Volumetric Flow Field Measurements of a Fixed-Wing Drone in Free-Flight

Vinicius Santos Costa[1,*], Jurij Sodja[1], Andrea Sciacchitano[1]

1: Delft University of Technology, Faculty of Aerospace Engineering, Delft, The Netherlands
*Corresponding author: V.M.SantosCosta@tudelft.nl

Keywords: 3D-LPT, OA-STB, Free-Flying, FSI, Forces from PIV.

ABSTRACT

A recent advancement in Lagrangian particle tracking, called Object-Aware Shake-The-Box (OA-STB), has enabled volumetric measurements to be taken with an object obstructing the field of view of the cameras (Wieneke & Rockstroh, 2024). With OA-STB, complex interactions unique to the flow around free-flying objects can be studied under more realistic conditions than what was previously achievable. To develop a robust strategy for free-flying experiments with OA-STB, measurements were taken on a 1.3m wingspan and 0.977m length model aircraft in levelled free-flight. This was done inside a wind tunnel where the drone was manually controlled and stabilized by an autopilot. Images were acquired using the twelve camera Omni-PIV setup (Hendriksen et al., 2025) with a volumetric illumination equal to 70 x 80 x 70 cm^3. By moving the drone inside the illumination, multiple measurements at different positions were stitched together, giving a measured volume of $0.99m^3$ within a 0.3 chord radius of the port side of the drone. Three aspects of the experiment are evaluated to assess the capabilities of free-flying measurements and to find areas of improvement. First, the model tracking provided by OA-STB is assessed. The best approach for tracking the drone is ascertained to be using time resolved STB. A partitioned camera system is used to account for the initial camera obstruction. By tracking both surface reflections and markers placed on the model, a positional uncertainty of $\pm 1.6mm$ and a rotational uncertainty of $\pm 0.28^\circ$ is found. The resulting tracking contained high frequency noise which is filtered out using a Butterworth filter. Secondly, the control strategy is evaluated using the in-flight kinematics of the model. Within a measurement of $4.7s$, the attitude varied by up to 0.5° from the mean in all components. Similarly, repeated measurements had deviations of the order of 1°. Significant improvements are required in the control to study specific conditions, as the obtained time average flow field is a mix of many different aerodynamic states. Lastly, an assessment was made of the resulting flow field data. Despite the averaging over multiple states, no discontinuities are observed in the flow field. The obtained data shows the potential of the technique, in this case enabling the characterization of the vortical structures around the wing and the propeller slipstream interaction with the tail.

1. Introduction

Over the last few years, significant developments have been made in the processing of 3D Lagrangian particle tracking (LPT) data. One of these, is an approach that accounts for camera obstructions due to the presence of solid objects inside the measurement volume. This approach,

commonly referred to as "Object-Aware" LPT (OA-LPT), was first introduced by Wieneke & Rockstroh (2024) for static objects, and later expanded to dynamically moving objects by Rockstroh & Brinkeman (2025). In OA-LPT, the obstructing object is equipped with fiducial markers whose positions can be obtained from the camera images. The fiducial positions are then used to register a CAD model of the object using the Iterative Closest Point (ICP) registration algorithm (Hendriksen et al., 2024). Finally, the registered CAD model allows for the determination of depth maps for each camera, which defines the line-of-sight distance of each pixel. Using the depth maps, particles are only reconstructed using cameras that can see the particle, thereby maximizing the number of cameras for reconstruction and minimizing the formation of ghost particles.

An area where the OA-LPT approach is particularly promising is in the determination of the flow field around free-flying objects. Past research in this area has predominantly focussed on visualizing structures inside the wake (Martínez Gallar et al., 2020) due to the limitation of object obstruction. In addition, precise positional control is difficult to achieve in free-flight, requiring the position of the model to be separately tracked for data analysis. OA-LPT removes these limitations by simultaneously accounting for the model obstruction and determination of model's position. By further combining OA-LPT with the state-of-the-art Shake-The-Box (STB) algorithm (Schanz et al., 2016), reconstruction of densely seeded 3D flow fields around free-flying objects can be achieved. This enables the characterization of (dynamic) forces in free-flight.

The necessity for measurements on free-flying models, as opposed to sting mounted ones, arises from the need to measure more realistic flight conditions to validate numerical simulations. In some applications, such as gust encounters, sting mounted models capture a different aerodynamic response than what an object in free-flight would experience (Sodja & De Breuker, 2022). This occurs by virtue of the fact that a free-flying model moves when a change in force is applied, unlike rigidly mounted models. While OA-LPT facilitates the data processing of free-flying measurements, the execution of these experiments remains a challenging task. Not only does it require for the acquisition system to image the model from all sides, but it also requires precise control of the model to ensure that it remains within the measurement domain. The control system must additionally not interfere with the flow imaging system and vice versa. Both aspects limit the available hardware and control strategies.

To develop a robust strategy for performing free-flying experiments, OA-LPT measurements were taken on a commercially available RR Easystar 3 model aircraft in levelled free-flight conditions inside a wind tunnel. The goal of these experiments was to find the best practices, limitations and validate the data acquisition and processing approaches. The lessons learned during these measurements will be used to improve the setup to enable for more complex measurements to be taken. In the present work, the experimental setup and data processing steps are described in detail. Herein, limitations of the chosen object-registration approach are highlighted. At the end, the obtained results of the flow field around the port-side of the drone is presented and discussed.

2. Experimental Setup

In this section, the experimental setup, drone model and control system are described. Experiments were performed inside the Open Jet Facility (OJF, 2.85 x 2.85 m^2 open test section) of the TU Delft. Measurements were taken at two free stream velocities, namely at 9 m/s and 10 m/s in levelled flight conditions. As the drone was in free-flight, the different measurements contained some degree of dynamic motion due to flight dynamics and adjustments of the control system. At 9 m/s the port half of the model was imaged in its entirety by stitching multiple separate measurements together. At 10 m/s, only the port side wing was imaged, along with approximately 1.5 chord up- and downstream. Figure 1 shows two photos of the drone in free-flight during the measurements. An overview of the different components used in the setup can be seen in Figure 2. Below follows a more detailed description of the components.



Figure 1. Photos of the Easystar 3 model aircraft in free flight during experiments. In (a), the model is shown in levelled flight prior to turning on the HFSB. In (b), the model is shown during a pitch-up manoeuvre.

The flow is seeded with 0.5 mm helium-filled soap bubbles (HFSB, Scarano et al. (2015)) generated with an in-house seeding rake situated in the settling chamber. The rake consists of 400 bubble generators, but only 200 were used to lower the particle density to improve tracking. A fraction of the generators (approx. 40) are non-functional, leading to gaps in the measurement. To evenly spread the HFSB, cylindrical turbulence generators were placed upstream of the rake. This led to a turbulence intensity of 1.43% inside the measurement volume. The measurement volume is illuminated using five LaVision flashlight-300 LED's, three placed on the bottom and two from the top, leading to an illuminated volume of approximately 70 x 80 x 70 cm^3 in X, Y, Z respectively.

For image acquisition a second version of the omni-PIV setup introduced by Hendriksen et al. (2025) was used. The omni-PIV setup consists of twelve cameras (LaVision Imager CX3, 5 MPx,

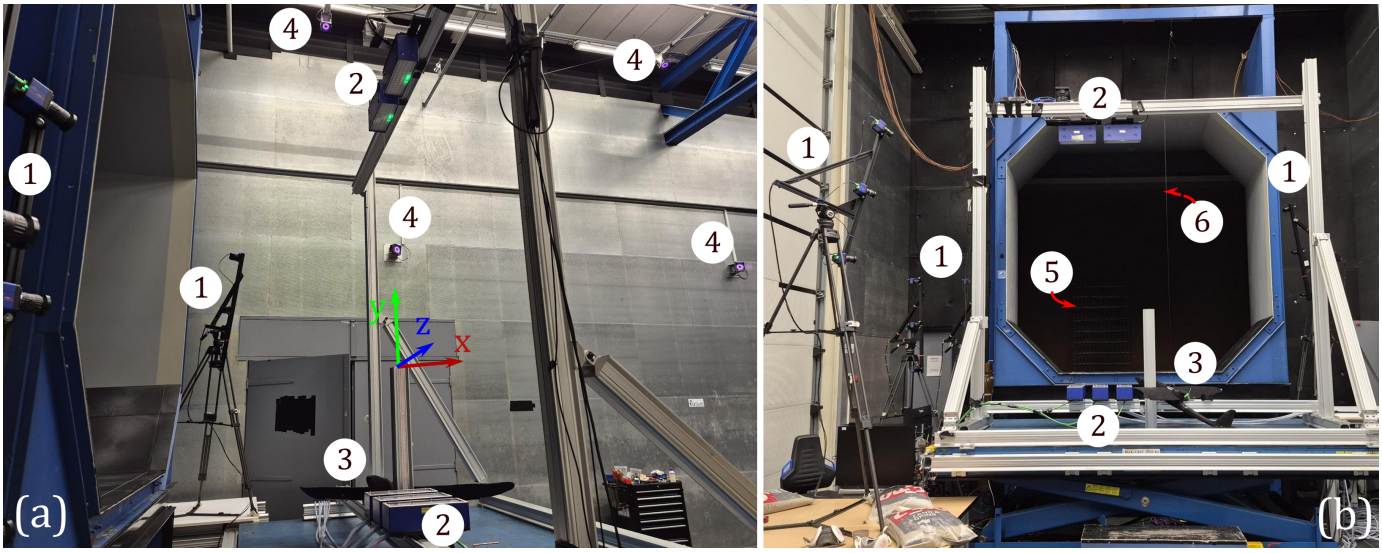


Figure 2. Photos showing a side view (a) and a back view (b) of the experimental setup. The equipment used in the experiment is numbered, with (1) indicating the camera tripods, (2) the LED's, (3) the model aircraft, (4) the OptiTrack cameras, (5) the seeding rake and (6) the nylon tether. Also shown is the coordinate system.

2.7 μm pixel pitch) recording in double pulse mode at 215.8 Hz, to obtain two-pulse STB tracks (Novara et al., 2023). Recordings were taken in 2x2 binned pixel mode, thereby reducing the effective resolution and pixel pitch to 1.25 MPx and 5.4 μm respectively. The cameras are mounted in sets of three on tripods, where they are arranged and pre-focused to measure at a $2m$ distance from the object using $12mm$ lenses at an f-stop of 8. This gives a measurement volume of 1.5 x 1.2 x 1.5 m^3 ($0.81Px/mm$) for a single tripod. By combining four tripods, the entire volume around the model is always in view by at least five cameras.

A commercially available foam model aircraft called the RR Easystar 3 is used. This model has a $1366mm$ wingspan (b), $977mm$ length (l), a $171mm$ mean aerodynamic cord ($\bar{c}$) and weighs 7.53 N. The model is tethered to the walls with 0.3mm thick nylon lines to initiate flight and to keep it from impacting against the tunnel walls. These lines were given sufficient slack to ensure free-flight inside the measurement volume.

The drone is painted black using Musou black to reduce reflections and equipped with two types of randomly distributed markers. The first are white paint dots that are used for object registration (OR) in OA-LPT, each around $2mm$ in size. Over 100 dots are spread over the model. Secondly, twelve square reflective tape markers of about 1 x 1 cm^2 are placed on the starboard side of the model outside of the PIV measurement volume. These are used to track the model using an OptiTrack motion capture system installed in the ceiling of the wind tunnel test section.

The model was controlled manually and stabilized by an autopilot. For this, a Pixhawk 4 flight controller was installed and connected to the open-source Paparazzi software. Here, an autopilot code refurbished from a previous experiment by Hwang et al. (2023) was used. The code uses

an Incremental Non-linear Dynamic Inversion based controller to stabilize the aircraft. Briefly explained, this controller attempts to maintain a constant zero angular acceleration for the drone. To do so, it takes IMU and Optitrack data as inputs to determine and correct for the angular acceleration. More details on the tuning and control theory can be found in the article of Hwang et al. (2023). The original autopilot had a position holding function which was not functional during these experiments, therefore requiring manual flight. It additionally used an airspeed sensor that was not used, as it was expected that the bubbles would interfere with the sensor.

3. Data Processing

Data acquisition and processing is performed using the Davis 11.3 software from LaVision GmbH. To support this, in-house codes were created in python and tecplot for further processing and visualization. The different processing steps are described below.

3.1. Calibration

The cameras were geometrically calibrated using a planar calibration plate with a pinhole model (Dawson-Howe & Vernon, 1994). Ideally, this would require a plate encompassing the entire field of view (FOV) of all cameras. Due to the impracticality of mounting and moving such a large plate, a smaller A2 sized plate was used. The plate was moved over the entire FOV of all cameras to ensure that a converged calibration was obtained. It was always positioned such that at least two different tripods view the plate, to link the different tripods together, requiring over 80 images for each geometric calibration. The tripods used in the experiment were not the most robust camera mounting, leading to small changes in position over time. To account for this, the geometric calibration was repeated at the start of each measurement day.

To further refine the geometric calibration, a volume self calibration (VSC, Wieneke (2008)) was performed multiple times during the measurements. It was done both at the start and in between measurements by taking images of just tracer particles without the model inside the flow. With this data, the particle image characteristics were further obtained by using the Optical Transfer Function (Schanz et al., 2012). Due to the non-robust mounting of the tripods, sub-pixel drifts in camera positions were seen over less than an hour of image acquisition. Once the model is registered, the VSC was repeated with depth maps in presence of the object. This was used to reduced the pixel uncertainties to below 0.1 for all tested cases.

3.2. Object Registration

For the Object Registration (OR) the structural markers have to be separated from flow tracers. For all cases, the movement of the model is much smaller than the displacement of the HFSBs. By taking the minimum pixel intensity over three to five consecutive images, most of the HFSBs are eliminated leaving images that only contain reflections, structural markers and some remnant intensity from the HFSB. To further improve image quality, a sliding minimum spatial filter (7 px) and normalization with respect to the first image (smoothed over 300 px) is used in conjunction with manual masking of the model.

When directly applying particle reconstruction to these images, hundreds of ghosts are created. In order to only track the motion of the surface markers and steady reflections, the first images of each image pair was extracted and processed via time-resolved STB (TR-STB). Before applying TR-STB, the data processing is split into a partitioned system using camera groups to account for object obstruction (similarly to Hysa et al. (2024)). One group is made for the bottom and one for the top side of the model. The reconstructed tracks are then filtered by track length and particle velocity, only allowing slow moving (<0.5 m/s) and long tracks (>10 time steps) to exist. Since the model is moving, the intensity of the structural markers and steady reflections can significantly change between consecutive images. Therefore, a track reconnection pass is used to ensure that as many markers are captured as possible. This process is visualized in Figure 3.

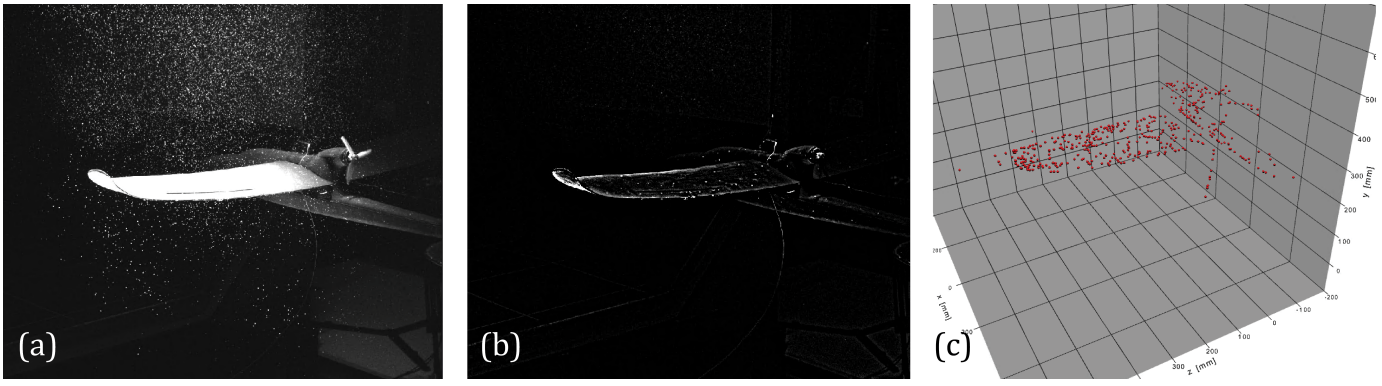


Figure 3. Process used to obtain the structural tracks. In (a) the raw image of one of the cameras is shown using a greyscale for the pixel intensity (0 to 300 counts). Image (b) shows the aforementioned image processing, without manual masking. Image (c) then shows the tracked markers and reflections, using time resolved STB.

The resulting particle tracks contain a mix of stationary reflection particles (referred to as SRPs) and real structural markers. The SRPs mainly form in spots where the paint deteriorates over time due to bubble collisions, such as the leading edge (LE), and along sharp outlines of the model, such as the trailing dege (TE). To start OR, the CAD is first manually placed inside the point cloud of the first time step. This step is crucial, as when it is skipped the ICP registration tends to fail due to registration ambiguities. Afterwards, ICP is started for each time step with a starting search radius of $50mm$ to account for model movement and rotation, and iterated 10 times while linearly

decreasing the radius to $2mm$. An example of the resulting position and attitude of the model over the duration of 1000 images in one of the levelled test cases be seen in Figure 4.

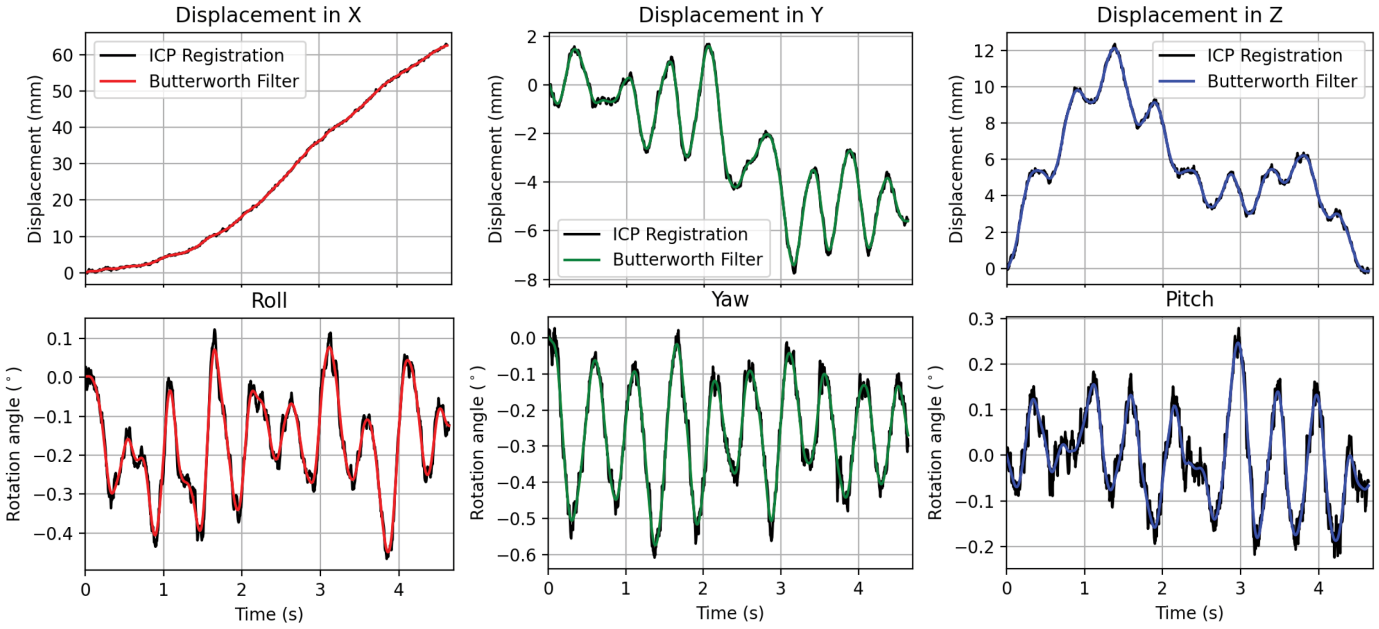


Figure 4. Rotations and displacements of the model over 1000 images obtained from ICP. The displacements and rotations are all given with respect to the position of the model in the first time step of the OR. Also shown is the filtered data using a 2nd order 5Hz forward and back Butterworth filter.

For all cases it was observed that the time-resolved ICP registration has high frequency noise in both the position and attitude of the drone, as can be seen in Figure 4. This noise likely originates from errors in the reconstruction of particle position, mismatches between the physical model and the CAD model and ambiguities in the registration of the object to the point cloud. To filter the high frequency noise, a low-pass forward and back Butterworth filter was used. Other filters likely suffice, here the choice for the Butterworth was made due to its maximally flat characteristic.

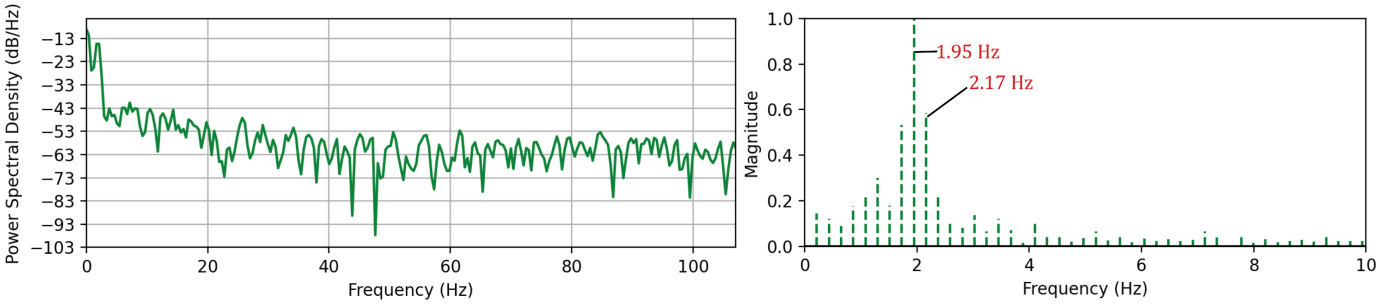


Figure 5. Spectral analysis of the yaw motion shown in Figure 4. On the left, the PSD can be seen whereas on the right the FT of the first 20 frequencies.

The order and the cut-off frequency of the Butterworth filter are chosen by analysing the power spectral density (PSD) and Fourier transform (FT) of the movements and rotations. For all rota-

tional components of all measurements a relatively large peak was found around 1.9 to 2.3 Hz, after which a sharp decrease was present. An example can be seen in Figure 5 for the yaw motion shown in Figure 4. Given that this peak is always present, it is likely associated with the response of the control system. A second order filter with a 5Hz -3dB cut-off ensures that this frequency is unattenuated, while eliminating most of the high frequency noise. The filtered position and orientation of the model over time are then applied for every time step to rotate and translate the CAD model such that the depth maps can be computed.

3.3. Objected Registration Uncertainty

As shown by Hendriksen et al. (2024), ICP registration can obtain sub-millimetre accuracy between the position of the structural markers and the CAD model of the object. For a wing, Hendriksen found an average mismatch below $0.3mm$. Assuming a similar uncertainty and that the markers are evenly distributed, this theoretically allows for a sub-millimetre positional uncertainty and a rotational uncertainty up to 0.05° for the drone. The latter was calculated assuming a mismatch of $0.3mm$ at the root and $-0.3mm$ at the tip of the wing ($0.6mm$ across half the span). However, this is only valid for rigid models where the CAD model accurately matches the real model.

To ensure that the CAD matches reality as closely as possible, the drone was 3D scanned using a Quantum^E FaroArm® scanner after the measurements. This scanner has achievable accuracies that are below μm, but it is expected that the actual accuracy of the scan was lower compared to free-flight. This is due to the fact that the model is made of relatively flexible foam, which is expected to have deformed slightly during flight. Deflections are expected to be worst at the tip of the model due to the integrated lifting load. Additionally, the control surfaces were scanned in a neutral position. The deflections of the surfaces were not recorded at the time of the image acquisition, therefore this leads to an additional mismatch between model and CAD. Lastly, the model experienced a couple of crashes during testing which introduced some dents on the surface. As the scan was made at the end, these dents might not have been present throughout.

Another source of uncertainty was observed to be caused by the TR-STB settings used in the model tracking. Lowering the minimum intensity in the TR-STB leads to both more SRPs and structural markers being detected. When tracking with a high intensity (>200 counts), less than half of the structural markers (<60) are tracked along with some SRPs (ranging from 10 to 90). Furthermore, tracks are constantly lost due to changing marker intensity as aforementioned. Registering the object under this condition (with <100 points), leads to a high noise (up to 1° between time steps) and high uncertainties ($\pm 0.5^\circ$) in the rotational components. The latter was determined by registering the object multiple times using a slightly different subset of points, which resulted in a bias difference up to $\pm 0.5^\circ$ for the rotational components.

Both the uncertainty and noise amplitude are reduced by lowering the intensity of the TR-STB

tracking, thereby obtaining more points. Using a very low intensity (20-30 counts) results in hundreds of SRPs (200-1000) being tracked while also reconstructing most of the structural markers (>100). With this approach, in between 200 to 700 points are used for the OR. The impact of the TR-STB tracking intensity for one measurement is shown in Table 1 and Figure 6.

Table 1. Root mean square error (RMSE) between the raw ICP motion and the filtered motion for the two cases shown in Figure 6. The values presented here are indicative of the differences that are obtained when using a high (>200 counts) and a low (20 counts) TR-STB intensity for the OR, but are highly case dependent.

Motion RMSE	X (mm)	Y (mm)	Z (mm)	Roll (°)	Yaw (°)	Pitch (°)
High Intensity	2.19	0.87	0.69	0.20	0.18	0.23
Low Intensity	0.24	0.12	0.11	0.02	0.02	0.02

Based on the RMSE values it can be seen that the noise is around a factor 10 higher when using a high TR-STB intensity for tracking and registration. The same also holds for the uncertainty which is reduced by roughly 10 times to about ($\pm 0.05°$) when using a different subset of points. Note that these values are highly case dependent and therefore difficult to generalize.

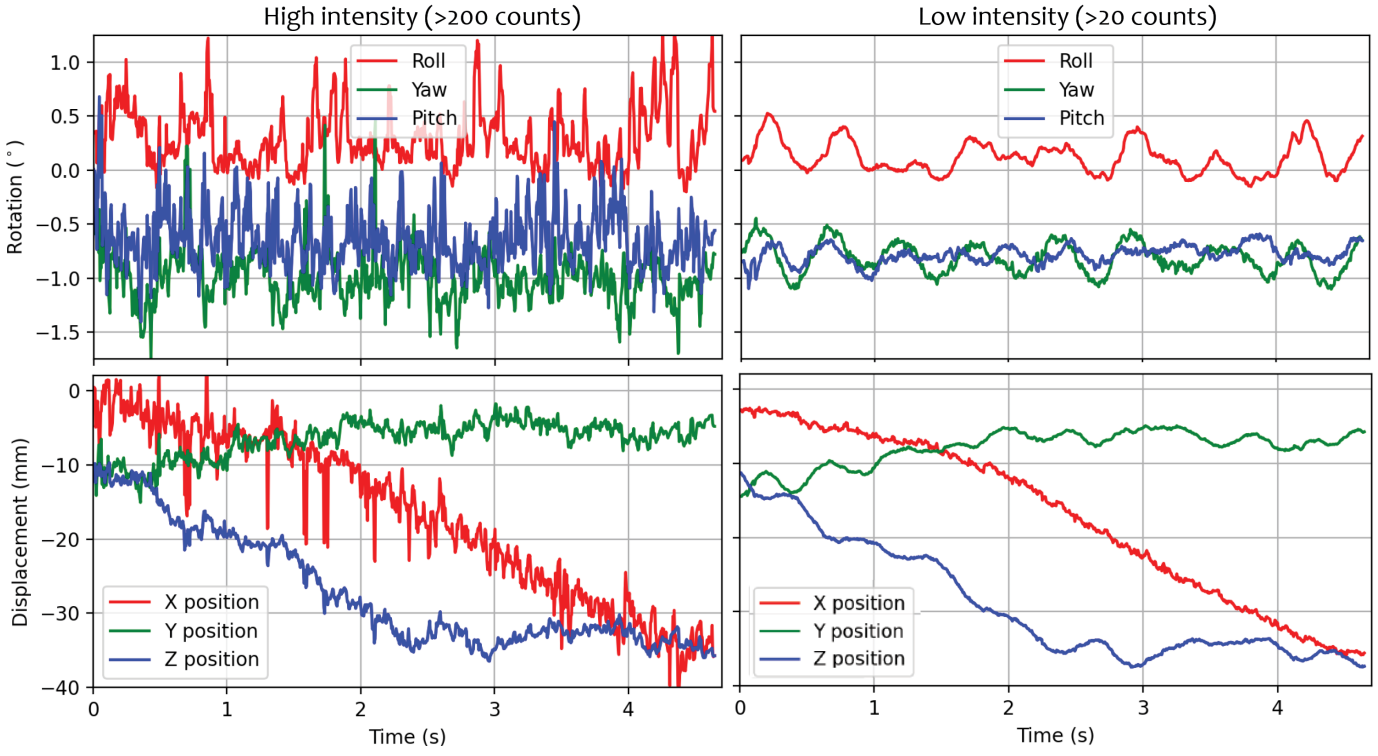


Figure 6. Comparison between the ICP results using a high (>200 counts, left) and a low (>20 counts, right) TR-STB tracking intensity over 1000 images. On the top the rotational components are shown and on the bottom the translational components. All components are given with respect to an arbitrary initial position of the CAD model.

The observed difference can be attributed to the number and distribution of the markers, which was previously also reported by Hendriksen et al. (2024). Having more tracks for the registration

adds constraints which reduces the registration ambiguity and the impact of random positional errors is mitigated. However, when using a low TR-STB intensity and accepting only tracks that are 500+ time steps (half the measurement time, roughly 50-100 tracks) an almost identical motion is obtained as when using all 200-700 tracks. This filtering includes both structural markers and very steady SRPs. The main difference between this case and the high TR-STB intensity case is that the points are more spread out, including dimly illuminated regions of the model. Given the fact that both low intensity approaches yielded the same motion, the unfiltered 200-700 tracks were used for the final registration to ensure that there are always sufficient markers for ICP.

The resulting registrations using the low TR-STB intensity have an average mismatch between the scanned CAD model and the tracked particles around $1.6mm$. Therefore a positional uncertainty of around $1.6mm$ is estimated, along with a rotational uncertainty of $0.28°$. This uncertainty is indicative of the accuracy of the mean value, as from Figure 4 and Figure 6 it can be seen that a much higher precision is obtained. From the reconstructed particle tracks it can indeed be seen that there was a mismatch near the tip, as gaps appear in regions where some cameras are obstructed by the tip. An example can be seen in Figure 7

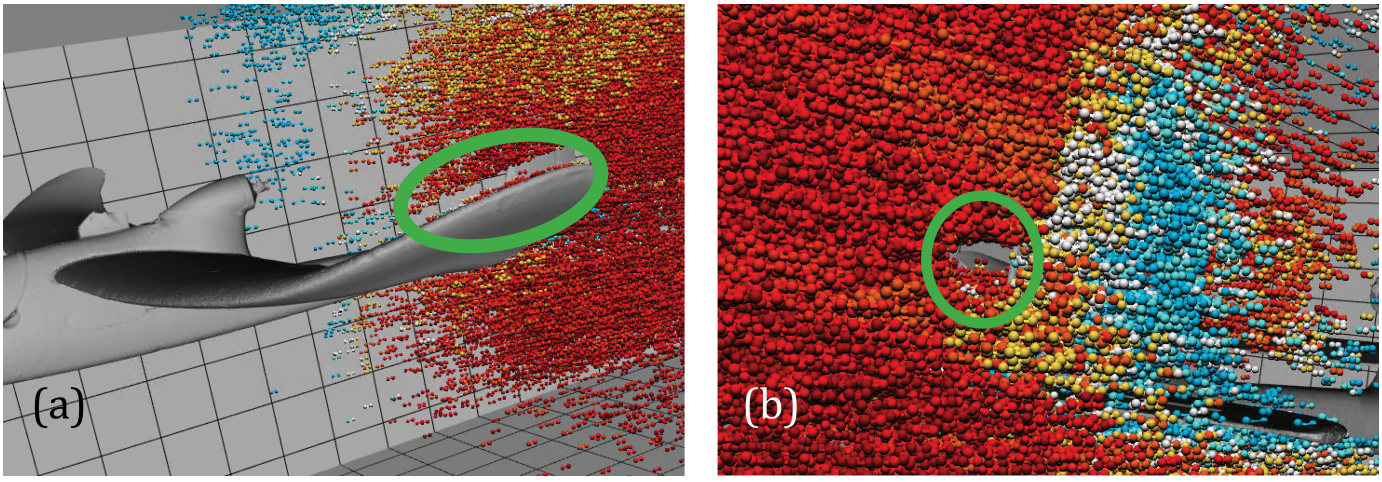


Figure 7. Gap in the measured particles caused by a misalignment of the tip of the CAD model with the real object. In (a) a front view can be seen, whereas in (b) a back view is shown of this gap.

3.4. Binning

With the depth maps obtained, two-pulse object aware STB (OA-STB) is used to reconstruct the flow tracers. Images are pre-processed using a minimum over time subtraction filter (5 images), a sliding minimum spatial filter (7 px) and a normalization with respect to the first image (smoothed over 300 px). Due to model movement, some structural markers and reflections remain after these operations. The reflections cause SRPs to form up to 10mm away from the surface. However, the vast majority of the SRPs are present within 3mm of the surface. In this region their number is in

the same order of magnitude as the number of real particles, leading to erroneous velocity vectors. By masking away any particles within 3mm of the surface of the model, the majority of SRPs are eliminated while not removing many real particles. An attempt was made to use a median filter, but to remove these SRPs requires a filter so strong that it also removes the tip vortex.

Next, the position of the obtained tracks for each time step are translated and rotated in space to a stationary CAD frame of reference. This is done to account for the model motion during the measurements. The rotations and translations are applied using a 4x4 transformation matrix ($T(t)$) that transforms from the registered model position to the CAD frame of reference. The matrix first rotates along X (ϕ), then Y (ψ) and finally around Z (θ) and is given by

$$T(t) = \begin{bmatrix} cos(\theta)cos(\psi) & cos(\theta)sin(\psi)sin(\phi) - sin(\theta)cos(\phi) & cos(\theta)sin(\psi)cos(\phi) + sin(\theta)sin(\phi) & X \\ sin(\theta)cos(\psi) & sin(\theta)sin(\psi)\sin(\phi) + cos(\theta)cos(\phi) & sin(\theta)sin(\psi)cos(\phi) - cos(\theta)sin(\phi) & Y \\ -sin(\psi) & cos(\psi)sin(\phi) & \cos(\psi)cos(\phi) & Z \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{1}$$

In this reference, the model is set at the average attitude it had with respect to the freestream during all measurements for a given condition. The tracks for all time steps are then combined together to obtain a time averaged particle field. At 9 m/s the port side of the model was imaged during 8 different measurements of 1000 snapshots each, and one of 2000 snapshots. Of these, 2703 snapshots were not included for the final binning due to the drone having a high velocity (>$0.5m/s$) or acceleration (> $1m/s^2$). Between 3000 to 13000 particles were tracked within 0.3 $\bar{c}$ of the model in each snapshot after masking away all particles within $3mm$ of the model. This gave an average of 7627 particles per time step. Most particles were tracked over the wing and in the slipstream of the propeller, as multiple measurements focussed on these regions. For 10 m/s, the flow was only imaged over the wing during 3 measurements of 1000 snapshots, giving an average of 9140 tracks per time step.

Next, the tracks are binned using a second order polynomial as described by Agüera et al. (2016). Bins of 36 voxel with 75% overlap are used. This corresponds to $3.77cm$ bins in physical space, or a vector spacing of $0.95cm$. From these, additional flow field properties such as spatial derivatives of the velocity vector and vorticity are obtained. For the 9 m/s case, a visual representation of the combined measured volume containing bins with >100 particles can be seen in Figure 8.

3.5. Circulation Based Lift Reconstruction

The Kutta-Joukowski theorem (Anderson, 2017) was used to obtain the lift over sections of the wing, similarly to the approach of Mertens et al. (2023). Here it is assumed that the flow field is quasi-steady, thereby neglecting any acceleration effects. The lift per unit span (L') is given by

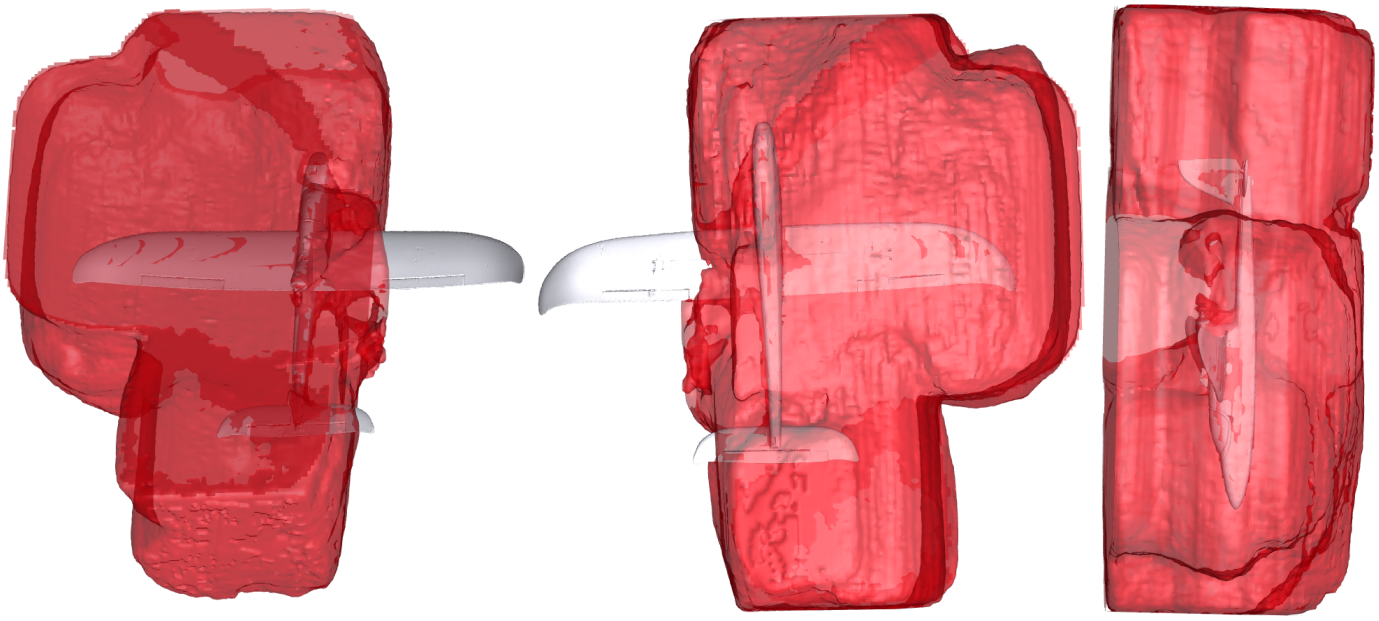


Figure 8. Top-back (left), bottom (middle) and side (right) view of the volume in which 100+ particles were detected over 7703 snapshots for the 9 m/s measurements. This volume is roughly 0.99 m^3.

$$L' = \rho U_\infty \oint_c \boldsymbol{u} \cdot \boldsymbol{ds} \tag{2}$$

where ρ is the fluid density, U_∞ the free stream velocity and the last term represents the contour integral of the velocity. Note that while it is originally derived for steady, inviscid and irrotational flows, Kutta-Joukowski theorem has been shown to work for experimental airfoil data and even for highly separated flows (see Lee & Su (2012) and Sharma & Deshpande (2012) for example). In the original theory, the calculated lift force is independent of the contour taken around the airfoil. For real flow fields, larger contours lead to a reduction in lift due to dissipation effects. Furthermore, experimental data contains noise leading to variations in the computed lift for different contours. To account for both effects, the lift was computed from multiple contours ranging from a distance of 0.17 $\bar{c}$ to 0.55$\bar{c}$ from the airfoil. Square contours aligned with x and y were used to simplify the mathematical implementation of the integral.

4. Results

4.1. Flight Kinematics

To assess the quality of the measurements that were combined together, this section describes the movement and rotations of the drone during the measurements. First, the mean aerodynamic angles are given and secondly the unsteady motion of the drone is discussed.

Aerodynamic Angles and Roll During Flight

As the flight data was not logged, the mean attitude of the model was an unknown of the measurement. Using the OR and the angle of the freestream, the angle of attack (α), sideslip (β) and roll (ϕ) angles were reconstructed. Mean results and their uncertainty are shown in Table 2. These results include the effective α and β due to the V,W velocity components of the drone.

Table 2. Mean aerodynamic angles and roll angle of the drone during measurements.

U_∞	$\alpha(^\circ)$	$\beta(^\circ)$	$\phi(^\circ)$
9 m/s	4.46 ± 0.63	1.28 ± 0.67	-0.47 ± 2
10 m/s	3.40 ± 0.63	1.48 ± 0.67	-0.91 ± 2

For α, the angle that chord line of the wing makes with the freestream was used. From $Z/(0.5b) = 0.83$, where Z is the distance along the midspan, the wing has both a sharp increase in dihedral and twist. This region was ignored when computing the average α. The total uncertainty of α is a combination of the OR uncertainty ($\pm 0.28^\circ$), the uncertainty in the chord line angle ($\pm 0.11^\circ$) and the uncertainty in the angle of the incoming flow ($\pm 0.24^\circ$). The uncertainty in the chord line was caused by the model being damaged and therefore not having a well defined LE and TE.

To estimate β, a line was drawn between the midpoint of the nose and the the propeller. The angle of this line with respect to the incoming flow is β. Here the total uncertainty is again a combination of the OR uncertainty ($\pm 0.28^\circ$), the uncertainty of where the midpoint of the propeller and nose are ($\pm 0.12^\circ$) and the uncertainty of the free stream angle ($\pm 0.27^\circ$).

Lastly, ϕ is the least accurate as the only straight line from which ϕ could be obtained was using the trailing edge from $Z/(0.5b) = 0.05$ to 0.3 as after $Z/(0.5b) = 0.3$ the wing has a non-zero dihedral angle. This line has a mean angle of $-0.47 \pm 0.3^\circ$. However, there was no horizontal reference anywhere in the images to estimate the roll angle with respect to the ground. The uncertainty in the roll is expected to be within the uncertainty of the initial angle of the calibration plate. This is estimated to be within 2° based on the position of the cameras for multiple calibrations.

As expected, α for 9 m/s is higher than for 10 m/s. The drone had the same mass and experienced the same air density ($\rho = 1.21 kg/m^3$) for the measurements of both airspeeds. Therefore the dynamic pressure was $49.09 Pa$ and $60.74 Pa$ for 9 m/s and 10 m/s respectively. From this, the 9 m/s cases needs a C_L that is 24% higher. Assuming a linear lift curve slope, the estimated value for α at 10 m/s is 3.38°, which almost perfectly matches with the measured value of 3.40°. This small difference can easily be explained by the measured difference in roll and yaw angle. Based on this, the relative difference between α for both cases can be assumed to be accurate.

Notable is that the drone flew on average at both a β and ϕ. This can be explained by damages the drone acquired early on. The tail of the drone was twisted and bent slightly towards the port

side. To compensate for this it likely needed to fly at some β and ϕ. As will be shown later, results indicate that the rudder had a large deflection. This probably compensates for the β and the warping of the tail.

Overall, the determination of the aerodynamic and roll angles is adequate to obtain an idea of how the drone was flying. While the uncertainty is relatively large compared to their values, these were extremely small angles. However, the uncertainty can be halved by optimizing the OR and the method used to obtain the angle from the model. In addition, synchronizing the autopilot data to the PIV acquisition can further mitigate the uncertainty.

In-Flight Motion

Figure 9 shows the computed aerodynamic and roll angles of the 9 m/s measurements. The angles vary by at most 1.5° from case to case excluding cases 6 and 7. Furthermore, within a given measurement, variations of $\pm 0.5^\circ$ are observed from the mean value of that measurement. The resulting flow field is therefore not an entirely accurate representation of the flow field at the mean aerodynamic angles, but a combination of different conditions centred around this mean. For the 10 m/s measurements, similar trends were observed as cases 1 to 4 of the 9 m/s measurements.

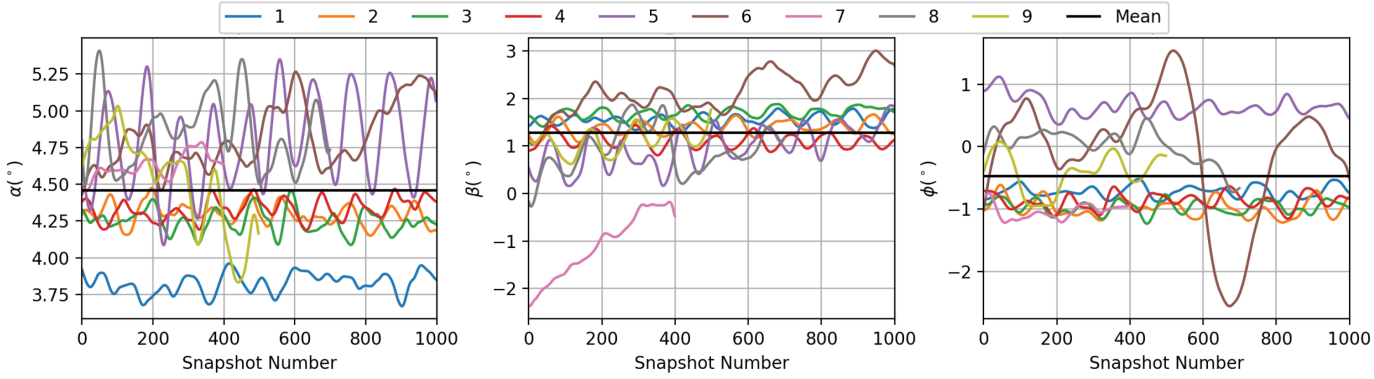


Figure 9. Computed α, β and ϕ of all 9 m/s cases during image acquisition. The numbers indicated in the legend represent the different cases. For clarity, the x axis is indicated with number of snapshots (recorded at $215.8 Hz$ rate).

Additionally, as the resulting flow field was stitched together from the different measurements, the flow field that was obtained has a spatial variation in mean aerodynamic conditions. For applications where precise angle control is required to characterize an interaction, this would be insufficient. These oscillations can however be improved upon. As can be seen from the most stable measurements (1 to 4), variations of $\pm 0.25^\circ$ were attainable with the current autopilot and manual flight combination. A fully functional autopilot would likely be able to provide a lower degree of oscillation with a higher degree of repeatability.

As previously discussed, the rotational components had a peak frequency around 1.9 to 2.3 Hz in

the PSD and FT. From Theodorsen theory (Theodorsen, 1949) the reduced frequency (k) of the flow is given by

$$k = \frac{\pi \cdot f \cdot \bar{c}}{U_\infty} \tag{3}$$

where f is the frequency of the motion. In the 9 m/s measurements this results in a k between 1.1 to 1.4. For these values the flow can not be assumed to be quasi-steady, and is instead described by unsteady aerodynamics. The measured particles at each time step are therefore affected by the history of the flow field, and lagging behind the aerodynamic angles in their response.

The drone was also constantly translating. For cases 1 through 4, velocities are within $4cm/s$ for all components. On the other cases, velocities are within $15cm/s$ with some instances reaching up to $30cm/s$. Instead, accelerations are mostly within $0.6m/s^2$, with the exception of cases 7 to 9, which reached peaks of $1m/s^2$. The magnitudes are relatively consistent between cases and can mainly be attributed to the $1.9 - 2.3Hz$ rotational motion. With a 0.768 kg mass the instantaneous change aerodynamic forces reached up to $0.46N$ to $0.77N$.

4.2. Reconstructed Flow Fields

In this section, a detailed description of the obtained flow field is given, demonstrating the capabilities of the measurements. Despite the relatively wide spread of aerodynamic angles of the binning data, a lot of details about the flow around the drone in free flight are observed. First, velocity fields are presented. Afterwards, vortical structures are discussed.

Velocity Fields

In Figure 10 two planes showing the U component of the velocity are presented, highlighting the impact of the propeller on the velocity over the wing. Starting on the left plane at $Z/(0.5 \cdot b) = 0.25$, some basic features of the flow over a wing can be seen. This includes the suction peak over the tail and the main wing, the wake of the wing and the stagnation point. The suction peak of the tail is on the lower side, indicating a negative lift force. The aircraft centre of gravity (c.g.) is set to be in front of the centre of pressure (c.p.) such that the drone is statically stable. To balance this, the tail needs to provide a negative lift force. Interestingly, it can also be seen that in this design the wake of the wing directly impacts the tail.

On the right plane at $Z/(0.5 \cdot b) = 0.11$ significant changes can be seen. This plane sits at a 10 mm distance from the tip of the propeller. A clear increase can be observed in the velocity on the top side of the wing which enhances the local lift. While the bottom side velocity also increases, the increase is lower than that of the top side. The wake of the wing is also much larger, further

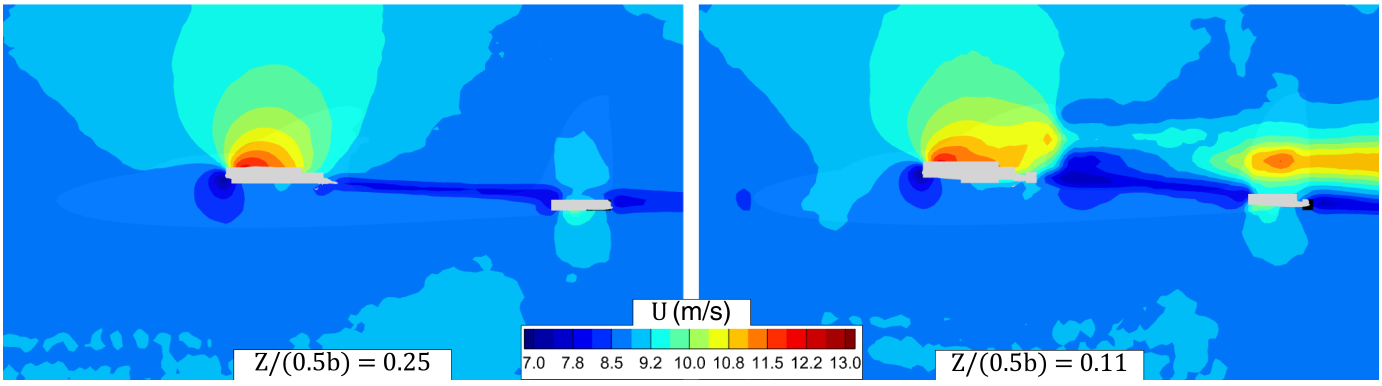


Figure 10. Slices showing the U component of the velocity at $Z/(0.5 \cdot b) = 0.25$ (left) and 0.11 (right). The tip of the wing has been hidden in the visualization of the model.

indicating the higher lift value. Above the tail, part of the slipstream of the propeller appears. The velocity on the bottom side of the tail is similar to the other plane.

Next, Figure 11 shows the U and W components of the velocity in the midspan of the drone, highlighting the slipstream and the flow over the fuselage. As expected, the propeller is ineffective at accelerating the flow near the hub. In addition, the lower part of the slipstream has an average lower velocity than the upper part. This can be explained by the fact that the lower part is behind the wake of the fuselage, so it experiences a lower momentum fluid inflow. Behind the tail, the slipstream of the propeller is slowed down significantly, by more than $3m/s$.

From the W component it can be seen that the propeller was rotating clockwise when looked at from the back. Swirl velocities in the slipstream reach up to $3.5m/s$. Despite the hub having a much lower U velocity than the edges, the swirl velocity is only lower over the very centre of the hub. The interaction between the slipstream and the vertical tail causes the swirl to be lowered, as most of the negative swirl velocity is diminished, while the positive velocity is slightly attenuated. This indicates that the vertical tail provides a force to the flow pointing towards the port side (in positive Z direction). The drone itself therefore experiences an opposite force.

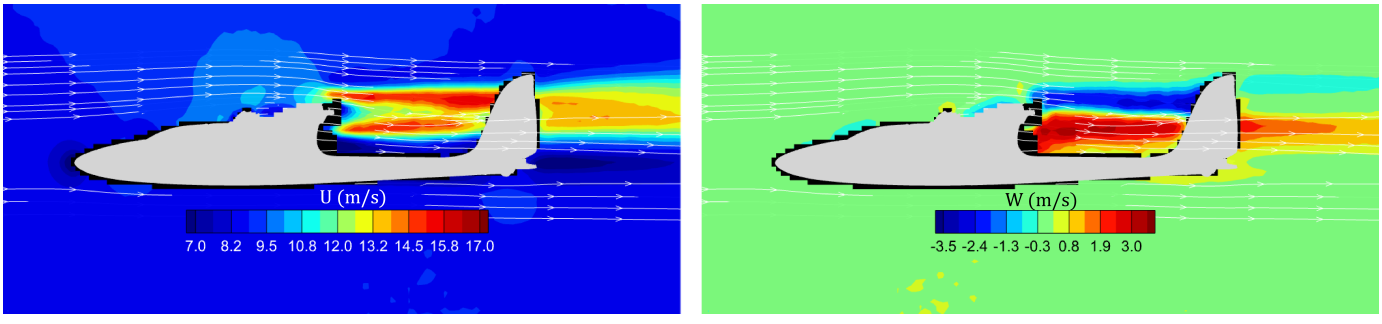


Figure 11. Slices showing the U component of the velocity at $Z/(0.5 \cdot b) = 0.25$ (left) and 0.11 (right). The tip of the wing has been hidden in the visualization of the model.

Lastly, Figure 12 shows the development of the slipstream at 4 different streamwise planes, situ-

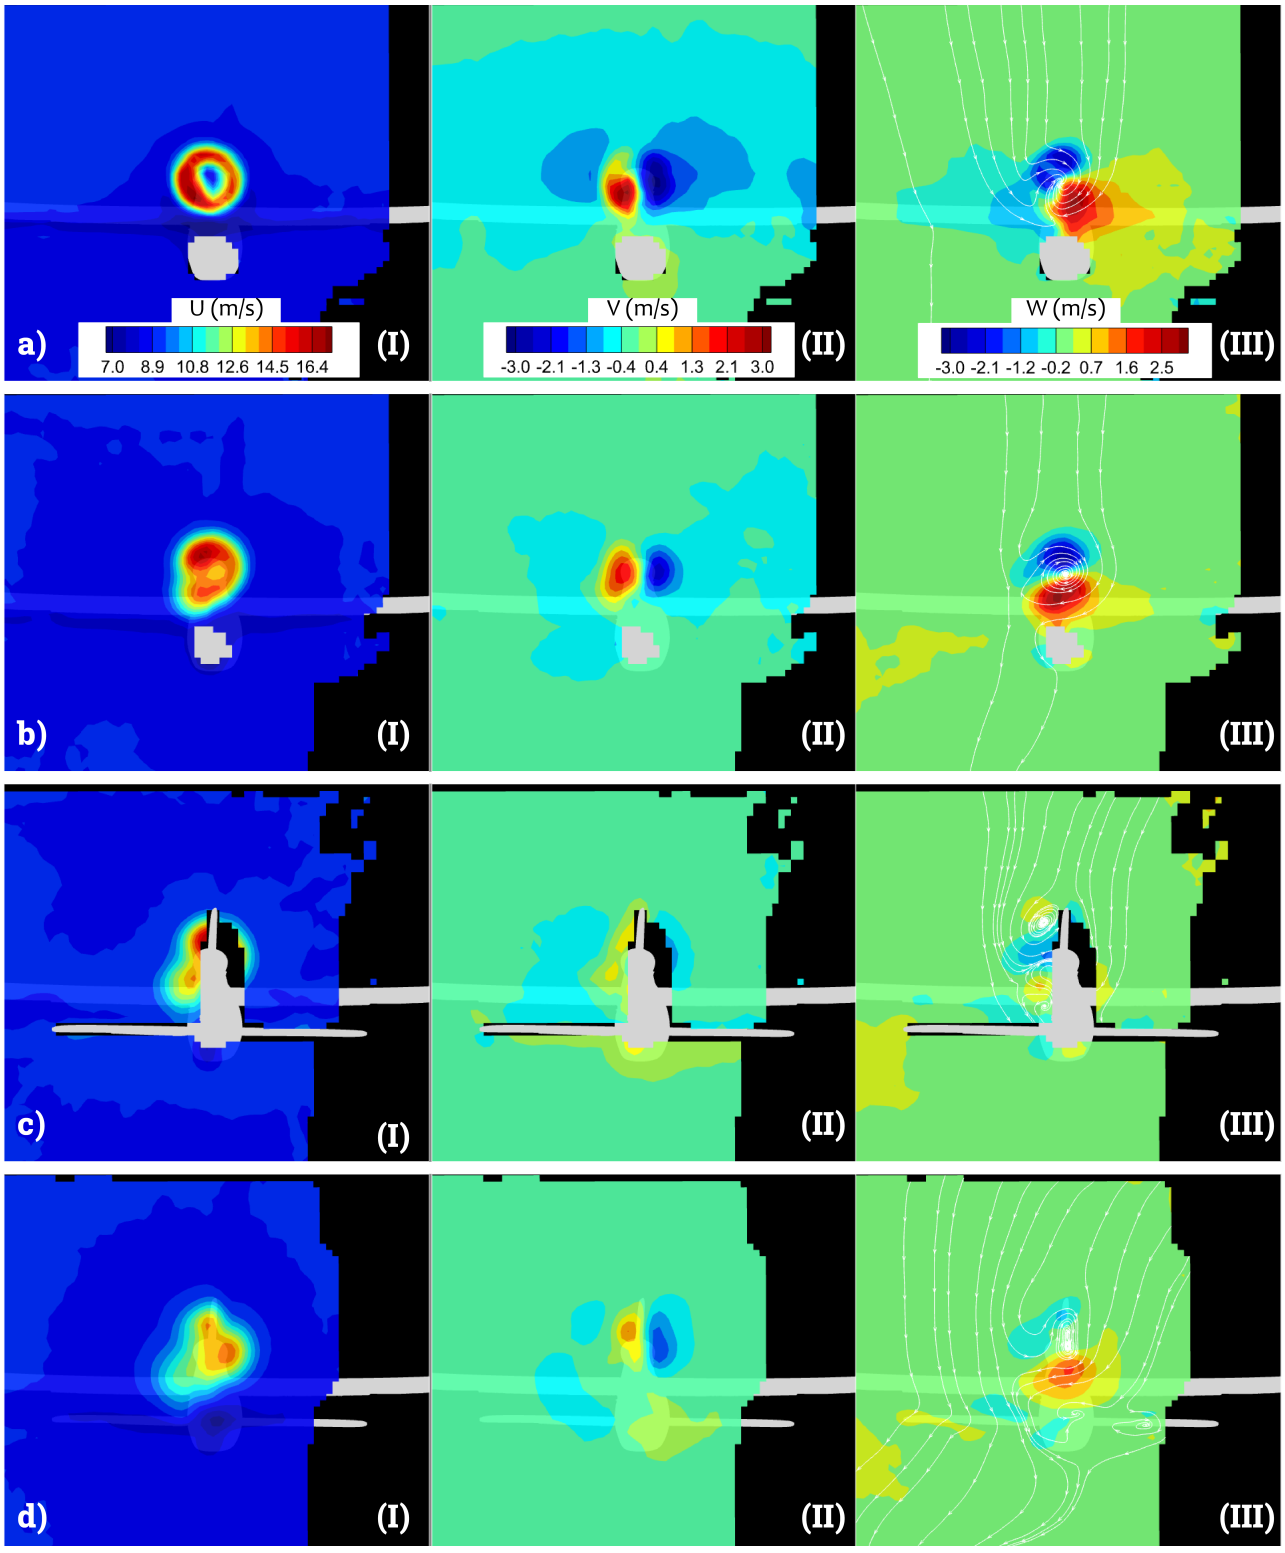


Figure 12. Wake planes showing the development of the propeller slipstream. The columns show contours of the U (I), V(II) and W(III) velocity components. Row a) shows the slipstream 40 mm behind the propeller hub at X/l = 0.56, Row b) at X/l = 0.80, row c) over the tail at 0.95 and row in the wake at X/l = 1.12. In addition, in-plane streamlines of the V,W velocity components are shown in the third column to indicate vortical structures.

ated at X/l = [0.56, 0.80, 0.95, 1.12]. Here X=0 is at the tip of the nose and l is the length of the drone. Briefly summarized, it can be seen that very near the propeller, the propeller affects the flow far into the wings. As you go further downstream to x/l = 0.8, the slipstream becomes a more self contained vortex. When it hits the vertical tail, a complex interaction takes place that forms different vortical structures on the port side of the tail, with at-least 3 large vortices being present. Unfortunately, the model obstructed the view of 10 of the 12 cameras over a small area next to the starboard side of the vertical tail, so the flow here has not been reconstructed. It is expected that other vortical structures form here. Lastly, in the wake a the propeller slipstream is deformed after the interaction. In addition, the slipstream vortex is no longer visible in the streamlines, with a elevator vortex forming behind the tail. This is further visualized in the next part.

Overall, the obtained velocity fields appear to be smooth, with no obvious jumps present due to the different mean conditions in which the measurements were taken. Given that measurements took place under small angles, this can be expected from the relatively linear aerodynamic behaviour that takes place over most of the model. Velocity fields at 10 m/s are not shown here, as they are very similar to those at 9 m/s over the wing.

Vortical Structures at 9 m/s

To visualize the vortical structures that are measured, isosurfaces of Q-criterion coloured by the streamwise vorticity (ω_x) are used. Both were non-dimensionalized using the mean aerodynamic chord and the freestream velocity according to

$$Q^* = \frac{Q \cdot \bar{c}^2}{U_\infty^2} \tag{4}$$

$$\omega_x^* = \frac{\omega_x \cdot \bar{c}}{U_\infty} \tag{5}$$

A top view of the structures over the model can be seen in Figure 13, whereas a back right view can be seen in Figure 14. Here many different vortical structures were captured. Starting with the expected structures, the tip vortex along with the propeller hub vortex are clearly represented. The hub vortex and tip vortex are both blue, rotating in the same direction as the propeller. A small gap can be seen on the hub vortex at the tail due to the lack of tracks in the starboard side.

Instead, the propeller tip vortex has not been clearly captured by Q^*. In a time resolved sense, it appears as three helical vortex tubes wrapping around the hub vortex. In a time averaged sense, if the acquisition frequency does not match with the propeller frequency, it should appear as a cone encircling the hub vortex. Instead, the propeller tip vortex appears as a broken structure in the figures, with it being more open on the starboard side. It is not clear why this occurred, as the propeller frequency was never an exact match or multiple of the acquisition frequency.

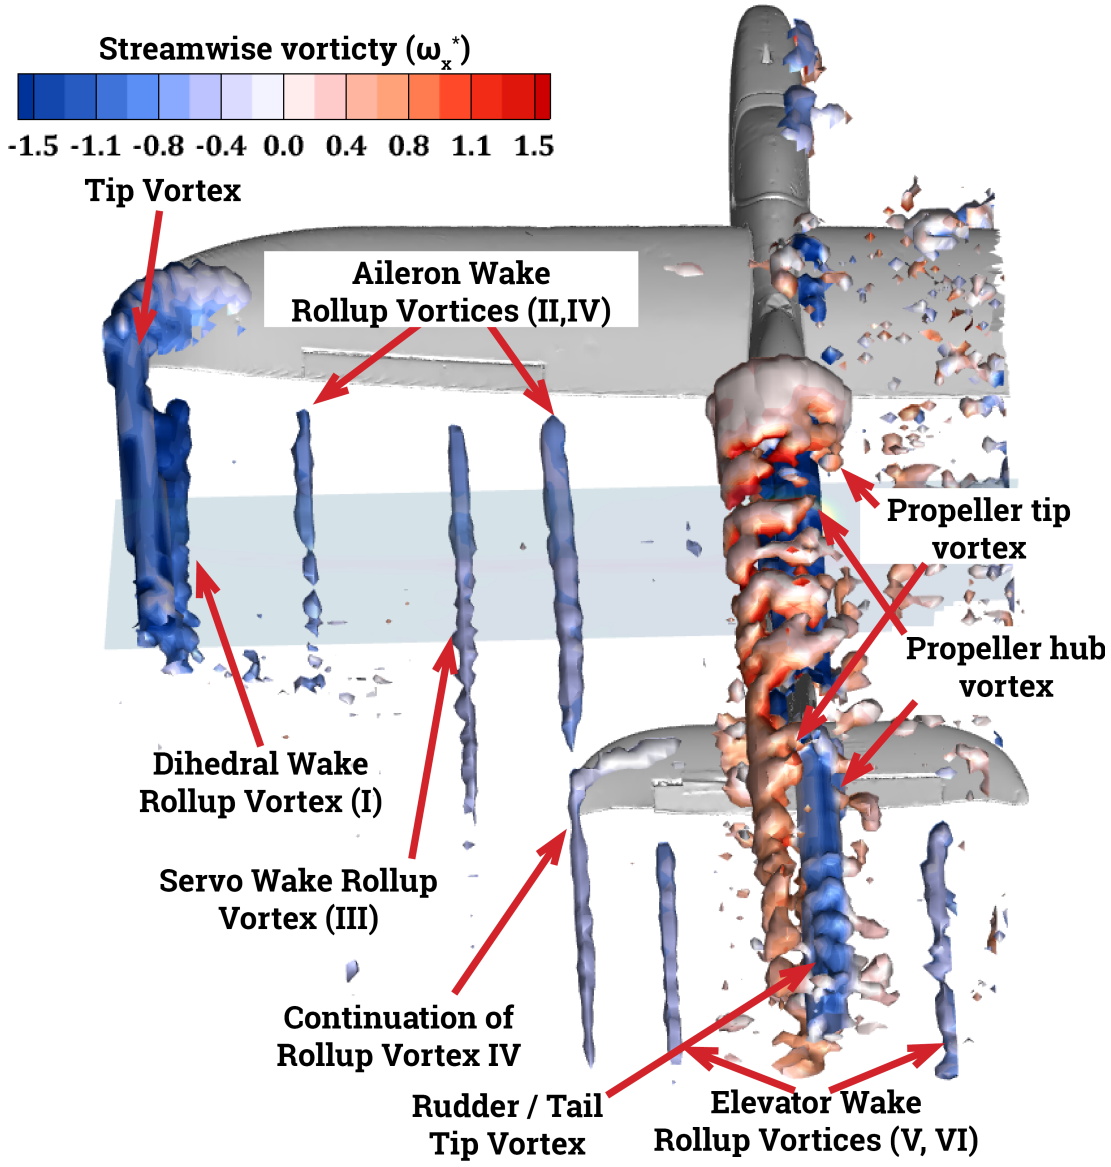


Figure 13. Top view of the vortical structures present in the flow-field. Visualized using isosurfaces of Q* = 0.11 coloured by the streamwise vorticity (ω_x^*) at 9 m/s. The transparent slice shown here can be seen in Figure 15.

Additionally, the vorticity here should be mainly in ω_z and ω_y, but as can be seen it has a low level of positive (red) ω_x vorticity.

Along with these three main vortices, 7 additional wake-like vortical structures were observed. Interestingly, these all have negative ω_x vorticity, indicating a clockwise rotation when visualized from the back. First, underneath the tip vortex, a wake roll-up vortex (I) is observed that is caused by the sharp increase in dihedral angle. This can be explained by the lift force rotating clockwise, inducing a clockwise curl in the wake which starts to rolls up behind the tip.

Secondly, behind the aileron three wake rollup vortices (II, III and IV) are observed. Vortex III occurs right behind a servo that sticks out underneath the wing to control the aileron. Given the positive β angle that the drone flew with, the incoming streamlines are slightly bent inboard with respect to the chord line. The servo locally straightens the streamlines in the chord wise direction. When the top and bottom flow meet there is a discontinuity in the W velocity component which rolls up into Vortex III. Vortex II is likely associated with the downward deflection of the aileron.

The origin of Vortex IV is uncertain. A downward deflection of the aileron causes an opposite vorticity here than at the location of Vortex II. Indeed, a patch of red vorticity can also be observed next to it in the wake plane shown in Figure 15. Vortex IV stays coherent up until far into the wake of the model, as it extends all the way to the end of the field of view. Hints of a similar vortex were also observed on the starboard side of the wing with an opposing vorticity, but the data there was too noisy to draw a decisive conclusion. But if this vortex is also present there, it is likely associated with the geometry of the wing. At this point, the dihedral angle roughly starts increasing from zero, so it could have been partially caused by this.

The lack of a horizontal tail tip vortex is caused by the spatial resolution of the binning not being sufficient to visualize it. With 12 voxel bins with 75% overlap (0.31mm vector spacing) the tail tip vortex is captured, but the data becomes extremely noisy (not shown for the sake of brevity). Note that behind the tail, there is a gap between the horizontal tail tip and Vortex IV, so these likely do not directly interact.

Behind the elevator, two additional wake vortices appear (V, VI). These appear to be similar in origin as Vortices II and IV, but there is no obvious indication as to how they originate. If it was purely due to the elevator deflection, they should have had opposite signs. Vortex VI can also be observed in Figure 12 d) III, whereas V does not appear. Further analysis of the data is required.

Lastly, behind the rudder multiple vortical structures appear. In the middle, the hub vortex continues, albeit weakened in vorticity compared to before the tail. Right on top of the hub vortex, a strong clockwise rotating vortex is present. From the velocity data the observation was made that the tail exerted a force in the positive Z direction on the flow, which would indicate a rudder deflection in the same way. This is further supported by the vortex shown here having a blue vorticity, which would match that rudder deflection. Additionally, on the bottom of the tail there

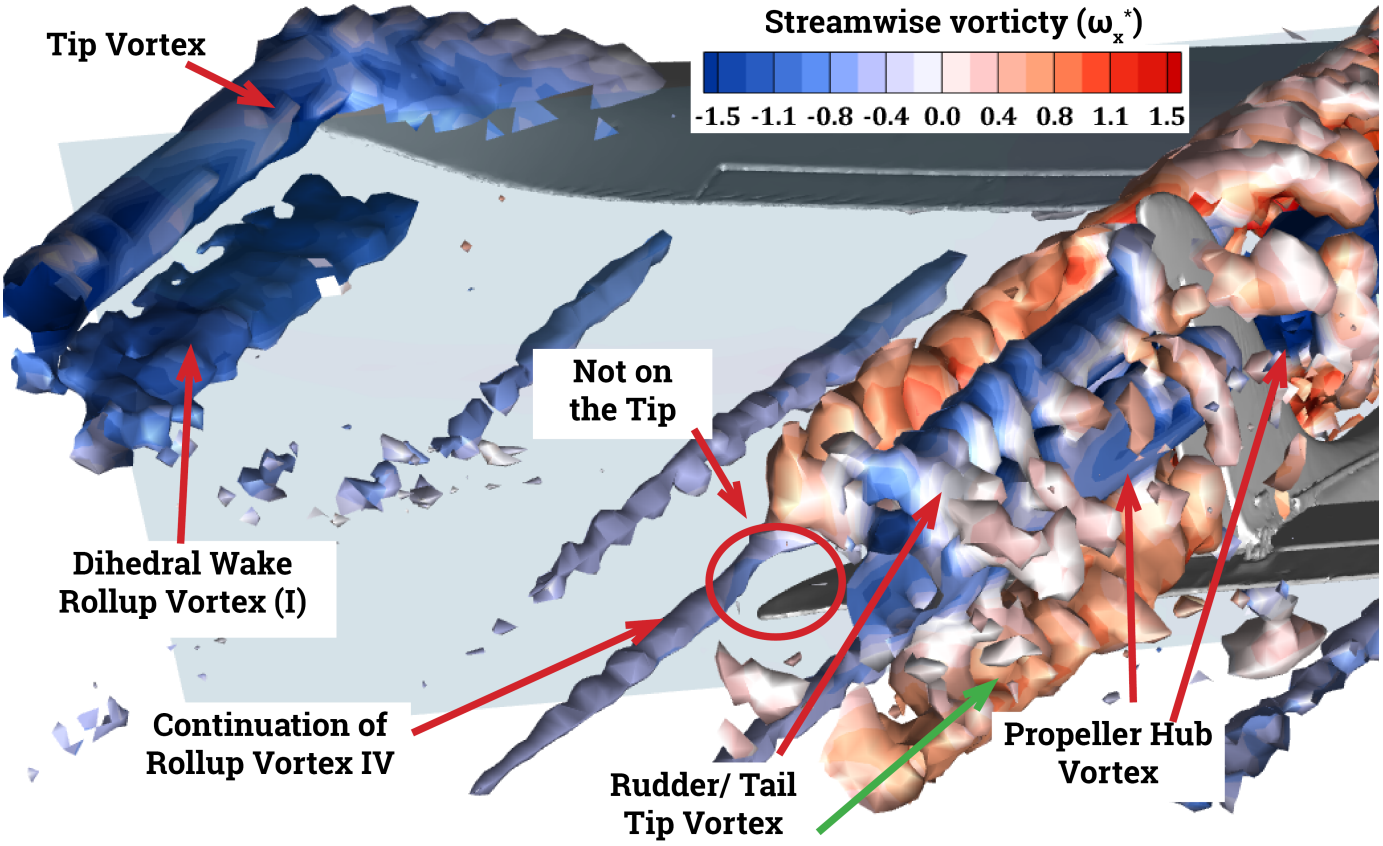


Figure 14. Back view of the vortical structures present in the flow-field at 9 m/s. Visualized using isosurfaces of Q* = 0.11 coloured by the streamwise vorticity (ω_x^*). The transparent slice shown here can be seen in Figure 15.

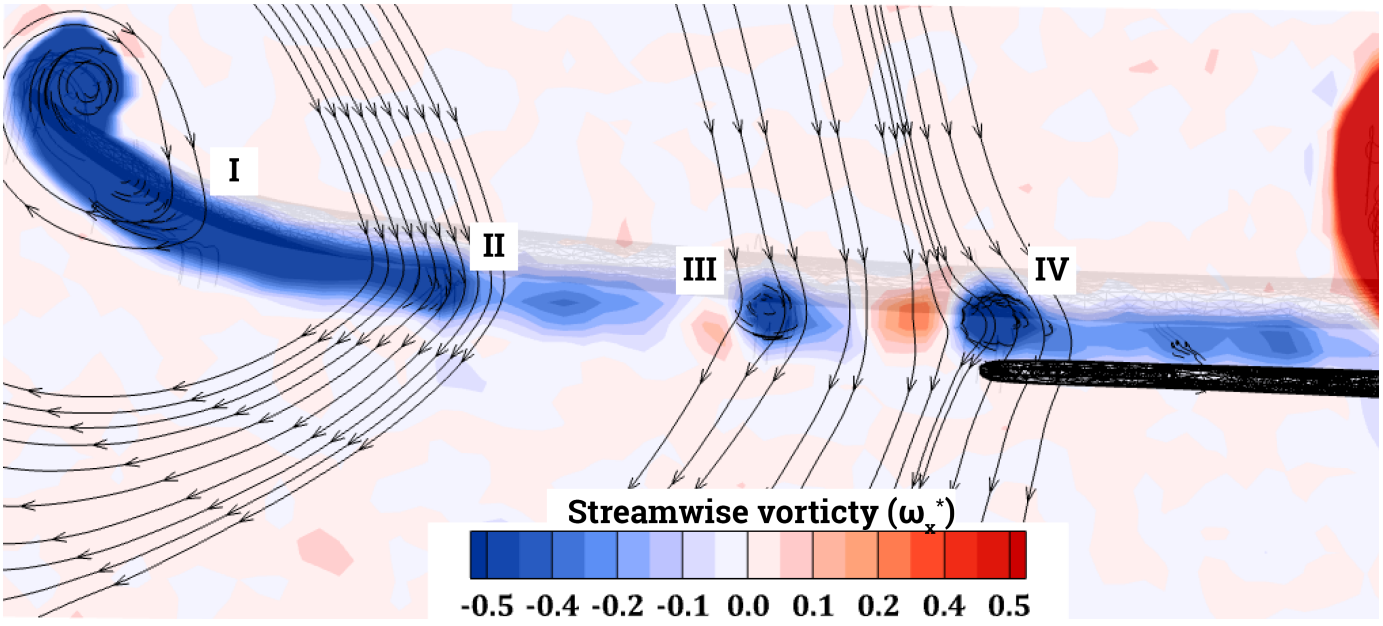


Figure 15. Streamwise vorticity (ω_x^*) on a wake place situated at X/l = 0.71. Additionally, some 2D streamlines are drawn on this plane. On the location of Vortices I-IV, 3D streamlines passing through the vortices are also drawn.

is an elongated structure with positive vorticity, indicated by the green arrow in Figure 14. This could be the other side of the rudder vortex, but it can also be a continuation of the propeller tip vortex, as that occurs in the same region. The most likely explanation is that it is a combination of the interaction of both, which in a time averaged sense could yield a chaotic blob of Q^*.

Vortical Structures at 10 m/s

In Figure 16 the vortical structures measured at 10 m/s are shown. The same structures were also observed for the 9 m/s case and are indicated with their numbers. Similar conclusions can therefore be drawn about their origins.

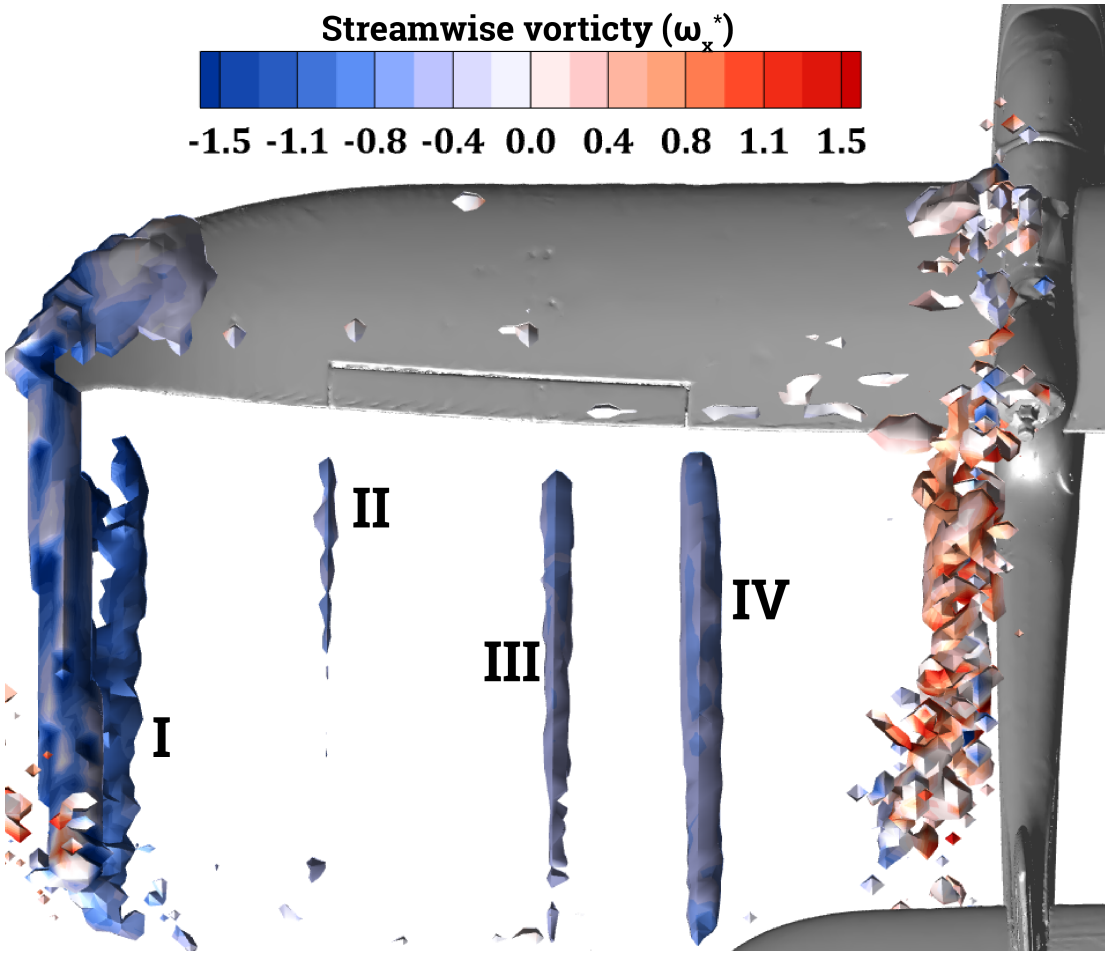


Figure 16. Top view of the vortical structures present in the flow-field at 10 m/s. Visualized using isosurfaces of Q* = 0.11 coloured by the streamwise vorticity (ω_x^*). The vortical structures previously described for 9 m/s are numbered accordingly.

4.3. Force reconstruction

Kutta-Jouwkowski theorem was used to compute the forces over the wing. An attempt was also made to estimate the forces over the tail (vertical and horizontal) but due to the highly viscous and three dimensional flow present there, the computed circulation varied significantly with the size of the integration contour. For these regions, a more advance control volume approach is required. In Figure 17, the lift per unit span results are shown for the 9 m/s and 10 m/s cases. In contrast, Figure 18 shows the C_L per unit span.

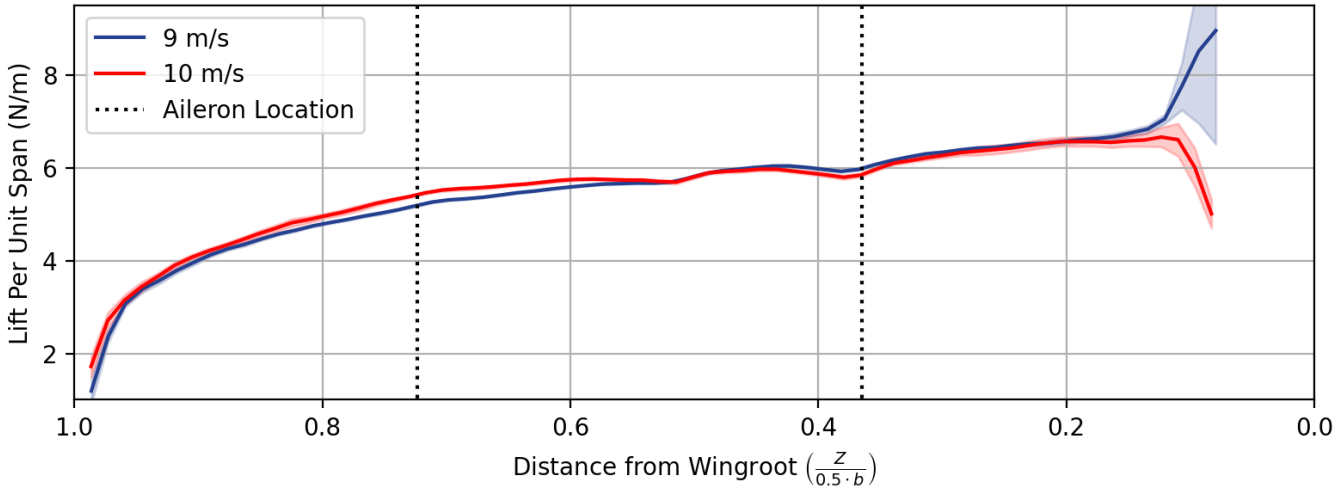


Figure 17. Lift per unit span (N/m) computed with Kutta-Joukowski. The shaded region shows the standard deviation of the different contours used to compute the lift.

It can be seen that the lift distribution after $Z/(0.5 \cdot b) = 0.125$ is almost identical between the two cases. This is consistent with theory for the lift of a 3D wing without twist. During levelled flight in the attached flow regime under the same loading conditions, it should be invariant as the C_{l_α} of the airfoil stays constant. With twist, the lift distribution is a linear combination of the $L_{\alpha=0}$ and L_α so by changing the α, the shape of distribution changes. This could potentially explain the small differences in L after the midspan, as the wing may have a small amount of twist.

Near the root stark differences can be seen. Whereas for the 9 m/s the propeller slipstream increases the local lift drastically, for the 10 m/s case the lift decreases. The likely cause is a change in the propeller slipstream shape or size between measurements, but as it was not measured for the 10 m/s case this can not be fully justified. Additionally, the values obtained near the root with Kutta-Joukowski are unreliable, therefore an analysis with a more adequate method is needed.

Neglecting the region before $Z/(0.5 \cdot b) = 0.125$ where Kutta-Joukowski is unreliable results in a total lift of $3.21N$ and $3.30N$ at 9 and 10 m/s respectively. This corresponds to 43% and 44% of the total weight of the aircraft respectively. This indicates that the measured data can be used to accurately determine the lift force distribution over the wing, but it necessitates a more advanced force calculation method.

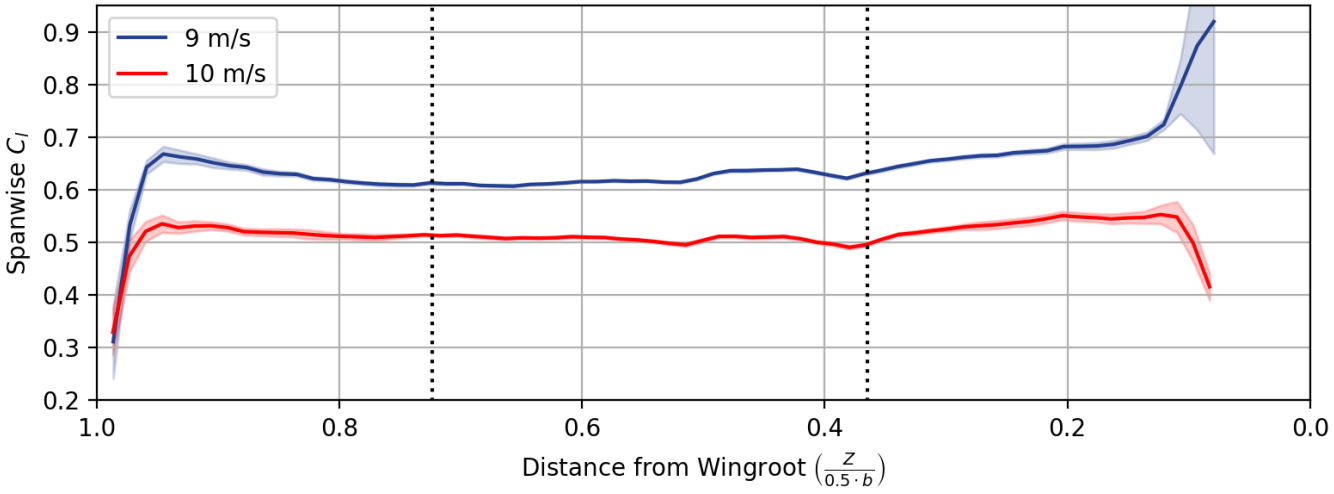


Figure 18. Spanwise C_l distribution computed with Kutta-Joukowski. Here C_l has been computed with the local airfoil chord value, not the $\bar{c}$. The shaded region shows the standard deviation of the different contours used to compute the C_l.

Next, from Figure 18 it can be seen that at 10 m/s the wing had a section C_l that was 18.7% lower than the C_l at $9m/s$. Using the measured $\alpha = 4.46°$ for 9 m/s and assuming linear C_{L_α} would indicate an α of $3.6°$, slightly higher than what was measured for 10 m/s. Purely based on the dynamic pressure, a decrease of 26% was estimated. A likely explanation for this discrepancy is the fact that the C_{L_α} is not linear due to the changing conditions near the root.

Lastly, the obtained force distribution can be related to the observed wake vortices. First, a bump can be seen in both the C_l and lift after the start of the aileron and ending at $Z/(0.5 \cdot b) = 0.517$ for both 9 m/s and 10 m/s. The start of the bump occurs after the aileron and Vortex IV, but its end point perfectly overlaps with Vortex III as indicated Figure 13. It is therefore likely associated with the presence of these two vortices. Furthermore, near the end of the aileron the slope of the lift curve has a small discontinuity, decreasing more sharply than before. This is also the location of vortex II. While such small changes in lift alone would not be expected to produce a visible vortex with the current spatial resolution, it is indicative that something else is going on in that region.

5. Conclusions

In the present work, volumetric flow field measurements are taken on a fixed wing drone in free, levelled flight. This was done inside a wind tunnel using object aware shake the box (OA-STB). As one of the first of these types of tests to ever be performed, the purpose of the test was to assess the capabilities of object aware measurements on objects in free flight and to find areas of improvement such that the experimental methodology can be advanced for more complex test cases in the future.

Three aspects of the experiment were evaluated. First, the use of OA-STB is discussed in detail and an assessment is made of its accuracy. To track the model in free flight, over 100 white markers were placed in a random distribution over the drone. The OR was found to be highly dependent on the tracking approach used to detect these markers. The best approach was found to be tracking the model using time resolved STB with a partitioned camera system to account for the initial obstruction from the object. Tracks were only accepted when within a limited velocity range, related to the expected velocity of the model, and filtered by track length. The most accurate reconstruction of the displacement and rotation of the model was obtained when tracking both the white markers and steady reflections over the model, yielding hundreds of points on the model. Using these settings, a positional uncertainty of $\pm 1.6mm$ and a rotational uncertainty of $\pm 0.28^{\circ}$ were postulated for the current measurements. The resulting time resolved OR contained some high frequency noise which was filtered out using a Butterworth filter before creating the depth maps.

Secondly, the control strategy was evaluated for the purpose of obtaining particle tracks to reconstruct the flow field around the model during flight. This was done through an analysis of the flight kinematics of the drone obtained from the OR. The drone was flown manually in the tunnel, supported by an autopilot to help stabilize the flight. The autopilot was found to respond at a 1.9 to 2.3 Hz frequency inducing rotations in the attitude of the drone up to 0.5°. At these frequencies, changes in the flow are described by unsteady aerodynamics. The magnitude of these rotations were very case dependent, as were the differences in aerodynamic and roll angles between different measurements. Under the most steady measurements, rotations of the order of 0.25° were observed, with a higher repeatability of the mean values. The final flow fields were obtained by combining different measurement together, therefore the reconstructed flowfield is an average of multiple conditions centred around the mean values. The mean values also have a high degree of uncertainty relative to the magnitude of the angle. Coupling between the flight control and the PIV acquisition system could help to improve this accuracy, enabling more accurate determination of the flow state. Overall, the conclusion can be drawn that fully autonomous flight that is linked to the PIV system is necessary to obtain more reliable measurements.

Lastly, the obtained flow field was studied to demonstrate the capabilities of the measurements, despite the inadequacies of the control method. A total volume of almost $1m^3$ was captured around the port side of the drone at 9 m/s, with a smaller volume over the wing at 10 m/s. From the velocity fields and isosurfaces of Q-criterion, intriguing features were revealed about the flow over the model. Despite the averaging over multiple conditions, no sharp jumps were observed in the reconstructed flow field. Basic features that are expected for this flow were clearly captured, such as the suction peak over the wing and tail, stagnation points, the tip vortex, the propeller hub vortex and propeller tip vortex. Additionally, the interaction between the propeller slipstream and the vertical tail was found to significantly warp the slipstream of the propeller. During the interaction, multiple different vortical structures form over the surface of the vertical tail. The

elevator was observed to be deflected towards the port side of the model, causing a force in the same direction on the flow. This resulted in a large tip and root vortex at the edges of the rudder. Lastly, multiple unexpected vortical structures were observed behind the wing, aileron and rudder. These were attributed to different characteristics of the flow.

To conclude, while the current work shows the potential for these measurements to study the flow field of objects in free-flight, significant improvements are required. The main area of relevance would be to study conditions which are difficult to capture in traditional wind tunnel tests, such as gust aeroelasticity. The main limitation shown here is the repeatability of test conditions, as this necessitates an improved control system. In addition, the object tracking method used here relies on the surface markers having sub-pixel displacement between frames, which might not be the case for more dynamic motions.

Acknowledgements

The authors would like to acknowledge the Micro Air Vehicle Lab of the TU Delft for all their assistance in setting up and flying the drone.

References

Agüera, N., Cafiero, G., Astarita, T., & Discetti, S. (2016). Ensemble 3d ptv for high resolution turbulent statistics. *Measurement Science and Technology*, *27*(12), 124011.

Anderson, J. D. (2017). *Fundamentals of aerodynamics*. McGraw-Hill.

Dawson-Howe, K. M., & Vernon, D. (1994). Simple pinhole camera calibration. *International Journal of Imaging Systems and Technology*, *5*(1), 1–6.

Hendriksen, L., Guerra, A. G., Sciacchitano, A., & Scarano, F. (2025). Omnidirectional particle image velocimetry. In *Ispiv.* Int. Symp. on Particle Image Velocimetry.

Hendriksen, L., Sciacchitano, A., & Scarano, F. (2024). Object registration techniques for 3d particle tracking. *Measurement Science and Technology*, *35*(12), 125202.

Hwang, S., Remes, B. D., & De Croon, G. C. (2023). Aosoar: Autonomous orographic soaring of a micro air vehicle. In *2023 ieee/rsj international conference on intelligent robots and systems (iros)* (pp. 5003–5010).

Hysa, I., Tuinstra, M., Sciacchitano, A., Scarano, F., van der Meulen, M.-J., Rockstroh, T., & Roosenboom, E. W. (2024). A multi-directional redundant 3d-lpt system for ship–flight–deck wind interactions. *Experiments in Fluids*, *65*(8), 126.

Lee, T., & Su, Y. (2012). Low reynolds number airfoil aerodynamic loads determination via line integral of velocity obtained with particle image velocimetry. *Experiments in fluids*, *53*(5), 1177–1190.

Martínez Gallar, B., Van Oudheusden, B., Sciacchitano, A., & Karásek, M. (2020). Large-scale volumetric flow visualization of the unsteady wake of a flapping-wing micro air vehicle. *Experiments in Fluids*, *61*(1), 16.

Mertens, C., Costa Fernández, J. L., Sodja, J., Sciacchitano, A., & Van Oudheusden, B. W. (2023). Nonintrusive experimental aeroelastic analysis of a highly flexible wing. *AIAA Journal*, *61*(7), 3062–3077.

Novara, M., Schanz, D., & Schröder, A. (2023). Two-pulse 3d particle tracking with shake-the-box. *Experiments in Fluids*, *64*(5), 93.

Rockstroh, T., & Brinkeman, R. (2025). Towards moving objects within the object-aware lagrangian particle tracking framework. In *16th International Symposium on Particle Image Velocimetry – ISPIV 2025, Tokyo, Japan.*

Scarano, F., Ghaemi, S., Caridi, G. C. A., Bosbach, J., Dierksheide, U., & Sciacchitano, A. (2015). On the use of helium-filled soap bubbles for large-scale tomographic piv in wind tunnel experiments. *Experiments in Fluids*, *56*(2), 42.

Schanz, D., Gesemann, S., & Schröder, A. (2016). Shake-the-box: Lagrangian particle tracking at high particle image densities. *Experiments in fluids*, *57*(5), 70.

Schanz, D., Gesemann, S., Schröder, A., Wieneke, B., & Novara, M. (2012). Non-uniform optical transfer functions in particle imaging: calibration and application to tomographic reconstruction. *Measurement Science and Technology*, *24*(2), 024009.

Sharma, S., & Deshpande, P. (2012). Kutta–joukowsky theorem in viscous and unsteady flow. *Experiments in fluids*, *52*(6), 1581–1591.

Sodja, J., & De Breuker, R. (2022). Design of an increased-fidelity aeroelastic experiment for wing response to gust excitation. In *Aiaa scitech 2022 forum* (p. 2418).

Theodorsen, T. (1949). *General theory of aerodynamic instability and the mechanism of flutter* (Tech. Rep.).

Wieneke, B. (2008). Volume self-calibration for 3d particle image velocimetry. *Experiments in fluids*, *45*(4), 549–556.

Wieneke, B., & Rockstroh, T. (2024). Lagrangian particle tracking in the presence of obstructing objects. *Measurement Science and Technology*, *35*(5), 055303.