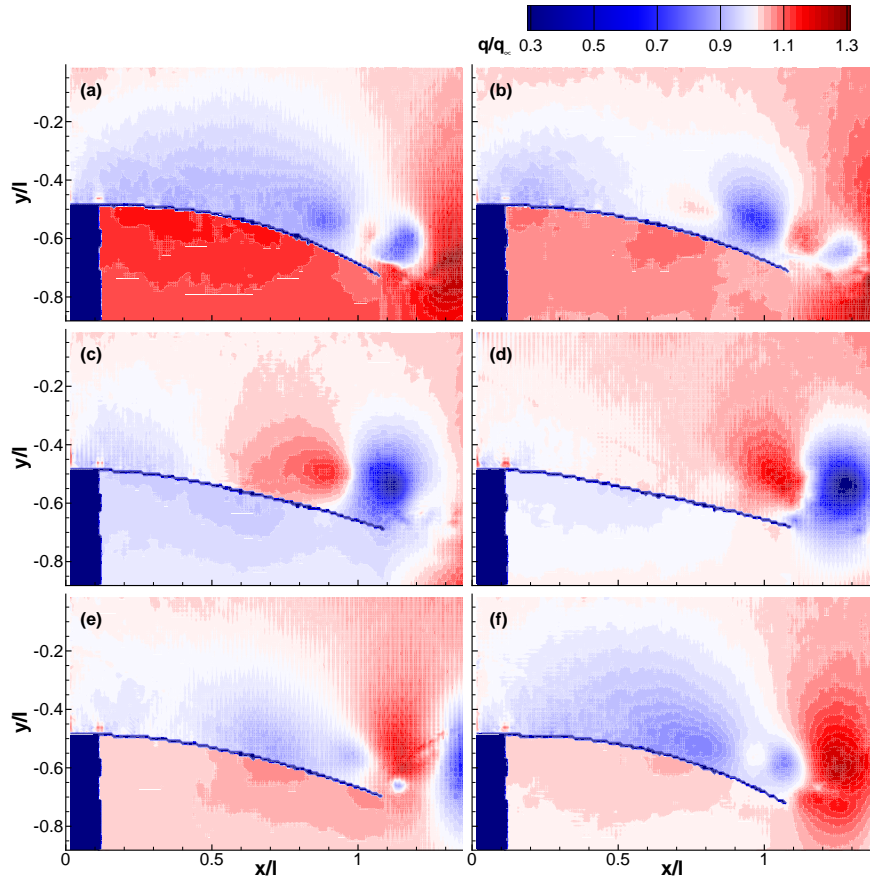
Fig. 5. Dynamic pressure results calculated from PIV data; (a)-(f): $t/T = 0.0 : 0.167 : 1.0$.

Near-wall turbulence characterization using 4D-PTV Shake-The-Box

Andreas Schröder[1], Daniel Schanz[1], Reinhard Geisler[1], Sebastian Gesemann[1] and Christian Willert[2]

[1] Institute of Aerodynamics and Flow Technology, DLR, Göttingen, Germany
[2] Institute for Propulsion Technology, DLR, Cologne, Germany
andreas.schroeder@dlr.de

ABSTRACT

A near-wall flow characterization including measurements of instantaneous wall-shear stresses of a turbulent boundary layer (TBL) developing along a flat plate with zero pressure gradients has been performed by using two advanced particle based measurement methods. The experiments were conducted in the 1m- wind tunnel facility of Göttingen at $U_\infty = 10$ m/s and a Reynolds number of $Re_\theta = 2{,}770$ corresponding to $Re_\tau = 960$. First, high-speed 2C-PIV was performed at two image magnification factors at 2 kHz and 4 kHz frame rates in order to obtain the overall statistical properties of the boundary layer profile together with a large time series of instantaneous 2C velocity vector fields in a streamwise, wall-normal plane. Single pixel line correlation applied to the particle image area close to the wall provided high spatial resolution velocity data down into the viscous sublayer. In a second step, the novel 4D-PTV technique Shake-The-Box (STB) was applied to a time series of particle images acquired with a typical tomographic high-speed camera set-up at 15.873 kHz. The STB measurement domain consists of a wall-bounded volume covering a stream- and spanwise area of 430 x 430 viscous units ($l^+ = \nu/u_\tau$) and a wall-normal extension of 32 viscous units. A comprehensive set of relatively dense Lagrangian track data was reconstructed from two time resolved sequences of 115,000 time steps each. The data enables an accurate and very high resolution measurement of the mean and rms-velocity profiles averaged in bins sized by a fraction of a viscous unit, of both components of the instantaneous wall shear stress (τ^+_w) and all components of the Reynolds stress tensor. Furthermore, the time-resolved 3D velocity vectors and corresponding gradient tensor have been interpolated onto a regular grid using the time series of irregularly distributed Lagrangian track data. With the present data coherent structures and their dynamics close to the wall can be investigated together with their role for various (rare) wall shear stress events. The STB method is proven capable of coping with strong velocity gradients close to walls and can also be extended to TBL flows with much higher Reynolds numbers.

INTRODUCTION

The role of turbulent flow features for the momentum transport near walls and corresponding (large) wall-shear stress events are prominent research topics for aerodynamics and technical flows. Large portions of the total aerodynamic drag are produced by skin friction along surfaces underneath turbulent boundary layers (TBL). Therefore, a better understanding of the related fluid dynamics is of high importance for advanced drag reduction strategies, high-lift-configuration design and the enhancement of flow- and separation control devices. Furthermore, advanced CFD validation procedures and recent LES wall-model developments requires accurate (and time-resolved) near-wall turbulence data at high Reynolds numbers provided by statistically converged data-sets. Nowadays, DNS of TBL flows is limited to moderate Reynolds numbers ($Re_\theta < 5{,}000$), so that appropriate measurement techniques that are able to deliver unsteady (or even time-resolved) three-component velocity information at industrially relevant Reynolds numbers at high spatial and temporal resolution, preferably at many points simultaneously, are highly welcome. All known measurement techniques are confronted with significant challenges in providing reliable data in close proximity to rigid surfaces due to strong mean and instantaneous velocity gradients. The use of accurate probe measurement techniques is limited due to its intrusiveness, coupling effects with the wall and spatial extensions (causing low-pass filtering effects) [1] and even modern miniaturized (pulsed) hot-wire probes or μLDA systems [2][3] are still suffering from spatial filtering effects close to the wall and are only capable of providing single-point information. So far on the experimental side only oil-film interferometry [4][5] is a widely accepted method for delivering the mean wall-shear stress magnitude and skin friction coefficient C_f value accurately. Nevertheless, in order to understand the (near-wall) boundary layer flow causing specific wall-shear stress events and producing drag a pure measurement of C_f would need to be complemented. In this sense a method is desired that is capable of providing both

both components of the unsteady wall-shear stress along with the friction velocity vector and at the same time provides multi-point 3C-velocity vectors in the near wall region in a non-intrusive manner, without biasing effects and at high accuracies. Therefore we will first have a look on recent developments of non-intrusive particle based planar and volumetric measurement methods.

PIV has undergone significant progress in the past three decades in terms of spatial and temporal resolution [6]. Within the recent years tomographic PIV [7] has matured into a reliable tool for investigating turbulent flows and delivering the complete 3D velocity gradient tensor [8]. However for all PIV based approaches the spatial resolution of the local instantaneous velocity measurement is limited by the finite size of the cross-correlation window, which acts as a low-pass filter biasing the measurement particularly in presence of strong velocity gradients. Two further basic limitations of PIV can be described by the dynamic velocity range (DVR) which is typically in the order of 1:100 and the dynamic spatial range (DSR), which is basically limited by the resolution of the employed camera sensor [9]. Several methods have been developed in order to overcome these limitations; nevertheless, only particle-tracking approaches have been found capable of delivering reliable results for at least the mean flow statistics near interfaces and walls [10] or in strong-shear flow. The first successful experiment using a 3D particle tracking method in the near wall region of a moderate Reynolds number turbulent channel flow at $Re_\tau = 1,400$ was based on digital holographic particle microscopy: A reconstruction of 3D particle distributions of a series of frame-straddled double images has been realized by Sheng et al. [11]. Results of mean and rms-profiles, fluctuation components of the wall-shear stress and spatial topologies of some velocity vector volumes in the direct vicinity of the wall have been presented. Nevertheless, the temporal resolution, a higher accuracy of the velocity estimation and statistical convergence for higher order moments are still missing.

Recently, the Shake-The-Box (STB) technique [12][13] has been developed, which is a 4D-PTV evaluation method for densely seeded flows capable of coping with ill-posed 3D particle reconstruction problems based on few camera projections by seizing the temporal information with predictive steps and applying an iterative particle reconstruction and image matching scheme (see Iterative Particle Reconstruction [14]). Within the resulting dense Lagrangian tracks the STB technique uses temporal fitting functions based on optimal Wiener filtering along all found particle paths. The parameters of an optimal Wiener filter are determined from statistical properties of the Lagrangian position, velocity and acceleration fluctuations along the reconstructed 3D particle positions and tracks for all three components separately. This optimal temporal filtering approach enables an accurate estimation of position, velocity and acceleration vectors and enhancing the DVR to values > 1:1000, when sufficient track lengths are provided. Therefore, STB is able to deliver accurate mean- and Reynolds stress values by bin-averaging (down to sub-pixel spatial resolution) and additionally provide the complete instantaneous velocity gradient tensor at a relatively high spatial resolution (comparable to a very well resolved tomographic PIV measurement) by using a proper interpolation scheme given with the "Flow-Fit" algorithm that was recently developed by DLR (brief description in [15]). Therefore it is expected that with STB valuable data for turbulence characterization with outstanding temporal and spatial resolution especially in (wall bounded) shear flow can be obtained.

In general, there is a strong need in fluid mechanics and turbulence research for defining scaling parameters which enable the description of mean and unsteady characteristics of turbulent flows in the form of laws, (self-similar) statistics and trends based on respective definitions of Reynolds numbers. In turbulent boundary layer (TBL) flows with zero pressure gradients (ZPG) mainly two scaling variants are used: The so called "outer scaling" is based on the boundary layer thickness δ, while the "inner scaling" is based on the friction velocity:

$$u_\tau = \sqrt{\frac{\tau_w}{\rho}} = \sqrt{\nu \frac{\partial u}{\partial y}}|_{y=0} \, .$$

Here τ_w is the mean wall-shear stress, ρ the density and ν the kinematic viscosity of the fluid. With the friction velocity u_τ a direct normalization of all measured velocities and length units can be realized. By the definition of a viscous length unit $l^+ = \nu/u_\tau$ the wall-normal position for instance is converted to a non-dimensional unit $y^+ = y/l^+$ (s.c. wall unit). Consequently, for all normalized parameters and statistics based on inner scaling an accurate estimation of the wall-shear stress and wall-position is mandatory. The mean wall-shear stress can be estimated by the velocity gradient within the viscous sub-layer $y^+ < 5$ in which a linear velocity profile can be assumed down to the wall. While accepted throughout the community, it should be noted that this scaling in principle only applies to the mean quantities. In practice however, the TBL unsteady flow may be nonlinear in the immediate proximity of the wall ($y^+ < 4$) due to high speed bursts or other extreme events, which is indicated as well by the present STB data. As pointed out by Brücker [16] the classical definition of the $y^+ = 5$ thickness of the viscous sublayer may need some revision in the context of unsteady near wall fluid dynamics of the TBL.

Today, there is still a need for improved measurement methods that can deliver the distribution of (instantaneous) wall shear-stresses and related skin-friction coefficients C_f together with near-wall velocity field information, especially at high Reynolds numbers. Furthermore, there is a strong necessity of accurately estimating the wall position [17] in order to calculate the wall normal velocity gradient du/dy correctly and perform accurate statistics for all profiles normalized by the

corresponding friction velocity u_τ. When using probe measurement techniques there are intrinsic difficulties and uncertainties of applying proper (optical) measurement techniques to achieve the relative probe position. Even with accurate position information the measured mean velocity profile and respective fluctuation component are still biased close to the wall due to intrusiveness and spatial filtering effects. A way to correct for biased probe measurement results applies assumptions of the linear profile or the self-similarity of the u-velocity PDF close to the wall [1]. Nevertheless, a direct non-intrusive and statistically bias-free measurement of all three velocity components near the wall and the wall-position itself would be preferable, especially because at high and aerodynamically relevant Reynolds numbers probe techniques reach ultimately their hardware limits. Based on results from DNS and a set of measurement data a formula has been built estimating the Reynolds number trend of the normalized RMS-value of the wall-shear stress to $\tau^+_{w,\,rms} = 0.298 + 0.018 \cdot \ln(\mathrm{Re}_\tau)$ [18]. The low-pass filtering effect of a limited hot-wire length of $l > 11 \cdot l^+$ produces a parallel shifted line at $\tau^+_{w,\,rms} = 0.240 + 0.018 \cdot \ln(Re_\tau)$ [19] for a series of high Reynolds number measurements. The authors correct for this effect by using low Reynolds number DNS values of the universal high-frequency part of the $\tau^+_{w,\,rms}$ distribution.

In recent DNS it has been observes that for increasing Reynolds numbers a growing fraction of the wall-shear stress is negative [20][21], which implies reverse flow events close to the wall, while Eckelmann wrote 1974 based on experimental investigations at relatively low Reynolds number turbulent channel flows using hot-wire and hot-film techniques [22]: "It can be stated with certainty, therefore, that there are no negative velocities near the wall".

Planar sensor array techniques which are mounted close to or at the wall surface like micro-pillars [23] or film-based methods [24] are delivering both components of the instantaneous friction velocities u_τ and w_τ with only slight intrusive interaction (micro-pillars cross through the viscous sublayer and surface films moves slightly), but need to be calibrated carefully, have a limited accuracy and show frequency response filtering properties. Both techniques would need a combination and synchronization with velocity measurement techniques in the flow region above the wall in order to gain insight into the near-wall flow dynamics and its effect on the measured wall-shear stresses.

For PIV methods the only way of resolving larger parts of the huge range of scales involved in high Reynolds number TBL flows is combining several cameras simultaneously in order to create large fields-of-view where low-frequency events dominate and additionally embed (µ)PIV resp. PTV systems with a large image magnification in order to view details of small scale flow structures and strong velocity gradients e.g. in a region close to the wall enabling the measurement of the mean [25] and instantaneous wall shear stress (slightly low-pass filtered) [26].

The above mentioned and other numerous efforts in developing and applying probe and particle based optical measurement methods in high Reynolds number turbulent shear flow are getting closer and closer to the desired goal, but they are (more or less) complementary and still limited at required resolutions either spatially, temporally or both. In the present work STB shall be introduced as a measurement method which offers the capabilities of delivering 3D non-intrusive, bias-free and high spatially and temporally resolved velocity (and acceleration) data for turbulent shear flow investigations. On the hardware side STB requires a typical multi-camera tomographic particle imaging system *and* time-resolution of the particle illumination and imaging with respect to the flow motion. So far appropriate time-resolved tomographic PIV experiments, which served as data for STB evaluations, have only been realized in low-speed water flows [12][15]. One way to overcome the hardware limitations in order to achieve particle tracking in high speed flows is the proposed multi-pulse STB approach (Novara et al. [27]). The other possibility is pushing the limits of existing hardware. Usually in relevant low-speed air flows (typically $U_\infty > 5$ m/s) the time-resolution of 3D PIV measurements using aerosol particles sized at $d_p \sim 1\,\mu m$ is limited to the frame-straddling method due to camera frame rate- and laser power limitations. In order to reach particle image frame rates of $f > 15$ kHz useful for particle tracking over many time-steps in low-speed air flows the resolution of existing cameras has to be reduced together with the size of the illuminated measurement volume. Today, high-speed CMOS cameras are available which deliver up to 25 kHz frame rates at 1 Mpx resolution (unfortunately often with relatively large pixel sizes). Together with pulse-burst lasers [28][29], which are capable of emitting pulse trains in the order of typically 10 to 100 laser pulses at very high frequencies (up to the MHz regime) and at laser pulse energies in the order of 5 to 200 mJ such cameras aligned in a tomographic set-up would enlarge the capabilities of STB for densely seeded Lagrangian turbulence investigations in a large variety of subsonic flows: small tracer particles in relatively large volumes could be illuminated sufficiently at several successive time steps enabling convergence of the STB algorithm even for high seeding densities [13], while pixel locking effects due to large pixel sizes of present high-speed cameras could be diminished using diffusor disks in front of the camera lenses.

Compared to such high cost solutions an efficient high-repetition 2C-PIV application for calculating turbulence statistics of a low Reynolds number turbulent boundary layer flow in air has been developed and performed by Willert [30] and will be applied in the following to characterize the overall boundary layer. In the second part of this paper the STB technique will be employed to provide, unsteady time-resolved 3-D data in the near-wall region.

EXPERIMENTAL SET-UP AND PROCEDURE

In the present study a high-repetition 2C-PIV implementation as described in [30] was applied at two large image magnification factors to characterize the turbulent boundary layer flow at $U_\infty = 10$ m/s free stream velocity covering the boundary layer flow over its whole thickness δ in y-direction in order to gain the overall statistical properties. Furthermore, STB has been applied to a thin wall-parallel volume at a high framing rate of 15.873 kHz in order to characterize near-wall turbulence and its spatial and temporal structures at a significant Reynolds numbers in air flow at very high spatial resolution. The measurement campaign was performed in the closed test section of the newly refurbished 1m-Wind Tunnel at DLR Göttingen with a cross-section of 740 x 1,000 mm² and a test section length of 3,000 mm. The boundary layer has been tripped by sandpaper stripes and zig-zag bands right after the contraction section. For both experiments DEHS particles with a mean diameter of ~ 1μm were generated by Laskin nozzles and introduced into the circuit wind tunnel enabling a homogenous distribution with adaptable seeding densities within the measurement volumes. The coordinate x refers to the streamwise direction while y and z respectively represent the wall-normal and spanwise directions. At a free stream velocity of $U_\infty = 10$ m/s the tripped boundary layer flow develops along $x = 2$ m at the lower wall of the test section to a zero pressure gradient turbulent boundary layer (TBL) (see Figure 1) with a Reynolds number based on the momentum thickness of $Re_\theta = 2,770$ corresponding to $Re_\tau = 960$, based on the friction velocity and boundary layer thickness δ_{99}. At the chosen stream-wise position the boundary layer thickness is estimated by the high-resolution profile 2C-PIV method at $\delta = 43.2$ mm and $\delta_{99} = 35.4$ mm with a skin friction velocity of $u_\tau = 0.418$ m/s and corresponding wall-unit size of $l^+ \sim 36.58$ μm (see Table 2 for further quantities). Then the viscous sub-layer height of $y^+ = 5$ corresponds to $y = 183$ μm. For the estimation of general measures two sets of in total ~100,000 particle images have been used for the statistics based on two large magnification and high-repetition 2C-PIV profile measurements.

For the corresponding set-up a roughly 5 mm wide measurement area was illuminated by a pair of relatively small sized externally modulated continuous wave lasers (*Kvant Laser*, Slovakia) with a combined output power of about 10 W at a wavelength of 520 nm. The non-collimated laser beam with a size of about 6×2 mm² was focused into a uniform 6 mm wide light sheet using a cylindrical lens ($f = 200$ mm). The resulting waist thickness was on the order of 200 μm before entering the wind tunnel glass panel from below (see Figure 1, left). Two imaging configurations were chosen. A 180 mm lens (*Zeiss Sonnar* T2.8/180) with extension tube imaged the complete boundary layer thickness ($m = 0.47$). Increased spatial resolution with $m = 1.1$ magnification factor was provided by a 400 mm lens (*Zeiss Makro Planar* T2.0/100 with two *Zeiss Mutar II* 2x teleconverters). The images were recorded with a CMOS high-speed camera (*PCO Dimax-S4*, 36GB) with a reduced field of view of 200 pixels width that allows the capture of up to 63,464 double-images at an image height of 1008 pixels. In order to keep the particle displacements at moderate levels around 30 pixels the laser pulse separation was set to $\Delta t = 80$ μs at $U_\infty = 10$ m/s and $m = 0.47$, and reduced proportionally at higher magnification (c.f. Table 1).

The acquired data was processed using a conventional 2-C PIV processing package featuring a coarse-to-fine resolution pyramid with intermediate image deformation (*PIVview2C, PIVTEC GmbH*, Germany). To obtain reliable mean velocity data and statistics within close proximity to the wall a high aspect ratio image sampling window of 64 pixels width and 12 pixels height was chosen. This corresponds to 1.5×0.28 mm² at $m = 0.47$ and 0.64×0.12 mm² at $m = 1.1$. For the latter the sample has an effective size of $17.5x^+ \times 3.28y^+$ at $U_\infty = 10$ m/s. The sample overlap was set at 75%.

Table 1: 2C-PIV image acquisition parameters

Magnification	m	-	**$m = 0.47$**	**$m = 1.10$**
Image size	(W x H)	pixels	200 x 2016	200 x 1008
Number of samples	N	-	31,738	63,464
Sample frequency	f_{acq}	Hz	2,000	4000
Laser pulse separation at 10 m/s	Δt	μs	80	40
Sequence length	T	s	15.87	15.87
BL turnover times at 10 m/s	$T\,U_\infty / \delta$	-	4,480	4,480
Eddy turnover times	$T\,u_\tau / \delta$	-	187.4	187.4

In a second step a STB measurement was performed at the same streamwise wall position and wind tunnel operating conditions. Illumination was realized using a *Quantronix DarwinDuo* Nd:YLF high-repetition laser with ~4 mJ energy per pulse at a repetition frequency of 15.873 kHz (two cavities operating at 7.937 kHz). The circular laser beam was shaped into an oval light profile, collimated, cut by a knife-edge and back-reflected onto itself by a mirror placed at the opposite side wall in order to increase the particle light scattering intensity for all in-line viewing directions. The rectangular shaped laser beam crossed the test section in spanwise directions and was aligned with the wall as close as possible. With a reduced resolution of 528 x 420 pixel four *PCO Dimax-S4* arranged in an in-line camera set-up below the wind tunnel enables a

high frame rate imaging of the illuminated particles in the measurement volume. Volumetric camera calibration was performed by using a 2D-glass target imaged at two planes with 2 mm separation. Volume-self-calibration [31] reduces the initial calibration error down to 0.1 px and parametrization of the OTF [32] was performed in order to account for present astigmatism effects at particle images during STB evaluation. At a magnification factor of ~1:3 and 11 µm per pixel, one pixel corresponds to 30 µm leading to a size of the measurement volume in x-z-directions of ~16 x 16 mm^2 (see illuminated area in Figure 1 (right)). This field is large enough to resolve a few near-wall low-speed streaks with an estimated average spanwise distance of $z^+ = 100$. In wall normal direction (y) the homogenous part of the light volume was about 1.2 mm thick which corresponds to an upper border of the 3D domain at $y^+ = 32$. Within the volume velocities up to $u = 7$ m/s corresponded to particle image shifts of ~13 pixels at the given 15.873 kHz frame rate, which is well suited for the STB tracking algorithm as most of the particles stay inside the volume for more than 40 time steps. For STB evaluation a predictor was used, which was gained from a STB pre-evaluation in form of a velocity profile. Additionally, its local rms-values for all three components have been used in order to regularize the particle search and shake radius (0.3 px + 5 times the local rms-component). Important values of the STB set-up are given in Table 2.

The aim of the present investigation is to acquire the full 3D velocity vector field of irregularly distributed particle tracks including the viscous sub-layer and the buffer layer with the region of the maximum stream-wise fluctuations $<u'u'>^+$ at $y^+ \sim 13$. At particle image densities between approximately 0.015 and 0.035 ppp (particles per pixel) up to 5,700 particles could be identified and tracked by STB for each time step. Of these particles approximately 700 particles are contained in the viscous sub-layer below $y^+ = 5$ and can be used to estimate the time-resolved skin friction velocity vector field in the measurement area of 16 x 16 mm or 430 x 430 viscous units for both stream- and spanwise components u_τ and w_τ using the dynamic viscosity µ and the near-wall velocity gradients du/dy and dw/dy respectively.

From ~ 230,000 time steps corresponding to 4,093 boundary layer turn over times U/δ or ~140 eddy turn over times u_τ/δ reasonable statistics of the near-wall mean velocity profile, together with all Reynolds stresses respectively skin friction velocity components, were then calculated from the volumetric Lagrangian track results from STB evaluation. The exact wall position was extracted from spatially highly resolved mean velocity profiles using bin averaging of the STB tracks in bin sizes of 1/13th l^+ at four sub-volume locations and extrapolating the linear profile between $2 < y^+ < 4.5$ to zero u-velocity. After estimation of the wall position span- and streamwise symmetry in a statistical sense for the relatively small volume can be assumed, so that wall normal 1D-profiles with a very high spatial resolution and high number of independent velocity data per bin can be extracted. The averaging process corresponds to 14.49 sec integration time at 15.873 kHz frame rate using ~8·10^8 data points.

Table 2: STB image acquisition parameters

Magnification	m	-	**$m = 0.36$**
Image size (each of four cameras)	(W x H)	pixels	528 x 420
Number of samples	N	-	230,000
Sample frequency	f_{acq}	Hz	15,873
Laser pulse separation ($1/f_{acq}$)	Δt	µs	63
Sequence length	T	s	14.49
BL turnover times at 10 m/s	$T\,U_\infty\,/\,\delta$	-	4,093
Eddy turnover times	$T\,u_\tau/\,\delta$	-	172.3

Figure 1: *Left:* Imaging setup for boundary layer profile measurements by high-repetiton-2C-PIV in the 1m-Wind Tunnel of DLR Göttingen. *Middle:* Collimated laser light volume with passe-partouts parallel to the wall *Right*: In-line camera high-repetition STB set-up with four *PCO Dimax-S4* cameras viewing through a glass insert of the wall.

BOUNDARY LAYER CHARACTERIZATION

As demonstrated in [30] the velocity profile and, in particular, the accompanying higher order statistics of a turbulent boundary can be obtained using long continuous PIV records of length $O(10^4-10^5)$. With a few modifications the methodology is extended for the present wind tunnel application. Two challenges had to be addressed: the distance to the measurement location is increased to about 550 mm while at the same time the free stream velocity was increased from 5 m/s up to 10 m/s. The latter results in a reduction of the viscous length scale which requires an increase in spatial resolution on the detector side. At the same time the higher velocity and higher magnification entails a proportionate reduction of the laser pulse separation.

Estimates of the mean and unsteady wall shear rate $\gamma = du/dy|_0$ were obtained using a single-line cross-correlation approach as described in [30] (i.e. the sampling window only has a wall-normal size of one pixel). The characteristic parameters for the boundary layer are summarized in Table 3. The measured data is normalized with inner variables using the traditional viscous scaling for velocity u_i^+ ($= u_i/u_\tau$) and length l_i^+ ($= l_i\, u_\tau/\nu$). In this sense the mean velocity profiles for u for both Reynolds numbers are shown in Figure 2 (left). The corresponding Reynolds stress variances $<u'u'>^+$, $<v'v'>^+$ and covariances $<u'v'>^+$, are provided in Figure 2 (right). Both pairs of plots also contain reference data from a DNS of a ZPG TBL provided by Schlatter et al. [18]. For the most part, the agreement between experiment and simulation is very good (the lines nearly coincide). Several discrepancies can be observed nonetheless: in the experiment the wake region above $y^+ > 200$ is less pronounced while the peaks of the streamwise variances are shifted outward by 1-2y^+. In the near wall region ($y^+ < 10$) the finite sized sample window results in significant departure from the predicted nearly linear velocity behavior, especially at the lower magnification.

Plots of the third and fourth moments of the streamwise velocity u are presented in Figure 3. While overall agreement between experiment and DNS is good, discrepancies do arise due to finite sample size and quite possibly the limited number of samples.

Table 3: Characteristic quantities of studied ZPG-TBL obtained with high-resolution 2C-PIV

	Symbol	Unit	$Re_\theta = 2770$
Tunnel free stream velocity	U_∞	m s^{-1}	10.0
Magnification factor (low / high res.)	m	-	0.47 / 1.1
Magnification on sensor (low / high res.)	m	µm pixel^{-1}	23.6 / 10.0
Measured free stream velocity	U_∞	m s^{-1}	10.09
Boundary layer thickness at 0.99 U_∞	δ_{99}	mm	35.4
Displacement thickness	δ^*	mm	5.75
Momentum thickness	θ	mm	4.16
Shape factor	H	-	1.384
Friction coefficient	C_f	-	0.00317
Wall shear rate, estimated	$du/dy|_0$	s^{-1}	11350
Friction velocity	u_τ	m s^{-1}	0.418
Viscous unit	ν/u_τ	µm	36.58
Wall unit on sensor (high / low res.)	ν/u_τ	pixel	1.56 / 3.68
Reynolds number	Re_δ	-	23515
Momentum Reynolds number	Re_θ	-	2768
Friction Reynolds number	Re_τ	-	961
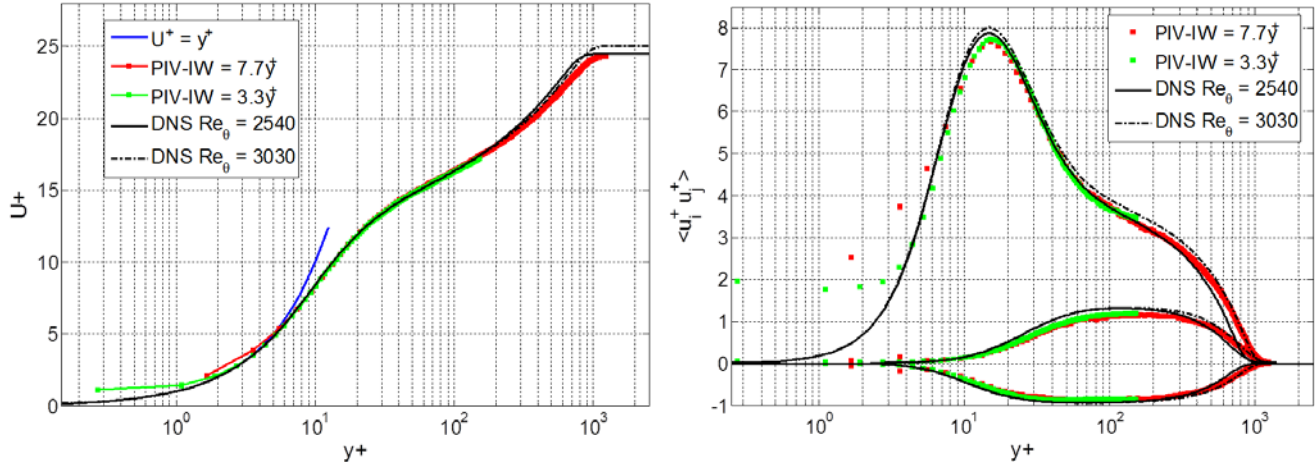
Figure 2: Mean velocity profiles at $U_\infty = 10$ m/s; (left) and inner-normalized profiles of the velocity variances from the high-resolution 2C-PIV measurements; top curve set represents $<u'u'>^+$, middle curve set $<v'v'>^+$, bottom curve set $<u'v'>^+$. Black lines DNS data from Schlatter et al. [18]
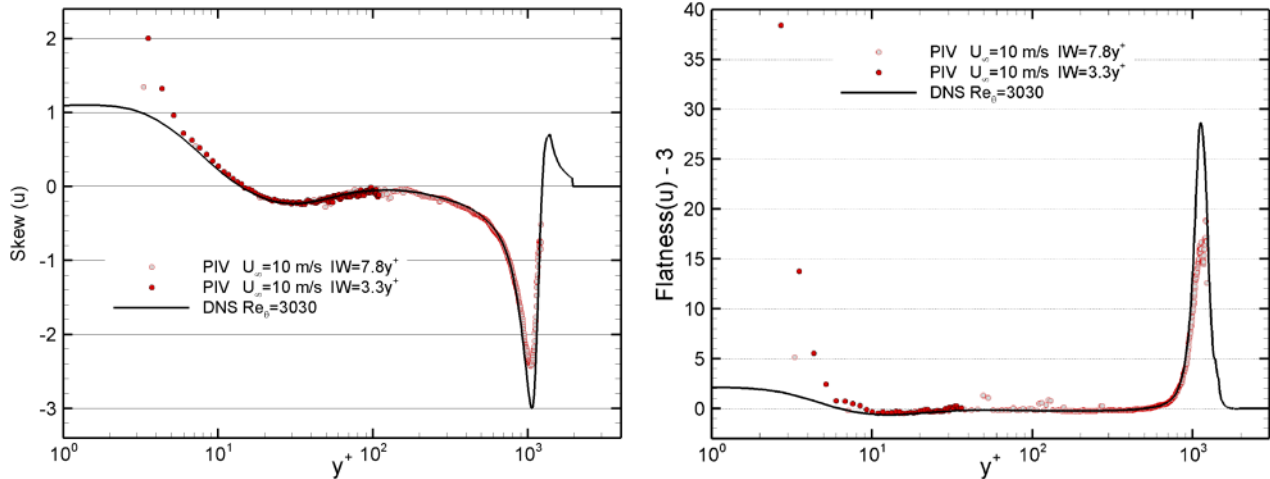

Figure 3: Profiles of the streamwise velocity skewness (left) and flatness (right) plotted versus y^+ (Flatness of 3 corresponds to a Gaussian distribution) ; black lines DNS data from Schlatter et al. [18]

NEAR-WALL STB RESULTS

An example of measured instantaneous velocity vectors distributed irregularly within a wall-parallel volume of approximately 16 x 1.4 x 16 mm³ according to 430 x 36 x 430 viscous units $(l^+)^3$ in x-, y- and z-directions evaluated by STB is shown in Figure 4. The velocity vectors are color coded by the *u*-velocity component and based on the application of an optimal temporal Wiener filter along 19 subsequent time steps of the reconstructed 3D Lagrangian particle tracks. The kernel size of the temporal Wiener filter is in the order of the Kolmogorov time scales η. The relatively dense Lagrangian tracks are built along two 7.27 s long time series of 115,000 particle images per set. Within the three projections of ~5,700 instantaneous velocity vectors onto the respective side planes of the measurement volume displayed in Figure 4 the typical features of a turbulent boundary layer flow in close proximity to the wall can be identified. High and low-speed flow regions elongated in flow directions are visible in the *x-z*-projection, while a spanwise meandering of these streaky regions with respective spanwise velocity components are as well detectable in the *y-z*-projection. Here in the top area of the projected volume (Figure 4, right) a spatially organized sweep (Q4-) event is combined with relatively large spanwise velocities, which, following the temporal development in the time-resolved series, can be generally addressed to be precursors for strong ejection (Q2-) events. Underneath these events the very low-velocity distribution (< 1 m/s) in the area

within and slightly above the viscous sublayer ($y^+ < 183$ µm) is more extended in wall-normal direction within low-speed streak regions and "compressed" to a thinner layer in high-speed streak or sweep regions. The x-y-projection shows the typical TBL u-velocity profile and respective strong instantaneous (wall-) shear gradients $du/dy|_{>0}$. As the particle position estimation using STB in experimental investigations is ideally in the order of less than 0.1 pixels which on top is filtered temporally (in the order of η) a very accurate velocity estimation close to the wall can be reached. This is particularly the case for the low-velocity particles near the wall that stay in the measurement volume for up to several hundred time-steps and the respective tracks are temporally highly oversampled with respect to the maximum acceleration values or smallest Kolmogorov time scales η, which are typically in the order of a viscous unit for TBL close to the wall [26].
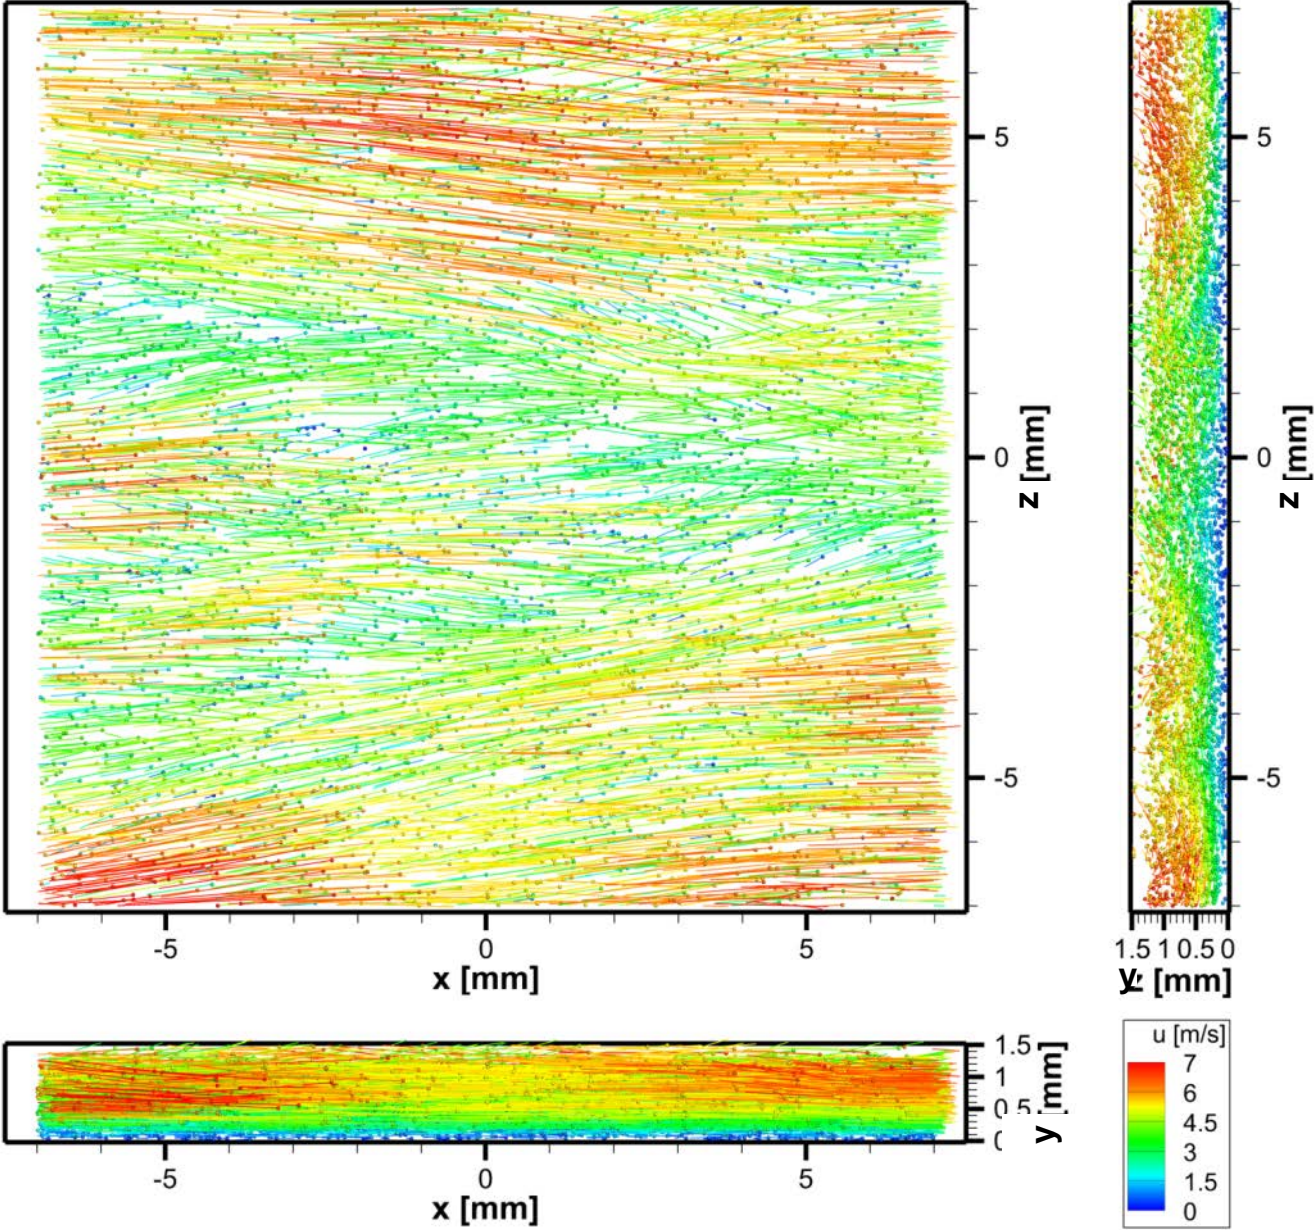
Figure 4: Instantaneous velocity vector distribution from optimal Wiener filtered Lagrangian tracks with ~5,700 particles in a volume of the TBL close to the wall projected onto x-z- (top-left), y-z- (right) and x-y- planes (bottom) of the STB measurement volume (u-comp of velocity color coded, y = 1.4 mm corresponds to ~36 viscous or wall units)

The total amount of data points ~8 x10^8 are available from the two STB runs for calculating a wall-normal 1D-profile in a bin averaging approach (projecting span- and streamwise extensions) including all components of the velocity vector and Reynolds stress tensor in bin sizes of 0.1 pixel resp. $1/13^{th}$ l^+. Depending on the intensity profile of the volumetric illumination and boundary effects of the STB evaluation approach the number of detected particle tracks varies along the respective wall normal row of bins. In Figure 5 (left) the number of samples collected in each bin along y^+ is displayed. In the region between y^+ = 1.5 and y^+ = 32 for both runs more than 1 x 10^6 samples are available, which are mostly statistically independent although there is a high time-resolution of the measurement. Due to the span- and streamwise projection onto a one-dimensional wall normal line along 14.49 s integration time these bin entries are statistically fully converged. In the area below y^+ = 1.5 and above y^+ = 32 boundary effects of the STB evaluation occurs mainly due to a reduction of the illumination intensity very close to the wall and some tracking issues due to out-of-focus particles above y^+ = 32. Nevertheless, here more than 1 x 10^5 entries could be collected.

A linear representation of the mean u-velocity profile along y^+ is given in Figure 5 (right) together with the scatter plot of the respective instantaneous u-velocity samples (blue scatter) of a small sub-set of the available data. It is clearly visible that almost all instantaneous velocity events fall in between the two black lines above and below the blue area which represent the boundaries of the STB track building approach.

The STB result of the converged U^+ velocity mean profile is given by the red squares along y^+ in Figure 6 in logarithmic (top left) and linear (top right) representation. The comparison with the black line given by DNS at Re_θ = 2540 [18] and with the green squares from measurement results of the high-resolution 2C PIV method show the advantages of the fully 3D particle tracking approach. The measured mean profile by STB follows exactly the DNS line down to ~ y^+ = 1.4 with only minor deviations which can be accredited to the slightly different Reynolds number and experimental boundary conditions (e.g. wall roughness). A direct comparison with the 2C PIV measurement confirms basically the low-pass filtering effect of the correlation based (even single-pixel line) approaches in comparison to the well resolved STB approach. The important statistic values for the TBL near-wall region characteristics determined by the STB bin averaging results are given in Table 3. Due to slight changes in temperature and pressure at the STB- compared to the 2C-PIV measurement the flow parameters changed slightly resulting in a Reynolds number of Re_τ = 929. This Reynolds number is based on the friction velocity of u_τ = 0.4042 m/s estimated by a linear fitting along the STB bin averaged mean u-velocity profile between y^+ = 2 and y^+ = 4.5, which can be considered to be statistically converged. The 2C-PIV friction velocity estimation of u_τ = 0.418 m/s has been obtained by the single-pixel line evaluation technique method and relies on the respective accuracy of the 2C-PIV profile measurement in the near wall region of the higher magnification factor displayed in Figure 2 (left). Accordingly, all the inner-scaled parameters differ slightly between 2C-PIV and STB results.
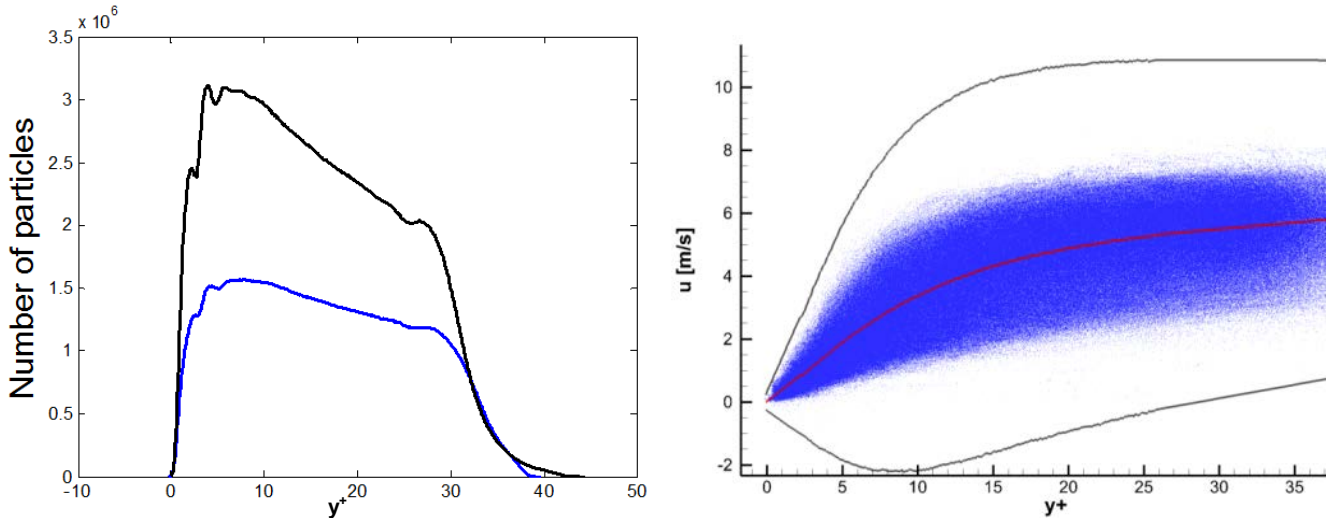
Figure 5: Number of particles with velocity vector information per bin size of 0.1 pixel resp. $1/13^{th}$ l^+ averaged over two runs (blue and black) with different particle image densities (ppp) along y^+ (left). Mean u-velocity profile (red) and scatter plot of instantaneous u-velocity entries (blue dots) of a sub-set of the available STB velocity data. Black lines indicate upper and lower track building boundaries of the STB evaluation scheme (right)

Table 3: Characteristic quantities of studied ZPG-TBL obtained with STB

	Symbol	Unit	$Re_0 = 2770$
Tunnel free stream velocity	U_∞	m s^{-1}	10.0
Magnification factor	m	-	0.36
Magnification on sensor (low / high res.)	m	μm pixel^{-1}	30
Measured free stream velocity	U_∞	m s^{-1}	10.00
Friction coefficient	C_f	-	0.003266
Wall shear rate, estimated	$du/dy|_0$	s^{-1}	10480
Mean wall shear stress	$\mu\, du/dy|_0$	kg m^{-1}s^{-2}	0.1933
U-rms wall shear stress (u-rms)/(mean)		-	0.396
Friction velocity	u_τ	m s^{-1}	0.4042
Viscous unit	v/u_τ	μm	38.39
Wall unit on sensor	v/u_τ	pixel	1.3
Friction Reynolds number	Re_τ	-	929

Nevertheless, maybe due to the strong reduction of the illumination intensity very close to the wall which is an intrinsic problem of a collimated tangential illumination mainly due to refraction along the surface a slight bias towards higher velocities can be detected in the mean velocity statistic below $y^+ = 1.5$ in Figure 6 (top-left). The expected spatial resolution of the STB bin averaging approach should be in the range of the positional uncertainty of the 3D particle reconstruction which is around 0.1 px for experimental data [15] or the used bin size. Therefore there is the suspicion that a systematic issue which correlates with the low illumination intensity is crucial: the signal-to-noise ratio of the particles very close to the wall is pretty low, which increases the trend of a pixel-locked representation of the particle position. This results in a 3D pattern of superimposed lines-of-sight of each camera pixel, which at their 3D knot distribution biased particle positions agglomerate and thus as well causes velocity biases (here towards higher velocities). Most probably this pattern biases the results in the near wall regions a bit more due to the fact that at low velocities the temporal filtering effect of such errors is further reduced. Such issues have to be addressed in future STB measurements of near wall TBL flow investigations by proper illumination strategies as the reachable spatial resolution resp. 3D particle positional- and related velocity estimation accuracy for well illuminated particles can be much higher [13].

The components of the Reynolds stresses show a very good agreement between the STB results of the $<u'u'>^+$ and the related DNS data down to sub- viscous units in the near wall region. Also the $<v'v'>^+$ profile of STB closely follows the DNS line, which at the same time lead to a nearly perfect alignment of STB and DNS data for the covariance $<u'v'>^+$ profile, with only a slight drift to lower measures for larger y^+ values. Note that the upper part of that profile is also well resolved by the 2C PIV approach and both experimental methods show the same trend, so that those small differences might be physical and relate to the wind tunnel conditions. Furthermore, the $<w'w'>^+$ profile from STB is also very close to the DNS data, but shows slightly lower values of the spanwise variances. Again the slight deviation could be accredited to the wind tunnel conditions.

In general so far no measurement technique is known to the authors which enables a more accurate and higher resolution profile of all components of the (mean) velocity vector and the full Reynolds stress tensor for such small viscous unit sizes without intrusiveness and band pass filtering effects within short measurement times. The use of PTV measurement tools for obtaining accurate profile measurements close to walls is motivated by the assessment done by Kähler et al. [33]. But especially the additional possibility to measure the full frequency range of all 3D velocity vectors and both components of the wall-shear stress fluctuation vectors at many points simultaneously in a relatively large volume resp. area at the wall with a high spatial resolution and without band pass filtering effects is a unique feature of the STB method. The analysis of the temporal resolution of the present data-set has not been explored yet. It would enable bin averaged 3D two-point space-time correlations for the determination of Taylor micro- and macro scales at various wall distances and e.g. the convection velocities, frequency analyses and the calculation of spectra of different velocity components or the wall-shear stress components. Unfortunately the full potential of the spatial accuracy of the STB technique at the very near wall region could not be reached in the present study for two reasons. Following the aforementioned problem of the near wall particle illumination and resulting bias issues below $y^+ = 1.5$ the accuracy of the STB variance measurement is also slightly lower resulting, for instance, in small positive values for $<u'u'>^+$ very close to the wall. Secondly, regarding the lower temporal filtering effect close to the wall a locally adapted optimal temporal Wiener fit length very close to the wall would be advantageous. Both, the illumination issue and temporal filtering adaptation can be easily solved by performing STB experiments in (TBL) water flows, which would allow the use of larger particle sizes and a full temporal resolution at lower frame rates and laser repetition frequencies at much higher Reynolds numbers.
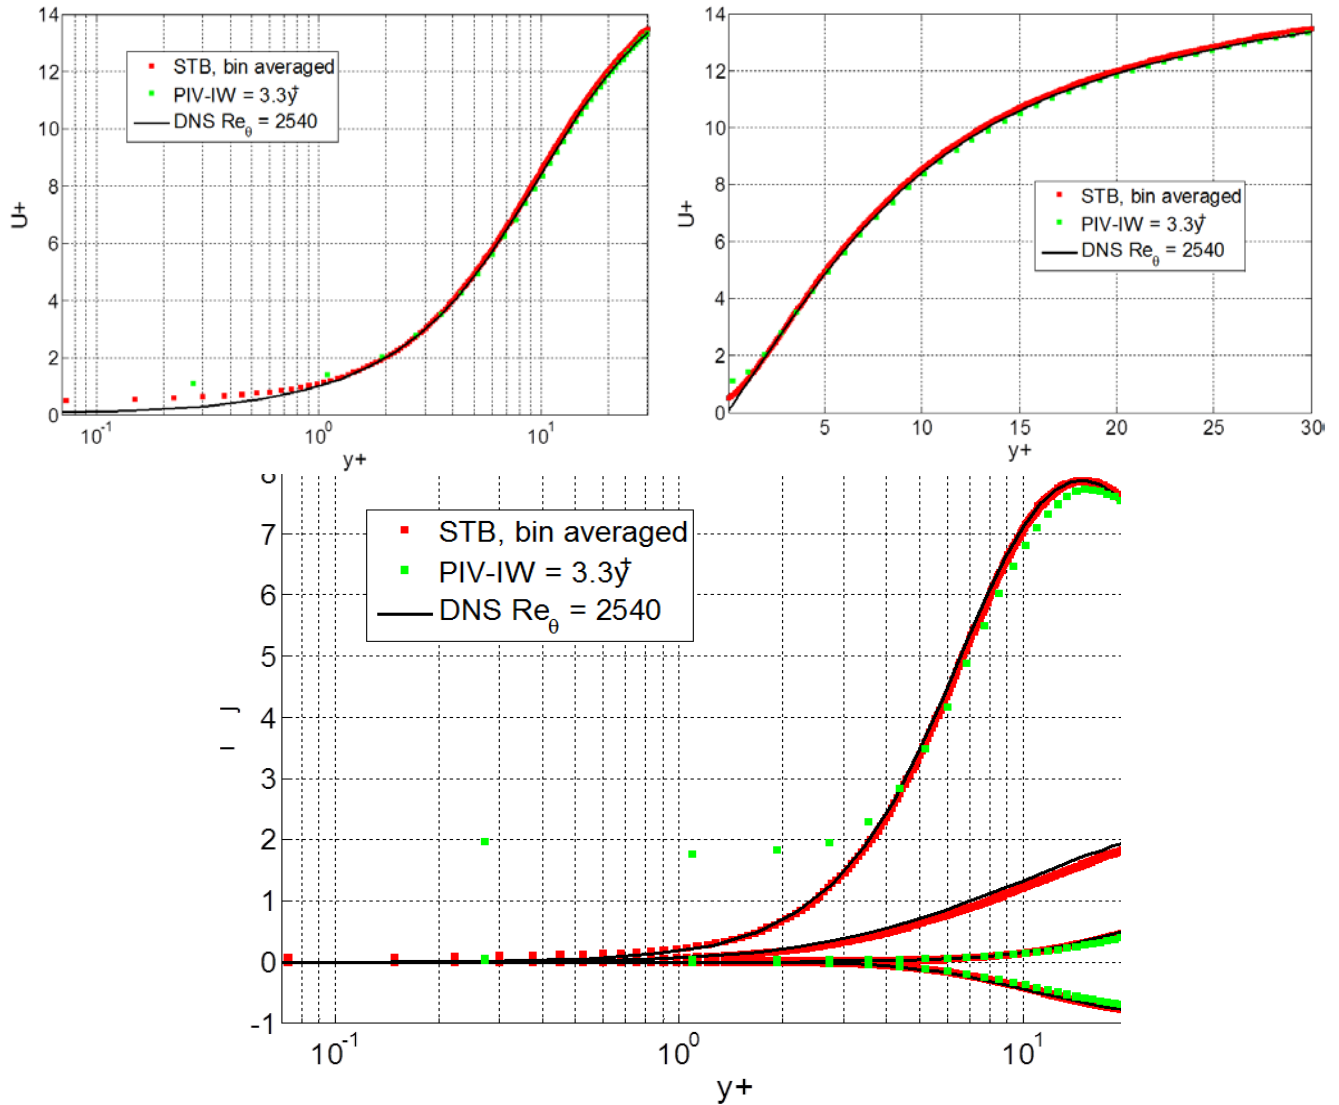
Figure 6: Inner-normalized mean velocity profiles U^+ against wall distance y^+ scaled logarithmically (left) and linear (right) and respective profiles of the velocity variances (bottom) at $U_\infty = 10$ m/s. At the velocity variances or Reynolds stresses (bottom) the top curve represents $<u'u'>^+$, middle curves set $<w'w'>^+$ (upper) and $<v'v'>^+$ (lower), bottom curve set $<u'v'>^+$. Red squares are STB results with 0.1 pixels or $1/13^{th}$ l^+ spatial resolution, green squares are from 2C-PIV (see Fig 2) and black lines are from DNS data at $Re_\theta = 2540$ [18].

The rms-value of the u-component of the wall-shear stress has been estimated to $\tau^+_{w, rms} = 0.396$ from the STB evaluated instantaneous velocity gradients $du/dy|_0$ of all particles between $y^+ = 2$ and $y^+ = 4.5$. This value fits quite well to 0.4 given by [35], while the formula given in [18] would result in $\tau^+_{w, rms} = 0.421$ for the present Re_τ. Looking on the corresponding probability density function of the fluctuation values of $\tau^+_w(u)$ in a linear representation estimated from the bin averaged STB results in Figure 7 (left) one can see a very good agreement for a respective DNS result given in Hu et al. [36] in Figure 7 (right). A possible explanation for the slightly lower rms-value compared to DNS data from [18][20][34][37] is visible in Figure 8. Here the same probability density function of the instantaneous $\tau^+_w(u)$ without subtraction of the mean value is displayed in a logarithmic representation. In comparison to DNS STB shows lower maximal values for the rare extreme events on the negative and positive ends of the respective distribution (Figure 8 (left)). Although a few negative values are present, which implies measured reverse flow in the respective STB interrogation area, the extreme values very close to the wall have not been considered by STB in the present investigation due to the mentioned illumination issues and chosen physical boundaries for gaining a reliable track building process (see black lines in Figure 5 (right)).
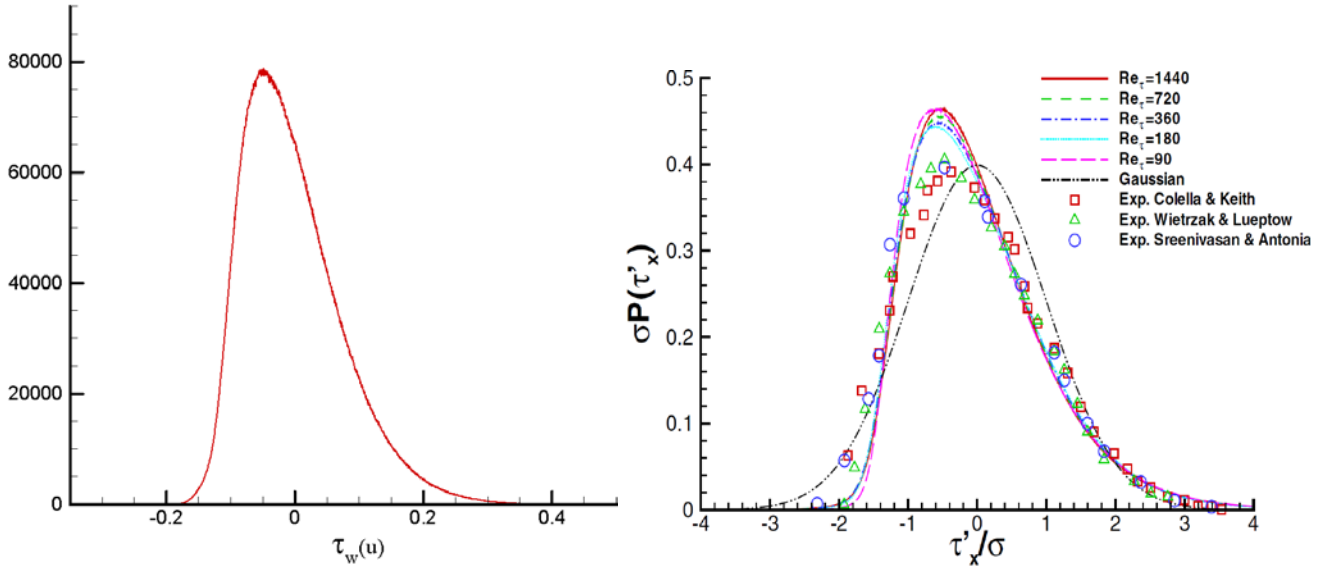
Figure 7: Probability density function (arbitrary units) of the *u*-component of the wall-shear stress fluctuations in a linear representation from STB (left) and DNS values at similar Reynolds numbers normalized by σ from Hu et al , Figure 10 [36] estimated from turbulent channel flow (right)
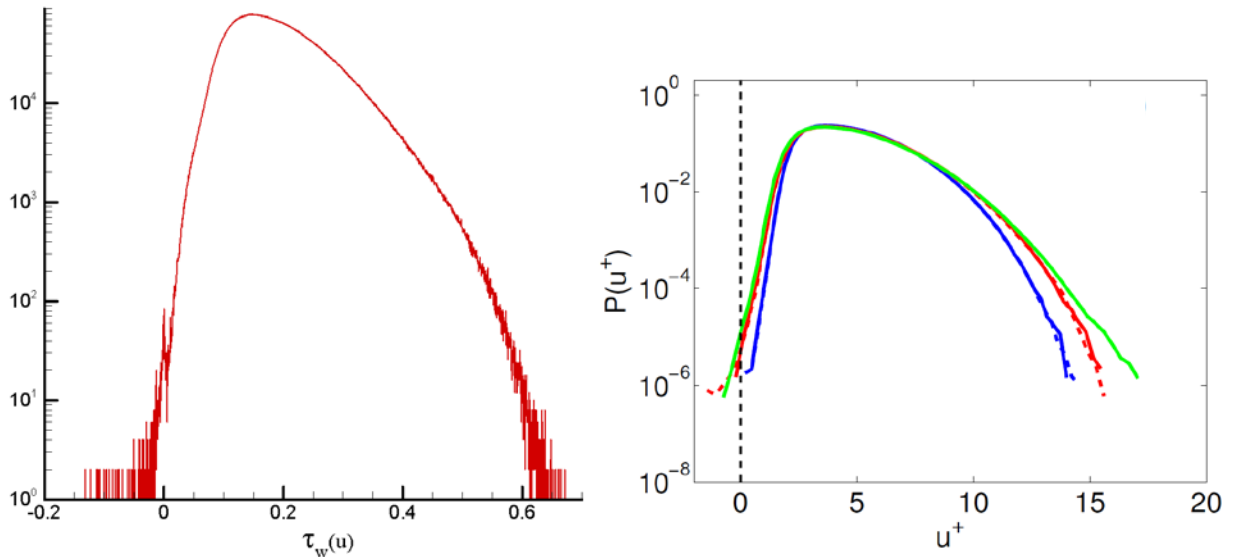
Figure 8: PDF of the wall-shear stress in a logarithmic representation showing rare negative values (reverse flow events) (left) and PDF of u^+-velocity fluctuations for three different Reynolds number (green line at Re_τ= 1000) from DNS data from Lenaers et al., Figure 4 [20]

It can definitely be stated that excluding the near wall position below $y^+ = 2$ for the estimation of the instantaneous wall-shear stresses decreases the number of rare and extreme wall-shear stress events, which can be confirmed by the similarity of the PDF of the present wall-shear stress estimation by STB in Figure 8 (left) and the shape of a PDF of the *u*-velocity fluctuations from a respective turbulent channel flow DNS [19] at a constant wall distance of $y^+ = 5$ displayed in Figure 8 (right). The number of negative (extreme) wall-shear stress events decreases with the wall distance and increases with the Reynolds number. A new computation of the present STB data including values closer to the wall would deliver a wider distribution of the wall-shear stress and more extreme events, but the main problem is again to enable a very accurate estimation of the velocity gradient even below $y^+ = 1$. The big challenge for measurements techniques in close proximity to surfaces is that every little uncertainty of the measured instantaneous velocity is significantly increased by calculating the

gradient $du/dy|_0$ when approaching the wall, because the value in the denominator approaches zero. In principle STB could provide the required high accuracies because a very good oversampling in time for those particles close to wall is given. For reaching the goal of calculating exact wall-shear stress values below $y^+ = 1$, beside the illumination issue, STB needs to be adapted with a local optimal temporal filtering approach close to the wall as already mentioned above.

CONCLUSIONS AND OUTLOOK

Two advanced particle based measurement methods, high-resolution high-speed 2C-PIV and STB, have been applied to a zero pressure gradient turbulent boundary layer (TBL) flow at $Re_\tau = 960$ and $Re_\tau = 929$ respectively. Both methods work very efficiently in delivering statistically converged and relevant data of mean and fluctuation velocity components for the overall characterization of the TBL within short measurement and evaluation times. The 2C-PIV method has the additional advantage of being simple and compact on the experimental side. While STB needs a high-power high-speed laser and four high-repetition frame rate cameras in a tomographic set-up, the evaluation time of only ~6 sec per time step on a simple multi-core computer for the STB track building is also very efficient for a 4D technique, mainly because of the limited amount of particles (3,000 to 6,000) to be reconstructed. Significant advantages of the STB 4D-PTV technique have been demonstrated in terms of accuracy, spatial and temporal resolution. STB delivers a well converged 3D data set of relatively dense Lagrangian tracks and related time resolved velocity vector volumes. The full Reynolds stress tensor has been determined based on a bin averaging process leading to a spatial resolution down to a fraction of a viscous unit (wall unit). The profiles of the mean velocity, components of the Reynolds stress tensor and the PDF of the instantaneous wall-shear stress distribution show very good agreement with DNS data at similar Reynolds numbers, in particular for the highly resolved STB data. Some adaptive steps on the hardware and evaluation side of STB could be identified to gain even better results for very near-wall velocity and instantaneous wall-shear stress measurements. With this first application of STB to a TBL flow in air the technique is proven capable to deliver valuable time-resolved 3D data in a non-intrusive manner. Further possibilities in extracting and analyzing Lagrangian and Eulerian properties of the near-wall dynamics are still open. In the future the application of the STB technique to higher Reynolds number TBL flows in water (beyond the possibilities of present DNS and comparable to those from earlier investigations [38]) shall further demonstrate the capabilities of STB in delivering accurate data for temporally and spatially highly resolved turbulence characterizations.

REFERENCES

[1] Alfredsson PH, Örlü R, Schlatter P "The viscous sublayer revisted-exploiting self-similarity to determine the wall position and friction velocity", Exp Fluids (2011) 51: 271-280

[2] Compton DA, Eaton JK "A high-resolution laser doppler anemometer for three-dimensional turbulent boundary layers", Exp Fluids (1996) 22: 111–117

[3] Fischer M, Jovanovic J, Durst F "Reynolds number effects in the near-wall region of turbulent channel flows", Phys Fluids (2001) 13, 6

[4] Tanner LH, Blows LG "A study of the motion of oil films on surfaces in air flow, with application to the measurement of skin friction", J Physics E, Scientific Instruments 9 (1976)

[5] Monson DJ "A nonintrusive laser interferometer method for measurement of skin friction", Exp Fluids (1983) 1: 15–22

[6] Raffel M, Willert C, Wereley S, Kompenhans J "Particle image velocimetry: a practical guide", Exp Fluid Mech (2007), Springer, Berlin

[7] Elsinga GE, Scarano F, Wieneke B and van Oudheusden BW "Tomographic particle image velocimetry", Exp Fluids (2010) 41: 933-947

[8] Schröder A, Geisler R, Staack K, Elsinga G, Scarano F, Wieneke B, Henning A, Poelma C, Westerweel J *"Eulerian and Lagrangian views of a turbulent boundary layer flow using time-resolved tomographic PIV"*, Exp Fluids (2011) 50: 1071–1091.

[9] Adrian RJ "Dynamic ranges of velocity and spatial resolution of particle image velocimetry", Meas Sci Technol (1997) 8(12):1393

[10] Kähler CJ, Scholz U, Ortmanns J "Wall-shear-stress and near-wall turbulence measurements up to single pixel resolution by means of long-distance micro-PIV," Exp Fluids (2006) 41: 327–341

[11] Sheng J, Malkiel E, Katz J "Using digital holographic microscopy for simultaneous measurements of 3D near wall velocity and wall shear stress in a turbulent boundary layer", Exp Fluids (2008) 45:1023–1035

[12] Schanz D, Schröder A, Gesemann S, Michaelis D and Wieneke B "Shake-the-Box: a highly efficient and accurate Tomographic Particle Tracking Velocimetry (TOMO-PTV) method using prediction of particle position", 10th Symp PIV (2013) Delft, The Netherlands

[13] Schanz D, Schröder A and Gesemann S "Shake-the-Box - a 4D PTV algorithm: accurate and ghostless reconstruction of Lagrangian tracks in densely seeded flows", 17th Symp Laser Tech (2014) Lisbon, Portugal

[14] Wieneke B "Iterative reconstruction of volumetric particle distribution", Meas Sci Technol (2013) 24, 024008

[15] Schröder A, Schanz D, Michaelis D, Cierpka C, Scharnowski S and Kähler CJ "Advances of PIV and 4D-PTV "Shake-The-Box" for Turbulent Flow Analysis – the Flow over Periodic Hills" Flow, Turbulence and Combustion (2015), 17p, Published online, May 24, DOI 10.1007/s10494-015-9616-2

[16] Bruecker C, "Evidence of rare backflow and skin-friction critical points in near-wall turbulence using micropillar imaging" Phys Fluids (2015), 27, 031705

[17] Örlü R, Fransson JHM and Alfredsson PH "On near wall measurements of wall bounded flows—the necessity of an accurate determination of the wall position", Prog Aero Sci (2010) 46:353–387

[18] Schlatter P and Örlü R "Assessment of direct numerical simulation data of turbulent boundary layers", J Fluid Mech (2010) 659: 116–126.

[19] Mathis R, Marusic I, Chernyshenko SI, Hutchins N "Estimating wall-shear-stress fluctuations given an outer region input", J Fluid Mech (2013) 715: 163-180

[20] Lenaers P, Li Q, Brethouwer G, Schlatter P and Örlü R „Negative streamwise velocities and other rare events", J Phys: Conference Series (ETC13) (2011) 318, 022013

[21] Cardesa J, Monty J, Soria J, Chong M. "Skin-friction critical points in wall-bounded flows," J Phys. (2014): Conf. Ser. 506, 012009

[22] Eckelmann H "The structure of the viscous sublayer and the adjacent wall region in a turbulent channel flow", J Fluid Mech (1974) 65:439–459

[23] Große S and Schröder W "Wall-shear stress patterns of coherent structures in turbulent duct flow", J Fluid Mech (2009) 633

[24] Amili O and Soria J "Wall shear stress distribution in a turbulent channel flow", 15th International Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal, 05–08 July, (2010).

[25] Cierpka C, Scharnowski S, Kähler CJ "Parallax correction for precise near-wall flow investigations using particle imaging", Appl Opt (2013) 52: 2923-2931

[26] de Silva CM, Gnanamanickam EP, Atkinson C, Buchmann NA, Hutchins N, Soria J, Marusic I " High-spatial range velocity measuremets in a high Reynolds number turbulent boundary layer", Phys Fluids (2014) 26, 025117

[27] Novara M, Schanz D, Kähler CJ, Schröder A „Shake-The-Box for multi-pulse tomographic systems: towards high seeding density particle tracking in high speed flows, 11[th] Symp PIV (2015), St. Barbara, CA, USA

[28] Wu PP, Miles RB "High-energy pulse-burst laser system for megahertz-rate flow visualization", Opt Lett (2000) 25, 1639-1641

[29] Thurow B, Jiang N, Samimy M, Lempert W "Narrow-linewidth megahertz-rate pulse-burst laser for high-speed flow diagnostics", Appl Optics (2004) 43, 26: 5064 - 5073

[30] Willert C H "High‑speed particle image velocimetry for the efficient measurement of turbulence statistics", Exp Fluids (2015) 56: 17p

[31] Wieneke B "Volume self-calibration for 3D particle image velocimetry" Exp Fluids (2008), 45: 549-556

[32] Schanz D, Gesemann S, Schröder A, Wieneke B and Novara M "Non-uniform optical transfer function in particle imaging: calibration and application to tomographic reconstruction" Meas Sci Technol (2013) 24, 024009

[33] Kähler CJ, Scharnowski S and Cierpka C "On the uncertainty of digital PIV and PTV near walls", Exp Fluids (2012) 52: 1641-1656

[34] Schlatter P, Örlü R, Li Q, Brethouwer G, Fransson JHM, Johansson AV, Alfredsson PH, Henningson DS "Turbulent boundary layers up to $Re_\theta = 2500$ studied through simulation and experiment", Phys Fluids (2009) 21(5), 051702

[35] Alfredsson PH, Johansson AV, Haritonidis J, Eckelmann H" The fluctuating wall-shear stress and the velocity field in the viscous sublayer, Phys Fluids 31 (1988) 31: 1016-1033.

[36] Hu ZW, Morfey CL, Sandham ND "Wall Pressure and Shear Stress Spectra from Direct Simulations of Channel Flow" AIAA Journal (2006) 44,7: 1541-1549

[37] Örlü R and Schlatter P "On the fluctuating wall shear stress in zero pressure-gradient turbulent boundary layer flows", Phys Fluids (2011) 23, 21704

[38] Fernholz HH, Finley PJ "The Incompressible Zero-Pressure-Gradient Turbulent Boundary Layer: an Assessment of the Data", Prog. Aerospace Science (1996) 32: 245-311

Towards high-resolution 3D flow field measurements at cubic meter scales

Daniel Schanz[1], Florian Huhn[1], Sebastian Gesemann[1], Uwe Dierksheide[2],

Remco van de Meerendonk[3], Peter Manovski[4], Andreas Schröder[1]

[1]German Aerospace Center (DLR), Institute of Aerodynamics and Flow Technology, Department of Experimental

Methods, Göttingen, Germany; [daniel.schanz@dlr.de]

[2]LaVision GmbH, Anna-Vandenhoeck-Ring 19, Göttingen, Germany

[3]Aerospace Engineering Department, Delft University of Technology, 2629, The Netherlands

[4]Defence Science and Technology Organisation (DSTG), Australia

* Correspondent author: daniel.schanz@dlr.de

Keywords: Lagrangian Particle Tracking, LED volume illumination, Helium filled soap bubbles

ABSTRACT

We present results from two large-volume volumetric flow experiments. The first of these, investigating a thermal plume at low velocities (up to 0.35 m/s) demonstrates the abilities and requirements to reach volume sizes up to and probably beyond one cubic meter. It is shown that the use of Helium filled soap bubbles (HFSBs) as tracers, combined with pulsed LED illumination yields high particle image quality over large volume depths. A very uniform particle imaging, both in space as well as in time enables using high particle image concentrations (up to 0.1 ppp), while still being able to accurately reconstruct the flow using Shake-The-Box particle tracking.

The experiment consisted of time-resolved volumetric flow measurements of a convectional plume within a volume of approx. 0.55 m³ (550 liters). The light yield needed for such a large scale measurement is realized by using HFSBs with 300 µm diameter as tracers and illuminating the measurement region using high-power, scalable arrays of white LEDs. Applying the Shake-The-Box algorithm, up to 275,000 bubbles could be tracked simultaneously. Interpolating the results on a regular grid (using 'FlowFit') reveals a multitude of flow structures. The setup can be scaled to larger volumes of several cubic meters, basically only being limited by the number and power of available LEDs and high-resolution cameras with sufficient frame-rate and pixel sizes.

A second experiment showcases the possibilities to reach higher flow velocities, while still measuring within a comparatively large volume, by applying high-speed imaging and advanced LED illumination. An impinging turbulent jet was investigated in volumes ranging from 13 to 47 liters, depending on the repetition rate of the camera system. The results show that even at a repetition rate of 3.9 kHz and flow speeds up to 17 m/s the tested system was able to deliver images that allowed for a reliable and accurate tracking of bubbles.

Again, the use of more LEDs would allow for larger volumes. New generations of high-speed cameras should enable the use of even higher flow speeds – thus enabling large-volume measurements in typical low speed wind tunnel experiments at high spatial resolution (provided enough HFSBs can be produced).

1. Introduction and Motivation

Time resolved volumetric flow measurements, using methods like TOMO PIV (Elsinga et al. 2006), 3D PTV (Maas et al. 1993, Malik et al. 1993) or Shake-The-Box (Schanz et al. 2016), are typically restricted to relatively small volume sizes in the order of $<= 200$ cm^3 (Scarano et al. 2015). This limitation stems from the small size of commonly used seeding material, in order to accurately follow the flow (~ 1 μm in air, $10 - 50$ μm in water). Currently available high-repetition laser systems, which are typically used as light source, do not provide enough intensity to allow for illumination of larger volumes - even for the larger particles used in water experiments. However, seeding particles whose density approaches that of the medium can be of much larger size, while still accurately following the flow (Melling 1997). Following this thought, neutrally buoyant Helium filled soap bubbles (HFSB) have been used in air to allow for large scale flow measurements. Applications range from traceline-visualizations (Pounder 1956) over large-scale 2D-PIV measurements (Müller et al. 2000, Bosbach et al. 2009) to three-dimensional tracking of single bubbles (Klimas 1973) and large-scale tomographic PIV of a convective flow (Kühn et al. 2011). Very recently, the feasibility of using HFSBs in wind tunnel facilities has been demonstrated (Scarano 2015).

While some the previous applications examined large measurement areas (Biwole et al. 2009, Klimas 1973), all of these were limited in particle number (e. g. tracing a few tens or hundreds of bubbles). The largest investigated volume that allowed the description of instantaneous flow structures was applied by Kühn et al (2011). A convection cell with a volume of approx. 56 liters was investigated using 2-pulse tomo-PIV. However, a large interrogation window size had to be chosen (48*48*24 mm), limiting the spatial resolution to large structures. The experiment reported by Scarano et al. (2015) was special in that it is the first application of HFBSs in a wind tunnel experiment. As Scarano et al. have shown, the production of enough bubbles to achieve a sufficient particle concentration within the measurement volume is a major topic for higher flow speeds. Due to the limitations in bubble number and due to the limits of the high-speed laser used for illumination, the volume size was restricted to 4.8 liters in this experiment and the interrogation windows were quite large (96x96x86 voxels).

The experiments discussed here avoid the problem of bubble production rate by operating in a closed chamber. LEDs provide a scalable light-source that is free of the typical artifacts of laser illumination (speckles, diffraction problems, etc.). HFSBs can be produced virtually mono-disperse and due to the light being reflected at the bubble surface (instead of scattered) uneven imaging between the different cameras is avoided. These features lead to a very high attainable image quality (as discussed in paragraph 2.4), which allows for a very reliable and accurate particle tracking using the Shake-The-Box (STB) algorithm. By applying a regularized interpolation scheme ('FlowFit', Gesemann et al. 2016) the locally highly accurate information is leveraged to maximize the spatial resolution (Schneiders et al. 2015).

2. Thermal plume investigation

2.1. Convection chamber

The experiments are performed in a cylindrical convection chamber with a height of 2.00 m and a diameter of 1.83 ,m (see Figure 1). Top and bottom plates are constructed of wood, the back wall is made of aluminum, and the transparent front window is acrylic glass of 1 mm thickness. Intransparent parts of the floor, walls and ceiling are painted black or covered with black adhesive film, in order to avoid scattered light and to improve the contrast of the particle images towards the background. The convection chamber is accessible from the back side through a door in the aluminum wall. The chamber is equipped with a circular perforated tube at the bottom to rinse it with pressurized air and remove seeding.

LED illumination enters through a plexiglas window of 1 m diameter in the ceiling, covered with a circular passe-partout (0.75 m diameter) that determines the width of the cylindrical measurement volume. The convective flow is forced by a standard 1500 W electric hotplate (188 mm diameter, Silva Homeline EKS 2121) that is placed a few centimeters below the measurement volume. It is covered by a black circular 250 mm diameter aluminum plate of 10 mm thickness that serves as a heat reservoir to keep the temperature constant over time.

Before conducting experiments, the hotplate is heated up for a few seconds to reach the desired temperature, and left some minutes until a uniform temperature distribution and a constant temperature is attained. Three temperature sensors are installed in the convection chamber to monitor the temperature of the heated aluminum plate, the air temperature 2 mm above the plate as a proxy for the maximum air temperature, and the ambient air temperature (see Figure 1 for positions of sensors). The accuracy of the absolute temperature measurement is estimated to be ±2°C, while the relative temperature differences, relevant for the convective flow, are accurate

within ±0.5°C. During the experiments, heat is provided by the aluminum plate, while the temperature was decreasing less than 0.5 K (see Figure 6).
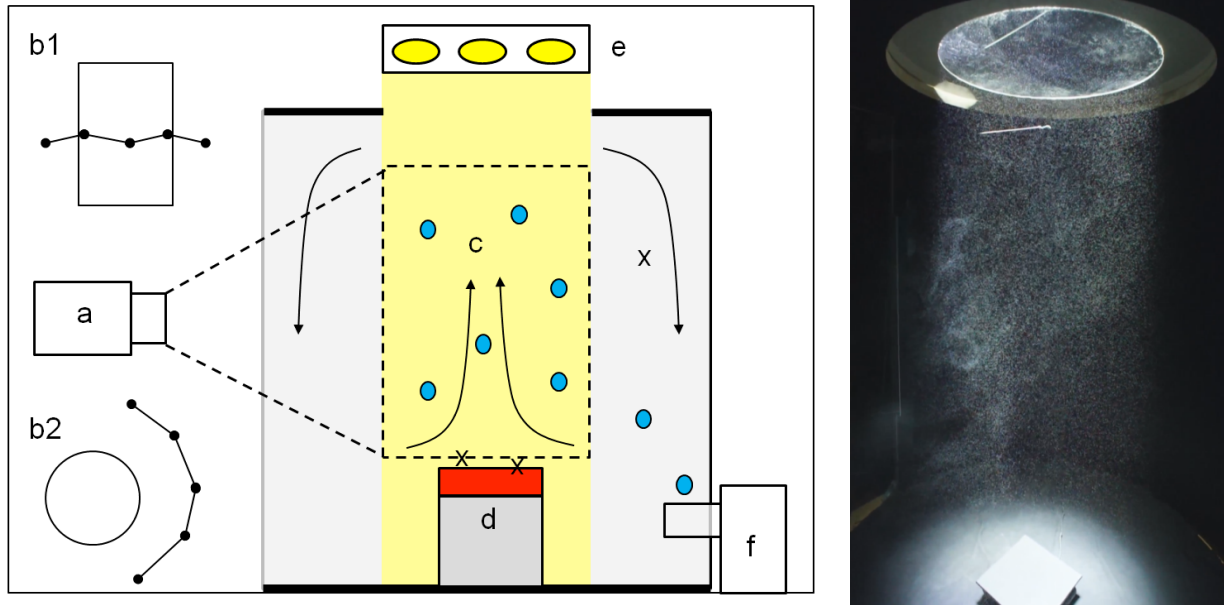

Figure 1: (left) Experimental setup of convection chamber: (a) cameras; (b1, b2) front and top view of camera configuration; (c) field of view (FOV), (d) hotplate; (e) LED array; (f) bubble generator; (x) positions of three temperature sensors

(right) Photograph of the interior of the convection chamber

2.2. Helium-filled soap bubbles

For PTV measurements, the flow is densely seeded with neutrally-buoyant helium-filled soap bubbles (HFSB) of 300µm diameter. They are produced by a bubble generator prototype of LaVision, based on the nozzle design presented by (Bosbach et al. 2009).

A nozzle consists of three concentric channels - providing helium, soap solution (ASAI 1035, Sage Action Inc.), and pressurized air, from the center outwards. It is covered by a cap with an orifice of 1.0 mm diameter.
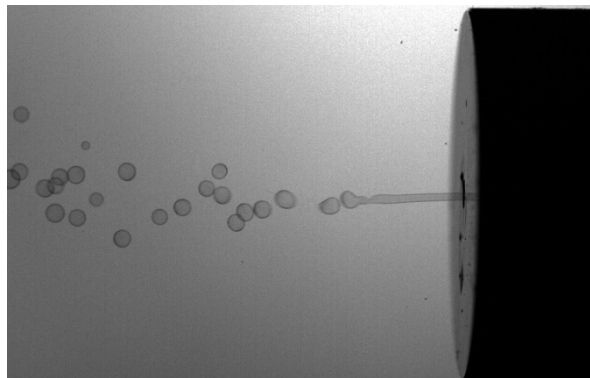
Figure 2: High-speed image (Exposure time $10\mu s$, recording rate 1.5kHz) of the seeding nozzle

The two inner channels produce a thin helium-filled soap tube that is transported through the orifice by the surrounding air flow and breaks up into a single chain of equally sized bubbles in the increasingly turbulent air flow (see Figure 2).

Six nozzles are operated in parallel with a bubble production rate of ~ 45,000/s each, counted with a high speed camera directed to the outlet of the nozzle. The nozzles are directly placed at the bottom inside the convection chamber to enable a high seeding density in the large volume. HFSBs are injected vertically close to the wall to avoid the influence of the momentum of the nozzles' jet towards the center where the hotplate is located. Before an experiment, the chamber was seeded for 30 seconds to reach a high particle concentration. A waiting period of around 135 seconds follows to reach a homogeneous spatial distribution of the seeding and let decay the motion induced by the nozzles' jet.

We adjust neutral buoyancy of the HFSBs by varying the flow rate of helium such that a zero settling velocity is attained. Careful experiments by Scarano et al. (2015) show that HFSB follow the air flow at high accelerations of $\sim 10^4$ m/s^2 in a wind tunnel even for a variation of the helium flow rate by a factor of two. Since in our convective flow, accelerations are much smaller (g~10^1 m/s, see results of STB), we expect the tracer to closely follow the air flow.

Soap bubbles burst, and the life time can be a limiting parameter especially for large convective flows that are typically slow and related to long time scales. Bosbach et al. (2009) estimate the life time of HFSB to be 1-2 min under similar conditions, e.g., room temperature and presumably relatively low relative air humidity in a lab.

2.3. LED light source

The measurement volume is illuminated by a LED array consisting of 7 standard collimated LED spotlights (Treble-Light, Power LED 20000) with an opening angle of 9° and 18,000 lm luminous flux at 170 W nominal electric power input each. One spotlight is composed of 48 LEDs with 3.5 W each. The LED array is located 1 m above the ceiling of the convection chamber and a passe-partout of 0.75 m diameter on the ceiling window defines the cylindrical measurement volume. In the experiment, the LED light source is synchronized with the camera system and is pulsed with a period of 3 ms at 29 Hz, corresponding to a 10 % duty cycle.
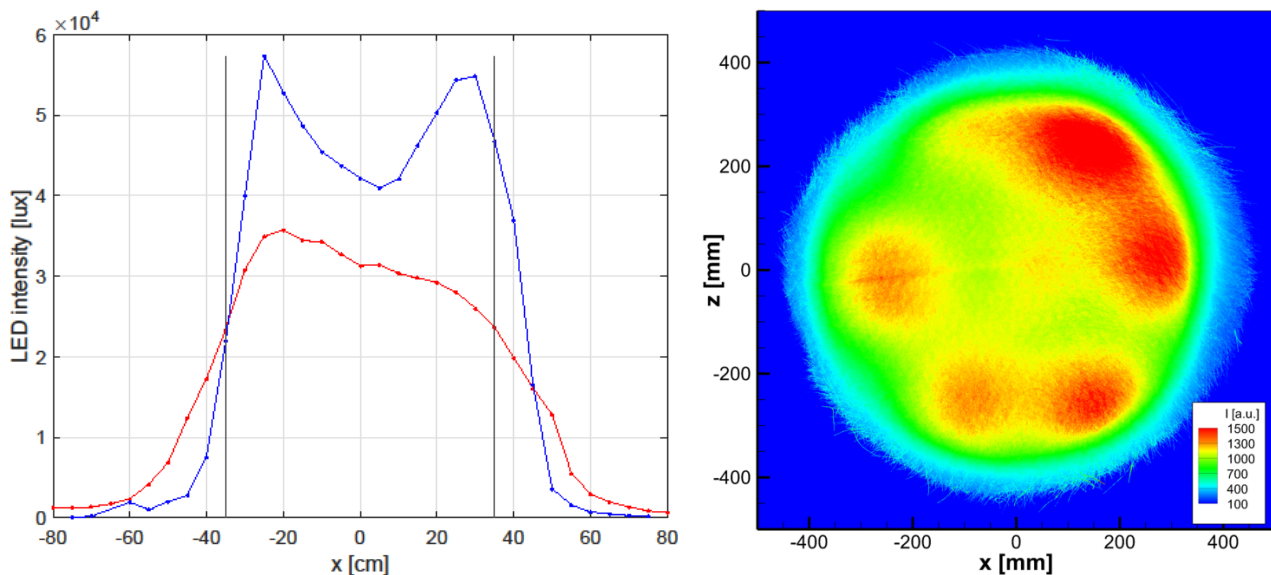
Figure 3: (Left) Intensity profile of continuous LED illumination across the measurement volume at a height of 110 cm (blue) and at the bottom of the convection cell (red) in the continuous mode. Black bars show the position of the passe-partout. (Right) Particle intensity as given by STB results, averaged over streamwise (y) direction

As a reference, we measure the horizontal profile of the continuous light intensity (light meter, Extech HD450) at the bottom of the convection chamber (red curve, Figure 3 left) and at a height of 1.10m across the measurement volume (blue curve). The illuminated volume is well defined by a sharp decay in light intensity, which helps to avoid light scattering from particles outside the measurement domain. An intensity dip in the center can be attributed to inhomogeneous distribution of LEDs in the light source and possible variations of intensity output between different LED arrays. Overall, the intensity of approximately $4.5 * 10^4$ lux corresponds to ~ 17,000 lm over the whole area.

Another view on the light intensity distribution is given in Figure 3 (right), which depicts a 2D ensemble average of the intensity of all particles as identified by the STB tracking (see paragraph 3) over a run of 500 images. The x- and z- direction of the measurement volume were discretized in 2×2-pixel bins and all particles located within such a bin are averaged (thus averaging over the streamwise direction of space). The different intensity of the LED spots can be easily seen. The spot located at x= 200 mm, z= 200 mm appears to have nearly double the intensity of the weakest one, located at x= -200 mm, z= 200 mm. These findings document that great care should be taken before assembling the illumination arrays in order to achieve homogenous lighting conditions.

Additionally it can be seen that due to the opening angle of 9° the illuminated region is widened from 75 cm at the top window to around 80 cm in the measurement volume. Given the height of

approx. 110 cm, the total volume comprises approx. 550 liters.

2.4. Camera system

The camera setup consists of five cameras (pco.edge 5.5 sCMOS, PCO) with a resolution of 2560×2160 pixels. They are arranged in a flat M-configuration with a small height difference between neighboring cameras of 15 cm (see Figure 1 and Figure 4). The cameras are placed on a circle around the convection chamber with a distance of 2.25 m to the vertical center line of the cylindrical measurement volume such that the cameras look perpendicularly through the front window.
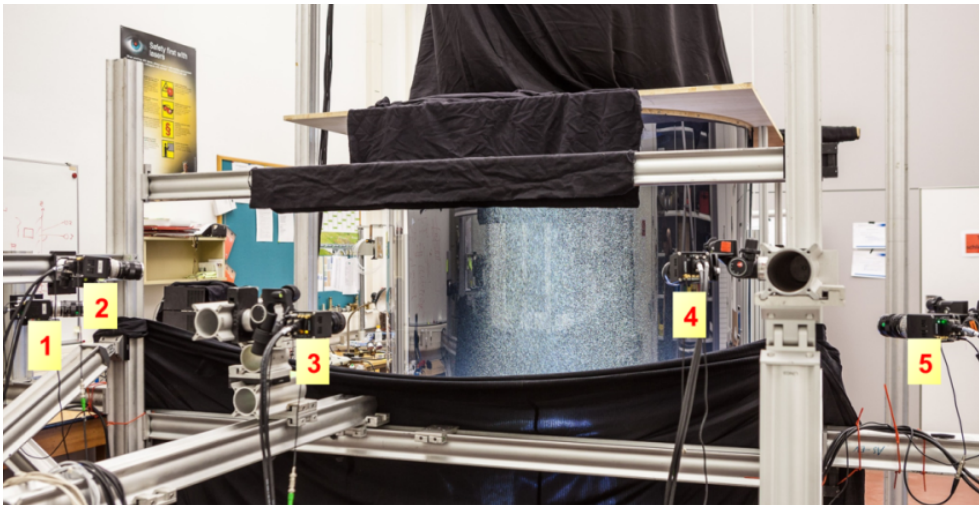
Figure 4: Photograph of the camera system and the convective cell.

The lines of sight of the outermost cameras have an angle >90° allowing for an accurate reconstruction of the particle position in all dimensions. The cameras are equipped with f=35mm lenses (Zeiss Contax) with the aperture set to $F_\#$ = 11, yielding sufficient depth of field to image the volume with 0.8 m diameter. The cameras are rotated by 90°, so their FOV has a width of 0.85 m and a height of 1.1 m to capture the vertically extended cylindrical volume.
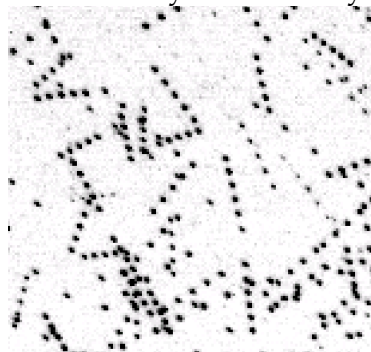
Figure 5: Unprocessed camera image (only background of 340 pixels subtracted, low seeding density, color

inverted). Sum of five consecutive images

The magnification is M = 0.016, corresponding to 0.4 mm/pix, so the observed velocities of up to 0.3 m/s result in a maximum particle shift of ~25 pix between two images. With this magnification, the maxima of the two glare points - reflections at two points of the bubble (Kühn et al. 2011, Scarano et al. 2015) - fall into one pixel, such that we obtain an isotropic circular particle image with a single peak. The image quality in general is very high. The production of the bubbles is a very stable process, leading to a monodisperse distribution of particle sizes at around 300 μm. Additionally, the absence of coherent light prevents effects of interferences or speckles. An example of the high image quality is given in Figure 5, showing a detail of a sample run at low seeding density and random flow. The sum of five consecutive images is shown, documenting the uniformity in particle imaging - both between the different bubbles, as well as in time for a single bubble.

For the 3D calibration of the cameras, a planar calibration target is set into the convection chamber and shifted to three positions, each 200 mm apart. Image acquisition, synchronization of cameras and the light source were controlled through the DaVis software (LaVision). The accuracy of the volumetric calibration was enhanced using Volume-Self-Calibration (Wieneke 2007); the particle imaging was calibrated, yielding a volumetrically resolved Optical Transfer Function (OTF, Schanz et al. 2013a).

All five cameras were connected to a single PC and the data was recorded directly to hard disc. The recording frequency of 29 Hz reflects the maximum write rate that was attainable. Connecting each camera to a single PC would allow for a repetition rate up to 100 Hz.

3. Data evaluation
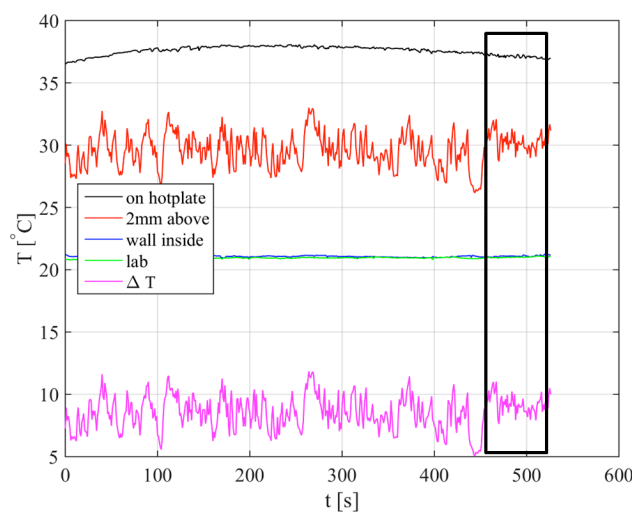
Figure 6: Temperature log for the different sensors, including a relaxation time. Measurement time (1000
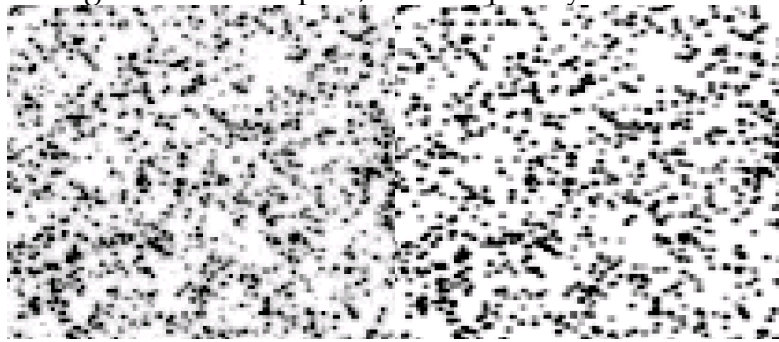
images) marked by rectangle

For this feasibility study a single measurement run was chosen from the available material, which features suitable conditions (high particle density, large particle shift) to assess how well the STB-evaluation copes with the current setup. The seeding density was found to be approx. 0.08 ppp in the center of the image (see Figure 7, left). The temperature difference between the air directly over the hot plate and the surrounding air was approximately 8°C (see Figure 6), which led to maximum velocities of around 0.35 m/s, corresponding to a maximum particle shift of nearly 30 pixels. 1000 images were recorded at a frequency of 29 Hz, of which 500 were evaluated using STB. Image preprocessing consisted of subtracting the smoothed minimum image and a constant of 50 counts.

3.1. Shake-The-Box processing

The DLR in-house STB algorithm was applied. For details of the method, please refer to (Schanz et al. 2016). The following parameters were applied to the current dataset: The number of triangulation iterations was set to 2 - using an allowed triangulation error of 1.0 pixel - followed by one triangulation iteration using a reduced set of cameras. Each of these was followed by five shake-iterations (n_1= 2, n_2=1, m = 5, ε = 1.0; see (Wieneke 2013)). For the initialization phase (the first four images), the number of triangulations was doubled. No help of a predictor in form of a vector field gained by TOMO-PIV processing was used in this case. As shown in Figure 8 (left), the number of tracked particles quickly increases with the number of processed images. After the initialization phase, around 53,000 tracks of length four are found. This number rises to 110,000 only three time-steps later and reaches 200,000 at time-step 18. For this first pass (going forward in time), the number of tracked particles saturates at just over 250,000 after about 80 images. From there on, the tracked particle number decreases, as bubbles burst and disappear from the tracking system. When reaching the end of the time-series, time is reversed and the algorithm walks backwards in time through the dataset. By doing so, known tracks that were not

immediately found within the first pass are extended and the first time-steps, where the tracking system was not yet converged in the first pass, are completely reconstructed.

Figure 7: (Left) detail from camera image; (right) reprojection of particle distribution as reconstructed by STB.

At maximum, around 275,000 particles are simultaneously tracked in the second pass. To the knowledge of the authors, this is an unprecedented number for particle tracking methods, which typically operate with hundreds or a few thousand particles within the same image.

Figure 7 compares the camera image to the virtual image, created by reprojecting all tracked particles. The high quality of the tracking process is documented in Figure 8 (right), which shows a statistic of the track length after pass 2. A very distinct peak can be seen at 500 images, showing that over 81,000 particles have been tracked over the whole time-series. These are particles slowly moving in the entrainment region. The rest of the tracks show lengths that are quite evenly distributed, reflecting the fact that many particles are transported at different speeds out of the volume. The computational effort, combined for both passes, is around 220 seconds per time-step. For comparison, the volume corresponds to a voxel space of 1960×2696×1960 voxels when using a voxel-to-pixel ratio of 1.0. Reconstruction and correlation times of such spaces of 10 teravoxel are very high.
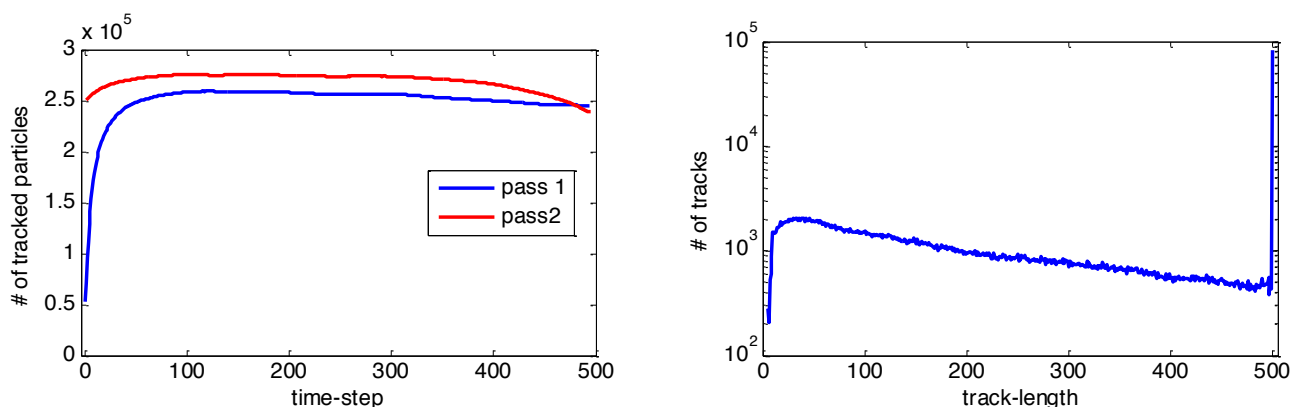
Figure 8: (Left) Development of number of tracked particles over time for both passes of STB; (right) track-length statistics after pass 2.

Following the second pass, the particle tracks are temporally filtered by means of an optimal Wiener filter, being represented by a series of 1D-B-Splines. On average, the particles are moved 0.116 px from their original position (0.046 px in x-, 0.040 px in y- and 0.083 px in z-direction) by the fit procedure. The velocity- and acceleration values are calculated as derivatives of the polynomial.

4. Flow field results

Figure 9 shows an instantaneous flow situation, depicted by ca. 275,000 tracks, whose velocity vectors are drawn for three consecutive time-steps and color-coded by streamwise velocity. Views from the front side and the top are provided. A large region of slowly moving particles, surrounding the thermal plume can be seen. The maximum velocity values are approx. 0.35 m/s, corresponding to a particle shift of nearly 30 pixels. The central regions can be better recognized in Figure 10, which shows the same tracks as Figure 9, albeit only a middle slice of 10 cm depth is shown, both for the x/y- and the z/y-plane. In general, the shape of the plume was varying in time quite visibly; in this time instant it can be seen that the plume is broadened in the z-direction, compared to the x-direction.

Vortical structures can already be identified by looking at the particle tracks; however a quantitative description is possible when quantities like vorticity or the Q-criterion are available. To this end, the discrete Lagrangian velocity- and acceleration information at the particle location is interpolated onto an Eulerian grid using the DLR in-house algorithm 'FlowFit' (Gesemann et al. 2016).
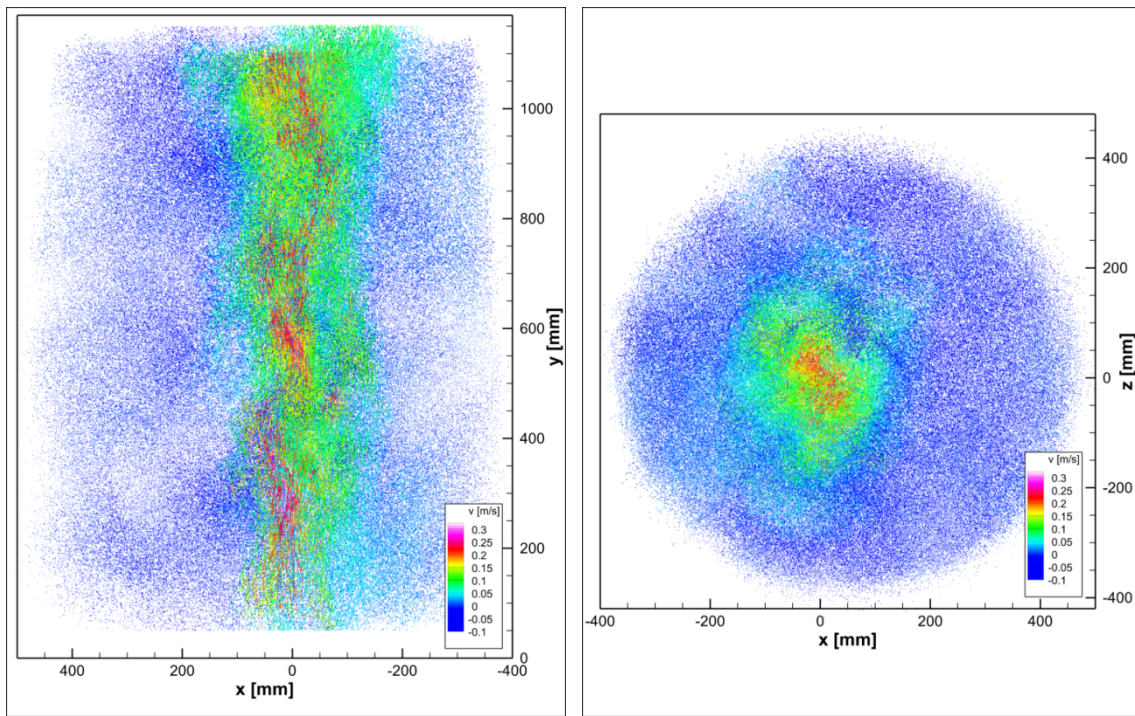
Figure 9: Front- and top view of the measurement volume, showing approx. 275.000 tracks for three consecutive time-steps, connected by the velocity vectors. Color-coding by streamwise velocity (v).
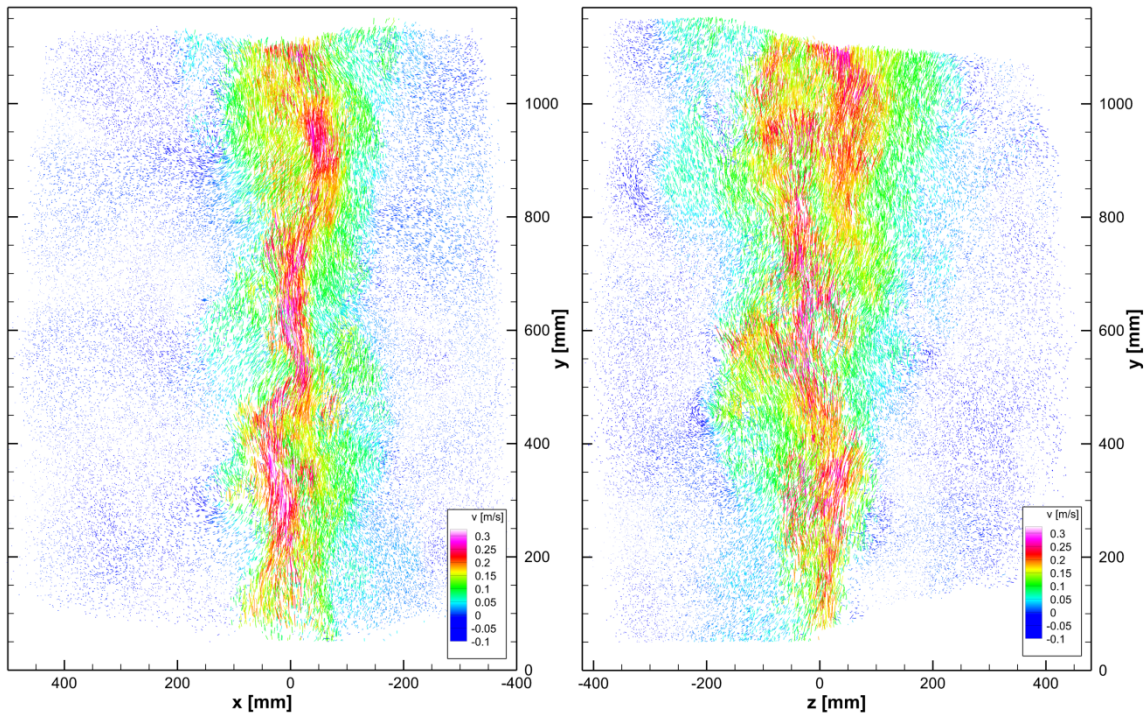
Figure 10 Front- and side view of the measurement volume, showing a middle slice of 100 mm thickness. Tracks shown for three consecutive time-steps, connected by the velocity vectors; color-coding by streamwise velocity (v).

The method models each component of the flow field as a weighted sum of three-dimensional and evenly spaced cubic B-splines. In order to evaluate this flow field on arbitrary coordinates, the weights have to be determined according to the known flow speeds at certain locations (being the particles with their velocity and acceleration). This results in a linear equation system where for each known flow speed at some particular position three equations are created. In addition to these equations based on the measurements other equations are used to regularize the equation system by penalizing non-zero curvatures (which is known in combination with spline fitting as "smoothing spline") and optionally (when the flow can be regarded as incompressible) by penalizing non-zero divergencies on a regular grid. This results in an overdetermined system where measurements and different kinds of regularizations can be weighted differently depending on how strong the smoothing effect should be, for example. This equation system is solved iteratively via the conjugate-gradient algorithm. The resulting flow field is then sampled on a regular grid including its spatial derivatives so that the derived values, such as vorticity or Q-criterion can be computed without numerical differentiation (Gesemann et al. 2016). For results of STB+FlowFit applied on Case D of the fourth international PIV Challenge, please see the DLR results in Kähler et al. (2016).
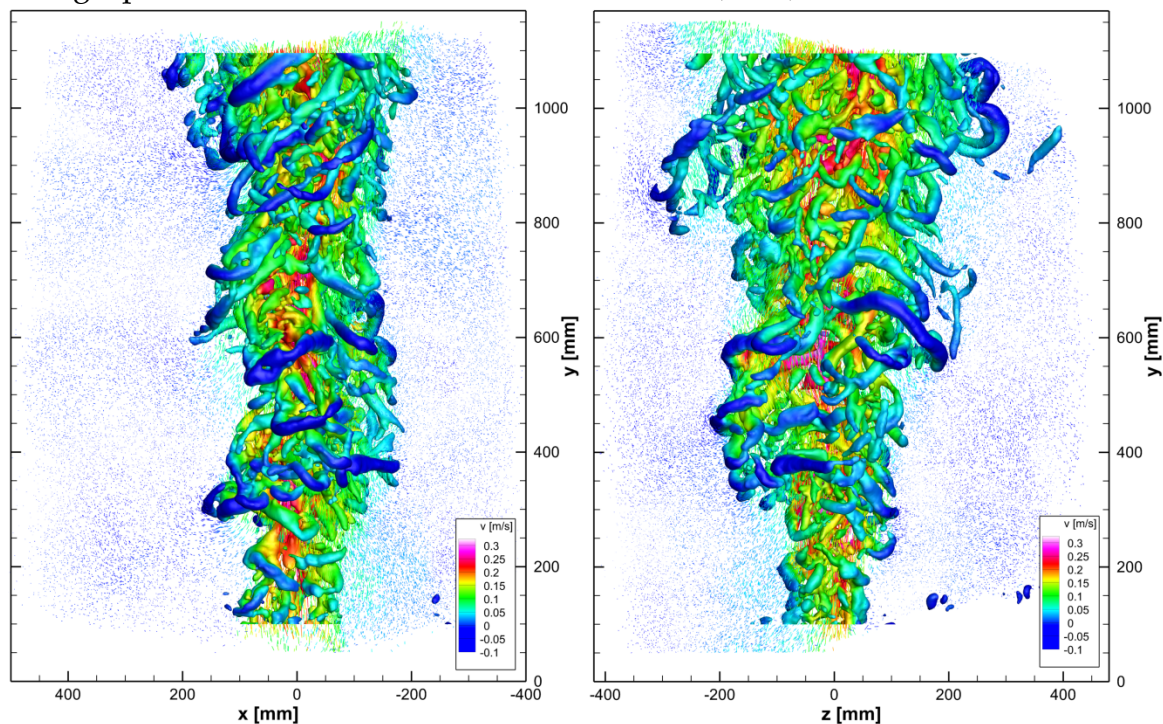

Figure 11: Isosurfaces of Q-criterion for one time-instant, color-coded by streamwise velocity. (Left) front view (looking from the cameras); (right) side view. Isosurfaces superimposed on the tracks as shown in Figure 10

FlowFit was applied to each of the 500 time-steps, taking velocities of the respective tracked

particles as data base. The B-spline system is setup, such that on average ten B-spline cells are present for every particle (0.1 *particles per cell*), leading to a spacing of approx. 8 mm between the cells. The compressibility constraint was applied (which is reasonable in the first approximation, given the low velocities and temperature gradients present), such that a penalization was put on divergency. 1500 iterations of the conjugate-gradient algorithm are applied to solve the equation system. The resulting continuous function is closely sampled on a grid with 3 mm spacing, ultimately leading to vector volumes of 266×333×266 vectors.

An example of the results gained by applying FlowFit to the STB track data is given in Figure 11, which shows isosurfaces of Q-criterion for the middle of the three time-steps shown in Figure 9 and Figure 10. The full amount and extent of the thermal flow structures becomes apparent. Long, undisturbed vortices are identified in the shear layers surrounding the center of the plume, while the central structures are smaller, but show equal strength. When looking at a time-series of such images a high temporal coherence is noticeable. The high quality of the tracking process translates directly into the quality of the Eulerian representation.

The high position accuracy, which is achievable due to the image quality, allows for an evaluation of particle acceleration (material derivative) - being the second derivative of space. Figure 12 shows tracks for the same time-steps as discussed before, color-coded by streamwise acceleration. Especially the detail view reveals that the acceleration and deceleration of particles drawn into flow structures is rendered smoothly by the tracking scheme. Just as for velocity, the acceleration can be interpolated onto an Eulerian grid using FlowFit. Figure 12 (right) displays the result of applying FlowFit to the particle accleration distribution in Figure 12 (left), with a penalization of rotation (therefore implying that viscosity effects plays a minor rule, which should hold for the small temperature differences present). Isosurfaces of streamwise acceleration indicate the regions of acceleration and deceleration above and below vortices (compare to Figure 11, right).
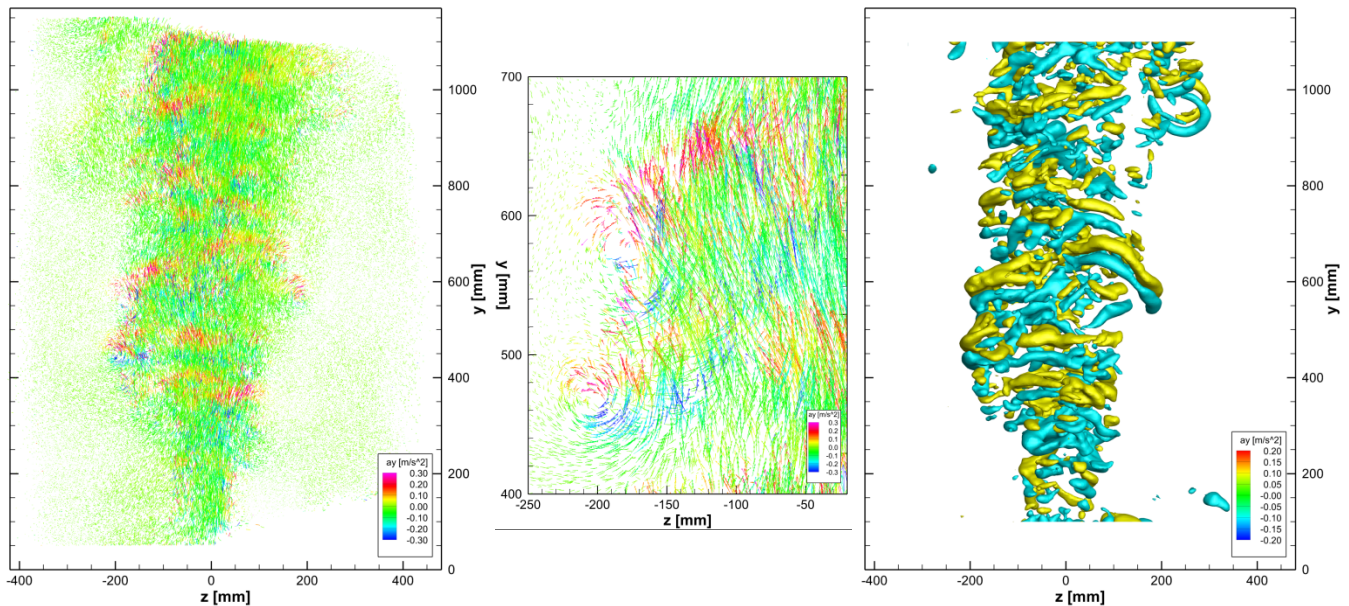
Figure 12: (Left) Tracks for three consecutive time-steps, color-coded by streamwise acceleration. Volume reduced to a middle slice of 200 mm; (Middle): Detail view from (left); (rigth) Isosurfaces of streamwise acceleration as extracted by FlowFit.

5. High-repetition measurement on an Impinging Jet

The results of the previous chapters demonstrate that combining HFSBs as tracer particles with LED illumination allows for a very accurate particle tracking in large volumes. In order to overcome the limitations in flow velocity, another experiment was created, using high-speed cameras and the latest generation of high-power LEDs to obtain enough light within the short pulse widths required at high flow velocities.

The experiment was set up in the same cylindrical chamber, in which the convection experiment was conducted (see paragraph 2.1). An air jet generated by a fan (PHYWE - 02742-93, upper and lower screen removed) with a nozzle exit diameter of $D = 11$cm and a variable exit velocity hits a flat acrylic glass plate at a distance of $H = 55$cm, $H/D = 5$, and at an angle of $\theta = 90°$. In the large measurement volume adjacent to the wall ($450×500×150$mm^3) the flow is seeded with helium-filled soap bubbles (HFSB) with diameters ranging from of ~ 300µm to ~ 500µm, depending on the air pressure applied on the generator (LaVision HFSB generator). Six high-speed cameras (PCO dimax) record particle images at different frame rates, ranging from $f = 1.25$kHz to $f = 3.9$kHz, depending on the flow velocity.
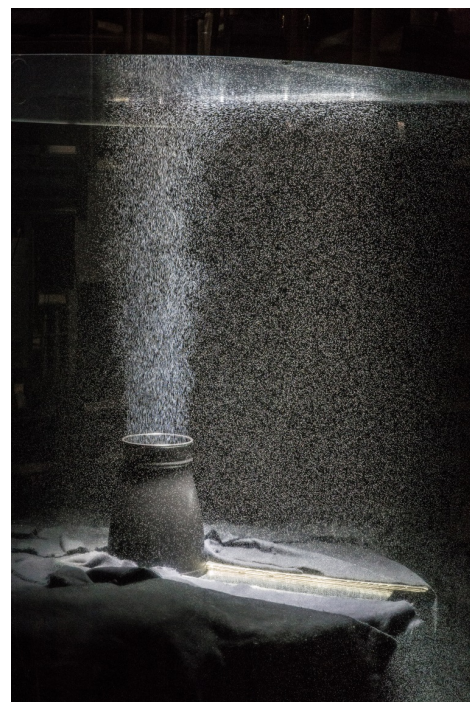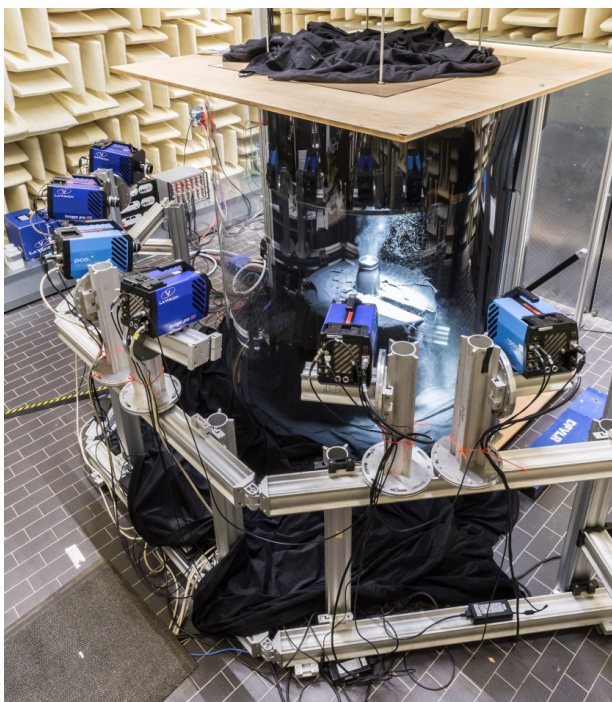
Figure 13: Photographs of the impinging jet experiement. Top: nozzle as viewed by the cameras sytem; down left: overview of the experimental apparatus; down right: nozzle in operation, emitting a turbulent jet, which impinges on the acrylic plate 55 cm above. The illuminated region can be recognized, with the light exiting through a passe-partout above the acrylic plate and being backreflected by a mirror below the nozzle.

The cameras are positioned in an in-line configuration and oriented in a way that lines of sight are tangential to the flat plate. The HFSBs are illuminated by two different pulsed LED arrays from above (through the acrylic glass plate). The central jet core is illuminated by a circular array of 150 high power LEDs, operated at 20 A (LaVision prototype); a double HARDsoft array of 42 LEDs each (operated at 90 A) is illuminating an area of approx. 20 cm in depth and 45 cm in

radial direction along the glass plate. The LEDs are operated at 10 % duty cycle. See Figure 13 for images of the setup.

Flow measurements at different jet velocities and aquisition rates were realized. The acquisition rate has to be increased accordingly, in order to restrict the particle shifts to no more than approx. 20 pixels. The usable resolution of the cameras sinks with the acquisition frequency, therefore reducing the imaged volume. At full resolution (which is in this case usable up to jet exit velocities of approx. 5 m/s) a volume of 47 liters could be reconstructed. At the highest repetition rate (3.9 kHz) the volume decreases to around 13 liters.

In order to ensure sharp particle imaging - avoiding temporal streaking - the pulse width – and therefore the available light – has to be reduced accordingly. For the shortest pulse widths an increase of the bubble size from 300 μm to 500 μm can increase the amount of reflected light reaching the cameras.

Operating at 1.25 kHz, using the full camera resolution, large numbers of bubbles (up to 190,000) could be successfully tracked with reliability similar to what has been demonstrated in the thermal plume experiment. Decreasing the pulse width leads to decreased signal-to-noise ratio, however even at a pulse-width of 27 μs the image quality allows for a very reliable tracking at seeding densities up to 0.045 ppp. Probably, even higher particle image densities could be processed. However, at high jet velocities (> 10 m/s) it was not possible to generate more bubbles, as likely a large number was destroyed by the blades of the fan.

Three exemplary cases were chosen for presentation within this text. Table 1 summarizes the main parameters of these cases. Very different jet velocities are presented, in order to document the differences in tracking properties, attainable volumes and resolvable scales.

Jet velocity [m/s]	Acquisition Rate [kHz]	Bubble size [μm]	Pulse width [μs]	Image size [px]	Volume [mm³]	Number of inst. tracked bubbles
1	1.25	~300	80	2016×2016	450×530×200	106,000
5.5	2	~500	50	1344×1808	370×530×200	118,000
16	3.9	~500	27	576×1728	180×530×140	40,000

Table 1: Parameters for volumetric measurements on impinging jet for different flow velocities
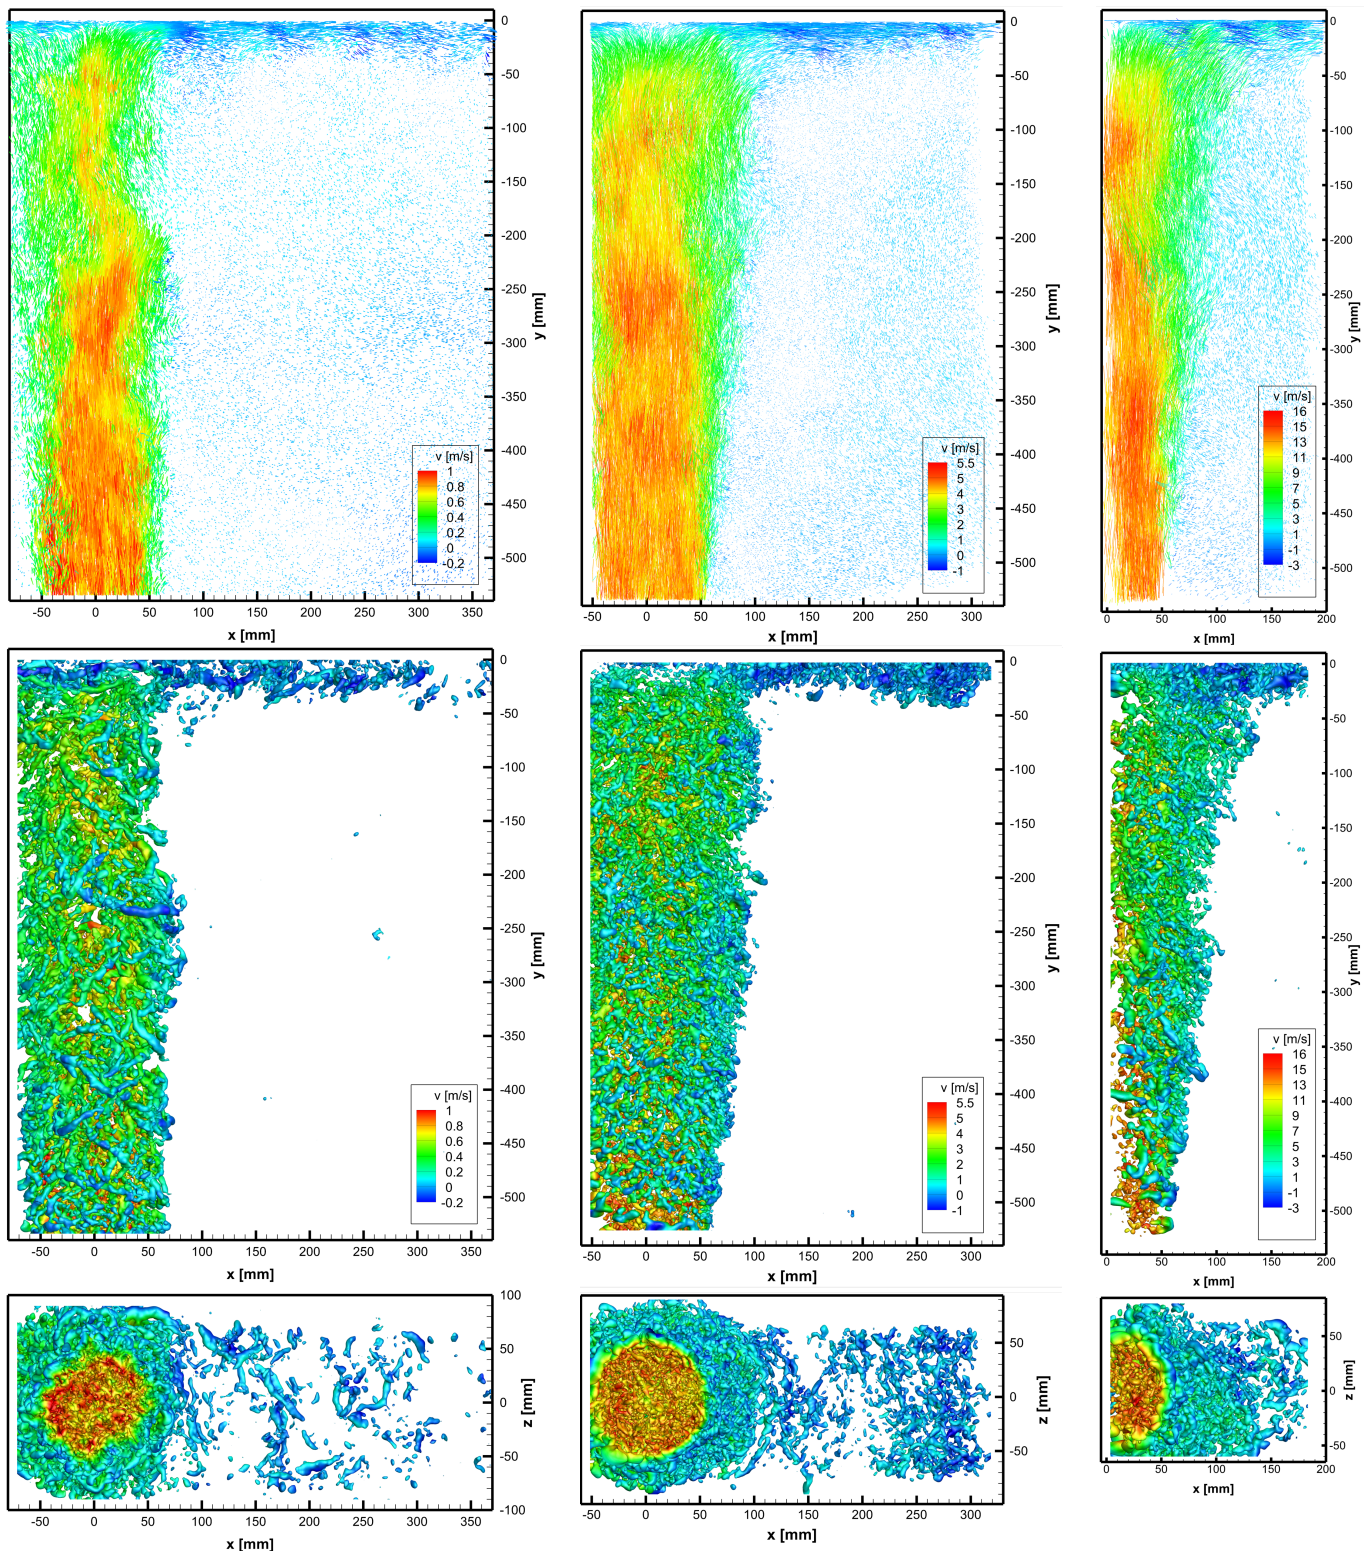
Figure 14: Results from STB + FlowFit evaluations for three flow velocities (from left to right 1 m/s, 5.5 m/s, 16 m/s) of an impinging jet. The top row shows velocity vectors of 9 successive time-steps from a 60 mm-slice in the center of the volume, color-coded by streamwise velocity. The second row displays isosurface of the Q-criterion, as extracted from FlowFit (from left to right: $Q = 700/s^2$; $Q = 25{,}000/s^2$; $Q = 150{,}000/s^2$) for the middle time-step. The third row shows the same isocontours, albeit looking upwards towards the impinging plate.

The following STB parameters were applied to all datasets: The number of triangulation iterations was set to 1, followed by one triangulation iteration using a reduced set of cameras. Each of these was followed by five shake-iterations ($n_1 = 1$, $n_2 = 1$, $m = 5$, $\varepsilon = 1.0$; see (Wieneke 2013)). The allowed triangulation error was set to 1.0 pixel. For the initialization phase (the first four images), the number of triangulations was doubled. The search radius for new tracks was chosen according to the expected particle shift. As soon as enough tracks are present to serve as predictor for the track identification, a constant search radius of 4 pixels around the predictor point was used. Outliers were identified using a neighborhood criterion (velocity difference larger than 8 times the rms). All cases quickly converged to a stable solution; from there on, the algorithm can quickly work through the time-series, as only few particles need to be newly triangulated (typically 1,000-3,000). As a reference, for the 3.9 KHz-case 2,500 images could be processed with two passes of STB overnight on a 20-core server.

Following a successful tracking of the particles, FlowFit (Gesemann et al. 2016) was applied to the results. In contrast to the thermal plume case, which is driven by density gradients, the full Navier-Stokes regularization of FlowFit (including the material derivative) can be used in this case. A closely spaced system of Cubic B-splines is setup (here around 0.04 particles per cell); the resulting equation system is iteratively solved using an LBFGS solver. The gained continuous function is closely sampled on a grid with 1 mm spacing, resulting in velocity volumes of 540 points in streamwise direction and varying width and depth, depending on the case.

Figure 14 shows examplary snapshots from the tracked particles and the FlowFit results from two perspectives. It can be seen that for all cases the tracking system was able to extract bubble trajectories at high numbers. Visual inspection shows no traces of obviously falsely tracked particles. Turning to the FlowFit results, for the lowest velocity (1 m/s, Figure 14 left) a multitude of flow structures can be identified using the Q-criterion. The flow coming form the fan is in a turbulent state, however large, elongated vortices can still be detected. Especially along the impinging plate well defined structures can be seen. At higher velocities (5.5 m/s, Figure 14 middle) the structures become smaller and higher in number (please note the increasing threshold value for the isosurfaces). While the larger structures can still be resolved – notably at the impinging plate – the small structures are most likely underresolved. When looking at a time series (not available here) the small structures start to flicker between the different time-steps. At the highest velocities (16 m/s, Figure 14 left) the structures become even stronger (again, note the isosurface threshold), as the turbulence increases. Still, the FlowFit method is able to extract temporally coherent large structures from the STB tracks; the small scales are not fully resolved in this highly turbulent dataset.

While the interpolated result might not resolve all scales, the particle tracks still contain this information. Particle statistics (PDFs, Bin ensemble averaging) should be free of any bias, as long as the temporal resolution is high enough.

The volumetric results from the impinging jet measurements will also be used for volumetric pressure reconstruction. Three reference microphones were integrated into the impinging plate (see Huhn et al. 2016 for details)

6. Conclusions and Outlook

Large-scale time-resolved volumetric flow measurements on a thermal plume and on an impinging jet in air were presented. In both cases Helium filled soap bubbles are used as flow tracers and are illuminated by an array of high power white LEDs.

The thermal plume investigation used a measurement volume of 550 liters. Up to 275.000 bubbles could be tracked simultaneously using the Shake-The-Box algorithm. Highly resolved vector volumes of velocity and acceleration, interpolated onto an Eulerian grid using the FlowFit method, show multitudes of flow structures. To the knowledge of the authors this investigation contains both the largest volume in which instantaneous flow measurements were successfully performed as yet and the largest number of tracked particles for PTV experiments.

These features can be realized mainly due to the high image quality. The bubble size distribution is very sharp around 300 μm (monodisperse) and all bubbles scatter the light very similarly to all cameras. The white, uncoherent light produced by the LEDs avoids effects like speckles or interference, leading to temporally very consistent particle images. All these aspects are beneficial for a reliable tracking process of STB and allow for the high particle concentrations of up to 0.1 ppp.

In order to demonstrate the possibilities of performing similar measurements at much higher flow velocities, enabling the operation in typical wind tunnel experiments, a second experiment using high-speed imaging of the flow of an impinging jet was carried out

Brighter LEDs of the latest generation, were operated in the kHz-range, allowing the tracking of bubbles at flow speeds of up to 16 m/s. While the volume was smaller in this case (ranging from 13 to 47 liters, depending on the repetition rate), solutions to enlarge the usable volume are foreseeable. The high-power LED arrays used for this investiagtion are still under development and therfore the availablity was limited. In the future the illumination of larger volumes, using large arrays of LEDs can be attained. Combined with new generations of cameras, that allow higher frame-rates at higher resolutions, the volume size and usable flow speeds can be further

expanded into regimes that are typical for low-speed wind tunnel experiments (e.g. 60 m/s). Applying such a setup to a wind tunnel measurement requires a much higher number of bubbles (Scarano et al. 2015), compared to the closed cell used in these experiments. However, bubble generators with 50 or more nozzles are just becoming available, possibly solving this problem in the near future.

References

Biwole PH, Yan W, Zhang Y, Roux JJ (2009) A complete 3D particle tracking algorithm and its applications to the indoor airflow study. Meas Sci Technol 20 115403

Bosbach, J., Kühn, M., Wagner, C. (2009) Large scale particle image velocimetry with helium filled soap bubbles, Exp. Fluids 46, 539-547

Elsinga G E, Scarano F, Wieneke B and van Oudheusden BW (2006) Tomographic Particle Image Velocimetry. Exp Fluids 41:933–947

Gesemann S, Huhn F, Schanz D and Schröder A (2016) "From Noisy Particle Tracks to Velocity, Acceleration and Pressure Fields using B-splines and Penalties", 18th Int Symp on the Application of Laser Techniques to Fluid Mechanics, July 4-7, Lisbon, Portugal

Gilet T, Scheller T, Reyssat E, Vandewalle N, and Dorbolo S (2007) How long will a bubble be? arXiv:0709.4412

Huhn F, Schanz, Gesemann S, Schröder A (2016) FFT integration of instantaneous 3D pressure gradient fields from Lagrangian particle tracking in turbulent flows. 18th Int Symp on Applications of Laser Techniques to Fluid Mechanics, (Lisbon, Portugal, July 04-07)

Kähler CJ, Astarita T, Vlachos PP, Sakakibara J, Hain R, Discetti S, La Foy R, Cierpka C (2016), Main results of fourth International PIV Challenge, Exp. Fluids 57:97, DOI: 10.1007/s00348-016-2173-1

Klimas P (1973) Helium bubble survey on an opening parachute flow field. J Aircraft 10:567-569

Kühn, M., Ehrenfried, K., Bosbach, J., and Wagner, C. (2011) Large-scale tomographic particle image velocimetry using helium-filled soap bubbles, Exp. Fluids 50, 929-948

Maas, H G, 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

Melling A (1997) Tracer particles and seeding for particle image velocimetry. Meas. Sci Technol 8:1406

Müller RHG, Flögel H, Schere T, Schaumann O, Markwart M (2000) Investigation of large scale low speed air conditioning flow using PIV. 9th Int Symp on Flow Visualization, Edinburgh, UK

Pounder E (1956) Parachute inflation process Wind-Tunnel Study, WADC Technical report 56-391, Equipment Laboratory, Wright Patterson Air Force Base. Ohio, USA, pp 17-18

Scarano, F., Ghaemi, S., Caridi, G., Bosbach, J., Dierksheide, U., Sciacchitano, A: On the use of helium-filled soap bubbles for large-scale tomographic PIV wind tunnel experiments, Exp. Fluids, 311 56:42, 2015.

Schanz D, Gesemann S, Schröder A (2016) Shake-the-box: Lagrangian particle tracking at high particle image densities. Exp Fluids 57:70

Schanz D, Gesemann S, Schröder A, Wieneke B, Novara M (2013a) Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. Meas Sci Technol 24 024009

Schanz D, Schröder A, Gesemann S, Michaelis D, Wieneke B (2013b) Shake-the-Box: a highly efficient and accurate Tomographic Particle Tracking Velocimetry (TOMO-PTV) method using prediction of particle position. 10th Int. Symposium on Particle Image Velocimetry – PIV13 (Delft, The Netherlands, July 1-3)

Schneiders JFG, Azijli I, Scarano F, Dwight RP (2015) Pouring time into space. 11th Int. Symposium on Particle Image Velocimetry – PIV15 (Santa Barbara, California, September 14-16)

Tobin, S.T., Meagher, A.J., Bulfin, B., Möbius, M., Hutzler, S.: (2011) A public study of the lifetime distribution of soap films, Am. J. Phys. 79(819), 819-824

Wieneke B (2013) Iterative reconstruction of volumetric particle distribution. Meas. Sci. Technol. 24:024008

CrossMark

RESEARCH ARTICLE

FFT integration of instantaneous 3D pressure gradient fields measured by Lagrangian particle tracking in turbulent flows

F. Huhn[1] · D. Schanz[1] · S. Gesemann[1] · A. Schröder[1]

Abstract Pressure gradient fields in unsteady flows can be estimated through flow measurements of the material acceleration in the fluid and the assumption of the governing momentum equation. In order to derive pressure from its gradient, almost exclusively two numerical methods have been used to spatially integrate the pressure gradient until now: first, direct path integration in the spatial domain, and second, the solution of the Poisson equation for pressure. Instead, we propose an alternative third method that integrates the pressure gradient field in Fourier space. Using a FFT function, the method is fast and easy to implement in programming languages for scientific computing. We demonstrate the accuracy of the integration scheme on a synthetic pressure field and apply it to an experimental example based on time-resolved material acceleration data from high-resolution Lagrangian particle tracking with the *Shake-The-Box* method.

1 Introduction

The pressure field in a turbulent flow reveals important coherent structures, e.g., vortex cores, and it is most relevant to determine the loads exerted by the fluid on a structure in the flow, e.g., a wing. Yet, spatially well-resolved pressure fields $P(\mathbf{x}, t)$ in three-dimensional unsteady flows are difficult to measure. Optical non-intrusive flow measurement techniques, in particular the novel Lagrangian particle tracking approach *Shake-The-Box* (Schanz et al. 2013; Schanz

2016), provide accurate and dense fields of the material acceleration in the fluid. Since, for the particle tracking approach, the material acceleration data are given at scattered particle positions, they are interpolated to a regular grid using physical regularizations of the interpolant function, as e.g., zero curl for the material acceleration field. The pressure gradient field $\nabla P(\mathbf{x}, t)$ is then derived from the Navier–Stokes equation for incompressible isothermal fluid and it is integrated in space. For this spatial integration, almost exclusively two methods have been used in the aerodynamic and hydrodynamic flow measurement community (van Oudheusden 2013): first, the direct numerical integration of ∇P along various paths in space with some spatial marching scheme based on finite differences,

$$P(\mathbf{x} + \Delta \mathbf{x}) = P(\mathbf{x}) + \nabla P \Delta \mathbf{x}, \tag{1}$$

and second, the numerical solution of the Poisson equation,

$$\Delta P = \nabla \cdot (\nabla P) \tag{2}$$

obtained by applying the divergence operator to the experimental pressure gradient field. The necessary Dirichlet or Neumann boundary conditions for (2) are determined from known flow conditions at the boundary or directly from the measured data set. Examples in which the first method was used include the experimental studies by Liu and Katz (2006), Dabiri et al. (2014), Tronchin et al. (2015); the second method has been applied, e.g., by Ghaemi et al. (2012), Ragni et al. (2012), Kat and Oudheusden (2012), Novara and Scarano (2013), Pröbsting et al. (2013), Neeteson and Rival (2015). Charonko et al. (2010) compare both methods for pressure reconstruction from 2D PIV measurements.

Certainly, the problem of reconstructing a two- or three-dimensional scalar field from its spatial gradient field, that is given as a discrete data set, is not limited to pressure reconstruction from experimental flow data, but has also

✉ F. Huhn
florian.huhn@dlr.de

1 Department of Experimental Methods, German Aerospace Center (DLR), Institute of Aerodynamics and Flow Technology, Bunsenstrasse 10, 37073 Göttingen, Germany

🌀 Springer

been adressed in several other scientific fields. In computer vision, Frankot and Chellappa (1988) were the first to propose a fast method to integrate the surface gradient of a three-dimensional body, obtained from images with the 'shape from shading' visualization technique. Their method enforces integrability and integrates the gradient field in Fourier space. Independently, Zhang (1996) introduces the very same method to reconstruct water surface elevations from optical surface slope measurements. He additionally considers the treatment of boundary effects in non-periodic domains. Both works consider two-dimensional problems, but the extension of the method to three dimensions is straightforward, as will be shown in the following.

This non-iterative integration technique in Fourier space is common in numerical approaches (e.g. Laizet and Lamballais 2009), but has surprisingly not been considered for pressure reconstruction in the aerodynamics community, so far. To the best knowledge of the authors, exclusively one group in medical imaging applies this fast integration technique to fluid flows (Wang and Amini 2005; Negahdar et al. 2013). They determine pressure gradient fields in blood flows, experimentally measured with magnetic resonance imaging, in order to obtain the pressure load on blood vessels. As a matter of low Reynolds numbers, the spatial complexity of the investigated blood flow is rather low, while we will consider multi-scale turbulent flows with a broad spectrum of spatial scales here.

The aim of this paper is to present the efficient integration method for spatial gradient fields and to show its applicability to 3D time-resolved pressure fields in turbulent flows. The gradient field data we use are based on accurate material acceleration fields that have become available recently through the advent of the high-density Lagrangian particle tracking (LPT) algorithm *Shake-The-Box* (Schanz 2016) in combination with the interpolation scheme *Flow-Fit* (Gesemann 2016).

The paper is organized as follows. In Sect. 2, we recall the principles of pressure reconstruction from flow measurements and present the integration method of the pressure gradient. In Sect. 3, we show examples of pressure reconstruction for a synthetic pressure field to validate the integration method, and in Sect. 4, we reconstruct the pressure field in an experimental turbulent jet flow. Finally, Sect. 5 summarizes the results.

2 Data and methods

2.1 Momentum equation

The momentum $\rho\mathbf{u}$ of a fluid parcel in an isothermal viscid Newtonian fluid evolves according to the Navier–Stokes equation

$$\rho\frac{D\mathbf{u}}{Dt} = \rho\left(\frac{\partial\mathbf{u}}{\partial t} + (\mathbf{u}\cdot\nabla)\mathbf{u}\right) = -\nabla P + \mu\Delta\mathbf{u} \tag{3}$$

with constant density ρ, constant viscosity μ, and the material acceleration $D\mathbf{u}/Dt$, i.e., the acceleration of a fluid element along its trajectory. Away from boundaries and strong shear flows, the pressure gradient force typically dominates the momentum balance,

$$\nabla P \gg \mu\Delta\mathbf{u}, \tag{4}$$

such that for flows with high Reynolds number, the viscous term is relatively small, and we get a good estimate of the pressure gradient field by

$$\nabla P \approx -\rho\frac{D\mathbf{u}}{Dt} = -\rho\,\mathbf{a}, \tag{5}$$

where we denote the material (or Lagrangian) acceleration with $\mathbf{a}$ from here on. In flows with low Reynolds number or in well-resolved measurements of boundary layers, the viscous term can be of the same order as the pressure gradient and has to be taken into account in these cases, for example, by a direct computation of the viscous term from highly resolved velocity data. For transonic steady flows, Oudheusden et al. (2007) describe a way to also include compressibility effects into the pressure reconstruction. For an overview over several formulations of the momentum equation for pressure reconstruction in different flow conditions, see the review paper by van Oudheusden (2013).

2.2 Measurement of Lagrangian acceleration

According to the term in brackets in (3), material acceleration could simply be obtained from time-resolved velocity fields as measured with particle image velocimetry (PIV). PIV is a robust method to obtain time-resolved velocity fields from pairs or sequences of particle images. However, this indirect composed measurement of material acceleration involves spatial and temporal derivatives of the velocity field (3). The derivatives enhance measurement noise in the velocity field, which can lead to noisy acceleration fields and, consequently, to noisy pressure gradient fields. Additionally, the spatial smoothing effect of the correlation window in the PIV technique may lead to an underestimation of the velocity gradient and to a bias of the derived material acceleration.

Lagrangian particle tracking (LPT) overcomes these problems by tracking individual fluid tracers, such that entire time-resolved trajectories of single particles $\mathbf{x_p}(\mathbf{x_0}, t_0, t)$ are known. While the well-known denotation particle tracking velocimetry (PTV) suggests the exclusive measurement of the fluid velocity, with LPT, the acceleration can additionally be determined as a point measurement at the position of individual tracers. Acceleration at particle

positions is obtained as $\mathbf{a}(\mathbf{x_p}(t)) = \mathrm{d}^2\mathbf{x_p}(t)/\mathrm{d}t^2$. The main source of error influencing the acceleration value is the uncertainty of the measured particle position, while no spatial smoothing is involved in the image processing.

LPT has long been used in statistical turbulence research at relatively low seeding densities (see e.g., Hoyer et al. 2005; Xu et al. 2007). Higher resolution has been reached by applying LPT to tomographic PIV data as a post-processing step in order to increase the accuracy of acceleration measurements (Schröder et al. 2011) and to reconstruct pressure fields (Novara and Scarano 2013). Only recently, the processable seeding densities in the particle images of direct LPT, measured in values of particles per pixel (ppp), increased substantially to values as large as 0.1 ppp (Schanz et al. 2014). Assuming a camera with a 1 Mpixel sensor, 0.1 ppp ideally corresponds to 10^5 simultaneously tracked particles, or an unstructured grid of roughly $50 \times 50 \times 50$ positions with accurate acceleration data. This spatially resolved acceleration field allows for a LPT-based pressure reconstruction in unsteady turbulent flows.

Here, we use the LPT algorithm *Shake-The-Box* by Schanz (2016). The strength of this algorithm is to use the temporal information of already established particle trajectories to predict the three-dimensional particle position in the next time step and identify the corresponding intensity peak in the camera images. The predicted particle position is then reprojected to the cameras using an estimated optical transfer function (OTF) (Schanz et al. 2013) and the particle position is iteratively corrected to optimally fit the intensity peaks in the camera images. Given the optimal particle position, the intensity signature of the particle in the camera images is subtracted from the images [iterative particle reconstruction, IPR (Wieneke 2013)], in order to reduce the complexity of the particle positioning problem. This procedure is crucial to handle high seeding densities. The discrete particle positions of an identified trajectory are fit by an optimal one-dimensional cubic B-spline curve, that takes the expected measurement noise into account. Temporal differentiation yields velocity and acceleration along the trajectory.

In a subsequent step, the acceleration data are interpolated to a Cartesian grid with the FlowFit algorithm (Gesemann 2016). Smoothing B-splines are fit iteratively to the scattered acceleration data. The B-splines are defined on a fine auxiliary grid with a selectable mean resolution given in particles per cell (ppc) with typical values of 0.1 ppc. The interpolant is smoothed by penalizing high frequencies in cells containing particles and in empty cells. Additionally, the curl of the acceleration field is penalized according to assumption (4) and tends to zero which supports integrability. Finally, the interpolant is evaluated on the output grid with a user-defined spatial resolution. A high spatial sampling resolution is beneficial for the Fourier transform in the pressure integration scheme below.

2.3 Integration scheme for pressure reconstruction

Following Frankot and Chellappa (1988) or Laizet and Lamballais (2009), we obtain the pressure field $P(\mathbf{x})$ by integrating the measured three-dimensional pressure gradient fields

$$\partial_x P = - \rho a_x \tag{6}$$

$$\partial_y P = - \rho a_y \tag{7}$$

$$\partial_z P = - \rho a_z \tag{8}$$

in Fourier space

$$\tilde{P}(\mathbf{k}) = \frac{\mathbf{k} \cdot \widetilde{\nabla P}}{\mathrm{i}|\mathbf{k}|^2} = \frac{k_x \widetilde{\partial_x P} + k_y \widetilde{\partial_y P} + k_z \widetilde{\partial_z P}}{\mathrm{i}(k_x^2 + k_y^2 + k_z^2)}, \tag{9}$$

and transforming back to normal space

$$P'(\mathbf{x}) = FT^{-1}\tilde{P}(\mathbf{k}). \tag{10}$$

For a short derivation of Eq. (9), see Ref. (Rocholz 2008). The tilde denotes a Fourier-transformed function, e.g., $\tilde{P}_x = FT(P_x)$, FT^{-1} is the inverse Fourier transform, and k_x, k_y, k_z are the components of the wave number vector $\mathbf{k}$. In (9), the separation of the curl-free longitudinal component of the vector field corresponds to a projection of the pressure gradient onto the k-vector $\mathbf{k}[\mathbf{k} \cdot \widetilde{\nabla P}]/|\mathbf{k}|^2$, and the integration in space corresponds to a division by $\mathrm{i}\mathbf{k}$.

Equation (9) has a singularity at $\mathbf{k} = \mathbf{0}$. In order to handle this, the amplitude for the constant component is set to zero, $\tilde{P}(\mathbf{k} = \mathbf{0}) = \mathbf{0}$. By this operation, the amplitudes of the constant component of the three pressure gradients are lost, e.g., $\widetilde{\partial_x P}(\mathbf{k} = \mathbf{0}) = \mathbf{0}$, but they can be reconstructed by adding linear planes to the pressure field (Zhang 1996)

$$P(\mathbf{x}) = P'(\mathbf{x}) + \langle \partial_x P \rangle x + \langle \partial_y P \rangle y + \langle \partial_z P \rangle z, \tag{11}$$

where $\langle \cdot \rangle$ is the spatial mean over the entire domain. Below, we will see that, in practice, we set boundary conditions at the periodic domain that account for the global linear pressure gradient. Due to small imperfections of the data in Fourier space (truncation errors etc.), the back transform in Eq. (10) generates a small imaginary part for the resulting pressure field. We only consider the real part and neglect the imaginary part. Finally, in order to obtain absolute pressure, the integration constant, a constant pressure offset P_0, obtained from additional measurements or from theoretical considerations at the boundaries, is added to the relative pressure field $P(\mathbf{x})$.

By construction, the integration scheme (9) satisfies the integrability constraint (Frankot and Chellappa 1988; Zhang 1996). The longitudinal component of a vector field is curl-free, i.e., partial derivatives commute, e.g.,

$$\partial_x \partial_z P(\mathbf{x}) = \partial_z \partial_x P(\mathbf{x}). \tag{12}$$

This is equivalent to the requirement that two different integration paths between two points in space yield the same result which imposes a smoothness constraint on the integrated field. In the presence of measurement noise, integrability is not necessarily fulfilled. For a direct path integration of the pressure gradient field, an alternative integration method (van Oudheusden 2013), the average over many paths is usually computed [using the mean or the median (Dabiri et al. 2014)] in order to reduce the bias of non-curl-free measurement noise. Instead, when using the FFT integration or the Poisson equation, the pressure gradient fields that are being integrated are curl-free by construction.

In order to put the integration scheme (9) in a more general perspective, we can look at the complete Navier–Stokes equation (3) in terms of the Helmholtz decomposition. According to the Helmholtz theorem, an arbitrary vector field $\mathbf{a}'$ decaying to zero at infinity can be decomposed into two components, one of which is divergence-free and the other component is curl-free (see, e.g., Griffiths 1999),

$$\mathbf{a}' = -\nabla\Phi + \nabla \times \mathbf{A}. \tag{13}$$

The curl-free component can be written as the gradient of the scalar potential Φ and the divergence-free component can be written as the curl of the vector potential $\mathbf{A}$. In the Navier–Stokes equation (3), the pressure gradient term is curl-free and the viscous term is divergence-free due to incompressibility, i.e., we can decompose the measured field of inertial forces $\rho\mathbf{a}$ into the pressure force and the viscous force by Helmholtz decomposition,

$$\mathbf{a}' = \rho\mathbf{a} = -\nabla P + \mu\Delta\mathbf{u} = -\nabla\Phi + \nabla \times \mathbf{A} \tag{14}$$

with $\Phi = P$ and $\mu\Delta\mathbf{u} = \nabla \times \mathbf{A}$. While the scalar pressure field P is given by (9), similarly, the vector potential is given by

$$\mathbf{A} = -FT^{-1}\left[\frac{\mathbf{k} \times \widetilde{\mathbf{a}'}}{i|\mathbf{k}|^2}\right] \tag{15}$$

where the cross-multiplication with the k-vector separates the divergence-free (transversal) component of $\mathbf{a}'$. In flows where the viscous force is of the same order as the pressure gradient force, this decomposition is a direct way to seperate the viscous force and to determine the contribution of the different forces. In more turbulent flows, the viscous force is typically small and the divergence-free component of the material acceleration is dominated by measurement noise. Therefore, we can use the concept of Helmholtz decomposition (13) to construct synthetic pressure gradient fields with non-zero curl component.

2.3.1 Boundary conditions for periodic domain

The Fourier transform assumes a periodic domain for the transformed fields, a condition that is typically not met by measurement data. A continuous extension of the field by mirroring the data has been proposed to minimize boundary artifacts due to non-periodicity of the data (Zhang 1996; Wang and Amini 2005). In the following, we discuss this approach and propose a new treatment of the boundaries that suppresses boundary artifacts while avoiding the need for excessive memory. The different boundary conditions discussed here are depicted in Fig. 1.

Continuous extension (CE) A periodic continuous extension avoids discontinuities at the boundaries of the field subject to the FT by reflecting the field about mirror planes along the boundaries. In (9) and (10), the Fourier transform acts on the pressure gradient field and on the pressure field. A crucial point is whether to construct continuous boundaries for ∇P or for P. Zhang (1996) proposes to continuously extend ∇P by reflecting it about the axes of the coordinate system such that
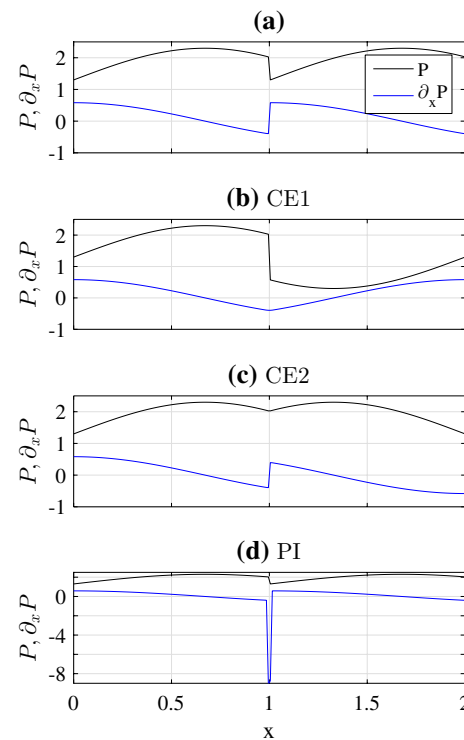
Fig. 1 Sketch of the different modified boundary conditions for the non-periodic pressure field P and pressure gradient field $\partial_x P$ (two periods of an arbitrary function are shown here for demonstration purposes). **a** Two periods of discontinuous data with period $L = 1$. **b** Continuous extension of $\partial_x P$ (*blue*) by mirroring, CE1, (Zhang 1996). **c** Continuous extension of P (*black*) by mirroring, CE2 (Wang and Amini 2005). **d** Representation of the discontinuity in pressure with an impulsive pressure gradient obtained from path integration, PI

$$\nabla P(x, y) = \nabla P(-x, y) = \nabla P(x, -y) = \nabla P(-x, -y). \quad (16)$$

We denote this choice of boundary conditions with a continuous extension of the gradient field as CE1. If this continuous gradient field is integrated in space, however, the resulting pressure field is not necessarily continuous at the boundaries (Fig. 1b). Wang and Amini (2005) enforce continuity of the resulting pressure field P instead (Fig. 1c), by extending the gradient field in a discontinuous way as

$$\partial_x P(x, y) = -\partial_x P(-x, y)$$
$$= +\partial_x P(x, -y) = -\partial_x P(-x, -y) \quad (17)$$

$$\partial_y P(x, y) = +\partial_y P(-x, y)$$
$$= -\partial_y P(x, -y) = -\partial_y P(-x, -y), \quad (18)$$

where the only difference to (16) is in the signs. These boundary conditions are denoted as CE2. In numerical experiments with a synthetic multi-scale pressure field (cf. Fig. 2), we find that the CE2 approach produces smaller residuals for the reconstructed pressure field than the CE1 approach.

Boundary gradients from direct path integration (PI) A disadvantage of the periodic continuous extension of the pressure field is the obvious increase of the domain size and the corresponding increase of required memory space by a factor of 8 in three dimensions due to the mirrored fields. Another way to match the boundary of the non-periodic pressure field to the periodic FFT computations is to consistently represent the high gradient of the pressure discontinuity in the gradient field (Fig. 1d). Assume a periodically extended pressure field with a discontinuity at the boundaries. The gradient of this field must have (infinitely) high values at the boundary, which are missing in the measured gradient data. The discontinuity in the pressure field can be accounted for by introducing an impulse at the boundary of the gradient field. The magnitude of the impulse is obtained by a direct path integration from boundary to boundary in the gradient field. For example, along a one-dimensional line in x-direction in the gradient field $\partial_x P$ with n sample points, the value of the boundary impulse can be expressed as

$$\partial_x P(1) = \partial_x P(n) = -\frac{1}{2} \sum_{i=2}^{n-1} \partial_x P(i) \quad (19)$$

The two boundary values in the measured $\partial_x P$ data are then replaced by the values from (19). $\partial_y P$ and $\partial_z P$ are corrected in the same way. This modification of the boundaries is denoted as PI boundary conditions. After this correction of the boundaries, the gradient field line-wise satisfies
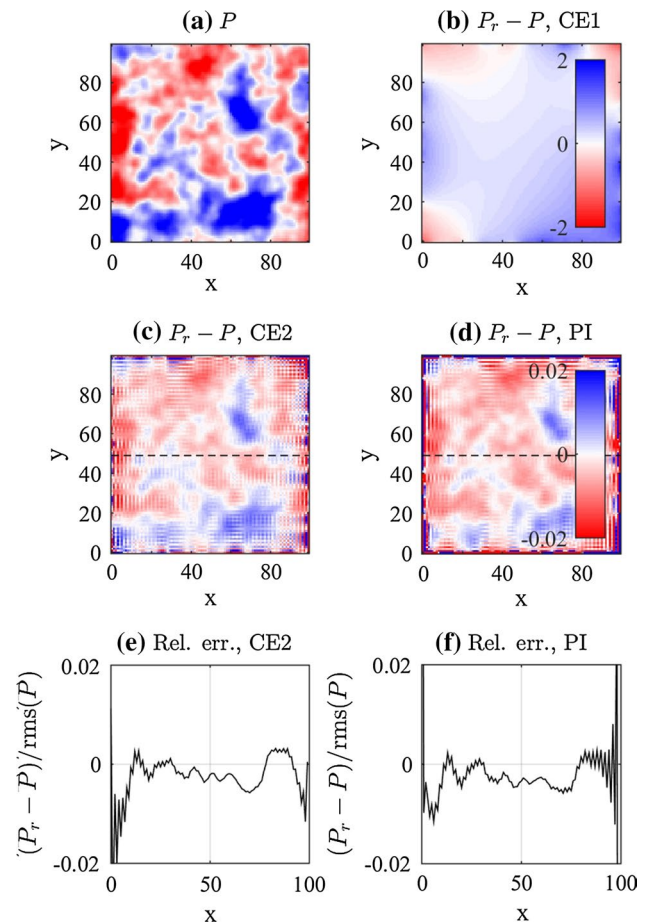
Fig. 2 a Synthetic three-dimensional pressure field (*central plane* shown) constructed as correlated noise. **b–d** Difference fields between reconstructed pressure field and ground truth data for different boundary conditions. **a**, **b** and **c**, **d** share the same *colorbar* respectively. **e**, **f** Relative error along a section (*dashed line*) for CE2 and PI boundary conditions

$$\int_0^L \partial_x P(x)\, \mathrm{d}x = 0, \quad (20)$$

also in y-direction and z-direction, which is equivalent with P being periodic with period L. The continuous extension approach CE2 also satisfies (20), while the approach CE1 does not.

2.3.2 Steps of the pressure reconstruction algorithm

Including the LPT measurement, we summarize the pressure reconstruction with the following steps:

1. Reconstruct particle trajectories from time-resolved particle images of at least three cameras using the STB algorithm (Schanz 2016).

2. Fit an interpolating continuous function consisting of 1D cubic B-splines to the trajectories and differentiate twice w.r.t. time to obtain the material acceleration at particle positions [*TrackFit*, (Gesemann 2016)].

3. Interpolate the material acceleration, given on an unstructured grid, to a fine Cartesian grid, using 3D cubic B-splines and penalizing the curl of the acceleration field [*FlowFit*, (Gesemann 2016)].

4. Neglect the viscous term in (3), assume constant density ρ and compute ∇P (5).

5. Modify ∇P by applying the CE2 or PI method for periodic spatial boundary conditions from Sect. 2.3.1, in order to avoid boundary artifacts.

6. Compute the Fourier transforms of ∇P, evaluate (9), set $\tilde{P}(\mathbf{k} = \mathbf{0}) = \mathbf{0}$, and transform back, using a fast-Fourier-transform (FFT) function.

7. Add a constant pressure offset P_0 to obtain absolute pressure.

Regardless of the exact way to obtain the pressure gradient on a Cartesian grid, be it a different measurement technique or the additional consideration of theoretical models, e.g., a compressible flow with varying density, the integration method is anyhow applicable, starting from step 5.

Once we have the pressure field, ground truth data are not always available to compare our experimental pressure results with. Yet, the random measurement noise that is uncorrelated in time can be filtered out in the temporal domain with a low-pass filter. Certainly, various sources of error are correlated in time or even constant, such that a validation against other measurements or models is in general highly desirable.

3 Synthetic pressure field

3.1 Curl-free pressure gradient

In order to validate the accuracy of the integration method, we create a three-dimensional synthetic pressure field on a cubic $257 \times 257 \times 257$ domain as ground truth data, compute the spatial gradient and reconstruct the pressure field again. A direct comparison of the original and the reconstructed pressure field reveals the accuracy of the integration. In order to generate the synthetic pressure field, uniformly distributed noise in Fourier space is correlated with a normalized correlation kernel of the form

$$f(\mathbf{k}) = \exp(-\alpha |\mathbf{k}|^{1/2}). \tag{21}$$

The slow decay of the kernel due to the exponent 1/2 generates a wide range of scales, mimicking turbulent pressure fields with large-scale gradients as well as small-scale structures related to vortices. The scale coefficient α controls

the spatial scale of the synthetic pressure field; by increasing α, the kernel decays faster and small scales are damped. The gradient of the periodic pressure field is computed in Fourier space by multiplying with $i\mathbf{k}$, then, a non-periodic $100 \times 100 \times 100$ subdomain is cropped in real space for the test. Figure 2a shows the cropped pressure field at the central z-plane. Pressure is reconstructed from the gradient fields, with and without accounting for periodic boundary conditions. With unchanged boundaries, the reconstructed pressure field significantly deviates from the original field (not shown). In the difference field, small boundary artifacts and large scale gradients up to 100 % relative error appear. When the gradient field is continuously extended according to CE1 (16), the deviation is still large. Figure 2b shows the difference between original and reconstructed pressure field. In contrast, when the boundaries are modified according to the CE2 method (17) or the PI method (19), the reconstructed pressure field agrees with the original field within an error of less than 1 %. Figures 2c, d show the vanishing difference between original and reconstructed pressure fields. Profiles along a section in x-direction show the small relative error (Fig. 2e, f). The CE2 boundary conditions have the advantage of smaller boundary artifacts than the PI method, to the cost of a 8 times larger mirrored volume that has to be integrated. On a 3.4 GHz standard PC, constructing PI boundary conditions and integrating pressure takes 0.5 s on 100^3 grid points, and 3.3 s for a 200^3 grid.

3.2 Pressure gradient with added non-zero curl component

Spatial fields of the material acceleration that are obtained from experiments typically have a non-zero curl which may be due to a significant contribution of the viscous term or due to measurement noise. In the data processing scheme presented in Sect. 2.3.2, we penalize the curl in the Flow-Fit interpolation scheme and therefore the FFT integration starts with a gradient field with vanishing curl. In general, however, it is interesting to quantify the sensitivity of the integration scheme to gradient fields with non-zero curl. Therefore, a pressure gradient field with a non-zero curl component is constructed as

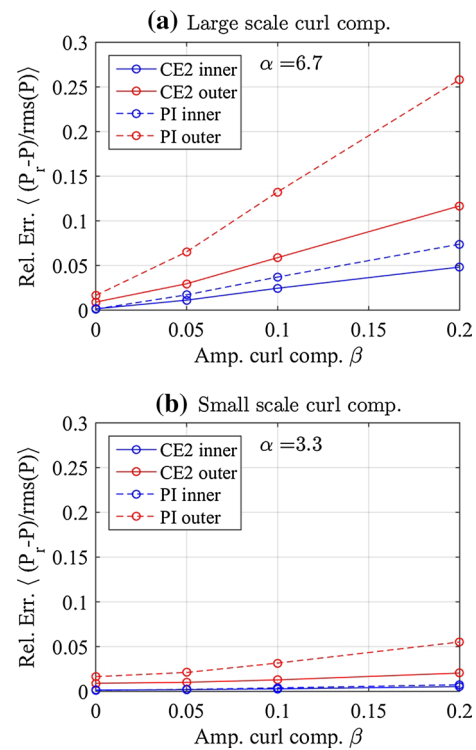$$\mathbf{a}_{nzc} = \nabla P(\alpha) + \beta'[\nabla \times \mathbf{A}(\alpha)]. \tag{22}$$

Both, the synthetic pressure field $P(\alpha)$ and the vector potential $\mathbf{A}(\alpha)$, are constructed as correlated noise as described above. The size of the spatial scales of the potential fields can be controlled via the parameter α in the correlation kernel, see Eq. (21). The amplitude of the added curl-component is given by the parameter $\beta' = \beta \, \text{std}(\nabla P)/\text{std}(\nabla \times \mathbf{A})$, such that for $\beta = 0.2$ the signal-to-signal ratio between the added curl-component $\nabla \times \mathbf{A}$ and the curl-free component ∇P is 20 %. Figure 3a
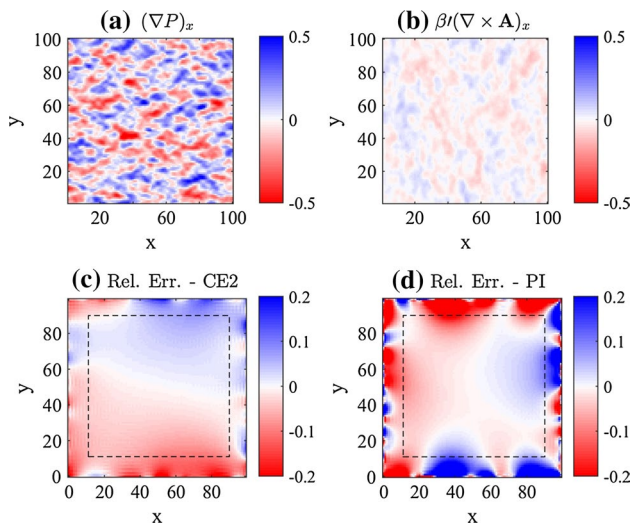
Fig. 3 Synthetic pressure gradient with non-zero curl. **a** Curl-free pressure gradient, **b** added curl-component with amplitude $\beta = 0.2$. Both fields have the same spatial spectrum with $\alpha = 6.7$. **c** Relative error $(P_r - P)/\mathrm{rms}(P)$ in the pressure reconstruction with CE2 boundary conditions and **d** PI boundary conditions. The *dashed square* splits the domain in interior and exterior domain (boundary with 10 % width)

Fig. 4 Relative error of pressure reconstruction in presence of non-zero curl. **a** Curl-free component and curl-component of the synthetic pressure gradient field have the same spatial spectrum with $\alpha = 6.7$ (cf. Fig. 3a). **b** The added curl-component has smaller scales ($\alpha = 3.3$) than the curl-free component ($\alpha = 6.7$)

illustrates the synthetic pressure gradient field and Fig. 3b shows the added curl-component with amplitude $\beta = 0.2$. Both fields have the same spatial spectrum with $\alpha = 6.7$. The pressure field is reconstructed as described in Sect. 2.3.2 using CE2 and PI boundary conditions, and the difference to the original pressure field is computed. Figure 3c, d shows the relative error for the reconstruction. In principle, the curl-component should have no effect on the reconstruction, since the pressure gradient field $\mathbf{a}_{nzc}$ is projected onto the longitudinal direction in Fourier space, such that only the curl-free component is integrated. However, the applied periodic boundary conditions are adapted for a curl-free field and therefore generate boundary artifacts when the curl is non-zero. In the interior of the volume (dashed square), the reconstruction error is within a few percent for both boundary conditions (CE2 and PI). In the boundary region, however, artifacts become large with relative errors up to 20 % for CE2 boundary conditions (Fig. 3c). For PI boundary conditions, the error exceeds 20 % (Fig. 3d).

The effect of non-zero curl on the pressure reconstruction with the proposed FFT scheme depends on the amplitude of the curl-component (controlled with parameter β) and on its spatial spectrum (controlled with parameter α). The amplitude β is varied and the relative error is quantified in the interior and exterior domain defined in Fig. 3 (dashed-line). Figure 4 shows the rms value of the relative error for $\alpha = 6.7$, i.e., the curl-component has the same spatial spectrum as the pressure gradient field and varying amplitude β. The CE2 boundary condition performs

better than the PI boundary condition, as already illustrated in Fig. 3c, d. In the interior of the domain, the reconstruction error of both methods stays below 8 % for a maximal amplitude of the curl-component of $\beta = 20\,\%$.

The error induced by imperfect boundary conditions is largely reduced, if the spectrum of the added curl-component is shifted towards smaller spatial scales, as one would expect for spatially uncorrelated measurement noise. Figure 4b shows the rms relative error for $\alpha = 3.3$. The error does not exceed 1 % in the interior, and 5 % at the boundaries for the tested values of the amplitude β. This result indicates that, in the tested range of values, the proposed FFT integration scheme is accurate within few percent for pressure gradient fields with a curl-component of moderate amplitude and with small dominant spatial scales.

3.3 Pressure gradient with internal boundaries

In a third case, we test the integration scheme on a synthetic pressure field around a body. This adds additional boundaries in the interior of the domain. For the determination of loads on a mechanical structure in the flow, as a wing, the flow field is typically measured in a volume around the structure, in order to reconstruct the pressure

field and integrate pressure over the surface of the structure (Tronchin et al. 2015; Ragni et al. 2012). Again, we generate a synthetic periodic pressure field, compute the gradient in Fourier space, and crop a non-periodic subset. In the center of the gradient field, an empty spherical region is defined, simulating a sphere in the pressure field which adds internal boundaries (Fig. 5a). Integration in Fourier space, only working on rectangular domains, may seem unapt to integrate the pressure field bounded by curved surfaces. For example, a pressure reconstruction that is directly computed on a surface following unstructured grid defined by particle positions has recently been proposed by Neeteson and Rival (2015). However, in our pressure gradient field, the empty region can be replenished with an interpolation by the FlowFit routine, such that the FFT integration can still be applied. Using a random subset of 5 % of the grid points as input data for the FlowFit (Fig. 5b), we interpolate the empty space with a smooth function (Fig. 5c). In the optimization, the curl of the field is penalized in the cost function, such that the result is nearly curl-free. In this way, no explicit boundary conditions have to be given at the curved interior boundary. Strictly speaking, the pressure gradient vanishes on the wall, since velocity and acceleration are identical zero at the wall. Here, however, we assume that the measurement point closest to the wall still is far from the wall in the microscopic sense and therefore experiences a non-zero pressure gradient.

Using PI boundary conditions at the exterior boundaries, the reconstructed pressure field closely matches the original pressure field. In Fig. 6, the pressure field is interpolated onto the surface of the sphere, and by visual inspection no difference can be seen between the two fields. The relative error is in the order of few percent. Figure 6c shows the relative error along an arbitrary circle on the sphere. Note that, strictly speaking, the pressure values on the surface of the sphere are extrapolations of the pressure field, since no gradient data are available exactly on the surface. When the load on the sphere is computed by integrating pressure over the entire surface of the sphere, the values between original and reconstructed pressure field differ by 2.1 % for the field shown.

4 Experimental pressure field

Our integration method for the pressure gradient field is now applied to an experimental data set. In the experiment, a transitional water jet emanates in vertical upward direction from a circular nozzle with diameter $d = 10$ mm at $Re = 5000$. Close to the nozzle, vortex rings form, grow, destabilize and decay further downstream. The LPT images are taken with three cameras at a frequency of 1 kHz. About 13,000 particles are tracked instantaneously by the *Shake-The-Box* method in a volume of roughly 28 cm^3, corresponding to a mean particle distance of 1.34 mm. The experiment is described in more detail in Schanz et al. (2013) and references therein.

For the pressure reconstruction, we follow the steps outlined in Sect. 2.3.2. The output of the *Shake-The-Box* LPT algorithm provides time-resolved acceleration fields at

Fig. 5 **a** Gradient of synthetic pressure field $\partial_x P$ with internal boundaries around a sphere, central *xy*-plane. **b** Random 5 % subset of the data points shown in a slice of the volume. **c** Curl-free interpolation in the spherical region and in an external buffer zone with FlowFit, central *xy*-plane
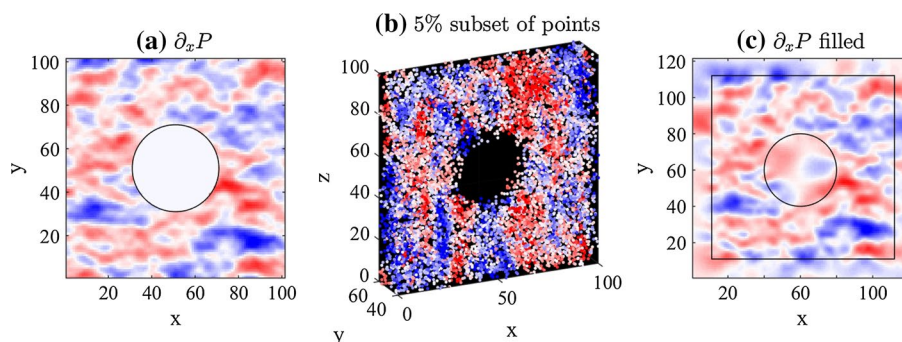
Fig. 6 **a** Original and **b** reconstructed pressure field on the spherical body. **c** Relative error $(P_r - P)/\sqrt{\langle P^2 \rangle}$ along a circle on the sphere (*dashed line*)
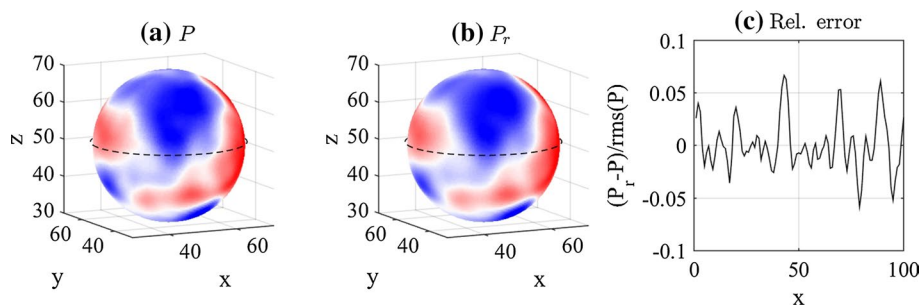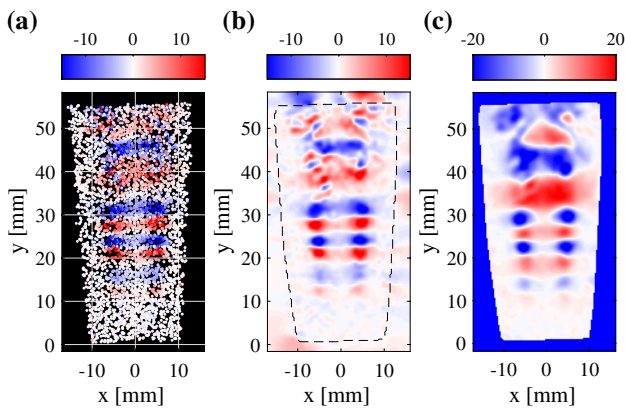
Fig. 7 **a** Particles color-coded with vertical acceleration a_y (m/s^2) in a slice $z \in [-3, 3]$ mm. **b** Smooth acceleration field a_y on a Cartesian grid after interpolation with FlowFit with surrounding buffer zone, $z = 0$ mm. **c** Resulting pressure field P (Pa) at $z = 0$ mm

instantaneous particle positions (Step 2). Figure 7a shows the distribution of tracer particles in a vertical section of the field of view for a fixed time frame. The particles are color-coded with the vertical component of acceleration a_y with a maximal magnitude of ∼20 m/s^2. Alternating regions of positive and negative acceleration can be seen that are associated with a chain of vortex rings. In step 3, the acceleration field is interpolated onto a Cartesian grid with the FlowFit routine (Fig. 7b). We choose a grid of size $111 \times 201 \times 111$ with a spatial resolution of $\Delta x = 0.3$ mm, which is roughly 1/4 of the mean inter-particle distance in the experiment, such that all details of the acceleration field are sampled. The domain of the Cartesian grid is chosen in a way that it entirely encloses the particle containing measurement volume (enclosed by dashed line). The interpolation function in the interior extends smoothly to the exterior buffer zone and it is curl-free in the whole domain. Small contributions of the viscous term to acceleration can be suppressed in this way, since the viscous term is divergence-free (curl-only) (14).

In step 4, we set the density to $\rho = 1000$ kg/m^3 and compute the pressure gradient ∇P. In step 5, the boundaries are modified according to the PI method, in order to account for periodic boundary conditions of the pressure field, cf. Eq. (19). In step 6, the relative pressure field is obtained after an integration in Fourier space. A section of the pressure field at the center plane of the jet is shown in Fig. 7c with the exterior region in dark blue. Each vortex ring in the jet can be seen in the pressure field as a pair of circular low-pressure regions. In upward (downstream) direction, the low pressure region inside consecutive vortex rings intensifies, indicating a growing rotation, since the pressure gradient around the pressure minimum mainly balances centrifugal forces in the rotating vortex. At the upper part of the domain, the vortex rings break up into less
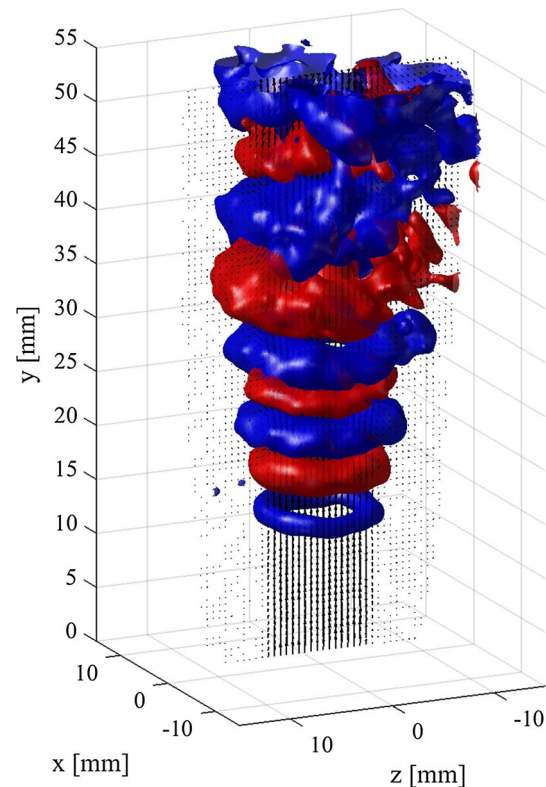

Fig. 8 3D pressure field (iso-surfaces) and central slice of the velocity field (vectors) of a transitional jet. Iso-surfaces are shown for $P = -5$ Pa (*blue*) and $P = +5$ Pa (*red*)

regular structures with an increasingly disordered pattern in the pressure field. This transition to turbulence can even better be seen in a three-dimensional representation of the pressure field. Figure 8 shows iso-surfaces of positive (red) and negative (blue) pressure regions that reveal the ring structure of the vortices and the increasing diameter of the growing vortex rings.

So far, relative pressure fields have been presented. These pressure fields lack the addition of a spatially constant pressure offset, the integration constant of the spatial integration. In the present experiment, no data of pressure boundary conditions are available, and the measurement volume barely reaches into regions of stationary flow where a pressure reference could be obtained from Bernoulli's equation (van Oudheusden 2013). Consequently, we cannot experimentally determine the pressure offset.

Nevertheless, each individual snapshot of the relative pressure fields is normalized by its zero spatial mean by construction, $\langle P \rangle_{x,y,z} = 0$, and we can address the question, whether such relative pressure fields are already smooth in time, or if a time-dependent pressure offset has to be added to reach the expected smooth temporal evolution of the pressure field. Figure 8 represents the temporal evolution of 50 consecutive relative pressure fields with
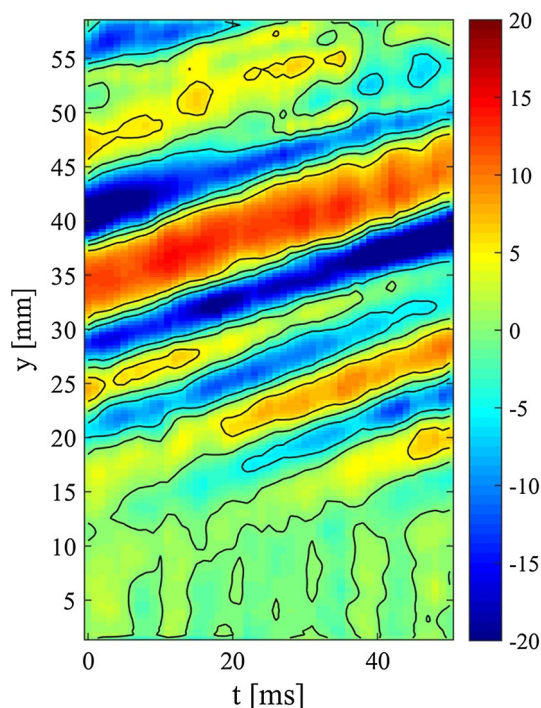
Fig. 9 Temporal evolution of pressure field $P(y, t)$ (Pa). Space–time plot of pressure along a *vertical line* intersecting the vortex rings

$\Delta t = 1$ ms in a space–time plot. The pressure profiles along a vertical line (y-direction) intersecting the vortex rings are plotted against time with contour levels (black) drawn at $[-5, 0, 5]$ Pa. Alternating low- and high-pressure regions, corresponding to the vortex ring pattern in Fig. 8, can be seen, and the slope displays the propagation velocity of the vortex rings. This diagram indicates an overall smooth evolution of the relative pressure fields, since the contour lines of the large-scale vortex rings evolve in time with small oscillations. For a strong time-dependent pressure offset, all contour lines would oscillate in phase. This points to a largely time-independent pressure offset (Fig. 9).

5 Summary

In this paper, we propose a new method for the spatial integration of the pressure gradient field in turbulent flows obtained from experimental Lagrangian particle tracking at high seeding density (*Shake-The-Box*). The pressure gradient field is integrated in Fourier space, simply including a Fourier transform of the pressure gradient, a multiplication and a back transform, resulting in a relative pressure field with zero mean. A single integration constant, the pressure offset, has to be added to the relative pressure fields in a final step. Before integration, the boundaries of raw data have to be modified to adjust to the periodic boundary

conditions inherent to the Fourier transform. We present two choices of periodic boundary conditions that result in numerical errors of less than 1 % in the test cases with a curl-free pressure gradient.

Prior to integration, the FlowFit interpolation scheme yields a representation of the scattered experimental acceleration data on a Cartesian grid. Due to a penalization, the FlowFit interpolation results in material acceleration fields with vanishing curl. In the case of significant curl in the acceleration field, the FFT integration scheme also works, since the field is projected onto the longitudinal (divergent) component in Fourier space, and only this component is spatially integrated. Yet, using the periodic boundary conditions presented here for non-curl-free acceleration fields leads to artifacts at the boundaries. These artifacts are reduced, if the curl-component has small spatial scales compared to the curl-free component, as it is typical for spatially uncorrelated measurement noise. Depending on the amplitude of the curl-component, the integrated field is accurate within few percent in the interior of the domain.

We also show that using the FlowFit interpolation scheme facilitates the treatment of internal boundary conditions, especially for curved interior boundaries, as e.g., the wall of a body inside the domain. The empty domain in the body can be filled with a smooth curl-free function, such that the integration, not limited to the FFT scheme, can be applied on a cubic domain ignoring internal boundaries. Similarly, external boundaries can be embedded inside the cubic domain by extending the domain with a buffer zone. This provides pressure data at the exact position of walls inside the measurement volume.

References

Charonko JJ, King CV, Smith BL, Vlachos PP (2010) Assessment of pressure field calculations from particle image velocimetry measurements. Meas Sci Technol 21:105401

Dabiri J, Bose S, Gemell BJ, Colin SP, Costello JH (2014) An algorithm to estimate unsteady and quasi-steady pressure fields from velocity field measurements. J Exp Biol 217:331–336

de Kat R, van Oudheusden BW (2012) Instantaneous planar pressure determination from PIV in turbulent flow. Exp Fluids 52:1089–1106

Frankot RT, Chellappa R (1988) A method for enforcing integrability in shape from shading algorithms. IEEE Trans Pattern Anal Mach Intell 10:4

Gesemann S (2016) From particle tracks to velocity and acceleration fields using B-splines and penalties. arXiv:1510.09034v1

Ghaemi S, Ragni D, Scarano F (2012) PIV-based pressure fluctuations in the turbulent boundary layer. Exp Fluids 53:1823–1840

Griffiths D (1999) Introduction to electrodynamics. Prentice Hall, Upper Saddle River

Hoyer K, Holzner M, Lüthi B, Guala M, Liberzon A, Kinzelbach W (2005) 3D scanning particle tracking velocimetry. Exp Fluids 39:923–934

Laizet S, Lamballais E (2009) High-order compact schemes for incompressible flows: a simple and efficient method with quasi-spectral accuracy. J Comput Phys 228(16):5989–6015

Liu X, Katz J (2006) Instantaneous pressure and material acceleration measurements using a four-exposure PIV system. Exp Fluids 41:227–240

Neeteson NJ, Rival DE (2015) Pressure-field extraction on unstructured flow data using a Voronoi-tesselation-based networking algorithm: a proof-of-principle study. Exp Fluids 56:44

Negahdar MJ, Kadbi M, Cha J, Cebral J, Amini A (2013) Noninvasive 3D pressure calculation from PC-MRI via non-iterative harmonics-based orthogonal projection: constant flow experiment. 35th annual international conference IEEE EMBS, Osaka, Japan

Novara M, Scarano F (2013) A particle-tracking approach for accurate material derivative measurements with tomographic PIV. Exp Fluids 54:1584

Pröbsting S, Scarano F, Bernardini M, Pirozolli S (2013) On the estimation of wall pressure coherence using time-resolved tomographic PIV. Exp Fluids 54:1567

Ragni D, van Oudheusden BW, Scarano F (2012) 3D pressure imaging of an aircraft propeller blade-tip flow by phase-locked stereoscopic PIV. Exp Fluids 52:463–477

Rocholz R (2008) Spatio-temporal measurement of short wind-driven water waves. Dissertation, Ruperto-Carola University of Heidelberg, Germany

Schanz D, Gesemann S, Schröder A (2016) Shake-the-box: accurate Lagrangian particle tracking at high particle densities. Exp Fluids 57:1–27

Schanz D, Gesemann S, Schröder A, Wieneke B, Novara M (2013b) Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. Meas Sci Technol 24:024009

Schanz D, Schröder A, Gesemann S, Michaelis D, Wieneke B (2013a) Shake the box: a highly efficient and accurate tomographic particle tracking velocimetry (TOMO-PTV) method using prediction of particle positions. 10th international symposium on particle image velocimetry, Delft, Netherlands

Schanz D, Schröder A, Gesemann S (2014) Shake the box—a 4D PTV algorithm: accurate and ghostless reconstruction of Lagrangian tracks in densely seeded flows. 17th international symposium on applications of laser techniques to fluid mechanics, Lisbon

Schröder A, Geisler R, Staack K, Elsinga GE, Scarano F, Wieneke B, Henning A, Poelma C, Westerweel J (2011) Eulerian and Lagrangian views of a turbulent boundary layer flow using time-resolved tomographic PIV. Exp Fluids 50:1071–1091

Tronchin T, David L, Farcy A (2015) Loads and pressure evaluation of the flow around a flapping wing from instantaneous 3D velocity measurements. Exp Fluids 56:7

van Oudheusden BW (2013) PIV-based pressure measurement. Meas Sci Technol 24:032001

van Oudheusden BW, Scarano F, Roosenboom EWM, Casimmiri EWF, Souverein LJ (2007) Evaluation of integral forces and pressure fields from planar velocimetry data for incompressible and compressible flows. Exp Fluids 43:153–162

Wang Y, Amini A (2005) Integrable pressure gradients via harmonics-based orthogonal projection. In: Christensen GE, Sonka M (eds) Information processing in medical imaging 2005. LNCS 3565. Springer, Berlin, pp 431–442

Wieneke B (2013) Iterative reconstruction of volumetric particle distribution. Meas Sci Technol 24:024008

Xu H, Oullette N, Vincenzi D, Bodenschatz E (2007) Acceleration correlations and pressure structure functions in high Reynolds number turbulence. Phys Rev Lett 99:204501

Zhang X (1996) An algorithm for calculating water surface elevations from surface gradient data. Exp Fluids 21:43–48

CrossMark

RESEARCH ARTICLE

Large-scale volumetric flow measurement in a pure thermal plume by dense tracking of helium-filled soap bubbles

Florian Huhn[1] · **Daniel Schanz**[1] · **Sebastian Gesemann**[1] · **Uwe Dierksheide**[2] ·
Remco van de Meerendonk[3] · **Andreas Schröder**[1]

Abstract We present a spatially and temporally highly resolved flow measurement covering a large volume (~0.6 m^3) in a pure thermal plume in air. The thermal plume develops above an extended heat source and is characterized by moderate velocities (U ~ 0.35 m/s) with a Reynolds number of Re ~ 500 and a Rayleigh number of Ra ~ 10^6. We demonstrate the requirements and capabilities of the measurement equipment and the particle tracking approach to be able to probe measurement volumes up to and beyond one cubic meter. The use of large tracer particles (300 μm), helium-filled soap bubbles (HFSBs), is crucial and yields high particle image quality over large-volume depths when illuminated with arrays of pulsed high-power LEDs. The experimental limitations of the HFSBs—their limited lifetime and their intensity loss over time—are quantified. The HFSBs' uniform particle images allows an accurate reconstruction of the flow using Shake-The-Box particle tracking with high particle concentrations up to 0.1 particles per pixel. This enables tracking of up to 275,000 HFSBs simultaneously. After interpolating the scattered data onto a regular grid with a Navier–Stokes regularization, the velocity

field of the thermal plume reveals a multitude of vortices with a smooth temporal evolution and a remarkable coherence in time (see animation, supplementary data). Acceleration fields are also derived from interpolated particle tracks and complement the flow measurement. Additionally, the flow map, the basis of a large class of Lagrangian coherent structures, is computed directly from observed particle tracks. We show entrainment regions and coherent vortices of the thermal plume in the flow map and compute fields of the finite-time Lyapunov exponent.

1 Introduction

Time-resolved volumetric flow measurements, using methods such as tomo-PIV (Elsinga et al. 2006), 3D PTV (Maas et al. 1993; Malik et al. 1993) or Shake The Box (Schanz et al. 2013b, 2016a), are typically restricted to relatively small volume sizes of the order of ≤200 cm^3 (Scarano et al. 2015). This limitation stems from the small size of commonly used seeding material, to accurately follow the flow (diameter range around ~1 μm in air, 10–50 μm in water). Currently available high-repetition rate laser systems, which are typically used as a light source, do not provide enough intensity to allow for illumination of larger volumes, even for the larger particles used in water experiments. However, seeding particles whose density approaches that of the medium can be of larger size, while still accurately being able to follow the flow (Melling 1997). In line with this thought, neutrally buoyant helium-filled soap bubbles (HFSBs) have been used in air to allow for large-scale flow measurements in the laboratory. Applications range from traceline visualizations (Pounder 1956) over large-scale 2D-PIV measurements (Müller et al. 2000; Bosbach et al. 2009) to three-dimensional tracking of single bubbles

✉ Florian Huhn
  florian.huhn@dlr.de

1 Institute of Aerodynamics and Flow Technology, Department of Experimental Methods, German Aerospace Center (DLR), Bunsenstr. 10, Göttingen, Germany

2 LaVision GmbH, Anna-Vandenhoeck-Ring 19, Göttingen, Germany

3 Department of Aerospace Engineering, Delft University of Technology, Kluyverweg 1, Delft, The Netherlands

⦾ Springer

(Klimas 1973) and large-scale tomographic PIV of a convective flow (Kühn et al. 2011). Recently, a feasibility study demonstrated that HFSBs are a promising seeding material for low-speed wind tunnels (Scarano et al. 2015).

At the uppermost end of the range of large measurement volumes, open air experiments in the atmospheric boundary layer reach spatial scales of tens of meters, using, for example, snow particles as seeding material illuminated with a searchlight (Hong et al. 2014) or centimeter-sized fog-filled soap bubbles illuminated by the sun (Rosi et al. 2014). These studies show promise as techniques for real-size flow measurements behind wind turbines, but are of minor applicability in a controlled laboratory environment. Some of the previous particle tracking experiments with HFSBs in the laboratory examined large measurement areas (Klimas 1973; Biwole et al. 2009), but were limited in particle number, often tracking only a few tens or hundreds of bubbles. The largest investigated volume that allowed the description of instantaneous flow structures was applied by Kühn et al. (2011). A convection cell with a volume of approximately 56 L was investigated with two-pulse tomo-PIV. However, a large interrogation window size had to be chosen (48 × 48 × 24 mm), limiting the spatial resolution to large structures. In the first application of HFBSs for a wind tunnel experiment, Scarano et al. (2015) have shown that for higher flow speeds (~30 m/s) the production of enough bubbles to achieve a sufficient particle concentration within the measurement volume is a major issue. Due to the limitations in bubble number and due to the limits of the high-speed laser used for illumination, the volume size was restricted to 4.8 L in this experiment and the interrogation windows for a tomo-PIV analysis were large (96 × 96 × 86 voxels). In a follow-up paper, the same group reaches a measurement volume of 12 L at a free stream velocity of 8 m/s. To increase the seeding density to a mean particle distance of ~10 mm, they present a piston-based seeding generator with an accumulation and release strategy (Caridi et al. 2016). At a free stream velocity of 5 m/s, Schneiders et al. (2016) present a 3D tomo-PTV measurement downstream of a surface-mounted low-aspect-ratio cylinder in an increased measurement volume of 6 L.

In this paper, we present an LPT experiment with helium-filled soap bubbles in a thermal plume with an unprecedented measurement volume of 560 L and a high spatial resolution (mean inter-particle distance d = 13.0 mm). The application of the Shake-The-Box (STB) algorithm allows for an accurate and time-effective reconstruction of particle paths at high seeding concentrations, providing dense velocity and acceleration fields of the flow. The experiment has been designed to apply STB LPT to a turbulent low-speed flow with a high seeding concentration of helium-filled soap bubbles. Since the seeding generator was limited to six seeding nozzles

at the time of the experiment, a closed flow chamber was chosen. The moderate thermal forcing of the flow allows for a precise control of the small flow velocity and the involved turbulent structures; this is in contrast, for example, to a fan that already induces complex turbulent structures on the blades. Hence, for the designed convection experiment with moderate velocities (0.35 m/s), the bubble production rate and the recording frequency are not limiting factors.

The paper is organized as follows. In Sect. 2, we present the experimental setup and in Sect. 3 we describe the data analysis. The results for the flow field are shown in Sect. 4. In Sect. 5, we present the results of a Lagrangian transport analysis in the thermal plume. Finally, Sect. 6 summarizes the results and discusses future perspectives of the measurement technique.

2 Experimental setup

2.1 Convection chamber

The experiments are performed in a cylindrical convection chamber with a height of 2.00 m and a diameter of 1.83 m (Fig. 1). The top and bottom plates are constructed of wood, the back wall is made of aluminum, and the transparent front window is acrylic glass of 1 mm thickness. Non-transparent parts of the floor, walls and ceiling are painted black or covered with black self-adhesive film, to avoid scattered light and improve the contrast of the particle images relative to the background. The convection chamber is accessible from the rear side through a door in the aluminum wall. The chamber is equipped with a circular perforated tube at the bottom to rinse it with pressurized air and remove seeding.

LED illumination enters through an acrylic glass window of 1 m diameter in the ceiling, covered with a circular passe-partout (0.75 m diameter) that determines the width of the cylindrical measurement volume. The convective flow is forced by a standard 1500 W electric hot plate (D = 188 mm diameter, *Silva Homeline EKS 2121*) that is placed a few centimeters below the measurement volume. It is covered by a black circular 250 mm diameter aluminum plate of 10 mm thickness that serves as a heat reservoir to keep the temperature constant over time. The circular hot plate is mounted on a support of ~0.4 m height above the floor and with a diameter of roughly 350 mm (Fig. 1b). The elevated position of the heat source allows for a vertical flow along the contours of the support toward the hot plate, feeding the thermal plume over the hot plate. As mentioned by Pham et al. (2005), the elevated position stabilizes the horizontal position and diminishes strong lateral oscillations of the thermal plume.
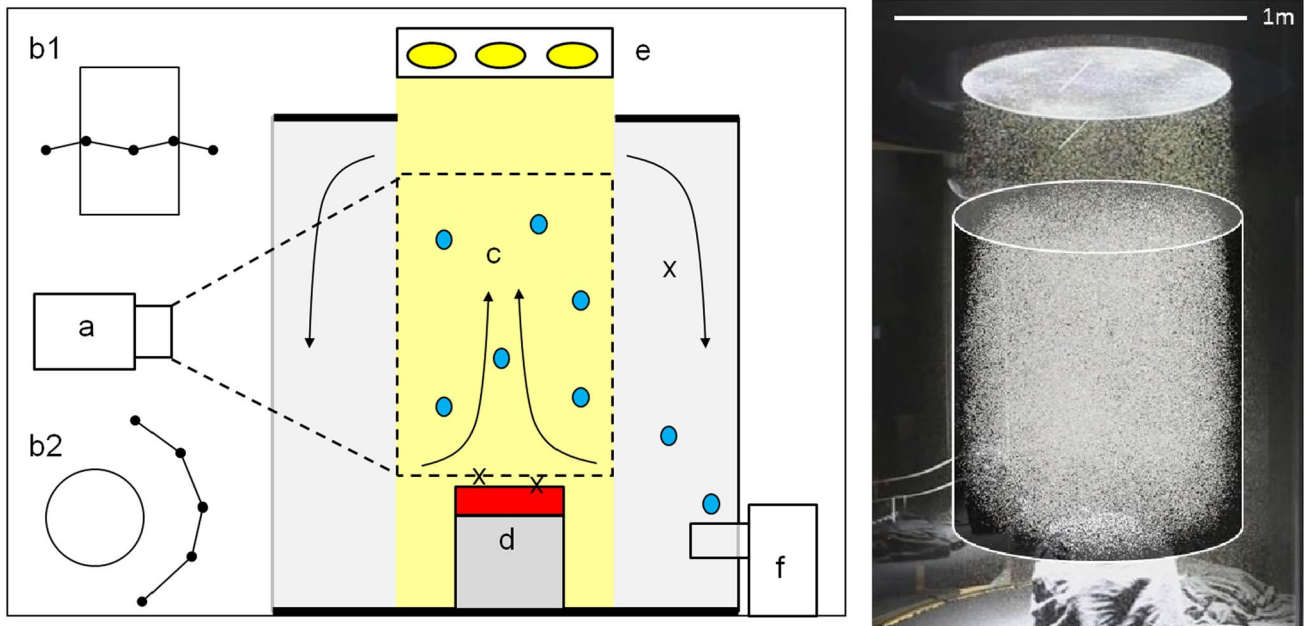
Fig. 1 (*Left*) experimental setup of convection chamber: (*a*) cameras; (*b1*, *b2*) front and top view of camera configuration; (*c*) field of view (FOV), (*d*) hot plate; (*e*) LED array; (*f*) bubble generator; (*x*) positions of three temperature sensors. (*Right*) photograph of the interior of the convection chamber with elevated heat source in the center and illuminated HFSB. In the cylindrical measurement volume, the reconstructed 3D particle cloud is shown as an overlay

Before conducting experiments, the hot plate is heated up for a few seconds to reach the desired temperature. After waiting for some minutes, a uniform temperature distribution over the hot plate is attained. Three temperature sensors (thermocouples *Omega 5SC-TT-KI-40-1M*, AD converter *Labjack T7 Pro*) are installed in the convection chamber to monitor the temperature of the heated aluminum plate, the air temperature 2 mm above the plate as a proxy for the maximum air temperature, and the ambient air temperature outside of the plume (see Fig. 1 for positions of sensors). We also measure the temperature in the laboratory and observe that during the experiment no measurable increase of the temperature inside the convection chamber compared to the laboratory can be observed (cf. Fig. 2). The accuracy of the absolute temperature measurement is estimated to be ±2 °C, while the relative temperature differences, relevant for the convective flow, are accurate within ±0.5 °C. During the experiments, heat is provided by the aluminum plate, while the temperature of the hot plate decreased by less than 0.5 K (cf. Fig. 2).

2.2 Helium-filled soap bubbles

For LPT, the flow is densely seeded with neutrally buoyant HFSBs of 300 μm diameter. They are produced by a bubble generator prototype of *LaVision GmbH*, based on the orifice-type nozzle design discussed by Okuno et al. (1993) and used by Bosbach et al. (2009). A nozzle consists of
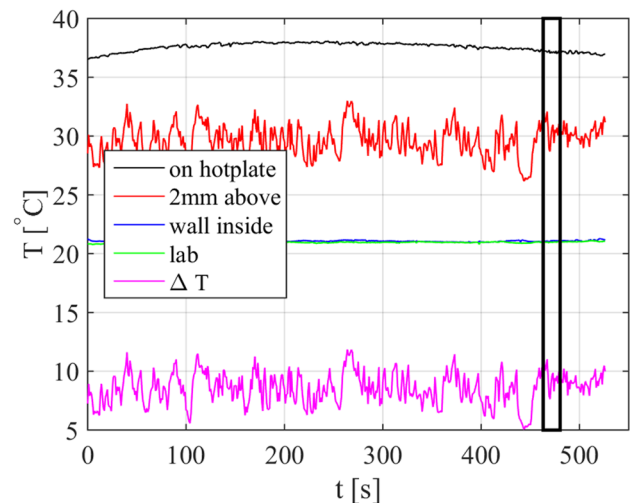

Fig. 2 Temperature log for the different temperature sensors. Measurement time window is marked by the *black rectangle*. ΔT is the temperature difference between air 2 mm over the hot plate and ambient air

three concentric channels—providing helium, soap solution (*ASAI 1035, Sage Action Inc.*), and pressurized air, from the center outward. It is covered by a cap with a small circular orifice. The two inner channels produce a thin helium-filled soap tube that is transported through the orifice by the surrounding air flow, subsequently breaking up into a single chain of equal-sized bubbles in the increasingly turbulent

air flow (see Fig. 3). Six nozzles are operated in parallel with a bubble production rate of ~45,000/s each. This production rate is estimated from high-speed camera images at the outlet of the nozzle. The nozzles are directly placed at the bottom inside the convection chamber to enable a high seeding density in the large volume. HFSBs are injected vertically close to the wall to avoid the influence of the momentum of the nozzle jets toward the center where the hot plate is located. Before an experiment, the chamber is seeded for 30 s to reach a high particle concentration. A waiting period of around 135 s follows, to reach a homogeneous spatial distribution of the seeding and to let the motion induced by the nozzle jets decay. We adjust the neutral buoyancy of the HFSBs by varying the flow rate of helium, such that, by visual inspection, a zero settling velocity is attained in calm air. Careful experiments by Scarano et al. (2015) show that HFSBs follow the air flow at high accelerations of ~10^4 m/s^2 in a wind tunnel even for a variation of the helium flow rate by a factor of two. Since in our convective flow, accelerations are much smaller ($a \sim 0.1$ m/s^2 see results of LPT), we expect the tracer to be able to closely follow the air flow.

Soap bubbles burst, such that the lifetime can be a limiting factor, especially for large convective flows that are typically slow and involve long time scales. Bosbach et al. (2009) estimate the lifetime of HFSBs to be 1–2 min under similar conditions as in this experiment, e.g., room temperature and presumably relatively low relative air humidity in a laboratory. Here, we further quantify the full lifetime distribution of HFSBs at these conditions, $T = 21$ °C and relative humidity of rh = 37%, considering a simple population model approach. The convection chamber is considered to be a well-mixed box with uniform seeding density, experimentally brought about by strong mixing with a continuously running fan. The number of HFSBs in the convection chamber is assumed to evolve as
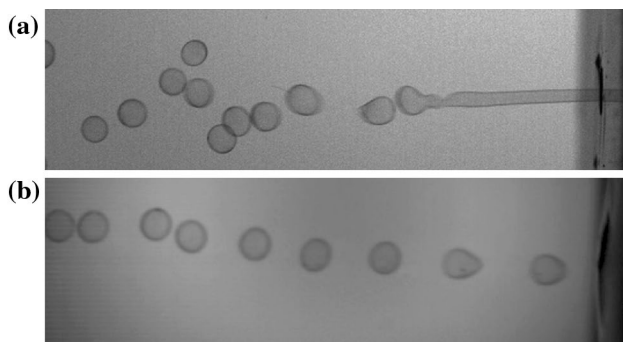
$$\frac{\mathrm{d}N}{\mathrm{d}t} = p(t) - \mathrm{d}(t), \tag{1}$$

with the production rate $p(t)$ and the death rate $d(t)$. The production rate $p(t)$ is a step function, controlled by switching on and off a single nozzle of the bubble generator. The death rate is given by
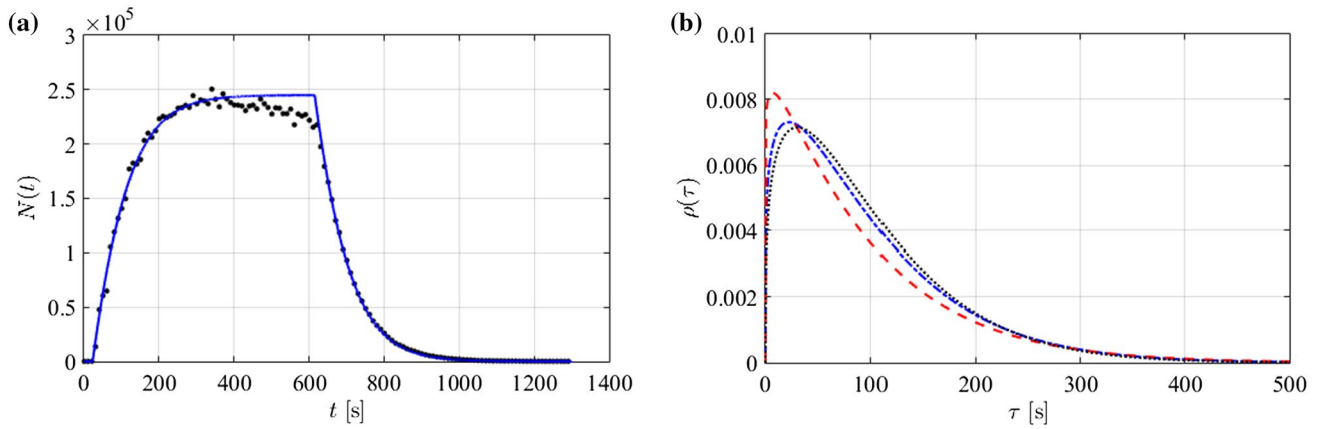
$$d(t) = \int_{-\infty}^{t} p(t-\tau)\rho(\tau)\mathrm{d}\tau, \tag{2}$$

with the probability density function of the lifetime $\rho(\tau)$, i.e., the death rate is given by integrating over all production rates in the past, weighted with the lifetime distribution. During the experiment in which the HFSB's lifetime is estimated, the particles are constantly produced with a single nozzle for 10 min, after which the production is switched off and images are acquired at a frequency of 1 Hz for a further 10 min. The number of HFSBs in the measurement volume $N(t)$ is determined by a particle detection algorithm applied to the images of the center camera. It turns out that $N(t)$ is well fit by the model when assuming $\rho(\tau)$ to be a Weibull distribution,

$$\rho(\tau) = \frac{k}{\lambda}\left(\frac{\tau}{\lambda}\right)^{-1}\exp\left[-\left(\frac{\tau}{\lambda}\right)^k\right], \tag{3}$$

with the scale parameter λ and the shape parameter k. The lifetime of soap films has been reported to follow a Weibull distribution with increasing decay rate due to an aging process (Gilet et al. 2007; Tobin et al. 2011).

Figure 4a shows the evolution of the particle number in the convection chamber and the fit function $N(t)$ obtained by integrating Eq. (1). The number of detected particles has a small dependence on the chosen intensity threshold; the curve in Fig. 4a is for a threshold of 150 gray value counts of the *pco.edge* sCMOS camera. It is expected that after an initial increase of the particle number, a plateau is reached, when production rate and death rate are equal with opposite signs such that the net change in particle number is zero. In Fig. 4a, a slight decrease of the plateau can be seen; this may be due to a non-constant production rate of the nozzle or inhomogeneous bubble quality, leading to earlier bursting of the initial bubbles. For the fit of the model to the data, data from $t \in [400, 620]$ s is excluded, leading to a better fit in the initial rise and the final decay of the curve. The blue curve in Fig. 4b shows the lifetime distribution $\rho(\tau)$ for the best fit. Curves for other intensity thresholds in the particle detection, 100 and 200 counts, are also shown. According to this Weibull distribution, the mean lifetime of HFSBs is $\tau = 97$ s which confirms the rough estimate of Bosbach et al. (2009). Beyond a simple mean lifetime value, the complete lifetime distribution shows that the lifetimes of HFSBs are widespread with a long tail, explaining

Fig. 3 High-speed images of the nozzle of the bubble generator. **a** Unstable working conditions with visible external breakup of the soap tube. **b** Stable working conditions where the soap tube breaks up internally and equally sized soap bubbles are produced

(a)

(b)

Fig. 4 **a** Temporal evolution of particle number $N(t)$ (*dots*) and optimized fit of the model Eqs. (1–3) (*blue line*). **b** Weibull distribution HFSB lifetime for intensity thresholds of 100 (*red, dashed*), 150 (*blue, dashed dotted*), and 200 counts (*black, dotted*). Parameters of *blue curve*: Shape parameter $k = 1.19$, scale parameter $\lambda = 103$ s, mean lifetime $\tau = 97$ s, standard deviation $\sigma_\tau = 82$ s

the observation that even after times as long as 5 min some HFSBs still remain.

Due to their large size, the intensity of the HFSB's particle images is high in the experiment (Fig. 5). Beyond their limited lifetime due to bursting, however, one also has to consider a significant intensity loss over time. The mean intensity decays to half the initial value in 10 min. Figure 5a illustrates this intensity loss as a mean over the whole ensemble of particles. In an experiment with an initial impulsive seeding for only 2 s and low resulting seeding density, the intensity of individual particle images was determined and averaged over the measurement volume. The number of particles (blue curve) also decays,

diminishing the size of the particle ensemble with practically vanishing particle number after 700 s. Figure 5b shows the full intensity distributions of HFSBs that narrow and shift to lower intensity with increasing time.

Shrinking of the HFSBs is a possible aging process responsible for the intensity loss, as helium diffuses out of the bubble. Another process that we observed qualitatively is the shift of the reflected spectrum toward smaller wavelengths. To the eye, old HFSBs appear bluish in contrast to the reddish fresh HFSBs. A thinning of the bubbles skin with time could be responsible for this spectral effect. The detected intensity would consequently be lower, since the camera's quantum efficiency curve drops quickly

(a)

(b)
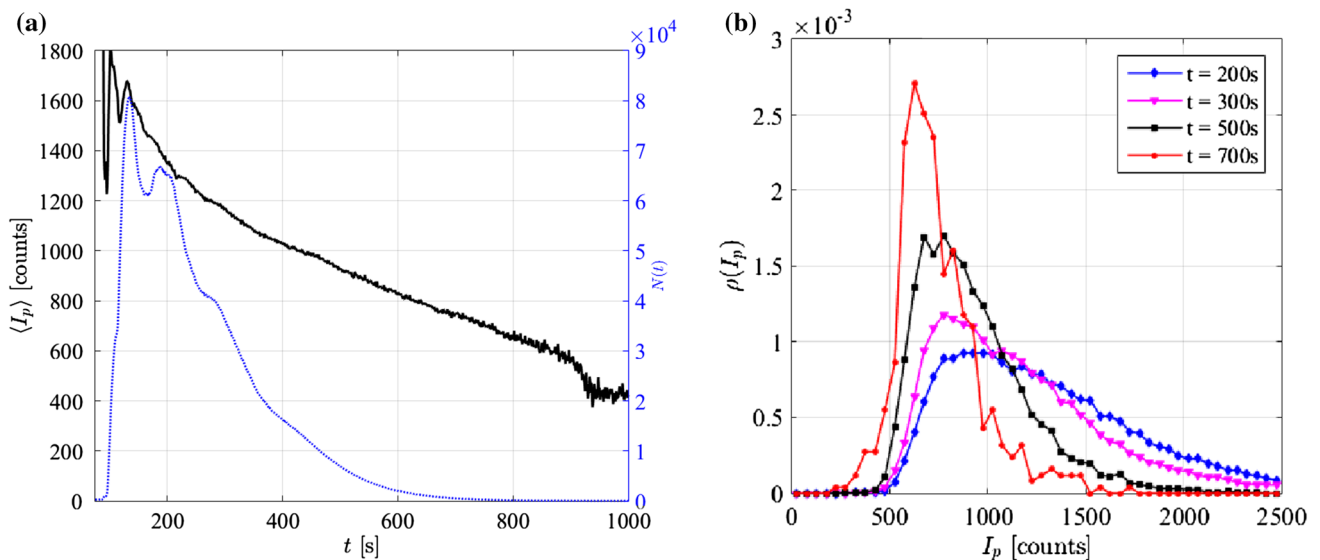
Fig. 5 Intensity loss of aging HFSB. **a** Decay of mean intensity of HFSBs (*black, solid*) and number of particles (*blue, dotted*) in an experiment with impulsive initial seeding for 2 s. **b** Probability distributions of particle intensity evolving with time

toward smaller wavelengths. Apart from aging processes, the quality of the soap solution influences the properties of the HFSBs. During the measurement campaign, the soap solution seems to change over time leading to less intense HFSBs. A thorough quantitative investigation of the physical reasons for the mentioned effects would require a specially designed experiment and is beyond the scope of this work.

In conclusion, given the time scales in the lifetime distribution and in the intensity loss, HFSBs can be recommended for measurements not lasting longer than ~2 min after bubble injection. For longer times, a continuous seeding is necessary to maintain the seeding density and the image quality. Before measurements, old and possibly recycled soap solution should be replaced by fresh solution.

2.3 LED light source

The measurement volume is illuminated by a white LED array consisting of seven standard collimated LED spotlights (*Treble-Light, Power LED 20000*) with an opening angle of 9° and 18,000 lm luminous flux at 170 W nominal electric power input each. One spotlight consists of 48 LEDs with 3.5 W each. The LED array is located 1 m above the ceiling of the convection chamber and a passe-partout of 0.75 m diameter on the ceiling window defines the cylindrical measurement volume. In the experiment, the LED light source is synchronized with the camera system and is pulsed with a period of 3 ms at 29 Hz, corresponding to a ~10% duty cycle.

As a reference, we measure the horizontal profile of the continuous light intensity (light meter, *Extech HD450*) at

the bottom of the convection chamber (red curve, Fig. 6) and at a height of 1.10 m across the measurement volume (blue curve). The illuminated volume is well defined by a sharp decay in light intensity, which helps to avoid light scattering from particles outside the measurement domain. Due to the opening angle of the collimated LEDs of 9°, the illuminated region is widened from 75 cm at the top window to around 80 cm in the measurement volume. An intensity dip in the center can be attributed to inhomogeneous distribution of LEDs in the light source and variations of intensity output between different LED arrays. Overall, the intensity of approximately 4.5×10^4 lx corresponds to ~17,000 lm over the whole area. A complementary view of the light intensity distribution is given in Fig. 6b, which depicts a 2D ensemble average of the intensity of all particle as identified by the STB particle tracking (see Sect. 3) over a run of 500 images. The x- and z- directions of the measurement volume were discretized in 2×2-pixel bins and all particles located within such a bin are averaged, thus averaging over the streamwise y-direction of space. The different intensity of the LED spots can be seen. The spot located at $x = 200$ mm, $z = 200$ mm appears to have nearly double the intensity of the weakest one, located at $x = -200$ mm, $z = 200$ mm. Possibly, the inhomogeneities are due to different types of LED arrays and different cable lengths in the experiment. On one hand, these findings document that great care should be taken before assembling the illumination arrays to achieve homogenous lighting conditions. On the other hand, the flow measurement results below indicate that the STB particle tracking algorithm can handle even the present inhomogeneities in the illumination. In sensitive cases, pretesting with the STB would also allow for a direct assessment of the three-dimensional intensity field
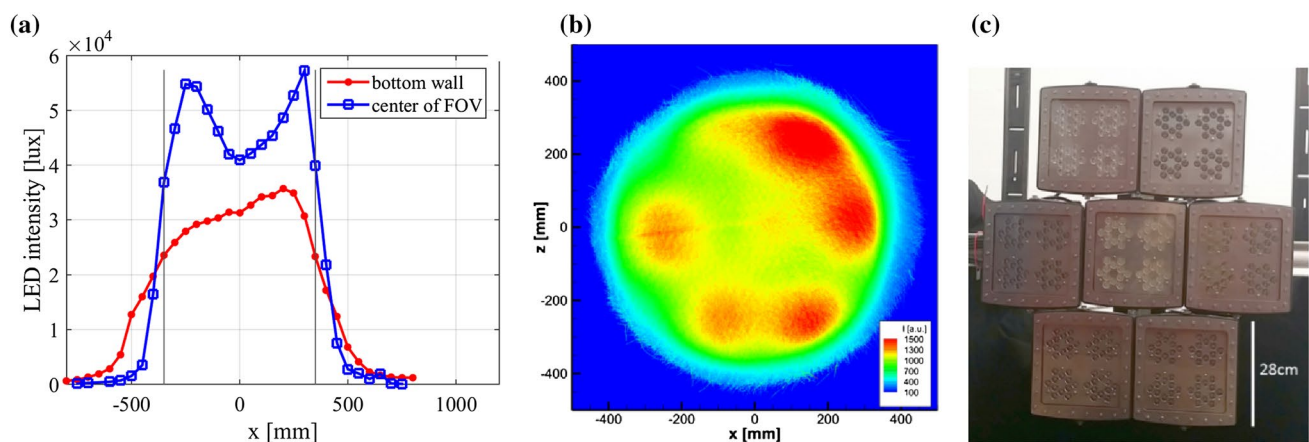
Fig. 6 a Intensity profile of continuous LED illumination across the measurement volume at a height of 110 cm (*blue squares*) and at the bottom of the convection cell (*red dots*) in the continuous mode. *Black bars* indicate the position of the passe-partout. **b** Horizontal profile of particle intensity as given by LPT results, averaged over streamwise (*y*) direction. **c** White LED array composed of seven LED lights with 48 3.5 W LEDs each

and an optimization of the illumination when setting up the experiment.

2.4 Camera system

The camera setup consists of five sCMOS cameras (*pco. edge 5.5 PCO*,) with a resolution of 2560 × 2160 pixels (6.5 μm pixel size). They are arranged in a flat M-configuration with a small height difference between neighboring cameras of 15 cm (see Figs. 1, 7). The cameras are placed on a circle around the convection chamber with a distance of 2.25 m to the vertical center line of the cylindrical measurement volume, whereby the cameras look perpendicularly through the front window. The lines of sight of the outermost cameras have an angle >90° allowing for an accurate reconstruction of the particle position in all dimensions. The cameras are equipped with $f = 35$ mm lenses (Zeiss Contax) with the aperture set to $F_\# = 11$, yielding sufficient depth of field to image the volume with 0.8 m diameter. The cameras are rotated by 90°, so their FOV has a width of 0.85 m and a height of 1.1 m to capture the vertically extended cylindrical volume.

The magnification is $M = 0.016$, corresponding to 0.4 mm/pix, such that the two glare points—reflections at two points of the soap bubble (Kühn et al. 2011; Scarano et al. 2015)—fall into one single peak of an isotropic circular particle image of ~3 pixel diameter. In general, the image quality is one of the best ever seen by the authors in comparable experiments. An example of the high image quality is given in Fig. 8, showing a detail of a sample run at low seeding density and random flow. The sum of five consecutive images is shown, documenting the uniformity in particle imaging—both between the different bubbles,
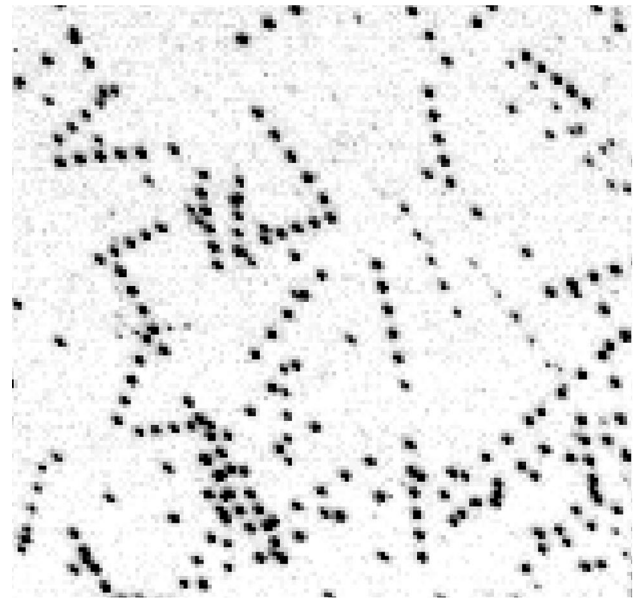
Fig. 8 Sum of five consecutive images at low seeding density with subtracted background (340 counts). Otherwise unprocessed camera image with inverted *colors*, showing the high quality of the HFSB's particle image

as well as in time for a single bubble. The high intensity of particle images due to the large diameter of the seeding particle (300 μm) and the absence of coherent light—preventing effects of interference or speckles—contribute to the image quality.

For the 3D calibration of the cameras, a large planar 950 mm × 760 mm calibration target is placed vertically in the convection chamber and manually shifted to three positions, each 200 mm apart. Image acquisition, synchronization of cameras and the light source were
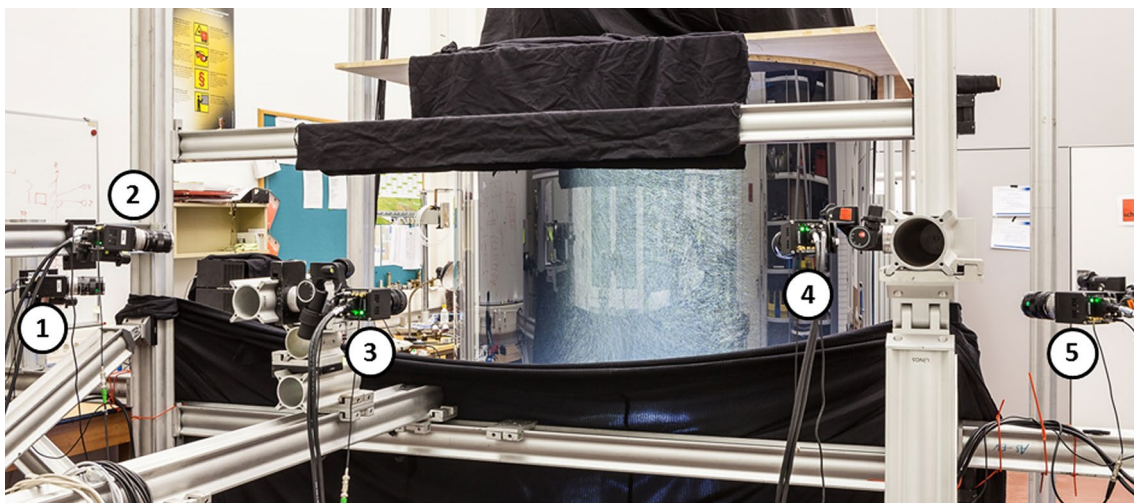
Fig. 7 Camera setup with five sCMOS cameras *pco.edge* around the convection cell. HFSB seeding in the convection cell appears white

controlled through the *DaVis*8.3 software (*LaVision*). The accuracy of the volumetric calibration was enhanced using Volume-Self-Calibration (Wieneke 2008); the particle imaging was calibrated, yielding a volumetrically resolved optical transfer function (OTF, Schanz et al. 2013a). All five cameras were connected to a single PC and the data were recorded directly to the hard disc. The recording frequency of 29 Hz reflects the maximum write rate that was attainable. Connecting each camera to a single PC would allow for a repetition rate up to 100 Hz at full frame resolution.

3 Data evaluation

For this study, a single measurement run with high particle density and large particle shift was chosen from the available material. The objective here is to assess the performance of the DLR STB algorithm under these conditions. The seeding density was found to be 0.08 ppp in the center of the image (Fig. 9). The temperature difference between the air directly over the hot plate and the surrounding air was approximately 8 °C (Fig. 2), leading to maximum velocities of around 0.35 m/s, and a maximum particle shift of nearly 30 pixels. 1000 images were recorded at a frequency of 29 Hz, of which 500 were evaluated using STB. Image preprocessing consisted of subtracting the spatially smoothed minimum image and a constant of 50 counts.

3.1 Shake-The-Box processing

The DLR in-house STB algorithm was applied. The main novelty of the LPT algorithm is to predict particle positions in the next time step by extrapolating already existing trajectories. By this procedure, the positioning problem and the problem of attaching particles to the right trajectory is greatly simplified. For details of the method, please refer to (Schanz et al. 2013b, 2014, 2016a). The following parameters were applied to the current dataset: The

number of triangulation iterations was set to 2—using an allowed triangulation error of 1.0 pixel—followed by one triangulation iteration using a reduced set of cameras. Each of these was followed by five shake-iterations [$n_1 = 2$, $n_2 = 1$, $m = 5$, $\varepsilon = 1.0$; see (Wieneke 2013)]. For the initialization phase (the first four images), the number of triangulations was doubled. No help of a predictor in form of a vector field gained by tomo-PIV processing was used in this case. As shown in Fig. 10a, the number of tracked particles quickly increases with the number of processed images. After the initialization phase, around 53,000 tracks of length four are found. This number rises to 110,000 only three time steps later and reaches 200,000 at time step 18. For this first pass (going forward in time), the number of tracked particles saturates at just over 250,000 after about 80 images. From there on, the tracked particle number decreases, as bubbles burst and disappear from the tracking system. When reaching the end of the time series, time is reversed and the algorithm walks backwards in time through the dataset. By doing so, known tracks that were not immediately found within the first pass are extended and the first time steps, where the tracking system was not yet converged in the first pass, are completely reconstructed. At the maximum, around 275,000 particles are simultaneously tracked in the second pass. To the knowledge of the authors, this is an unprecedented number for particle tracking methods, which typically operate with hundreds or a few thousand particles within the same image.

Figure 9 compares the camera image to the virtual image, created by reprojecting all tracked particles. The high quality of the tracking process is documented in Fig. 10b, which shows the statistics track length after pass 2. A distinct peak can be seen at 500 images, showing that over 81,000 particles have been tracked over the whole time series. These are particles slowly moving downwards in the recirculation region. The rest of the tracks have lengths that are evenly distributed, reflecting the fact that many particles are transported at different speeds into

Fig. 9 **a** Details of camera image. **b** Reprojection of particle distribution as reconstructed by STB
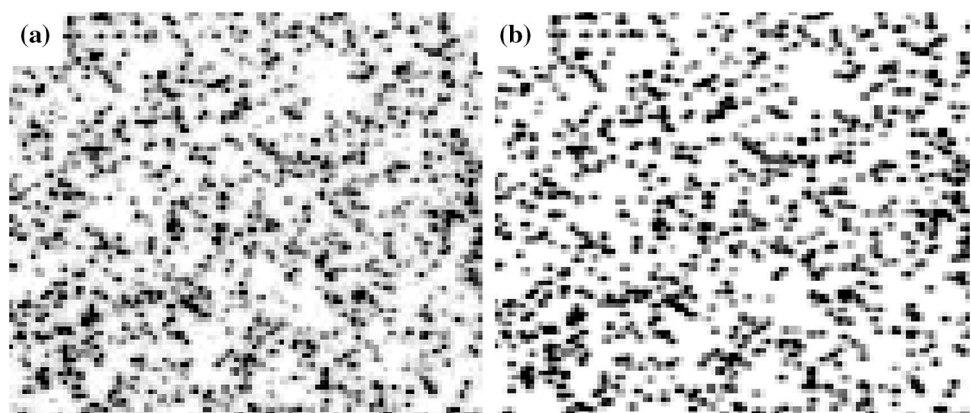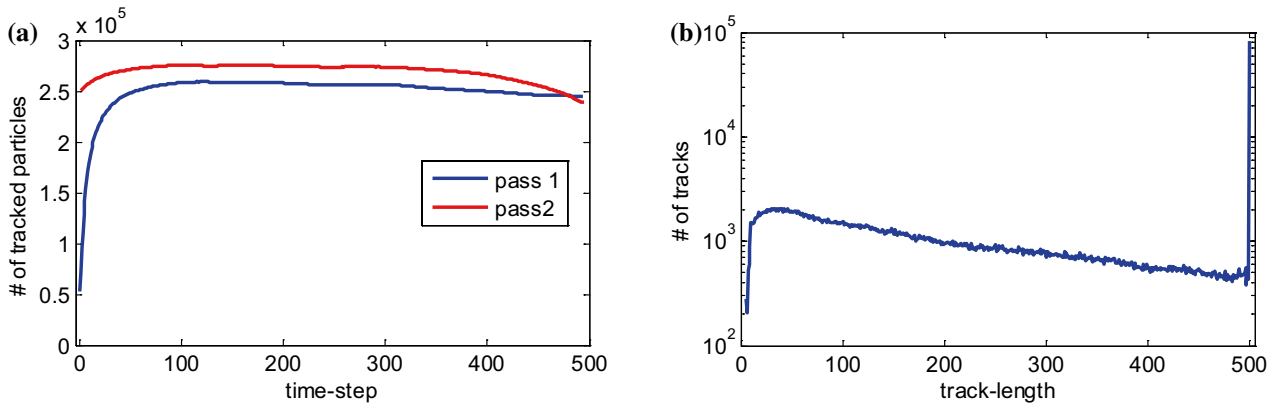
Fig. 10 **a** Development of number of tracked particles over time for both passes of STB. **b** Statistics of track length after pass 2

and out of the volume. The computational effort, combined for both passes, is around 220 s per time step on a 16-core Xeon server. A recently optimized version of the STB code further reduced this time to around 40 s for two passes on a 4-core i7 desktop computer.

Following the second pass, the particle tracks are fitted with a cubic B-spline curve with the TrackFit algorithm (Gesemann et al. 2016). On average, the particles are moved by 0.103 px from their original position (0.046 px in x-, 0.040 px in y- and 0.083 px in z-direction). The velocity values and acceleration values are calculated as derivatives of the B-spline curve. After fitting, a particle position accuracy of ~36 μm (0.086 px) and a velocity uncertainty of ~0.4 mm/s is estimated from the spatial spectrum of the tracks; this corresponds to a dynamic velocity range of ~875:1.

4 Flow field results

4.1 Mean field and velocity profiles

Based on a recording of 1000 consecutive time frames at a lower seeding concentration (~42,000 particles), the mean field of axial velocity is obtained by bin averaging the scattered velocity data. For the lower seeding concentration, a longer time series is available with potential advantages for the convergence of the statistics. Despite the relatively poor convergence of the statistics, some basic properties of the mean field can be observed and compared to another experiment (Pham et al. 2005). Figure 11a shows the radially averaged axial velocity. Directly above the heat source, fluid is accelerated by the buoyancy force, reaching a maximal axial velocity at $y^* = y/D \sim 1.5$. At larger y^*, axial momentum increasingly diffuses into the lateral quiescent
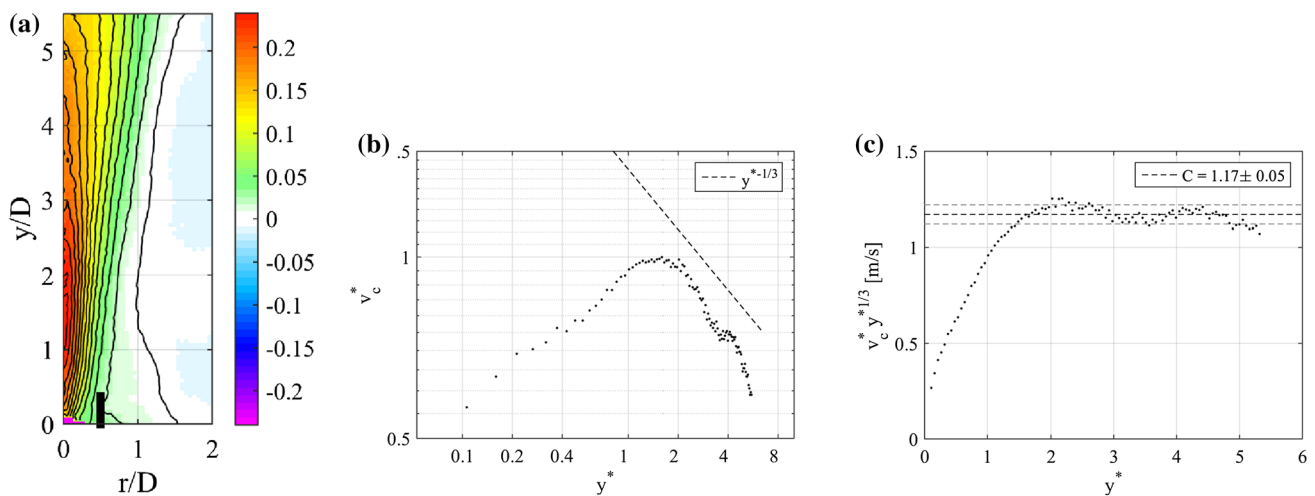
Fig. 11 **a** Mean axial velocity [m/s] averaged in circular bins with radius r. Contours are 0.02 m/s apart. *Black bar* indicates extension of the heat source. **b** Central profile of axial velocity v_c^* along the plume axis (circular bin with radius 0.1 D) with turbulent scaling regime. **c** Compensated axial velocity profile

fluid and the central axial velocity decreases. Figure 11b shows the axial velocity along the central axis of the plume. A simple box model of the plume assuming conservation of volume, momentum and density deficiency (equivalent of heat) (Morton et al. 1956; Pham et al. 2005; Plourde et al. 2008) leads to an expected −1/3 power scaling of the decaying axial velocity at the center $v_c^*(y^*) = Cy^{*1/3}$, where v_c^* is the axial velocity at the center normalized to its maximum. We observe an early initiation of the -1/3 scaling at $y^* \sim 2.0$ (Fig. 11b), equivalent to a fully turbulent plume. This is earlier than the transition at $y^* \sim 4$ reported by the experimental study of Pham et al. (2005). In their setup, the heat source is coplanar with the bottom wall, while in our setup, it is located on a support of height ~2D. Thus, air is entrained into the plume laterally from below the heat source and the inward flow around the edges of the heat source already generates turbulent instabilities, leading to earlier onset of turbulence. With a value of $C = 1.17 \pm 0.05$ (Fig. 11c), the involved constant is somewhat smaller than the value of 1.26 reported by Pham et al. (2005).

4.2 Acceleration statistics

The acceleration of particles decorrelates faster than the velocity and therefore statistics show a better convergence for the same data base of 1000 consecutive frames. Figure 12 shows the mean fields of acceleration in the thermal plume averaged in circular bins. With a bin size of 10 mm, most bins comprise between 10^3 and 10^4 data points, with only the central region having a number of data points below 10^3. While the mean of horizontal acceleration, $\langle a_x \rangle$ and $\langle a_z \rangle$, vanishes, the buoyancy force in positive y-direction is clearly apparent in the mean field of the vertical acceleration $\langle a_y \rangle$. The maximal value of $a_y = 0.12$ m/s^2 is reached at $y/D \sim 0.5$. This mean vertical acceleration corresponds to a mean temperature difference of $\langle \Delta T \rangle = T_0 \langle a_y \rangle / g \sim 3.6$ K (cf. Eq. 4), which seems realistic, given the temperature difference of 8 K between the heat source and the ambient air (cf. Fig. 2). A black bar at $r/D = 0.5$ marks the extension of the heat source at the bottom of Fig. 12b. The region of positive upward buoyant acceleration coincides with the edge of the heat source. The spatial distributions of fluctuations of acceleration are uniform for all three directions (Fig. 12d–f). Contours show a similar opening angle of the plume as the mean velocity field in Fig. 11a. The fluctuations of acceleration reach values of up to ~0.12 m/s^2 at the center at $y/D \sim 2.5$. This maximal magnitude of fluctuations is similar to the maximal mean upward buoyant acceleration.

Figure 13 shows the probability density functions (pdfs) of acceleration at selected points in the thermal plume. The grid of reference points is marked in Fig. 12b. Each subfigure in Fig. 13 represents a set of pdfs along a vertical

column in the thermal plume ($y/D = [1, 2, , 3, 4, 5]$) at a constant distance from the center ($r/D = [0, 0.5, 1.0]$), cf. Fig. 12b. The pdfs are plotted with a logarithmic y-axis. For a better overview, the normalized pdfs have been shifted by a constant factor of 10 and ordered such that the uppermost curves correspond to the uppermost point in the column. At the center ($r/D = 0$), the shift of vertical acceleration a_y toward positive values due to the buoyancy can clearly be seen, while the distributions in the two horizontal directions (x, z) have zero mean and are isotropic. Toward the top of the plume, the a_y-curve approaches the curves of the horizontal acceleration (a_x, a_z) and all three curves converge at the uppermost point, i.e., the developed turbulence leads to isotropic distributions of acceleration. At a distance from the center of $r/D = 0.5$, in the shear layer at the edge of the heat source, the standard deviation of the acceleration is smaller than in the center (Fig. 12e). However, higher extreme accelerations appear in the longer tails of the pdf (Fig. 13b). While the acceleration pdfs at the center follow closely a Gaussian distribution (bold black line in Fig. 13a, Gaussian distribution with standard deviation σ_{ax} and zero mean), in the shear layer the acceleration pdfs significantly deviate from a Gaussian shape. The distribution at the top of Fig. 13b shows an exponential decay. The horizontal acceleration pdfs in Fig. 13b fall on top of each other, while in the vertical acceleration (red curve), a positive mean value, i.e., a contribution of the buoyancy, can still be observed. At a distance from the center of $r/D = 1.0$, the acceleration pdfs are isotropic without any influence from the buoyancy forces. The non-Gaussian shape suggests that in the outer part of the thermal plume, the acceleration is dominated by turbulent motions (cf. e.g., La Porta et al. 2001; Rosi et al. 2014), i.e., by the pressure gradient.

4.3 Instantaneous flow results

Figure 14 shows an instantaneous flow situation depicted by ~275,000 particle tracks, whose velocity vectors are drawn for three consecutive time steps and color-coded by streamwise velocity. A front view and a top view of the entire measurement volume are provided. The convex hull around all particles has a mean volume of 560 L. A large region of slowly moving particles surrounds the rising thermal plume. Due to the confining boundaries of the convection cell, outside of the thermal plume the flow is slightly downwards on average. In the center of the plume, the maximum velocity values are ~+0.35 m/s. In the presented time instant, the plume is broadened in the z-direction, compared to the x-direction, despite the circular symmetry of the hot plate. Figure 15 focuses on a central slice of the particle trajectories, giving an impression of the spatial sampling of flow structures with the dense seeding.
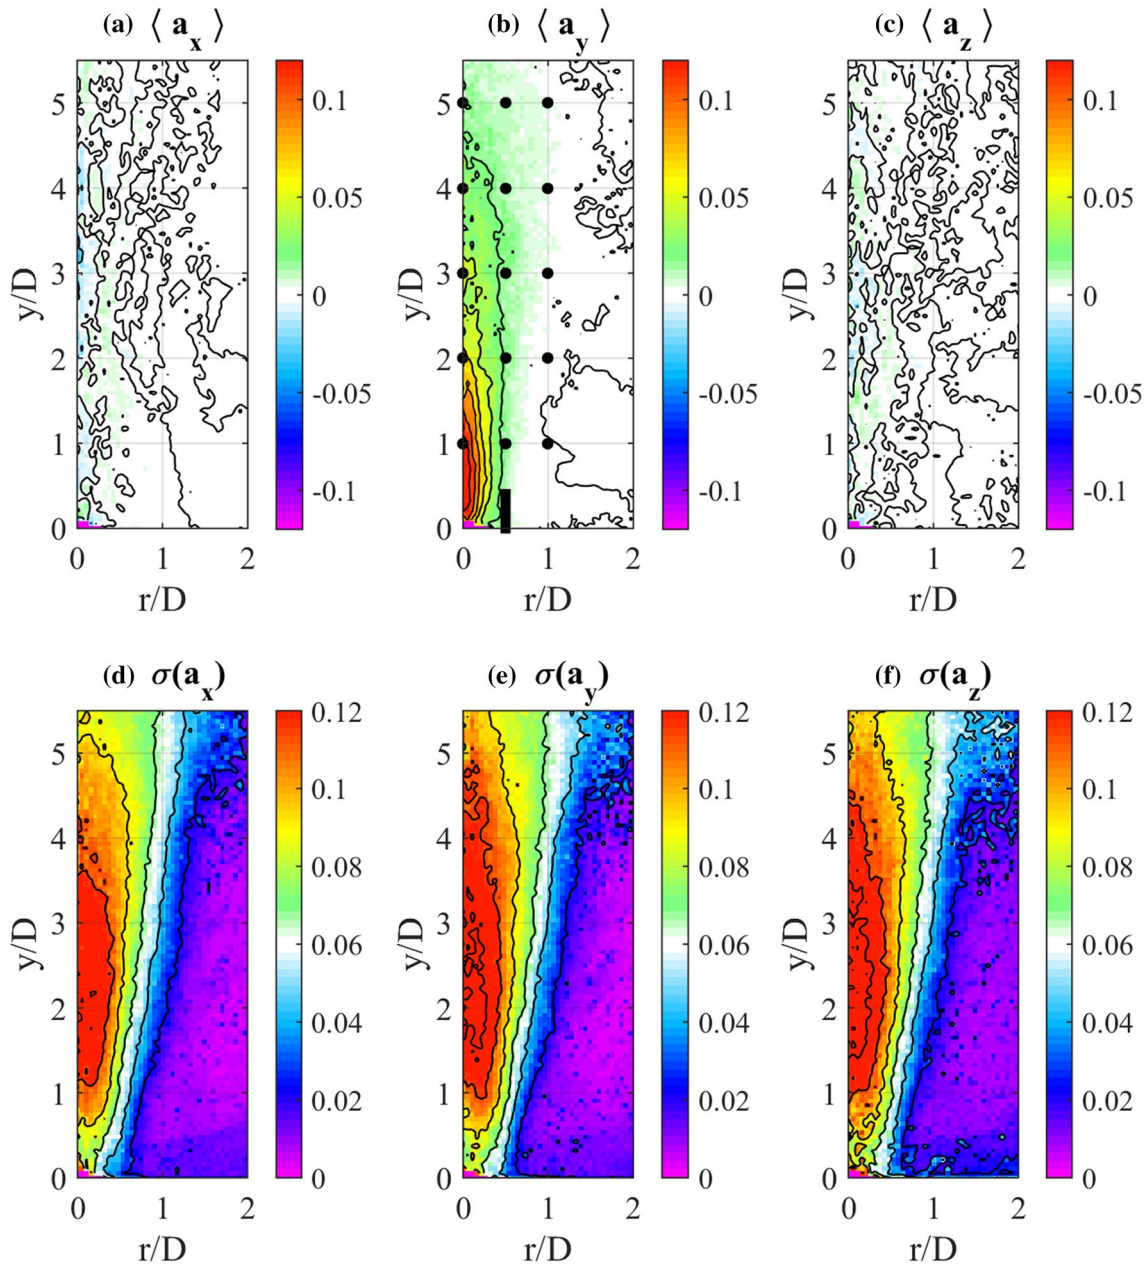
Fig. 12 **a–c** Mean and **d–f** fluctuations (standard deviation) of the particle acceleration. Bin size is 10 mm in both radial and vertical direction. Contours have a separation of 0.01 m/s^2

Locations of vortical structures can already be identified by looking at the particle tracks in Fig. 15; however, a quantitative description of vortices expressed as vorticity or using the Q-criterion would be desirable. To this end, the discrete Lagrangian information of velocity and acceleration at the scattered particle location is interpolated onto an Eulerian grid using the DLR in-house algorithm 'Flow-Fit' (Gesemann et al. 2016). In the FlowFit algorithm, the interpolating function is composed of a weighted sum of three-dimensional and evenly spaced cubic B-splines. To determine the weights, the interpolant is fitted to the measured velocity and acceleration values at scattered particle positions using velocity and pressure as fit variables. In the most advanced version of FlowFit used here, a full Navier–Stokes regularization is implemented. The cost function of this optimization problem includes six terms: the difference between data points and the B-spline interpolant for (1) velocity and (2) acceleration (fit it data); the penalization of high wavenumbers of the (3) velocity and (4) pressure field (smoothing); (5) the penalization of

**(a)** r/D =0.0   **(b)** r/D =0.5   **(c)** r/D =1.0
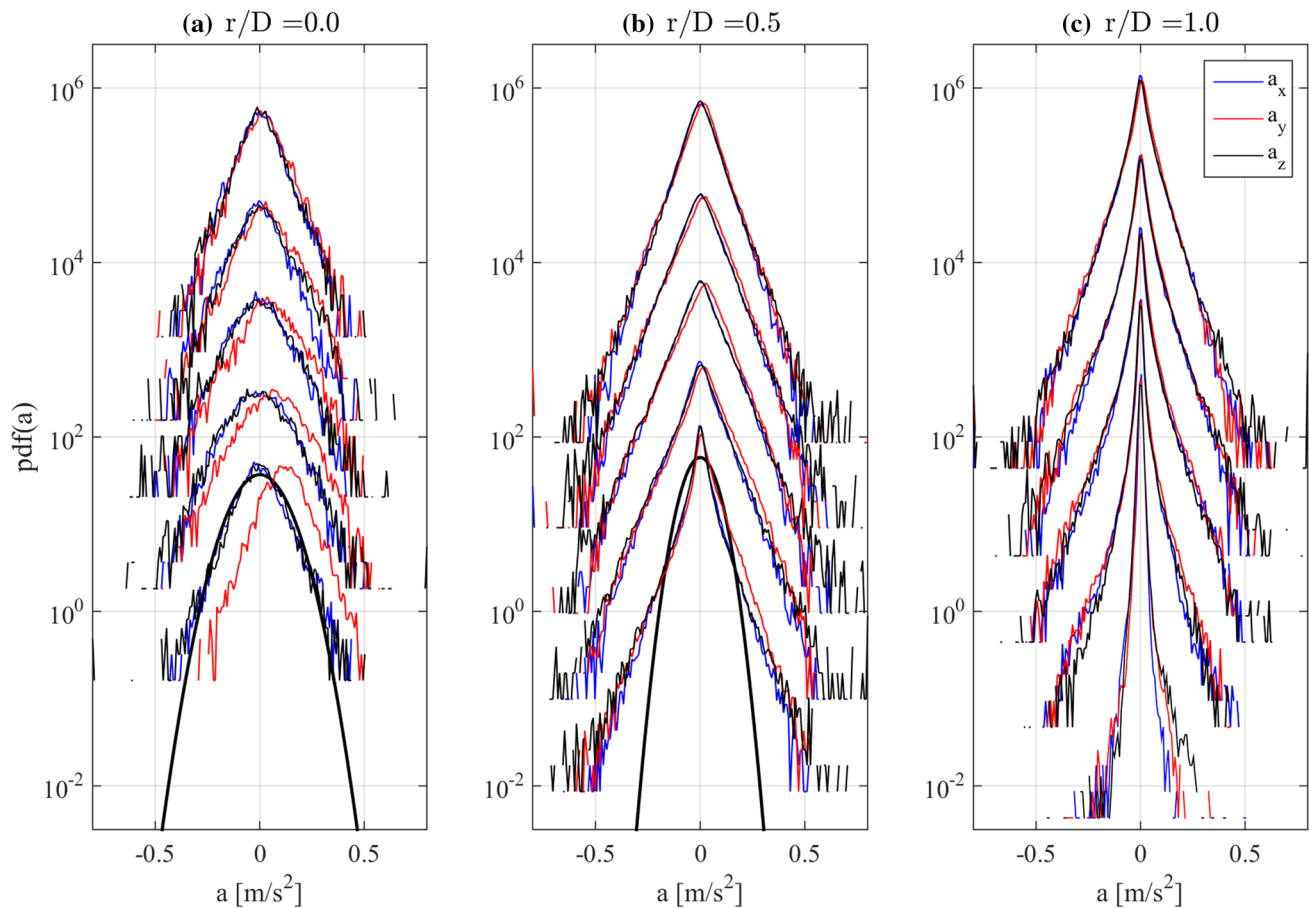
Fig. 13 Probability density functions of particle acceleration based on bin averaging over a circular bin with 50 mm width in both radial and vertical direction. The central position of the bins are ordered on a grid which is marked in Fig. 12b above. Pdfs are shifted such that the lowest corresponds to $y/D = 1$ and the uppermost to $y/D = 5$. The number of data points is **a** $n \sim 7000$, **b** $n \sim 130,000$, **c** $n \sim 280,000$. *Bold black curves* in **a** and **b** are Gaussian distributions for comparison

divergence of the velocity field, $\nabla \cdot \mathbf{u}$ (incompressible); and finally, (6) the penalization of $\nabla \cdot \partial \mathbf{u}/\partial t$. The last term being zero translates to a simplified coupling of the velocity and the acceleration field through the material derivative. The material derivative introduces a non-linearity into the cost function, leading to a non-linear least squares problem which is solved with the L-BFGS method. Once the interpolant is computed, it can be sampled on a regular grid, including its first and second spatial derivatives that are analytically obtained from the B-splines. Vorticity or the Q-criterion can thus directly be computed without numerical differentiation. The FlowFit interpolation was applied to each of the 500 time steps, taking velocities and accelerations of the respective tracked particles as input data. The B-spline system is setup, such that on average ten B-spline cells are present for every particle (0.1 particles per cell), leading to a spacing of ~6 mm between the cells. The flow was assumed to be incompressible which is reasonable to first approximation, given the small temperature

differences in the flow. The L-BFGS algorithm is iterated until convergence and the resulting continuous function is closely sampled on a grid with 3 mm spacing, ultimately leading to volumes of $254 \times 334 \times 254$ vectors. On a single core processor, the FlowFit interpolation takes about 2 h per frame. Performing the FlowFit interpolation for several frames simultaneously is embarrassingly parallel.

An example of the results gained by applying FlowFit to the STB track data is given in Fig. 16, which shows isosurfaces of the Q-criterion for the same time step shown in Fig. 15. The full amount and extent of the vortical flow structures in the plume becomes apparent. Long, undisturbed vortices are identified in the shear layers surrounding the center of the plume. Central structures are smaller, but show equal strength. Movies of the temporal evolution of the plume are available in the supplementary data [link to Suppl. Data]. When looking at a time series of such images, a high temporal coherence is noticeable, i.e., isosurfaces evolve smoothly in time. The Q-criterion
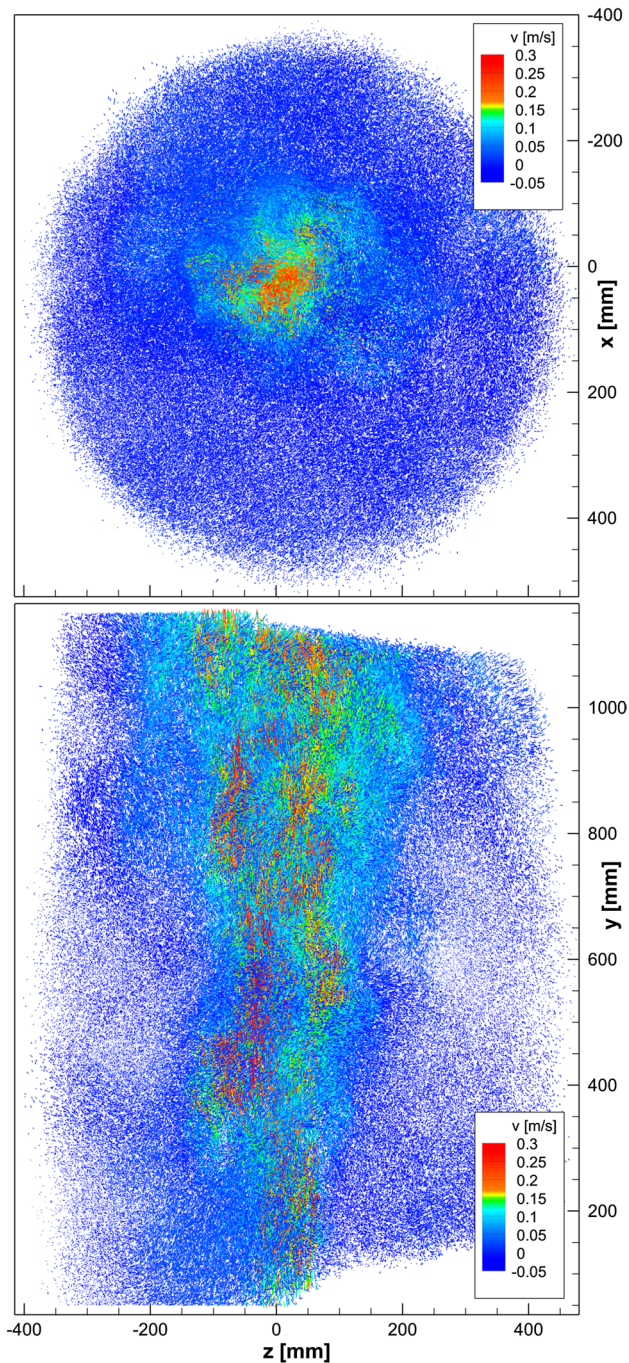
Fig. 14 *Top view* and *front view* of the entire measurement volume, showing ~275,000 particle tracks for three consecutive time steps, color-coded by streamwise velocity *v*. The inclined upper and lower bounds of the measurement volume (lateral view) are due to the opening angle of the cameras (∼ 25°)

is derived from the velocity gradient tensor. As such, the temporal coherence is a clear indication that the LPT measurement at high seeding density in combination with the FlowFit interpolation strategy is able to recover the vast majority of the present flow structures. The high quality of

the tracking process translates directly into the quality of the Eulerian representation.

The high position accuracy in the reconstructed trajectories, which is achievable due to the high particle image quality, allows for an evaluation of the particle accelerations; these are the second temporal derivative of the trajectory. Figure 16 shows tracks that are color-coded by streamwise acceleration. The acceleration varies smoothly in the measurement volume. Around vortices oriented with their axis in x-direction (out of plane), the centrifugal acceleration of the circular motion can be observed.

Just as for velocity, the acceleration can be interpolated onto an Eulerian grid using the FlowFit algorithm. Figure 17a displays an interpolation with only a smoothing constraint, but no further regularizations. The resulting field shows many isolated small structures that seem unphysical and point to a poor reconstruction of the acceleration field. In a strict sense, adding further regularizations to the interpolation is limited in the present flow by the unknown buoyancy term in the momentum equation (last term)

$$\rho_0 \mathbf{a} = -\nabla P + \mu \mathbf{\Delta u} - \rho_0 \mathbf{g} \Delta T / T_0, \qquad (4)$$

written here in the Boussinesq approximation (e.g., Kundu and Cohen 2008). The viscous term is often neglected in turbulent flows, since it is small compared to the dominant pressure gradient term. If also the buoyancy term was negligible, the acceleration field would be curl-free and could be regularized with the condition of vanishing rotation. However, a rough estimation of the buoyancy term with $\Delta T = 2$ K, a typical temperature fluctuation at 100 mm above the heat source, gives a value of ~0.07 m/s^2 for the buoyancy acceleration. This acceleration is not small compared to the magnitude of the total measured acceleration of ~0.20 m/s^2. Maximal temperature differences directly above the heat source may be even as high as $\Delta T = 8$ K (cf. Fig. 2), making buoyancy locally the dominant term in the momentum equation close to the heat source.

Nevertheless, when neglecting the unknown buoyancy, a joint interpolation of the velocity and acceleration fields with the complete Navier–Stokes regularization (pressure and viscous force) leads to a reconstruction of the acceleration field that seems smoother and more physical (cf. Fig. 17a, b). Isolated pieces of the isosurfaces become connected and structures of positive and negative centripetal acceleration adjacent to vortices appear elongated, as expected (cf. Figs. 17b to 16). Why and where is the reconstruction improved when applying a physically incomplete regularization? The elongated vortices are located in the shear layer, where ambient air is entrained and the temperature difference is lower, leading to a small buoyancy force. Furthermore, in the rotating

Fig. 15 Particle trajectories with a length of 11 time steps in a central slice with 100 mm thickness, color-coded with vertical velocity. Detailed view (*upper right*) corresponds to the region in the *black rectangle* with 41 time steps in a slice of 40 mm thickness
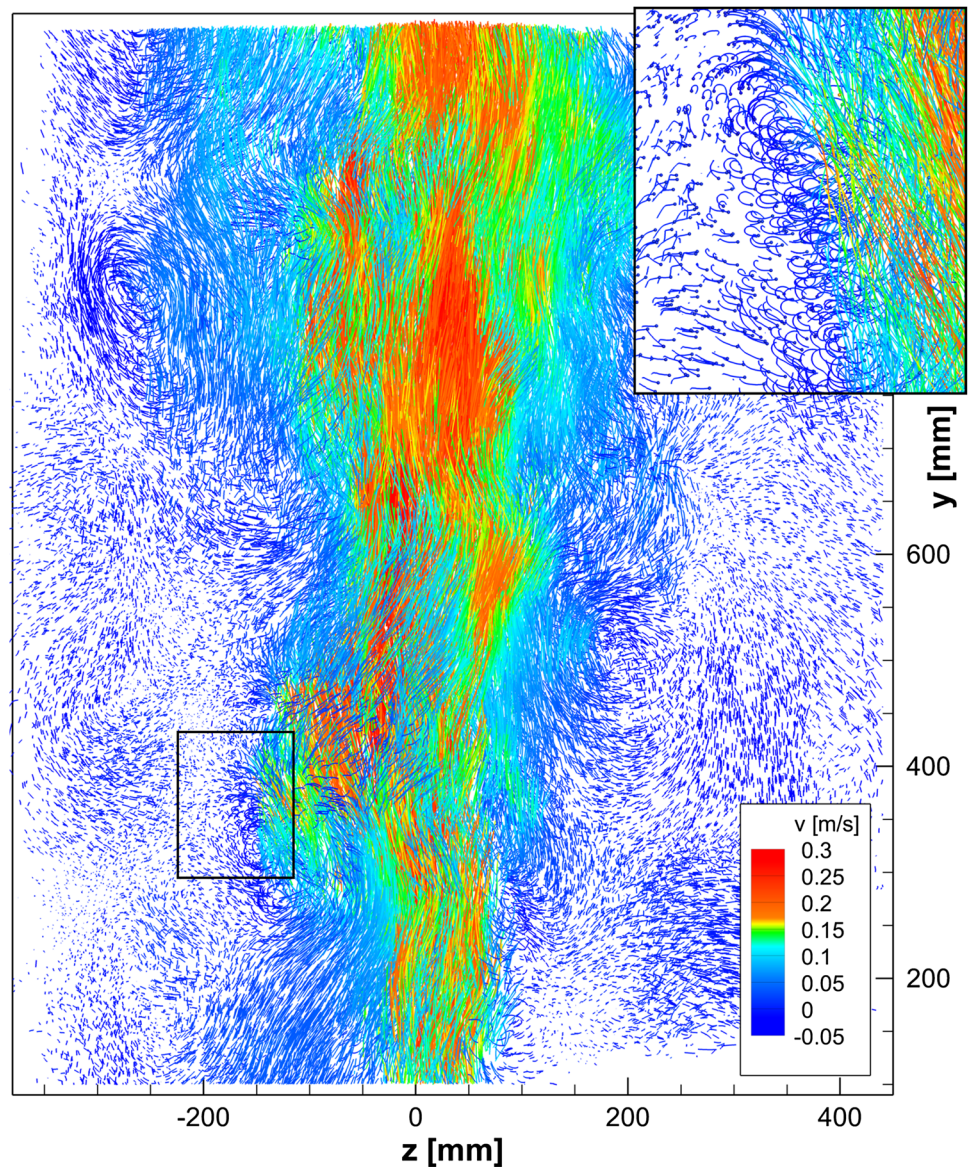

vortices, the centripetal pressure gradient force is strong. Therefore, applying the Navier–Stokes regularization without buoyancy improves the reconstruction of these vortex cores. In regions with a large buoyancy term, i.e., predominantly in central filamentous regions emanating from the heat source, the regularization without buoyancy might indeed bias the reconstructed acceleration field. Systematically investigating the local magnitude of the buoyancy force and the implications on the regularization scheme could be done on numerical data, but is beyond the intended scope of this paper.

In summary, the presented results demonstrate that detailed instantaneous fields of velocity and acceleration can be measured in large volumes when tracking HFSBs with the STB method.

5 Finite-time fluid transport and Lagrangian coherent structures

Flow measurements with the Lagrangian particle tracking (LPT) method STB for the first time offer a dense tangle of experimental particle trajectories that represent a direct measurement of the flow map. This flow map is the building block of the detection of a wide class of Lagrangian coherent structures (Peacock and Dabiri 2010; Peacock et al. 2015; Haller 2015). The following Lagrangian analysis of the experimental flow field is concerned with transport of fluid in space over finite times. In line with the Lagrangian frame of fluid motion, trajectories are tracked from an initial position at time t_0 to the final position at time $t = t_0 + T$ after a predefined time period T.

Fig. 16 Isosurfaces of Q-criterion (5 s^{-2}) color-coded by streamwise velocity and particle tracks (11 time steps, slice of 50 mm thickness) color-coded with streamwise acceleration
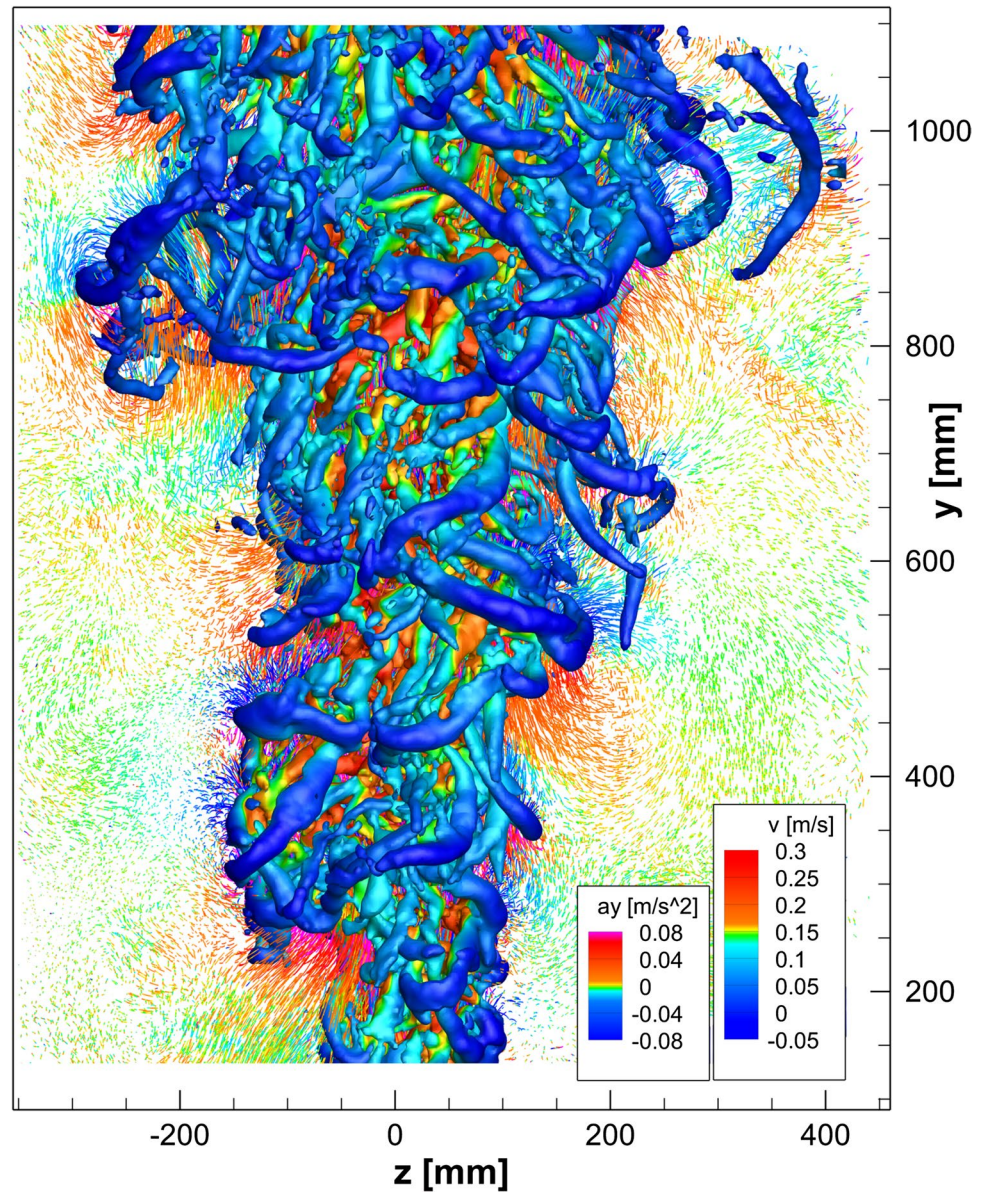
Properties of the fluid are advected along the trajectory, where the advective transport of heat, momentum and chemical tracers is of special interest in many applications. The key object for finite-time transport and the related Lagrangian coherent structures is the flow map $F_{t_0}^t(x)$ that maps initial fluid positions $x_0 = x(t_0)$ to their final position $x(t) = F_{t_0}^t(x_0)$ at a later time. Most flow data sets are given as a velocity field from experimental or numerical experiments, and the flow map must be computed by integrating $\partial_t x_p(t) = v(x)$ to obtain trajectories of fluid elements. With dense LPT, however, the flow map can directly be derived from measured trajectories. This idea was put forward by Raben et al. (2014) for two-dimensional flows and thereafter extended to

three-dimensional flows on unstructured grids by Rosi et al. (2015). In contrast to their work that focuses on simple vortical flows, the attainable spatial resolution of the flow map with the dense trajectory data in our experiment is significantly higher such that more complex flow structures can be represented. In the following example, we reconstruct a detailed flow map and analyze fluid transport in the thermal plume.

The flow map $F_{t_0}^t(x)$ is given as the concatenation of a series of Δt-flow maps

$$F_{t_0}^{t_n} = F_{t_{n-1}}^{t_n} \circ \cdots \circ F_{t_0}^{t_1} \tag{5}$$

with the time period $t_i - t_{i-1} = \Delta t$ (Brunton and Rowley 2010), where the time increment Δt is the
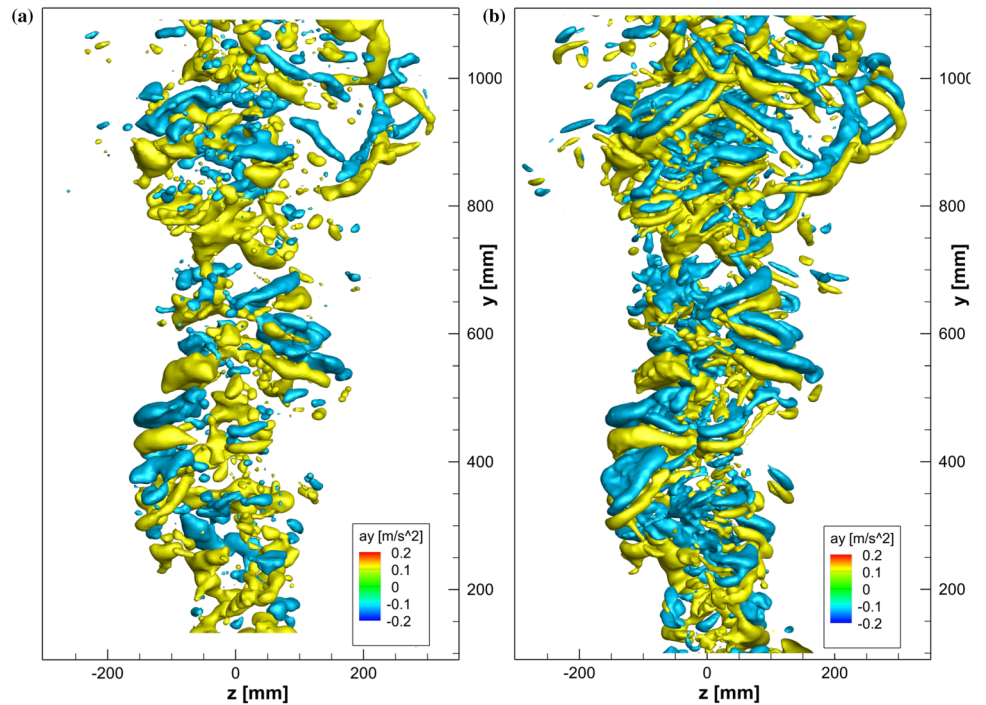
temporal resolution of the measured particle trajectories, $\Delta t = 34.5$ ms in the experiment. The composed flow map $F_{t_0}^{t_n}$ covers the time period $T = n\Delta t$. A Δt-flow map is given only at particle positions x_p and must be interpolated to a regular grid by the interpolation operator I_x,

$$F_{t_0}^{t_1}(x) = I_x x_p(x_p(t_0), t_1).\qquad(6)$$

For this interpolation, we use the FlowFit routine described above. On the level of the flow map, incompressibility of the flow corresponds to the condition $|\det \nabla F| = 1$. This condition can be used for a regularization of the field,

in the case of poor spatial resolution. Here, for a densely seeded flow, we refrain from using this regularization.

For the analysis of the flow of the thermal plume, the flow map is computed backward in time. This approach is linked to the question where the fluid originates from. The flow map is constructed on a regular grid and fluid particles are tracked back to their origin. Figure 18 shows the flow map at a central section of the plume at $z = 0$ mm. The original x-position of the fluid is shown in Fig. 18a, the original y-position is shown in Fig. 18b. The color-coded representation of original position can also be interpreted as the forward advection of particles carrying their original
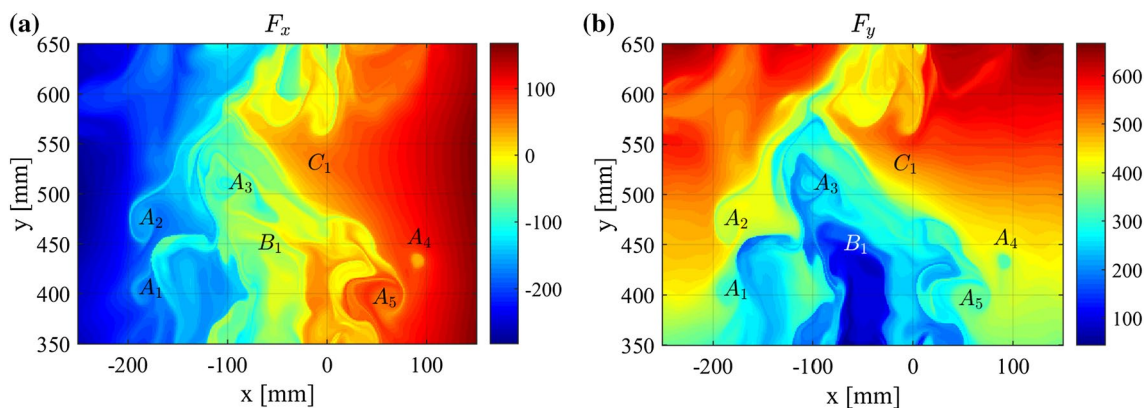
Fig. 18 Backward time flow map at $t_0 = 261, T = -35\Delta t$. *Colors* indicate where a particle—now located at position (x, y)—was at the earlier time $t = t_0 + T = 226$. Both spatial components F_x and F_y are shown. **a** F_x-component with vortical structures (A_{1-4}), and a prominent lateral intrusion of air from the right (C_1). **b** F_y-component with vertical intrusion of air from below (B_1)

position at time $t = t_0 + T = 226$ to their new position at time $t_0 = 261$, with $T = -35\Delta t$. In the arising pattern, several distinct coherent structures can be identified. Several circular tracer patterns induced by vortices are marked with the letter A. The colors in Fig. 18 show, for example, that the small circular fluid region A_4, now located at $(x, y) \approx (90, 440)$, was earlier located at $(F_x, F_y) \approx (50, 300)$. Hence, the fluid region originates from a lower, more central part of the plume. The vertical intrusion of air from below at the center of the thermal plume is characterized by low original y-position (Fig. 18b, blue, B1). Also lateral entrainment of air into the rising plume can be observed (cf. Pham et al. 2005, their Fig. 9). The region C_1 has high original x-position (Fig. 18a, orange/red), indicating that the fluid was drawn into the rising plume from the quiescent right side. These three coherent structures of finite-time mixing are relevant for the fluid dynamics of the plume, especially if they are combined with the momentum and temperature fields. Heat is transported from the thermal boundary layer of the hot plate upwards, and momentum is mixed laterally from the interior of the plume through the turbulent shear layer into the quiescent ambient fluid. This transport is not continuous in time but is generated by intermittent intrusions of air regions carrying the particular quantity which can be visualized by means of the flow map.

The deformation of the fluid over a predefined finite time gives rise to hyperbolic structures where fluid from different origins merge. These hyperbolic structures, a subclass of Lagrangian coherent structures (LCS) (Haller 2015), can be diagnosed from fields of the backward finite-time Lyapunov exponent (FTLE) (Haller 2001). The FTLE $\sigma(x, t, t_0)$ is a measure of stretching of the fluid over a finite time and is computed from the flow map $F_{t_0}^t(x)$ as

$$\sigma(x, t, t_0) = \frac{1}{|t - t_0|} \ln \sqrt{\lambda_{\max}\left(C_{t_0}^t(x)\right)} \tag{7}$$

where $C_{t0}^t(x) = \nabla F^T \nabla F$ is the Cauchy–Green strain tensor and $\lambda_{\max}$ is its largest eigenvalue. Figure 19 shows the backward-time FTLE field with the typical elongated structures that separate fluid from different origins. According to its definition, the FTLE reaches high values along lines of high gradients of the flow map (cf. Fig. 18). With its units of 1/s, the FTLE represents an average separation rate between neighbored fluid elements over finite time in the direction of maximal stretching or shearing. According to their different origins, the patterns in the flow map in Fig. 18 are surrounded by lines of high FTLE values. Furthermore, lines of high FTLE values coincide with filamentary structures in the flow map, since filaments are characterized by a high stretching rate in one direction. In summary, the flow map constructed purely from dense experimental trajectory data from our flow measurement
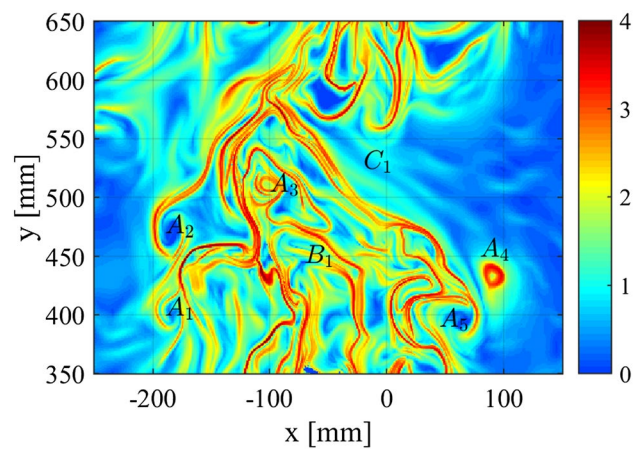
Fig. 19 Lagrangian coherent structures diagnosed from the backward-time FTLE field $\sigma(x, t, t_0)$ computed from the flow map in Fig. 18

allows for the identification of detailed transport patterns and opens the possibility to locate LCSs from the FTLE field and also from more involved LCS methods that rely on accurate trajectory data (Hadjighasem et al. 2016; Haller 2015).

6 Conclusions and outlook

A time-resolved volumetric flow measurement of a pure thermal plume with extended heat source was presented. Helium-filled soap bubbles with a quantified lifetime and intensity loss are used as flow tracers and are illuminated by an array of white high-power LEDs. In the measurement volume of 560 L, up to 275,000 bubbles could be tracked simultaneously using the STB algorithm. To the best knowledge of the authors, this investigation involves the largest number of tracked particles for LPT experiments so far. Evaluations using an updated version of the STB algorithm (Jahn 2017) allowed the processing of time series at even higher particle image densities, tracking up to 430,000 bubbles instantaneously in the present experiment. The results, spatially highly resolved vector volumes of velocity and acceleration, interpolated onto an Eulerian grid using the FlowFit method, show a multitude of well-resolved flow structures with a remarkable coherence in time. We also show that, due to the high spatial resolution, the experimental trajectories can directly be used to compute the flow map, to visualize Lagrangian coherent structures showing the fluid transport during the evolution of the thermal plume.

These achievements are possible mainly due to the high image quality in the experiment, i.e., small uniform particle images. The bubble size distribution is narrow around

the peak value of 300 μm (monodisperse) and all bubbles reflect the light to all cameras with similar intensity. The white, nearly incoherent light produced by the LEDs avoids effects like speckles or interference patterns, leading to temporally consistent particle images. All these aspects are beneficial for a reliable tracking process of STB and allow high particle concentrations of up to 0.1 ppp to be used.

The current study successfully demonstrates a large-volume flow measurement with flow velocities below 1 m/s. It is highly desirable to extend the general setup to higher flow velocities, enabling the operation in typical wind tunnel experiments with applications to boundary layer research, the car industry or sports aerodynamics. The three key factors needing to be addressed for a transfer of the presented measurement technique to low-speed wind tunnels are: (1) HFSB generators with an augmented production rate, i.e., increased number of nozzles, (2) high-speed cameras with a resolution of several mega-pixels and a frame rate of 1–10 kHz, i.e., data throughput of several Giga-pixels/s, and (3) pulsed collimated high-power LED arrays with a small opening angle of the emitted light. Modern equipment with these properties is available and allows for LPT measurements with high seeding concentration in faster flows. This has been shown in a recent experiment with an impinging jet in the same experimental chamber ($U = 16$ m/s, $V = 13$ L, 10 HFSB nozzles, frame rate 3.9 kHz) (Schanz et al. 2016b) and another still unpublished experiment at DLR Göttingen in a turbulent boundary layer in a wind tunnel ($U = 13$ m/s, $V = 180$ L, 100 HFSB nozzles, frame rate 2 kHz). Further studies report similar developments (Scarano et al. 2015; Caridi et al. 2016; Schneiders et al. 2016).

References

Biwole PH, Yan W, Zhang Y, Roux JJ (2009) A complete 3D particle tracking algorithm and its applications to the indoor airflow study. Meas Sci Technol 20:115403

Bosbach J, Kühn M, Wagner C (2009) Large scale particle image velocimetry with helium filled soap bubbles. Exp Fluids 46:539–547

Brunton SL, Rowley CW (2010) Fast computation of finite-time Lyapunov exponent fields for unsteady flows. Chaos 20:017503

Caridi GCA, Ragni D, Sciacchitano A, Scarano F (2016) HFSB-seeding for large-scale tomographic PIV in wind tunnels. Exp Fluids 57:190

Elsinga GE, Scarano F, Wieneke B, van Oudheusden BW (2006) Tomographic particle image velocimetry. Exp Fluids 41:933–947

Gesemann S, Huhn F, Schanz D, Schröder A (2016) From noisy particle tracks to velocity, acceleration and pressure fields using B-splines and penalties. In: 18th international symposia on applications of laser techniques to fluid mechanics, Lisbon, Portugal

Gilet T, Scheller T, Reyssat E, Vandewalle N, Dorbolo S (2007) How long will a bubble be? arXiv:0709.4412

Hadjighasem A, Karrasch D, Teramoto H, Haller G (2016) Spectral-clustering approach to Lagrangian vortex detection. Phys Rev E 93:063107

Haller G (2001) Distinguished material surfaces and coherent structures in three-dimensional fluid flows. Physica D 149:248–277

Haller G (2015) Lagrangian coherent structures. Annu Rev Fluid Mech 47:137–162

Hong J, Toloui M, Chamorro LP, Guala M, Howard K, Riley A, Tucker J, Sotiropoulos F (2014) Natural snowfall reveals large-scale flow structures in the wake of a 2.5-MW wind turbine. Nat Commun 5:4216

Jahn T (2017) Volumetric flow field measurement: an implementation of Shake-The-Box, Master thesis, DLR Göttingen/Georg-August-Universität Göttingen

Klimas P (1973) Helium bubble survey on an opening parachute flow field. J Aircr 10:567–569

Kühn M, Ehrenfried K, Bosbach J, Wagner C (2011) Large-scale tomographic particle image velocimetry using helium-filled soap bubbles. Exp Fluids 50:929–948

Kundu PK, Cohen IM (2008) Fluid mechanics. Academic Press, Elsevier

La Porta A, Voth GA, Crawford AM, Alexander J, Bodenschatz E (2001) Fluid particle accelerations in fully developed turbulence. Nature 409:1017–1019

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

Melling A (1997) Tracer particles and seeding for particle image velocimetry. Meas Sci Technol 8:1406

Morton BR, Taylor G, Turner JS (1956) Turbulent gravitational convection from maintained and instantaneous sources. Proc R Soc A 234(1196):1–23

Müller RHG, Flögel H, Schere T, Schaumann O, Markwart M (2000) Investigation of large scale low speed air conditioning flow using PIV. In: 9th international symposium on flow visualization, Edinburgh, UK

Okuno Y, Fukuda T, Miwata Y, Kobayashi T (1993) Development of three dimensional air flow measuring method using soap bubbles. JSAE Rev 14:50–55

Peacock T, Dabiri J (2010) Introduction to focus issue: Lagrangian coherent structures. Chaos 20:017501

Peacock T, Froyland G, Haller G (2015) Introduction to focus issue: objective detection of coherent structures. Chaos 25:087201

Pham MV, Plourde F, Kim SD (2005) Three-dimensional characterization of a pure thermal plume. J Heat Transf 127:624–636

Plourde F, Pham MV, Kim SD, Balachandar S (2008) Direct numerical simulations of a rapidly expanding thermal plume: structure and entrainment interaction. J Fluid Mech 604:99–123

Pounder E (1956) Parachute inflation process Wind-Tunnel Study, WADC Technical report 56-391, Equipment Laboratory, Wright Patterson Air Force Base. Ohio, USA, pp 17–18

Raben G, Ross SD, Vlachos PP (2014) Computation of finite-time Lyapunov exponents from time-resolved particle image velocimetry data. Exp Fluids 55:1638

Rosi GA, Sherry M, Kinzel M, Rival DE (2014) Characterizing the lower log region of the atmospheric surface layer via large-scale particle tracking velocimetry. Exp Fluids 55:1736

Rosi GA, Walker AM, Rival D (2015) Lagrangian coherent structure identification using a Voronoi tessellation-based networking algorithm. Exp Fluids 56:189

Scarano F, Ghaemi S, Caridi G, Bosbach J, Dierksheide U, Sciacchitano A (2015) On the use of helium-filled soap bubbles for large-scale tomographic PIV wind tunnel experiments. Exp Fluids 311(56):42

Schanz D, Gesemann S, Schröder A, Wieneke B, Novara M (2013a) Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. Meas Sci Technol 24:024009

Schanz D, Schröder A, Gesemann S, Michaelis D, Wieneke B (2013b) Shake-the-Box: a highly efficient and accurate Tomographic Particle Tracking Velocimetry (TOMO-PTV) method using prediction of particle position. In: 10th international symposium on particle image velocimetry—PIV13, Delft, The Netherlands, July 1–3

Schanz D, Schröder A, Gesemann S (2014) Shake-the-Box—a 4D PTV algorithm: accurate and ghostless reconstruction of Lagrangian tracks in densely seeded flows. In: 17th international symposium on applications of laser techniques to fluid mechanics, Lisbon, Portugal, July 07–10

Schanz D, Schröder A, Gesemann S (2016a) Shake-The-Box: Lagrangian particle tracking at high particle image densities. Exp Fluids 57:70

Schanz D, Huhn F, Gesemann S, Dierksheide U, van de Meerendonk R, Manovski P, Schröder A (2016b) Towards high-resolution 3D flow field measurements at cubic meter scales. In: 18th international symposium on applications of laser techniques to fluid mechanics, Lisbon, Portugal

Schneiders JFG, Caridi GCA, Sciacchitano A, Scarano F (2016) Large-scale volumetric pressure from tomographic PTV with HFSB tracers. Exp Fluids 57:164

Tobin ST, Meagher AJ, Bulfin B, Möbius M, Hutzler S (2011) A public study of the lifetime distribution of soap films. Am J Phys 79(819):819–824

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