

The VICCTOR algorithm for 3D Particle Tracking Velocimetry

Jean Le Bris^{1,*}, Maximilien Hebey², Benjamin Leclaire², Philippe Cornic³, Frédéric Champagnat³, Adam Cheminet¹

1: SPEC, CEA, CNRS, Université Paris-Saclay, CEA Saclay, Gif-Sur-Yvette, France

2: ONERA, Department of Aerodynamics, Aeroelasticity and Acoustics, Meudon, France

3: ONERA, Department of Information Processing and Modelization, Palaiseau, France

* Corresponding author: jean.lebris@cea.fr

Keywords: PTV processing, Two-Pulse PTV

ABSTRACT

We present a new two-pulse particle tracking velocimetry (PTV) algorithm that builds on recent advances in 3D particle reconstruction, matching, and position optimization. Unlike most current state-of-the-art approaches, such as Godbersen et al. (2024), which progressively recover the true particles while retaining a small proportion of outliers, we take a symmetric approach, operating at a high rate of detected real particles, and focusing most of our algorithmic choices on the gradual elimination of outliers. Initial estimates of particle positions at both time points are obtained by introducing a modified, large-scale version of the particle volume reconstruction (PVR) from Champagnat et al. (2014), which discretizes the image formation process onto a fine voxel grid. The positions and intensities are then globally optimized using an L-BFGS optimizer. Once the particles with low reconstructed intensity are discarded, the resulting clouds cover a large portion of the actual particles, although they contain a significant amount of noise (i.e., ghost particles). The F-VFC algorithm presented in Le Bris et al. (2025) is then used to match the particles from the two time steps while removing a large number of ghost particles from the particle clouds. This entire process is repeated on the residual images until convergence.

1. Introduction

In recent years, particle-based volumetric velocimetry methods have achieved significant performance improvements, enabling much higher maximum seeding densities than before. The time resolved aspect of Lagrangian Particle Tracking is key to its success: both the reconstruction and the tracking aspects benefit from it, as, to process a given time step, one can "guess" where the particle will be given its position and velocity prior to that. In the context of high speed flows, measurements have to be performed using the traditional two-pulse acquisition process, making prediction of the particle position at the second time-step much more difficult. To tackle this challenge, current state-of-the-art approaches put the emphasis on first obtaining reconstructed particle clouds of the highest quality possible. Reconstruction techniques use advanced methods which highly limit the number of ghost particles reconstructed (Wieneke, 2012; Jahn et al., 2021; Cornic et al., 2020). Without too many ghost particles, classical cross-correlation based approaches (Folki

(Cheminet et al., 2014), Particle Space Correlation (Novara et al., 2023)) suffice to recover the displacement field. The trade-off might be that some real particles are missed. From a symmetrical yet still mostly unexplored point of view, the user could choose to tolerate more ghost particles in the reconstructed particle clouds, with the benefit of recovering more true positives. With larger amounts of ghost particles, classical tracking approaches will fail, but recent advances in particle tracking allow to be more robust to ghost particles (Mercier et al., 2024; Le Bris et al., 2025). Whatever the approach, it is generally considered that iterative processes are key to recovering most true particles Yang et al. (2019); Lasinger et al. (2019); Novara et al. (2023). The general idea is that processing residual images will improve reconstruction quality through processing of previously missed information.

We first introduce the main steps of the VICCTOR algorithm (so-called for Voxel-grid Initalized reConstruction with Consensus Tracking and Outlier Removal), detailing each working principle as well as optimal operating conditions. The VICCTOR algorithm is then used on the two-pulse test cases from both LPT Challenges (2020 and 2025), available on the dedicated web portal www.onera.fr/flow-benchmarks (Leclaire et al., 2022), in Section 3. Finally, we assess VICCTOR's performances on experimental data acquired on the Giant von Kármán platform at CEA SPEC in Section 4, by comparing our results with velocities obtained by processing full time sequences with the Shake-The-Box algorithm.

2. The VICCTOR algorithm

Like state-of-the-art 3D PTV algorithms, the VICCTOR algorithm is iterative in nature. Its complete workflow is schematically represented in Figure 1. To explain how it works, we will first describe each of its main blocks, highlighting in particular how each one is adapted or modified to leverage its strengths and fit into the overall logic of the approach. Details will then be provided on how these blocks fit into the overall iterative process, converging toward the desired vector fields.

2.1. Individual blocks

2.1.1. Reconstruction: Large-Scale PVR-SMART

To obtain particle reconstructions at the two instants, we use a modified version of the Particle Volume Reconstruction (PVR), a tomographic reconstruction approach initially introduced by Champagnat et al. (2014). The main idea of the PVR approach, which is kept in this version, is to use a different weight matrix W_{ij} as traditionally considered (e.g. by Elsinga et al., 2006), built by discretizing the equation modelling image formation. This equation preserves the pointwise nature of the particles, and models the diameter of their images on the cameras (due to diffraction-limited optics) via a Point Spread Function (PSF). Another difference to traditional methods is that for op-




Fig. 1. Overview of the full VICCTOR algorithm. Blue stars $\star$ indicate particles from time step $t=0$, and red triangles $\triangle$ indicate particles from time step $t=1$

timal performance, the 3D voxel grid is set at a finer voxel-to-pixel ratio, of $v/p = 0.5$. Starting from images thresholded to decrease the impact of noise, the reconstruction is initialized using the Multiplicative Line-Of-Sight approach (Atkinson and Soria, 2009), followed by SMART (Simultaneous Multiplicative Algebraic Reconstruction Technique) iterations using the PVR matrix, the latter choice stemming from SMART's parallel structure, adapted to a GPU implementation. At the time of writing of their article, due to computational time constraints, Champagnat et al. (2014) advised the use of 25 iterations, as often recommended for SMART-based methods.

Back then, volumes considered in experiments were on average smaller than what is today considered as standard, so that the fine discretization with $v/p = 0.5$ remained computationally tractable. As a landmark, such a grid now leads to a total of 34 resp. 85 billion voxels for cases 3 resp. 2 (largest domain, with ratio of intersection over union volumes viewed by the cameras of $I/U = 0.46$) of the recent 2nd LPT Challenge. Besides, another change of nowadays practice compared to 10 – 15 years ago is the drastic increase in maximal seeding density; together with our present objective to maximize the amount of found particles at each step, translating at the initialization phase by a moderate value of the image threshold (estimated as the mean of background noise), this leads to the initial MLOS pruning keeping a quite large amount of the voxel grid. Considering e.g. that 30% of the voxels could be kept, which is realistic for experiments at high density, this leads to a size for the SMART problem of 305 Gb for the above mentioned 2nd Challenge case 2, which is well above both the maximum global memory of modern GPUs, and even of the RAM of powerful CPU workstations.

To avoid having to store the kept voxels on CPU RAM, or even worse, on disk, which would make the computational times totally unrealistic, we devised a dedicated global strategy, logically termed Large-Scale PVR-SMART. Its principle consists in replacing the first iterations by modified, so-called "on-the-fly" iterations, where the residual images are successively updated, allowing the voxel intensities to not be stored. Each of these iterations in fact recomputes the problem from the start (i.e. MLOS); this is in line with the philosophy of GPUs, which balance a moderate memory capacity by a very large number of simple but fast computing cores, so that duplications can be considered. The number of these iterations is found by counting the number of voxels whose intensity is above a minimal threshold. When this number is compatible with storage in the GPU global memory.

These standard iterations run logically much faster than the initial ones, so that the former guideline of a total of 25 iterations can easily be extended to several hundreds with no significant increase in computational time; in practice, for the VICCTOR algorithm, we found that 100 iterations guaranteed a high enough degree of convergence. To then obtain candidate 3D particle positions and intensities from the SMART results, which consist of grey levels positioned at the coordinates of their corresponding voxel centres, we perform a three-point 3D Gaussian fit over the local intensity maxima. This fit, which can also be implemented efficiently in a parallel way, is simply used

to provide a first guess of the particles' positions and intensities, as these will be optimized in the next phase (section 2.1.2).

Alongside the above mentioned parameters (image threshold, total number of SMART iterations), another parameter of interest is the estimated size of the image PSF. In the PVR model (Chamagnat et al., 2014), consistent to widely used PIV/PTV image formation models, it is taken as a Gaussian integrated over the pixel surface, thus parameterized by its standard deviation σ_{PSF} (which can be obtained by calibration from the images).

Note that, while the PVR approach uses voxels which are half of the size of a back-projected pixel, in the sequel, we will however refer to voxels in a more traditional sense, i.e. with size of a back-projected pixel, as a characteristic scale for error metrics.

2.1.2. Optimization: Global Particle Cloud Optimization (GPCO)

Consistent with the characteristics of the SMART reconstruction, and of the three-point Gaussian fit to obtain particle candidates out of the voxel grid, it can be expected that the obtained particle distribution will contain close to all correct particles, but with large errors in their positions and intensities, and exhibit a large number of ghosts. To tackle these issues, we then implement a processing step which we term Global Particle Cloud Optimization (GPCO). Similar to the sub-voxel refinement step introduced by Cornic et al. (2020), GPCO relies on a global optimization of the positions and intensities of particles, but used in a different way and with added components. However, whereas Cornic et al. (2020) considered the global optimization after having matched particles between the two instants, in order to refine their subvoxel position, here we choose to perform this optimization right after the reconstruction. In fact, as the particle distribution contains a large amount of ghosts, we precisely use this step both to refine the position of the true particles, and to filter out some of the ghosts.

The global optimization is built around criterion $\mathcal{C}$:

$$\mathcal{C}(\Theta) = \sum_c \sum_x \left\| I_c(x) - \sum_p E_p h(x - \mathcal{F}_c(\mathbf{X}_p), \sigma) \right\|^2 \quad (1)$$

where $I_c(x)$ is the gray level of pixel x of the image acquired by camera c , $\mathcal{F}_c$ its projection function obtained from calibration, h is the PSF, modeled in the exact same way as in the PVR-SMART reconstruction (i.e. Gaussian of standard deviation σ integrated over the pixel surface), and $\mathbf{X}_p$ and E_p are respectively the position and intensity of particle p . Notation Θ refers to the variables to be minimized, i.e. collects $\mathbf{X}_p$ and E_p for all $p = 1 \dots P$ particles. Note that, although this will not be done in the following, the PSF of each particle can also be added to the parameter Θ . $\mathcal{C}(\Theta)$ is thus the global residual of the problem, since it is the squared sum of the residual images, i.e. the difference between the raw images, and the images formed by the projection of the particle cloud

estimate. Additionally to the first difference mentioned above, i.e. the step (among the different algorithmic phases) at which the optimization is performed, another novelty compared to Cornic et al. (2020) is the choice of a robust norm instead of the standard L2-norm in criterion (1)

In practice, starting from the initial set Θ provided by PVR-SMART and Gaussian fit, criterion (1) is minimized iteratively, here with the L-BFGS algorithm. It has been our observation that during the iterations, the intensities of ghost particles decrease, while that of the true particles increase, whatever the seeding density. This leads us to the second ingredient of GPCO, consisting in thresholding the particle distribution obtained after optimization, based on the intensity. This operation is in fact the first “ghostbuster” step of the VICCTOR algorithm: a well-chosen intensity threshold - hereafter referred to as “GPCO-threshold” - will indeed allow filtering a first amount of ghosts, while keeping a large majority of true particles. This is why GPCO is not only an optimization of the particle’s positions and intensities, but an optimization of *the whole particle cloud itself*.

2.1.3. Matching: F-VFC

To match particle clouds obtained at the two instants using the previous steps, we use the F-VFC matching algorithm, which we recently introduced (Le Bris et al., 2025). Similar to GPCO, this is not only a pure matching step, but also the second “ghostbuster” of the approach, as the specificity of F-VFC is to be robust to a large proportion of ghost particles, and to be highly efficient in filtering them out. We here briefly recall its working principle, and refer the reader to Le Bris et al. (2025) for more details.

F-VFC encapsulates a modified version of the Vector Field Consensus (VFC) method introduced in the computer vision domain by Ma et al. (2014). As this approach has been designed to be robust to a large proportion of outliers, we initialize it with a naive first guess, built by considering all pairing possibilities within a predefined search radius R_S . This is consistent with our global philosophy, since this guarantees that all true vectors (i.e. inliers) are present in the initial set (provided R_S is chosen according to the expected maximum displacement), notwithstanding the cost of a huge number of outliers. F-VFC then exploits the fact that, while outliers will have random orientations and magnitudes, and thus obey a uniform probability distribution, inliers will exhibit a spatial consistency and thus can be modeled as a Gaussian distribution around the local mean of the displacement field, denoted by $\bar{u}$. Given these modeling parameters, F-VFC then iteratively computes the probability of each vector to be either an inlier or an outlier, using the Expectation-Maximization algorithm, and by progressively refining the estimation of $\bar{u}$, sought on a basis of Gaussian functions (one per vector).

In short, F-VFC is an iterative process that looks at maximizing the rate of inliers kept, while discarding most outliers. As mentioned in Le Bris et al. (2025), its success also relies on a double-pass structure, using different stopping criteria depending on the pass. In the VICCTOR context, we

keep this exact same structure, except when we perform the very last step of the whole algorithm, on the last residual images (see section 2.2). The fact that F-VFC is now part of a larger algorithmic chain does also not change the way its optimal tuning parameters can be found; as described in Le Bris et al. (2025), aside from the search radius R_S , which depends on the maximum displacement expected, these parameters can be chosen easily based on minimum physical knowledge, with the most influential of them, the width of the Gaussian kernel used for $\bar{u}$, being based on seeding density.

2.2. Iterative scheme and convergence

As depicted in Fig. 1, once the three preceding operations -reconstruction, optimization and matching-, have been performed on the acquired images, the VICCTOR algorithm then enters a loop structure, where these steps are repeated (with slight modifications which we will highlight below), but on residual images. Exploiting these residual images, which are simply built as the difference between raw images and images obtained by projection of the current particle cloud estimate (similar to the way we build the GPCO criterion, see section 2.1.2), is a traditional way of ensuring to retrieve as many true particles, especially at high seeding densities (see Novara et al., 2023). The use of such a step exploiting residual images might appear paradoxical, since our method is targeted to operate at high Recall (maximum number of found particles) throughout. While our other algorithmic choices (PVR-SMART, F-VFC) are indeed in line with this strategy, in practice, they cannot maintain an optimal (i.e. unity) Recall, and will rather maintain a high value. Besides, as e.g. outlined by Novara et al. (2023) as well, their interplay can also slightly decrease Recall. Indeed, after the two PVR-SMART reconstructions, the missed particles have mostly no reason to be the same between the two instants, so that filtering out particles left unmatched by F-VFC will decrease Recall. This is why we similarly rely on iterations on the residuals.

Once the residual images have been formed, (with all negative values set to zero), a reconstruction step is used to find new particle candidates. These new particles are then appended to the current particle clouds, which are then optimized with the GPCO step. Here, it is important to emphasize that previous optimization information of the already available particles is only used as guess value, since these particles are re-optimized within their new set. Then, matching with F-VFC is applied, also with no preliminary knowledge, i.e. without exploiting the matches found for the initial particle set; here as well, we found that this way of functioning guaranteed the optimum performance. When matching is completed, new residual images are generated and the process is repeated. While, for nearly all iterations, we use F-VFC in the standard way described in section 2.1.3, an exception is made for the last iteration, for which we chose to slightly modify its tuning. Indeed, in highly turbulent flows, where consensus can be hard to find, it has been found that the double pass F-VFC tends to discard a small proportion of true vectors. At this point though, since most particles have been already been paired, there are close to no ghost particles in the dataset.

We thus found that running the last VICCTOR iteration with only the first F-VFC pass, and then simply removing all least probable duplicate vectors was efficient in recovering the few missing vectors.

In practice, this iterative loop is convergent, and does not add a high computational cost. Indeed, the residual images are less dense than the original input images, making particle reconstruction easier. PVR-SMART can still be expected to reconstruct close to all correct particles existing in the residual images, although with much fewer ghost particles. This, in turn, makes both the optimization and the matching step much easier. This convergence can be observed on the number of vectors obtained, which, after enough VICCTOR iterations, tend to be nearly constant or slightly oscillate. In a majority of cases, exploiting this behaviour is sufficient to stop the computation and obtain the final result. In other, more exceptional cases, it has been observed that the number of vectors could remain roughly constant over the iterations but the performance still increases: in such situations, monitoring the corresponding Eulerian interpolated velocity fields has been sufficient to diagnose if they were still affected by outliers, or if they could be deemed sufficiently converged.

2.3. Guidelines

We provide in Table 1 several guidelines for parameter tuning of the different blocks of the VICCTOR algorithm. These guidelines were established over extensive testing over several synthetic datasets. We also provide in the table the exemple of the values used for the datasets of the 2nd LPT Challenge presented in the following section.

3. LPT Challenges

We now challenge VICCTOR using the 3D two-pulse synthetic data from the first (Leclaire et al., 2021; Sciacchitano et al., 2021) and the second LPT Challenge. As can be seen on the portal hosting the data and results (<https://www.onera.fr/flow-benchmarks/>, see Leclaire et al., 2022), performance metrics in terms of particle precision and recall, vector precision and recall, as well as the mean squared error of position and displacement are documented for each published result. Although we participated in the second challenge in 2025, regarding the first challenge, which dates back to 2020, we simply used the portal's feature allowing us to upload a result, obtain an evaluation, and publish it so that it could be added to the results tables and graphs. This was also done in 2025, as shown by the information available on the portal. The first LPT challenge included a density scan ranging from 0.005 ppp to 0.160 ppp. The second challenge included a similar scan, ranging from 0.025 ppp to 0.250 ppp. Images with added Mie scattering effects, as well as variations in the ratio between the volume seen by all 4 cameras and the volume seen by at least 1

	Parameter	Guideline	Value for 2nd LPT Chal- lenge
Image pre-processing	Threshold	Mean value of back-ground noise	50
Large-Scale PVR-SMART	# of iterations	100	100
Model PSF	σ_{PSF}	Estimated PSF	0.6
GPCO	# of L-BFGS iterations	Check convergence of the criterion	200
	Threshold	Tune at lowest seeding density if possible, or set at minimum grey value of particle peak	175
F-VFC	Search Radius R_S	Depending on the maximal estimated inter-frame displacement	22 voxels
	Other parameters	Detailed in Le Bris et al. (2025)	

Table 1. Suggested optimal values for VICCTOR’s user parameters.

camera (denoted “I/U,” for Intersection/Union), were also generated for $\text{ppp} = 0.08$. As seen on the web portal for the First challenge, and in Fig. 2 for the Second, the VICCTOR algorithm reliably reconstructs and matches particles quasi perfectly up to a seeding density of 0.200 ppp. Very few False Positives and Negatives remain, and the obtained position RMS error is remarkably low, comparable to the DLR implementation of Two-Pulse Shake-The-Box (Novara et al., 2023), which can be pinned on the accuracy of the GPCO step. Furthermore, in the case of the I/U sweep at 0.080 ppp, also shown in Fig. 2, it can be observed that, while VICCTOR maintains very high Precision and Recall at all I/U values, it obtains a lower or even much lower position (and displacement) RMS error than the other tested algorithms. This as well is probably to be ascribed to the efficiency of the GPCO step, and validates the philosophy of VICCTOR, which was designed to minimize particle localization errors. As shown by the complete set of curves, it seems that this very low localization error presently comes at the cost of slightly less found particles and slightly more ghosts than for the competitors, although Precision resp. Recall are systematically maintained above 0.94




Fig. 2. Results from the ppp (top) and I/U (bottom) sweeps of case 2 of the 2nd LPT Challenge. Figure appearing on the related web portal <https://www.onera.fr/flow-benchmarks/>, as of 01/21/26.

resp. 0.92. This observation could motivate directions for possible future developments.

While we here mostly report on the ppp and I/U sweeps of case 2, note that the Challenge web portal <https://www.onera.fr/flow-benchmarks/> also hosts further results which may help the reader compare VICCTOR to other state-of-the-art algorithms. The same case 2 of the Second LPT Challenge also investigated the impact of Mie scattering at 0.08, both at perfect and complex I/U ratios (resp. 0.99 and 0.46). For both these I/U values, VICCTOR keeps the lowest position error when Mie scattering is activated, with also either the best or second best performances in terms of Precision and Recall. This confirms its overall robustness, as its building model (1) cannot account for different Mie regimes among the cameras, since each particle is represented by a unique intensity E_p , and the Point Spread Functions h of all cameras are taken identical, with same standard deviation σ .

4. Experimental test case

In this section, we assess the performances of VICCTOR on a von Kármán flow experiment. To do so, pairs of images are extracted from a full time-resolved acquisition. We then run the VICCTOR algorithm on these pairs of images and assess the obtained displacement field with respect to results provided by the LaVision Davis 10.2 implementation of Shake-the-Box (STB) on the full

time-resolved dataset. Because STB benefits from the knowledge of the full time-resolved acquisition, we assume we can consider its results as a reference for our comparison.

4.1. Experimental setup

A turbulent flow is generated inside the CEA SPEC's Giant von Kármán platform (Cheminet et al., 2022; Geneste, 2023). The platform has a radius of 0.5 m and a height of 0.8 m. The scale of these dimensions allows us to optically access the smallest scales of the flow, even down to scales smaller than those of Kolmogorov. The flow is generated by two counter-rotating, face-to-face impellers. The turbines act as a fluid pump, setting the fluid in motion Debye (2019); Dubrulle (2019). We rotate the turbines at a frequency $f = 0.1$ Hz. The fluid used is water, maintained at a constant viscosity $\nu = 1.0 \times 10^{-6} \text{ m}^2 \cdot \text{s}^{-1}$ by keeping it at a constant temperature of 20 °C using cooling circuits located above and below the two turbines. The Reynolds number of the resulting flow is then $\text{Re} = \frac{2\pi R^2 f}{\nu} = 1.57 \times 10^5$ and its Kolmogorov scale is $\eta = 0.18$ mm. To seed the flow, we use polystyrene beads manufactured by Spherotech, which are virtually monodisperse: their diameter ranges from 5.0 to 5.9 μm with an average of 5.33 μm . The density of these particles is 1.05 kg/m³, and their Stokes number in the flow, which we estimate using the Kolmogorov time scale ($\tau_\eta = 0.02$ s), is 9.1×10^{-5} . This ensures that we have passive tracers.

With a 30 mJ pulse high speed Nd-YLF laser, we illuminate a $50 \times 44 \times 6 \text{ mm}^3$ volume located at the centre of the volume (both in the horizontal and vertical directions) of a longitudinal plane of the von Kármán mixer. The laser has its beam waist located 900 mm from the 8 mm diameter output. To image the measurement volume, we use four Phantom Miro M340 cameras, equipped with CMOS sensors of 2560×1600 pixels with 10 μm pitch, resulting in a voxel size of 30.2 microns. Cameras 1 and 4 (resp. 2 and 3) are placed at a $\pm 67.5^\circ$ (resp. $\pm 45^\circ$) angle relative to the laser. We observe the particles through forward scattering, which means that the perceived intensity of the particles varies widely with the different angles of observations of the cameras. ND8 filters are positioned in front of Cameras 2 and 3 to homogenize the average intensity on all four cameras. Cameras are equipped with an $f = 200$ mm lens, set at an aperture number $f_\# = 11$. Taking into account the magnification of the optical setup, which we estimated to be around 0.37 Geneste (2023), this leads to theoretical particle image diameters $d_\tau \approx 2$ pixels (i.e. theoretical $\sigma_{PSF} = 0.5$ pixel). As shown by sample images in Fig. 3, despite the filters, the different observation angles among the two groups of cameras lead to different global aspects, as well as to inhomogeneities in the image density. This can be understood by the fact that cameras 2 and 3's lines of sight have a larger intersection with the illuminated volume than cameras 1 and 4. This also leads to a ratio between the intersection and union volumes viewed by the four cameras of $I/U = 0.7$, which is significantly smaller than 1, similar to case 2 of the 2nd LPT Challenge (see section 3). Therefore, similar to this case, and alongside the always more difficult nature of experimental




Fig. 3. (Top) The Giant von Kármán facility (left) and experimental set-up (right) at CEA SPEC. The four cameras observe a $50 \times 44 \times 6 \text{ mm}^3$ volume at the center of the von Kármán. (Bottom) Sample particle images from cameras 1 (Left) and 2 (Right), rescaled between 0 and 150 counts.

images, the present images will bring further characterizations of how VICCTOR behaves in the presence of Mie scattering and when images contain particles which are not viewed by all cameras, both these elements acting as model errors. Another reason why we emphasize the presence of these experimental difficulties is that, consistent with all results of the I/U and Mie explorations of case 2 of the Second LPT Challenge (not only VICCTOR's), they still concur to date to lower the maximum seeding density at which a satisfactory final accuracy is still obtained. This explains in particular why the following images will investigate a somewhat smaller ppp range than other experiments (e.g. with LEDs, and perfect I/U ratios), i.e. up to roughly 0.07 ppp.

Each measurement sequence is made of 3226 consecutive time-steps acquired in a time-resolved way with an acquisition frequency of $f_{\text{acq}} = 1.2 \text{ kHz}$. The acquisition frequency was tuned so as to get an average image particle displacement of 5 pixels. This leads to a total of 24 time steps per Kolmogorov timescale. Calibration was done using a LaVision 3D double-sided dual-plane calibration plate.

To provide a reference for our experimental evaluation of the VICCTOR algorithm, we obtained time-resolved PTV data through the Shake-The-Box (STB) algorithm in LaVision Davis 10.2 software. After calibration and volume self-calibration, images are pre-processed with a spatial filter that removes the local sliding minimum (computed over a square of side 11 pixels), in order to reduce background noise, before running STB. In order to ensure that STB results are the most trustworthy, we ran an extensive parameter study on the 2D detection threshold for IPR, which holds the most impact on the outcome results, while all other parameters of STB were set at the value recommended by the Davis user manual. The typical gray level value of the maxima of particle images was found to be in the range of 200 – 300 counts in the center of the particle images, but much lower on the edges, due to the above mentioned geometric and Mie scattering effects. Accordingly, we found that for thresholds varying from 30 to 65 counts, we obtained very similar results, although with slightly varying reconstructed number of trajectories, with the number increasing when decreasing the threshold. This variability will appear as an error bar on the es-




Fig. 4. Close up on a 50×50 pixel zone of images obtained on the GVK platform at varying densities (camera 4), estimated to be (from left to right): 0.019 ppp, 0.054 ppp, 0.080 ppp.

estimated number of trajectories reconstructed by STB. Note that STB results presented hereafter in the visualizations are those obtained with a threshold of 40 counts.

To estimate the seeding density of the images, we used the number of particles found in the STB reference, as this was more precise than a simple count of local maxima. Because cameras angles differ between cameras 1/4 and cameras 2/3, apparent seeding densities also differ widely between them. For instance, at the highest seeding density tested, we measured values around 0.057 ppp for cameras 2 and 3 and 0.080 ppp on cameras 1 and 4 (as seen in Fig. 4), leading to the average value 0.067 ppp. In the following, except mentioned otherwise, the indication of seeding density in ppp will refer to the average over the four cameras.

4.2. Processing parameters

For each sequence at a specific seeding density, after having obtained our reference via the STB processing, we then single out seven different pairs of neighboring instants, to process them with the VICCTOR algorithm. To do so, we first process the calibration plate images with our in-house calibration methods and perform a self-calibration derived from Cornic et al. (2016), in order to obtain a calibration model for the volume seen by the 4 cameras. The measured reprojection errors are estimated around 1.7 pixels before the self-calibration, and 0.11 pixels after, with a minimum error of 0.012 pixel and a maximum error of 0.31 pixels. Then, we use the VICCTOR algorithm with the following parameters:

- Reconstruction with PVR-SMART: while the theoretical σ_{PSF} is of 0.5 pixel, observation of the images and PSF calibration using a dedicated approach (by a fitting process on sample particle images) rather shows that the effective PSF is smaller, this being an effect of the signal-to-noise ratio, leading to the edges of the erf function to be within the image noise for the least bright particles. Consequently, we chose to set $\sigma_{PSF} = 0.35$ pixel, according to our calibration.

- GPCO step: we set the number of L-BFGS iterations to 300. Using reconstructions at the lowest densities as in section 2.3, we estimated the adapted value of the GPCO to be of 50 counts. Interestingly, this value is quite close to the optimal range of thresholds for the STB processing.
- Matching with F-VFC: we simply adapted the search radius R_S as a function of the expected maximum displacement, and the width β of the Gaussian kernel in the local mean flow $\bar{u}$ using the seeding density estimation.

4.3. Results




Fig. 5. Number of tracked particles in a reference sub-volume, averaged over the seven different time steps for each tested run. These number have been computed over a 250×250 pixel window. We show results from VICCTOR, and STB reconstructions. We indicate with a red bar the margin of error for the STB estimate, obtained from results with 2D particle detection thresholds of 30 and 65. We also indicate an estimate of the seeding density (in ppp) for each run, above its corresponding bar plot: this seeding density was computed using results from STB with a 2D particle detection threshold of 40.

First, to evaluate performance, we compare the number of particles reconstructed by VICCTOR and the STB baseline model, respectively. Given that the geometry of our setup yields an I/U ratio of approximately 0.7, and since the Davis 10.2 user manual does not specify whether only the intersection volume, or also parts of the union (and if so, which ones), are reconstructed, we focus on a subregion that, when projected, corresponds to a 250×250 -pixel window in the images. This allows for a more fair comparison, as it is certain that VICCTOR and STB reconstruct at least the intersection. Figure 5 shows the corresponding average number of particles generated by each algorithm. The values above the histograms also indicate the corresponding seeding density, calculated from the STB results. We observe that, up to the highest seeding density tested,

VICCTOR produces particle counts very similar to those of STB. As mentioned in section 4.1, the red error bars were obtained by varying the IPR's 2D particle detection threshold in STB. In total, at the highest seed density tested, STB tracks approximately 135,000 particles in the measurement volume. We then compare the corresponding flows obtained by each algorithm. As seen in Fig. 6, where we show linear interpolations of the displacement fields, our two-pulse algorithm reconstructs very faithfully the flow field. We find that, up to ppp 0.067, our algorithm yields a flow field that resembles the reference STB results, with very little noise added. This points to the success of the VICCTOR algorithm when it comes to dealing with high seeding densities in experimental two-pulse applications and is all the more remarkable because here we compare it with a measurement exploiting the full temporal context.

To further investigate the performances of the VICCTOR algorithm, we now compare, in Fig. 7, the raw, projected and residual images for one of the two pulses of camera 2, at the lowest and highest seeding densities shown here. In the left column (raw images), blue dots identify the projections of the particles found in VICCTOR's reconstructions, confirming that a large majority of the particles are retrieved, and thus the overall performance of the algorithm. Comparison between the three series of images then allows to identify some limits and thus possible future developments. We spotted two different types of situations, which are highlighted respectively in red and green in the raw and residual images. Red circles identify isolated particles (in the sense that their images do not seem to overlap with another particle's) which are missed by the algorithm. Within the three ones that were singled out in the figure, the particle on top illustrates a first distinct situation. Indeed, it only spreads over one pixel, i.e. has a very small effective PSF. Except in the context of time-resolved measurements, these particles are known to be extremely complicated to recover: since they spread over a single pixel, residual self-calibration errors may result into the crossing of epipolar lines outside the single illuminated pixel, meaning here that the particle can never be reconstructed by the initial PVR-SMART step. The bottom two missed particles, on the contrary, exhibit images which spread over several pixels, which is usually a favorable situation for their reconstruction. Therefore, is it probable that these particles are only viewed by a subset of the cameras and not all four of them, which points towards one of the possible future developments of VICCTOR, i.e. to account as well for particles viewed by only three or even two out of the four cameras. The other category of bright pixels remaining in the residual images, of which examples are identified by the green circles, rather seem to correspond to zones with overlaps of particle images. The fact that particles reconstructed by VICCTOR do not lead to a full extinction on the residual images then probably has to be ascribed by inaccuracies in the estimation of their intensities, and thus to the effect of Mie scattering. This observation could as well motivate some of the further development steps for VICCTOR.




Fig. 6. Snapshots of linear interpolations of the velocity magnitude obtained with VICCTOR (Left) and STB (Right) for three different seeding densities (Top) 0.015 ppp (Center) 0.053 ppp (Bottom) 0.067 ppp.

5. Conclusion

We have proposed a two-pulse PTV algorithm called VICCTOR (short for “Voxel-grid Initialized reConstruction with Consensus Tracking and Outlier Removal”), which is based on a ‘ghostbuster’




Fig. 7. Sample raw (left), synthetic image generated with the reconstructed particles and their intensity (center) and residual images (right), at the lowest (top) and highest (bottom) densities tested (camera 2). In the raw images, blue dots $\bullet$ correspond to projected positions of particles reconstructed by VICCTOR. We also single out with red circles some particles in the raw images that have not been reconstructed and remain in the residuals, as well as in green circles some particles that were reconstructed but possibly with insufficient intensity, as some bright pixels remain close to their positions in the residuals.

type of strategy—that is, aiming to recover as many real particles as possible in the early stages, and then gradually filtering out the associated ghosts. The voxels are first reconstructed at both time points using the PVR-SMART approach, then subjected to a three-point Gaussian fit to provide an initial estimate of their location and intensity. These estimates are then refined using global optimization: after convergence, we also take advantage of the fact that ghost particles have a (much) lower intensity than real particles, and we filter out a portion of them using an appropriate threshold. These two steps, optimization and filtering, are grouped under the acronym GPCO (Global Particle Cloud Optimization), as they improve the quality of the particle distributions as a whole. Pairing and additional outlier removal are then performed by the recently introduced F-VFC algorithm (Le Bris et al., 2025). This is followed by an iterative loop, with the above steps repeated until the number of vectors converges on the residual images.

The robustness of the VICCTOR algorithm to very high seeding densities was confirmed on data from the First and Second LPT Challenges. With seeding densities up to 0.200 ppp, we recov-

ered close to all particles with very few false positives, while equaling or improving the position RMS error, depending on the cases, of state-of-the-art 3D-PTV methods. Finally, the VICCTOR algorithm was tested on experimental data acquired on the Giant von Kármán platform at CEA SPEC. Results were compared to reference results obtained using the Shake-the-box algorithm on full time-resolved sequences. We found that, up to the highest seeding density tested (around 0.067 ppp), VICCTOR obtained displacement fields very similar to those of STB, and a very similar number of particles.

The excellent performance achieved by the VICCTOR algorithm in all tested cases leads us to believe that it is a state-of-the-art solution in the field of 3D PTV algorithms, with the added benefit of reduced localization error on the reconstructed particles. Its relatively small number of user parameters, all of which can be estimated with very little prior knowledge, as well as its robustness against very high seeding densities and real experimental noise, are all highly desirable characteristics.

Conflicts of interest

The authors acknowledge that there are no known competing financial interests or personal relationships that have influenced the work presented in this article.

Acknowledgments

This work was funded through ANR EXPLOIT, grant agreement no. ANR-16-CE06-0006-01 and ANR BANG grant agreement no. ANR-22-CE30-0025 . The authors would like to thank Pierre Bragança, Christophe Cuvier, Jean-Philippe Laval and Jean-Marc Foucaut from Laboratoire de Mécanique des Fluides de Lille (LMFL), as well Damien Geneste, Ariane Barlet, Vincent Padilla, Benjamin Musci, Cécile Wiertel-Gasquet, Bérengère Dubrulle and François Daviaud from CEA Paris-Saclay, CNRS, for their help and expertise throughout the several months of experimental measurements on the GVK Platform.

References

- Atkinson, C. and Soria, J. (2009). An efficient simultaneous reconstruction technique for Tomographic Particle Image Velocimetry. *Experiments in fluids*, 47(4):553–568.
- Champagnat, F., Cornic, P., Cheminet, A., Leclaire, B., Le Besnerais, G., and Plyer, A. (2014). Tomographic PIV: Particles versus blobs. *Measurement Science and Technology*, 25:084002.
- Cheminet, A., Geneste, D., Barlet, A., Ostovan, Y., Chaabo, T., Valori, V., Debue, P., Cuvier, C.,

- Daviaud, F., Foucaut, J.-M., Laval, J., Padilla, V., Wiertel, C., and Dubrulle, B. (2022). Eulerian vs Lagrangian irreversibility in an experimental turbulent swirling flow. *Physical Review Letters*, 129.
- Cheminet, A., Leclaire, B., Champagnat, F., Plyer, A., Yegavian, R., and Le Besnerais, G. (2014). Accuracy assessment of a Lucas-Kanade based correlation method for 3D PIV. In *17th International Symposium on Applications of Laser Techniques to Fluid Mechanics*, Lisbonne, Portugal.
- Cornic, P., Illoul, C., Cheminet, A., Le Besnerais, G., Champagnat, F., Le Sant, Y., and Leclaire, B. (2016). Another look at volume self-calibration: calibration and self-calibration within a pinhole model of Scheimpflug cameras. *Measurement Science and Technology*, 27(9):094004.
- Cornic, P., Leclaire, B., Champagnat, F., Le Besnerais, G., Cheminet, A., Illoul, C., and Losfeld, G. (2020). Double-frame tomographic PTV at high seeding densities. *Experiments in Fluids*, 61(2).
- Debue, P. (2019). *Experimental approach to the problem of the Navier-Stokes singularities*. Theses, Université Paris Saclay (COMUE).
- Dubrulle, B. (2019). Beyond Kolmogorov cascades. *Journal of Fluid Mechanics*, 867:P1.
- Elsinga, G., Scarano, F., Wieneke, B., and Oudheusden, B. (2006). Tomographic Particle Image Velocimetry. *Experiments in Fluids*, 41:933–947.
- Geneste, D. (2023). *Des évènements extrêmes de dissipation aux évènements extrêmes de dispersion, une empreinte des singularités : Approche expérimentale*. Theses, Université Paris-Saclay.
- Godbersen, P., Schanz, D., and Schröder, A. (2024). Peak-CNN: improved particle image localization using single-stage CNNs. *Experiments in Fluids*, 65(10):153.
- Jahn, T., Schanz, D., and Schröder, A. (2021). Advanced iterative particle reconstruction for Lagrangian Particle Tracking. *Experiments in Fluids*, 62.
- Lasinger, K., Vogel, C., Pock, T., and Schindler, K. (2019). 3D Fluid Flow estimation with Integrated Particle Reconstruction. *International Journal of Computer Vision*, 128(4):1012–1027.
- Le Bris, J., Leclaire, B., Cornic, P., Champagnat, F., Musci, B., and Cheminet, A. (2025). Consensus based tracking for 3D-PTV at high seeding densities. *Experiments in Fluids*.
- Leclaire, B., Cornic, P., Champagnat, F., Fabre, E., and Calmels, F. (2022). A web portal for automatic performance evaluation of Lagrangian Particle Tracking and Data Assimilation algorithms. In *LX Laser 2022-20th International Symposium on the Application of Laser and Imaging Techniques to Fluid Mechanics*.

- Leclaire, B., Mary, I., Liauzun, C., Péron, S., Sciacchitano, A., Schröder, A., Cornic, P., and Champagnat, F. (2021). First challenge on Lagrangian Particle Tracking and Data Assimilation: datasets description and planned evolution to an open online benchmark. In *Proceedings of 14th International Symposium on Particle Image Velocimetry*, volume 1. ILLINOIS Tech/Paul V. Galvin Library.
- Ma, J., Zhao, J., Tian, J., Yuille, A., and Tu, Z. (2014). Robust Point Matching via Vector Field Consensus. *Image Processing, IEEE Transactions on*, 23:1706–1721.
- Mercier, B., Thomas, L., Tremblais, B., and David, L. (2024). A robust pairing method for two-pulse particle tracking velocimetry based on coherent point drift. *Measurement Science and Technology*, 35(6):065301.
- Novara, M., Schanz, D., and Schröder, A. (2023). Two-Pulse 3D particle tracking with Shake-The-Box. *Experiments in Fluids*, 64.
- Sciacchitano, A., Leclaire, B., and Schröder, A. (2021). Main results of the first Lagrangian Particle Tracking challenge. In *Proceedings of 14th International Symposium on Particle Image Velocimetry*, volume 1. ILLINOIS Tech/Paul V. Galvin Library.
- Wieneke, B. (2012). Iterative reconstruction of volumetric particle distribution. *Measurement Science and Technology*, 24:024008.
- Yang, Y., Heitz, D., and Mémin, E. (2019). Lagrangian Particle Image Velocimetry. In <https://athene-forschung.unibw.de/128737> <https://www.unibw.de/ispiw2019/abstracts-paper>, pages 1–9.