Validation and Uncertainty Analysis of Stereo Time-Resolved PIV Measurements for Boundary-Layer Transition Research

Jenna L. Eppink*

NASA Langley Research Center, Hampton, VA 23681

An uncertainty analysis was performed on the stereo time-resolved particle image velocimetry measurements that were obtained in the boundary-layer of a swept wing model. Several different approaches were undertaken to estimate the uncertainty of the PIV measurements and to understand how that uncertainty propagates through to the results of interest. It was found that the stereo registration error, which is often neglected in stereo-PIV uncertainty analysis, can cause significant errors in the measured amplitude of the stationary crossflow instability due to the large spanwise velocity gradients that occur in the flow. The self-calibrations do not always adequately correct for the stereo-registration error. A simple approach is demonstrated for estimating the error due to a misalignment of the calibration target and laser plane. This approach was found to be very successful at predicting the error in stationary crossflow amplitude caused by the misalignment. The time-resolved PIV results are also compared to previously-acquired hotwire data. The TRPIV technique is successful at measuring the same types of unsteady instabilities as the hotwire, and thus does not appear to have any adverse effect on the transition process due to the introduction of particles into the flow. It was also found that a larger interrogation window size is beneficial for reducing the noise of the TRPIV results.

I. Introduction

Stereo time-resolved particle image velocimetry (TRPIV) measurements have been performed on a swept-wing boundary layer to study the transition characteristics in the presence of 2D steps. TRPIV measurements are advantageous when compared to typical single-point, single-velocity component measurements (such as hotwires), since one can obtain instantaneous planes of data, while resolving all three velocity components, at significant frequencies. However, the use of PIV in transition research is fairly novel, thus, it is desirable to validate the results through an extensive uncertainty analysis, as well as comparison with a separate, well-proven, measurement technique. Thus, the purpose of this paper is two-fold. The first purpose is to attempt to quantify, or at least bound, the uncertainty in the mean and time-resolved stereo-PIV measurements that were performed. The second purpose is to compare the PIV measurements with the hotwire (HW) measurements that were also acquired for several of the cases under consideration.

The uncertainty in a PIV measurement is dependent on many factors, including particle image size, seeding density, in-plane displacement, and camera noise. Additionally, stereo-PIV measurements include an additional potential source of uncertainty known as registration error, which is due to the misalignment of the calibration target and the actual measurement plane. This misalignment could occur due to either a systematic or random problem in the alignment procedure of the calibration target with the presumably constant laser sheet position. Additionally, the laser sheet position may shift from the set or desired location due to changes in temperature affecting the optics. Apart from the uncertainty due to flow fluctuations, the PIV community has only recently begun to address uncertainty quantification in planar PIV measurements. The uncertainty for stereo-PIV measurements includes a combination of the planar PIV uncertainty for each camera, as well as the uncertainty in the stereo calibration function, caused by the registration error. There are several different approaches to PIV uncertainty quantification. Initially, attempts to quantify the PIV

*Research Aerospace Engineer, Flow Physics and Control Branch, M.S. 170, AIAA Member

uncertainty typically followed *a priori* approaches, which attempt to model the measurement chain. The uncertainty associated with each of the different sources (such as particle size) is typically estimated and then propagated through the chain to obtain a total velocity uncertainty.[1] Typical values for the random errors using this technique have been found to range from approximately 0.05 to 0.1 pixels.[2,3] Unfortunately, these *a priori* approaches based on simplified models cannot account for the many error sources that contribute to the overall measurement uncertainty. Recently, several *a posteriori* methods have been developed.[4–6] Among these, the direct methods[4,5] estimate the planar PIV measurement uncertainty from the measured velocity field and the particle images. While these *a posteriori* approaches have been found to more accurately estimate the random uncertainty, they do not take into account the calibration errors, which can be a major error source particularly for stereo-PIV configurations. Researchers tend to neglect the registration error, claiming that the self-calibration procedure will correct for any misalignment. The self-calibration essentially compares the dewarped particle images from the two cameras and attempts to correct for the disparity between the two images by shifting and/or rotating the initial calibration plane to line up with the laser plane. However, as shown by Beresh et al.,[7] a converged self calibration with small disparity does not always sufficiently correct for the stereo calibration error. Bhattacharya et al.[8] recently presented a method for stereo-PIV uncertainty quantification using the direct *a posteriori* methods to estimate the planar uncertainties and including the effect of the multicamera calibrations and stereo reconstruction on the uncertainty. This method will be considered in the following sections.

Quite a few studies have been performed comparing PIV data to HW data, as well as DNS data.[9–13] These studies have largely been performed with a standard (10 to 15 Hz) mono-PIV system and typically compare mean and turbulence quantities for various turbulent flows. In general, the mean measurements agree well, but the various studies tend to disagree as to whether PIV over- or underestimates the fluctuations measured using an HW. Foucaut et al.[9] found that the measurement noise introduced throughout the PIV chain is important when comparing the PIV wavenumber power spectra to the HW spectra. The interrogation window size was found to play an important role in minimizing the noise that was introduced. The study by Gericke and Scholz[10] compared results for a mono- and a stereo-PIV configuration to HW measurements in the turbulent wake of an air outlet. Their results showed that the fluctuation component measured by the stereo-PIV system resulted in significant errors due to high frequency noise. There have not been quite as many studies comparing the time-resolved PIV technique to HW measurements. Wernet[14] compared mono-TRPIV data obtained at 10 kHz to HW results from measurements performed in a hot jet. The power spectra for the two techniques agreed very well up to the HW cut-off frequency of 4 kHz. For the most part, the axial turbulence intensities also agreed well with the HW results. Dróżdż and Uruba[15] compared HW and TRPIV statistics in a turbulent boundary layer for a mono-PIV setup. The mean velocities agreed well for all PIV estimation methods and Reynolds numbers studied. The fluctuating velocities agreed fairly well away from the wall, but these results varied depending on the PIV estimation method and interrogation window size. The PIV fluctuating velocity also departed from the HW measurements near the wall, and the disagreement increased with increasing Reynolds number. It is clear from these studies that while PIV and HW data can agree well under certain circumstances, the results often depend on the situation and the choice of PIV processing parameters such as window size. Thus, it is important to quantify the uncertainty for a given setup and PIV processing approach.

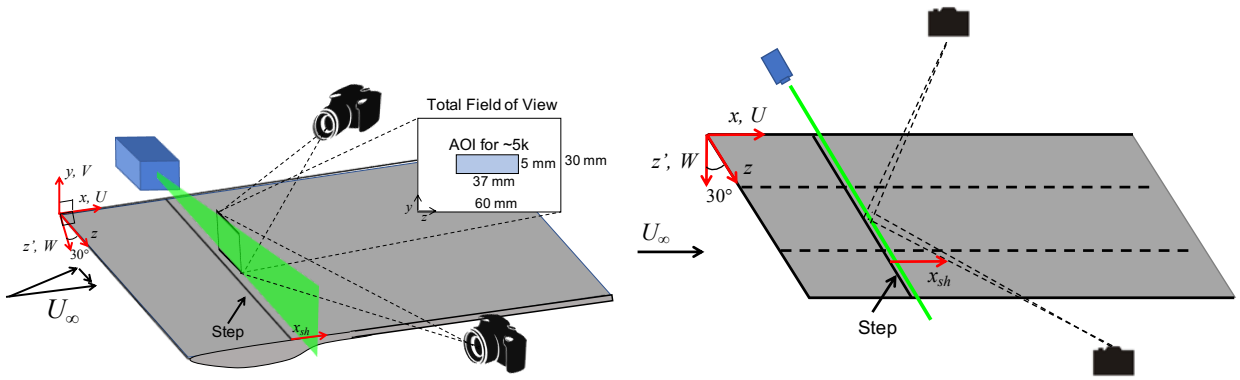
PIV is typically considered a nonintrusive measurement technique. However, in most cases, it is necessary to seed the flow with particles. It is unknown how the introduction of particles will affect the boundary-layer transition process. Since the measurement of boundary-layer instabilities using TRPIV (or even standard PIV) is only recently being done,[16,17] it is important to compare a well-established technique (HW) to the PIV measurements in order to validate the PIV technique for these types of measurements and to gain some idea of their quantitative value. Some limited comparisons have been performed up to this point. Serpieri and Kotsonis[18] show a good comparison between the tomo-PIV and hotwire measurements of the stationary-crossflow instability for a single plane. Eppink[16,19] shows good agreement between the mean boundary layer profiles, stationary crossflow, and unsteady measurements (including velocity spectra) performed using HW and PIV. The present work is a more thorough comparison of the mean and fluctuating velocities measured using TRPIV and HW, including uncertainty analysis. The questions that we wish to answer include: is PIV able to measure the same type of transition instabilities and mechanisms without fundamentally altering the transition process due to the addition of particles? If so, how good, quantitatively, are the PIV measurements of the instabilities and the mean flow? In addition to addressing these questions, we also wish to share some of the challenges and lessons learned with the current setup in order to emphasize the level of care that is

required to accurately perform these types of measurements.

II. Experimental Setup

The experiment was performed in the 2-Foot by 3-Foot Low Speed Boundary-Layer Channel at the NASA Langley Research Center. The tunnel is a closed circuit facility with a 0.61-m high by 0.91-m wide by 6.1-m long test section. The tunnel can reach speeds up to 45 m/s ($Re' = 2.87$ x 10^6/m) in the test section. Freestream turbulence intensity levels, $Tu = \frac{1}{U_\infty}\sqrt{\frac{1}{3}(u'^2 + v'^2 + w'^2)}$, were measured using a crosswire in an empty test section to be less than 0.06% for the entire speed range of the tunnel, and less than 0.05% for the test speed of 26.5 m/s. This value represents the total energy across the spectrum, high-pass filtered at 0.25 Hz. Thus, this tunnel can be considered a low-disturbance facility for purposes of conducting transition experiments.[20]

The 0.0127-m thick flat plate model consists of a 0.41-m long leading edge piece, swept at 30°, and a larger downstream piece (see Fig. 1). The model is 0.91 m wide (thus, spanning the width of the test section) and 2.54 m long on the longest edge. The downstream or leading edge pieces can be adjusted relative to each other using precision shims to create either forward-facing or backward-facing 2D steps of different heights, parallel with the leading edge. A 3D pressure body along the ceiling was designed to induce a streamwise pressure gradient, which, along with the sweep, causes stationary crossflow growth. A second purpose of the ceiling liner was to simulate infinite swept-wing flow within a midspan measurement region of width 0.3 meters. This was achieved by designing the liner such that the C_p contours were parallel with the leading edge within the measurement region.



(a) Perspective View **(b) Top View**

Figure 1: Sketch of PIV setup.

All measurements were performed at a freestream velocity of 26.5 m/s ($Re' = 1.69$ x 10^6/m). The current experiment utilized a single leading-edge roughness configuration consisting of discrete roughness elements (DREs) with a diameter of 4.4 mm. The DREs were applied approximately 50 mm downstream of the leading edge, with a spanwise spacing, λ_z, of 11 mm and were approximately 20 μm thick. The spacing of the DREs (11 mm) corresponds to the most amplified stationary crossflow wavelength calculated for the baseline case with no step. For more details of the experiment setup, refer to Eppink.[21]

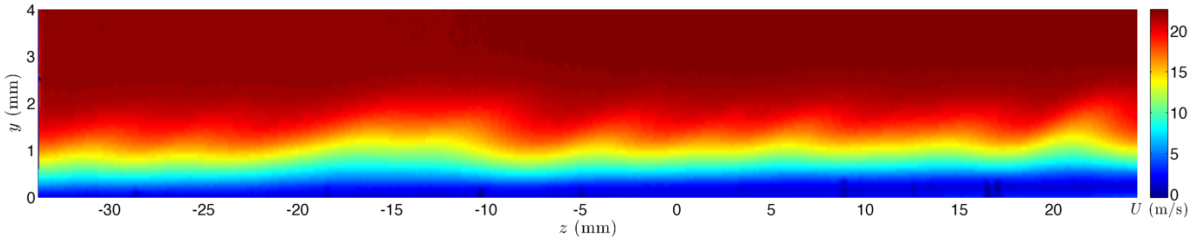
Initially, HW measurements were performed for the no-step and backward-facing step cases. The data were acquired at 25 kHz and were AC-coupled (high-pass filtered at 0.25 Hz) and low-pass filtered at 10 kHz. The hotwire was mounted on a traversing system, which allowed for several different types of boundary-layer scans to be performed. For more details on the hotwire measurement procedure, see Eppink et al.[22]

A high-speed double-pulsed Nd:YLF laser provided the laser sheet for the PIV measurements (see Fig. 1). The laser sheet was set up parallel with the leading edge and the step. The two high-speed 4-megapixel cameras that were used to acquire the TRPIV measurements were placed downstream of the step. One was placed on the outboard side of the test section at approximately 30° to the laser sheet, and the second camera was placed on the inboard side (in backward scattering) at an angle of approximately 45° to the laser sheet (Fig. 1). To achieve the desired field of view and resolution, 300 mm lenses were utilized, resulting in a

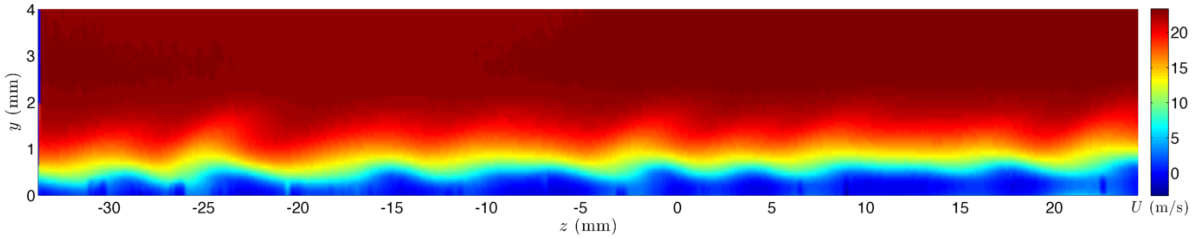
total possible measurement area of approximately 60 mm x 30 mm. For the majority of the measurements, the area of interest was reduced to approximately 60 mm x 8 mm to obtain an acquisition rate of 2 kHz. This area allowed acquisition of approximately five wavelengths of the stationary crossflow instability in a single frame, while still acquiring approximately 30 points (using 75% overlap and 16x16 pixel interrogation size) inside the boundary layer. For the mean flow measurements, data were acquired starting at the step and moving downstream at approximately 1 mm increments. Five hundred image pairs were acquired at each location. For the time-resolved measurements, data were acquired at a faster rate of 4.7 kHz, for which the area of interest was necessarily reduced further to approximately 37 mm x 5 mm, allowing acquisition of three wavelengths of the stationary crossflow instability. For these measurements, 10,000 image pairs were acquired. The PIV data were all processed using DaVis 8.2. The vector calculation was performed in multiple passes with decreasing window size. The first pass was computed with a window size of 48x48 pixels. This was followed by two iterations at either 16 x 16 pixels or 32 x 32 pixels, both with 75% overlap. These results will be compared in the time-resolved section of this paper. The cameras and laser were all mounted on the same traversing system, which allowed measurements at multiple locations with relative ease. An oil-based fog machine generated the seeding, which was introduced downstream of the test section. The particle size is estimated to be approximately 1 μm. Using the approach of Mei,[23] we can estimate a low-pass frequency cut-off of the particles, assuming a certain particle diameter and density ratio. If we assume the particle density is approximately that of water, and the particle diameter ranges from 0.5 to 2 μm, the estimated low-pass frequency cutoff ranges from 14 to 227 kHz. Thus, the particles in this flow should be able to respond to the maximum frequencies encountered in the present experiment, in the 800 to 1500 Hz range. The stereo-PIV calibrations were performed by placing a target with known dot size and spacing on the model, aligned as closely as possible with the actual laser sheet position. The traversing system was then moved downstream at 0.5 mm increments, and calibration images were obtained at three stations. These images were processed using a polynomial calibration to obtain the mapping functions for each camera.

III. Uncertainty Analysis

In this section, results from two different cases will be utilized for the uncertainty analysis. These consist of a plane from the no-step case, and a plane taken approximately 3 mm downstream of the 1.7 mm forward-facing step (FFS). The time-averaged streamwise velocity for both cases is shown in Fig. 2. Notice that there is a small positive slope in the contours of these images, which is likely due to a slight misalignment of the calibration target with the model surface. In the following, this slope was removed before the computation of the spanwise averages and mean perturbation profiles.



(a) No step



(b) Downstream of 1.7 mm FFS

Figure 2: Mean streamwise velocity for both cases studied.

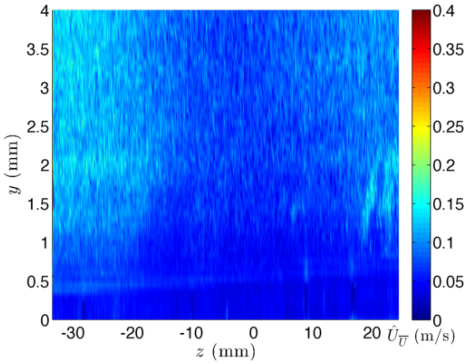
A. Uncertainty Analysis from Existing Techniques

First, we consider the standard approach of approximating the mean uncertainty of the measurements, following the propagation methods of Sciacchitano and Wieneke.[24] The mean uncertainty is not exactly a measure of the uncertainty of the PIV technique itself, but, rather, is dependent on the measured fluctuations. The mean uncertainty depends on the standard deviation of the measured quantity and the effective number of independent samples (Eq. 1). Since there are typically real fluctuations in the flow, the standard deviation will include effects from both the real fluctuations and the measurement error. The effective number of independent samples (N_{eff}) is computed from Eq. 2, where ρ is the autocorrelation coefficient.

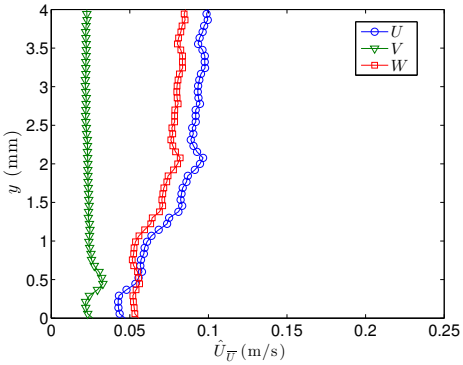
$$\hat{U}_{\overline{x}} = \frac{\sigma_x}{\sqrt{N_{eff}}} \tag{1}$$

$$N_{eff} = \frac{N}{\sum_{n=-\infty}^{+\infty} \rho(n\Delta t)} \tag{2}$$

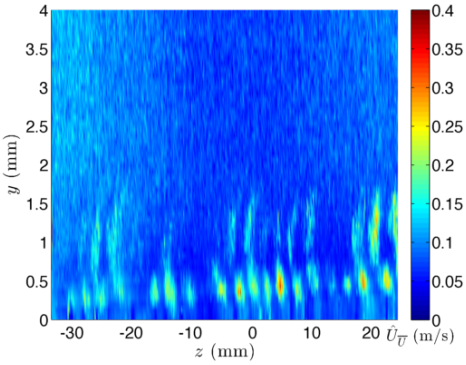
Since these results were acquired using the time-resolved system, there is a period of time over which the samples are correlated, thus, $\sum_{n=-\infty}^{+\infty} \rho(n\Delta t) > 1$, and N_{eff} is smaller than the total number of samples, N, meaning the uncertainty is larger. The results for the mean uncertainty of the U component across the whole measurement plane is shown in Figs. 3a and 3c for the no-step and forward-facing step cases, respectively. The pattern seen in Fig. 3c is due to the spatial variation of both the standard deviation of U and N_{eff}, which are both modified in this case due to the very strong stationary crossflow vortices. Spanwise-averaged profiles of the uncertainty of all three velocity components are shown in Figs. 3b and 3d. The mean uncertainty of U and W are the largest, remaining around 0.1 m/s throughout the boundary layer. The mean uncertainty of V is much smaller, remaining at values near 0.002 m/s.



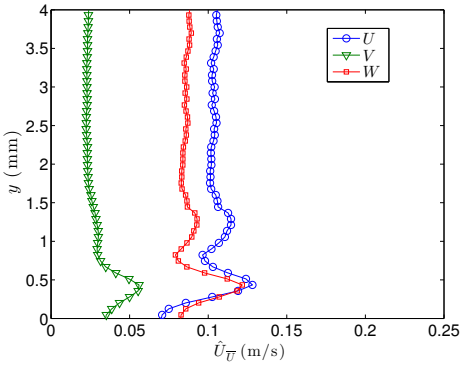
(a) No step, mean U uncertainty



(b) No step, spanwise-averaged mean uncertainty



(c) 1.7 mm FFS, mean U uncertainty

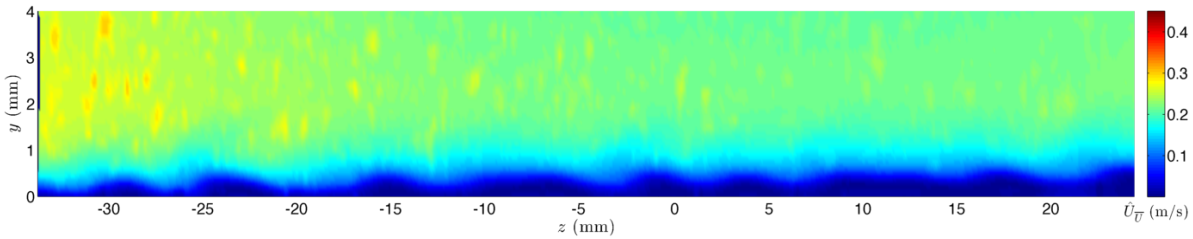


(d) 1.7 mm FFS, spanwise-averaged mean uncertainty

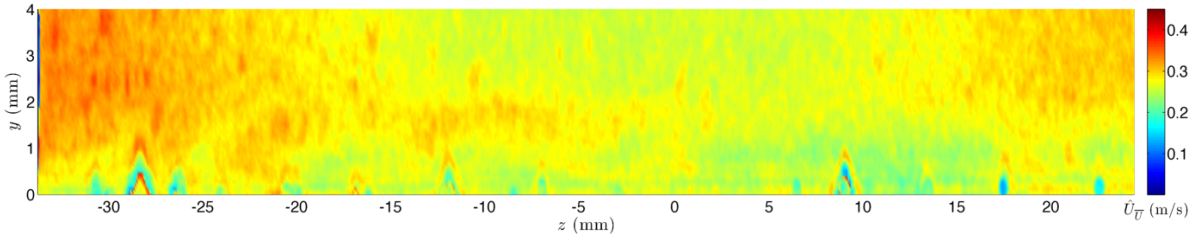
Figure 3: Mean uncertainty for both cases.

Next, the approach of Bhattacharya et al.[8] is used to estimate the uncertainty. This approach takes into account the effect of the stereo registration errors, which are incorporated into the analysis as uncertainties in the camera angles. These camera angle uncertainties are estimated from the measured disparity field of each image, which is used to estimate the change in the laser sheet location. This is then propagated through as an uncertainty in the calibration coefficients, which is used to compute the uncertainty in the camera angles. The camera angle uncertainties are then used to compute the uncertainty of the velocity. This portion of the velocity uncertainty is referred to as the angular uncertainty. The total uncertainty is then computed by combining the angular uncertainty and the planar uncertainty, which is computed using the correlation statistics approach[4] that is built into DaVis.

Results were computed for just the angular uncertainty, as well as the total uncertainty. The streamwise velocity uncertainties for the 1.7 mm step height case are shown in Fig. 4. For these results, we introduced a constant offset of 0.5 mm in the out-of-plane direction, in addition to the random disparity, so we could study how well the bias error would be predicted with this technique compared to the techniques presented later. Notice that in Fig. 4a, the angular uncertainty has a similar shape to the mean velocity (Fig. 2b). This occurs because the angular uncertainty depends on the local velocity field. The angular uncertainty is low near the wall and increases to around 0.2 to 0.25 m/s towards the edge of the boundary layer. The planar uncertainties, on the other hand, are significantly larger near the wall, but smaller away from the wall. This is more evident in Fig. 5, which shows spanwise-averaged results for both cases and for all three velocity components. With both angular and planar uncertainties considered, the total uncertainty is still relatively low, remaining at approximately 0.35 m/s or lower throughout the boundary layer.

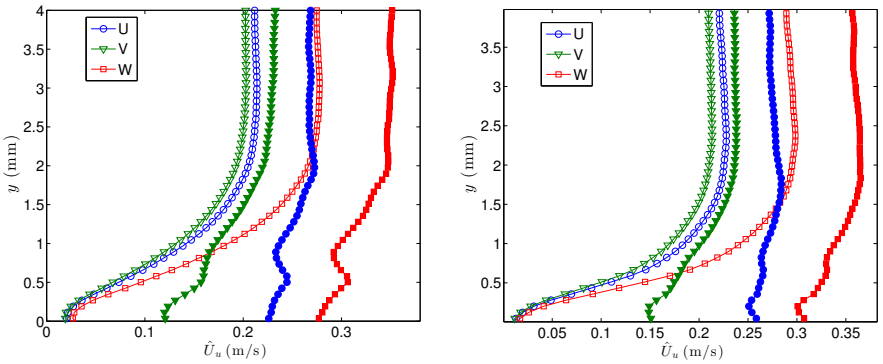


(a) Angular uncertainty only.



(b) Total uncertainty.

Figure 4: U-velocity uncertainties computed using the approach of Bhattacharya et al.[8] for the 1.7 mm FFS case.

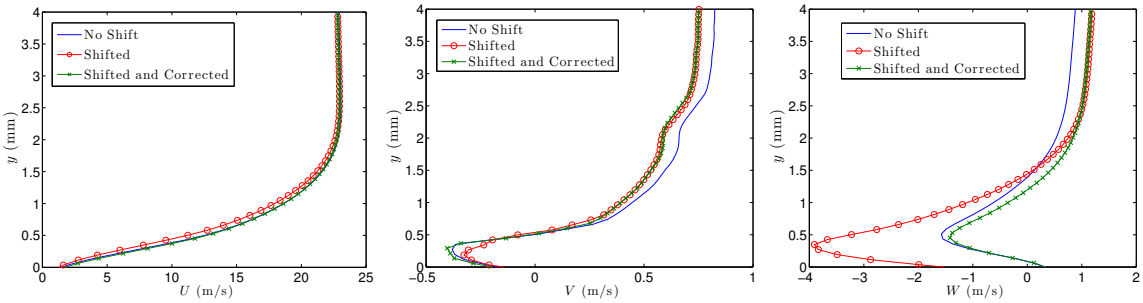


(a) No step **(b) 1.7 mm FFS**

Figure 5: Spanwise-averaged velocity uncertainties computed using the approach of Bhattacharya et al.[8] for both cases, including angular uncertainties only (empty symbols) and total uncertainties (filled symbols).

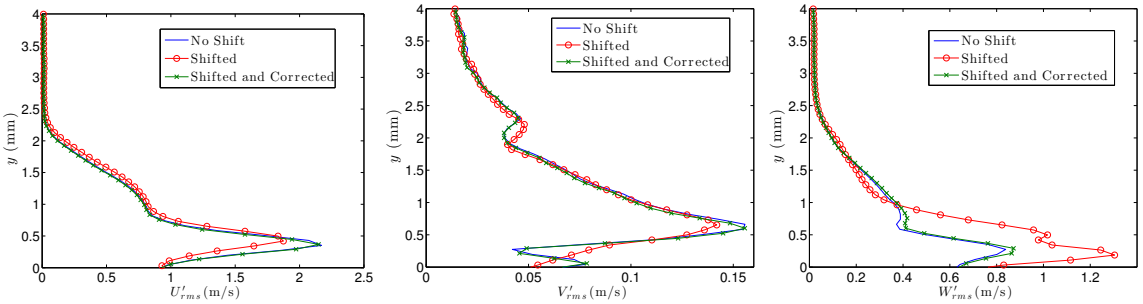
B. Effect of Small Vertical Misalignment

One of the major challenges of the current setup is the extreme sensitivity to very small (and often, unintended and undetected) changes to the setup. Because of the high magnification and long focal length of the current setup, any small undetected change to the camera or lens position will be magnified at the measurement plane. These issues can be especially problematic since a self-calibration typically assumes the camera positions have not changed, and only accounts for a change in laser sheet position. The biggest problem in this regard with the current setup was an apparent change in the relative vertical camera positions. This shift was evident by examining the dewarped particle images for both cameras and comparing the wall location, which was evident due to reflection across the surface. Typically, if the data were obtained immediately before or after a careful calibration was performed, the wall locations matched well. However, over time, and typically due to movement of the traversing system, the relative vertical locations could shift at the measurement plane by as much as 10 pixels (which translates to approximately 0.1 mm). While this may not seem like a large shift, the discrepancy was apparent in the results due to the large velocity gradients near the wall. To test the effect of a known shift, the initial calibration was performed, followed by a second calibration where the system was translated vertically by 0.1 mm, and only the images from camera one were retaken. Then, the calibration mapping function was computed using the original camera two images, resulting in a vertical discrepancy between camera one and camera two. The PIV images were processed with both calibrations, and the results are shown in Figs. 6 and 7 for the 1.7 mm FFS case. Fig. 6 shows spanwise-averaged mean profiles of each of the velocity components for the original (blue line) calibration, and the shifted (red diamond) calibration. Discrepancies are evident near the wall for all velocity components, but it is especially apparent for the W-velocity component. Fig. 7 shows the mean spanwise perturbation profiles for each velocity component. The stationary crossflow (or perturbation) profiles are essentially a measurement of the spanwise modulation of the mean flow caused by the stationary crossflow vortices. The steady disturbance velocity profiles (U', V', W') are first computed by subtracting the spanwise-averaged profile from the velocity profile at each spanwise (z) location. Then, the perturbation profiles are calculated from the wavelength spectra of the respective steady disturbance velocity (U', W', V') across the span (integrated across a wavelength range of 8 to 20 mm). Again, the shifted calibration causes some differences to the mean perturbation profiles for every velocity component, but especially the W component.



(a) U-profile **(b) V-profile** **(c) W-profile**

Figure 6: Mean profiles for the vertical offset calibrations downstream of the 1.7 mm FFS.



(a) U-perturbation profile **(b) V-perturbation profile** **(c) W-perturbation profile**

Figure 7: Mean perturbation profiles for the vertical offset calibration downstream of the 1.7 mm FFS.

Short of performing calibrations for each and every measurement plane (which was not feasible given the extra time that would be required), the best solution we found to this problem was to artificially shift the particle images to match the wall location. Since this is a very small shift and the calibration does not vary greatly over this small distance, this should not cause any major additional error. To test this hypothesis, we first applied a vertical shift to the particle images from camera one so that the dewarped wall location (using the shifted calibration) would once again match camera two. Then, we processed the data using the shifted calibration, and compared the results to the unshifted calibration results. These results are also included in Figs. 6 and 7 as the"shifted and corrected" results (shown as green xs). In general, the corrected results match the original results very well, though there are some small discrepancies evident, particularly in the freestream for the V and W-mean profiles. This discrepancy, however, could be due to small differences in the calibrations themselves, not necessarily related to the vertical shift. In fact, this is likely the case, since the velocity in the freestream would be expected to match since there is no vertical velocity gradient. Overall, the results illustrate that the approach of shifting the images so that the wall-location matches for both cameras is successful at recovering the correct velocities.

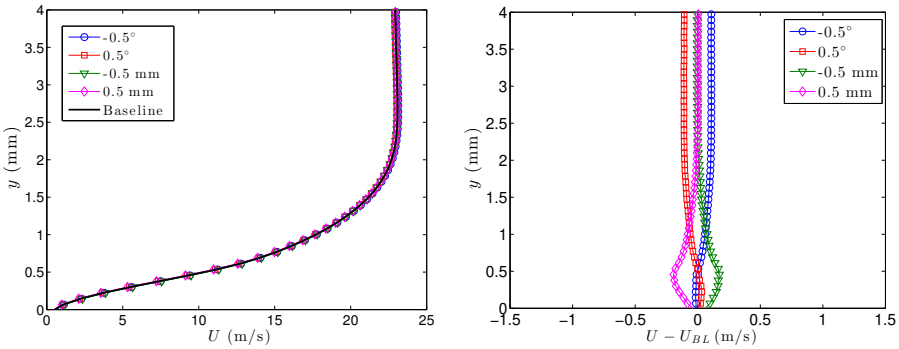
C. Effects of Registration Errors

One of the main efforts undertaken in this work was to assess the effect of the stereo calibration registration error on the final results since this is believed to be one of the major possible error sources. As mentioned earlier, stereo registration errors can come about due to either poor alignment of the calibration target with the known laser plane location, or due to a shift in the laser plane location. This effect was evaluated by purposefully misaligning the calibration target by a known amount ($\pm$ 0.5 mm) in the out-of-plane direction, as well as $\pm$ 0.5° in the horizontal angle. These "bad" calibrations were then used to bound the uncertainty due to the target misalignment. These results will also be used to evaluate the effectiveness of the self-calibration procedure.

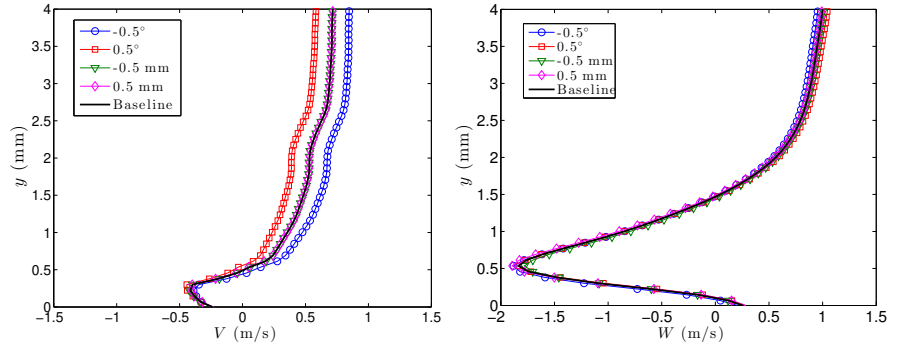
It was first necessary to determine an adequate initial target alignment, which was not simple with the current setup. The target was first aligned manually as closely as possible with the measurement plane, and a three-plane polynomial calibration was performed. Data that were acquired with the current setup were then used to compute a self-calibration using this initial calibration. The planar and angular offsets estimated by the self-calibration were then used to precisely adjust the position of the calibration target, and a new calibration was performed and the calibration quality assessed again using the data. Only a few iterations of this approach were required to achieve alignment within ± 0.02 mm and ± 0.06° degrees, according to the self-calibration. This calibration is referred to as the baseline calibration throughout this paper. Once the baseline calibration plane was determined, we then performed 4 additional calibrations by artificially perturbing the location of the calibration target. The spatial offsets of ±0.5 mm were achieved artificially within the calibration computation by changing the initial plane location. The angular offsets of ±0.5° were performed by modifying the calibration coefficients of the baseline calibration to achieve the desired angular rotation about the center of the plane ($z \approx$ -5 mm). These offsets of ± 0.5 mm and ± 0.5° were chosen to bound the uncertainty that would be introduced by our initial calibration target placement. From the offset found from the self-calibrations, we found that we were consistently able to place the target within these offsets. In addition, self-calibrations were performed for each of the four offset calibrations to assess the effectiveness of the self-calibration at correcting the known offsets. One initial observation is that the offsets found by the self-calibration procedure were often not correct, since it tended to add an out-of-plane shift and angular displacement at the same time to account for the simple out-of-plane shift. Additionally, for the angle-misaligned calibrations, the self-calibration tended to include a small vertical angle adjustment in addition to the horizontal angle adjustment.

Several cases were investigated, which for the most part produced similar results. The case shown here corresponds to the measurements performed approximately 3 mm downstream of a 1.7 mm FFS. This case is interesting because there are some very strong spanwise and wall-normal velocity gradients, particularly near the model surface. Results for the five different calibrations are shown in Fig. 8 for the spanwise-averaged mean profiles for each of the velocity components. The velocity components presented here are not the directly measured components, but are transformed into the tunnel coordinates, U and W, as illustrated in Fig. 1. The offset calibrations have a very small effect on the mean U-profiles. The ±0.5 mm calibrations are the worst near the wall (Fig. 8b), but collapse back to the baseline profile near the edge of the boundary-layer. The ±0.5° calibrations agree well near the wall but are off near the freestream by about 0.1 m/s. The V-profiles (Fig. 8c) all collapse very well, except for the ±0.5° calibration, which is off by approximately 0.05 m/s near the freestream. The W-profiles collapse fairly well as well, though there are some small discrepancies (around 0.05 m/s) near the freestream for the ±0.5° calibrations.

Since one of the primary measurements of interest for this case is the stationary crossflow amplitude, we also need to know how the stereo registration errors affect the measured stationary crossflow amplitude. The mean perturbation profiles are shown in Figs. 9a-9c. The U-perturbation profiles (Fig. 9a) all collapse very well except for the -0.5 mm calibration, which results in a reduction in max amplitude of approximately 20% compared to all the other cases. Similarly, the V-perturbation profiles (Fig. 9b) collapse very well, except for the -0.5 mm case, which is approximately 35% less than the other cases in peak amplitude. Except for the -0.5 mm case, the W-perturbation profiles (Fig. 9c) agree fairly well, though not as well as the U and V-perturbation profiles. The -0.5 mm calibration, in this case, results in an amplitude approximately twice as large as the other cases. It is interesting that despite the fact that the spanwise-averaged mean profiles for the -0.5 mm case agree well with the baseline calibration, the calibration registration error still results in such a large change in the mean perturbation amplitudes. This shows that the misalignment of the laser sheet and the calibration target can result in increased (or decreased) spanwise modulation of the measured velocity, while not affecting the overall average velocity. It is also interesting that the +0.5 mm calibration results are not similar to the -0.5 mm results, since one would expect a shift in either direction to result in similar errors. Overall, these results illustrate that a small misalignment of the target and the laser plane can result in significant errors in the stationary crossflow amplitude, even though only minor errors occurred in the spanwise-averaged mean profiles. For configurations such as the present setup, with high magnification and sensitivity to misalignment, researchers should be aware of these effects on their specific setup and the measurements of interest before drawing conclusions based on their quantitative results.

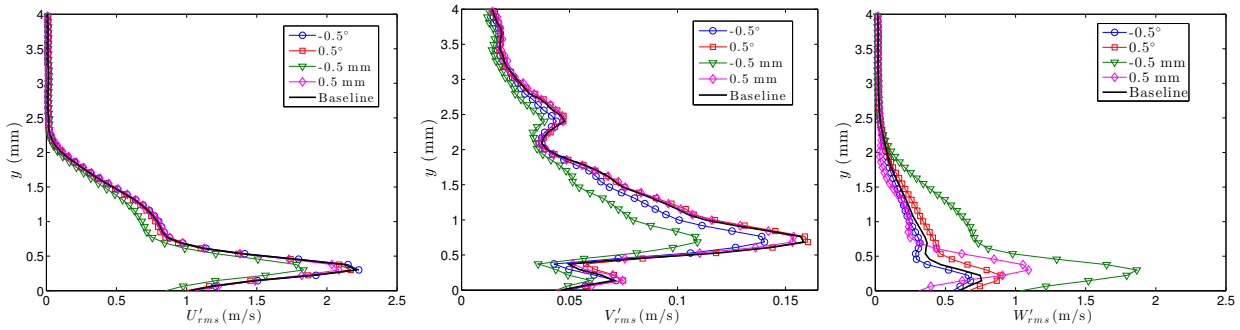


(a) U-profile **(b) U-profile with baseline U profile subtracted**



(c) V-profile **(d) W-profile**

Figure 8: Spanwise-averaged profiles for the five different calibrations downstream of the 1.7 mm FFS.



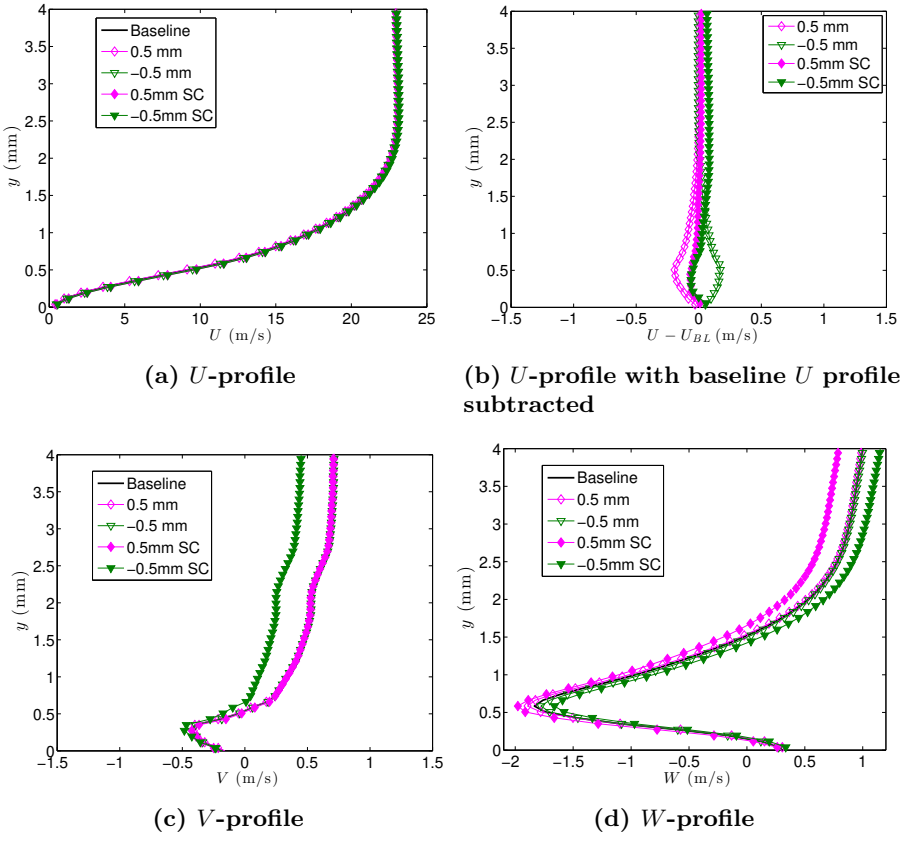
(a) U-perturbation profile **(b) V-perturbation profile** **(c) W-perturbation profile**

Figure 9: Mean perturbation profiles for the five different calibration downstream of the 1.7 mm FFS.

D. Self-Calibration Performance

As mentioned, many PIV users rely on self-calibrations to correct for misalignments between the calibration plate and laser plane. This may be sufficient in many cases, but particularly for challenging setups such as the current work, it is important to assess the effectiveness of the self-calibration. Thus, self-calibrations were performed for each of the 4 offset calibrations to see how well the self-calibration could correct for the known offset. The self-calibrations were performed with 500 particle images using a sum-of-correlation approach. The results for the ±0.5 mm offset cases, with and without the self-calibration, are shown in

Fig. 10. The results for the ±0.5° cases are not shown here, but are qualitatively similar to the results for the ±0.5 mm cases. The spanwise-averaged mean profiles are plotted in Figs. 10a-10d. For the U-profile (Figs. 10a-10b), the self-calibrations help to reduce the error near the wall. However, for the -0.5 mm case, the self-calibration results in more error in the freestream. Similarly, for the V-profile (Fig. 10c), the -0.5 mm self-calibration results in more errors, both near the wall and in the freestream. For the W-profiles (Fig. 10d), the self-calibrations for both cases result in discrepancies throughout the boundary layer, particularly near the freestream. It was mentioned earlier that the self-calibration for the ± 0.5 mm cases attempts to correct the calibration plane location by not only shifting it in the out-of-plane direction, but also by rotating it. This is obviously incorrect, since we know that the offset was purely a planar shift, but the self-calibration process has no way of knowing that, and in a typical situation the user would not know whether or not this is correct. The large differences in the spanwise-averaged mean profiles are likely due to this problem with the self-calibration. The proposed solution, and the approach adopted for the current work, is to use the self-calibration to get an idea of the actual physical offset, and then physically correct for that offset as much as possible. It is possible within the calibration procedure to correct for a shift offset by setting the initial location of the calibration plane. However, it is not easy to correct for a rotation offset, though, as in the present work, this can be done by modifying the calibration coefficients. If the rotation offset is small enough, this may not be a problem, and one could proceed or perform the self-calibration, if desired, after correcting for the shift offset. Each case is unique and may not be as sensitive as the present case to small misalignments and self-calibration errors, so it is important to quantify the effect for a given setup.



(a) U-profile

(b) U-profile with baseline U profile subtracted

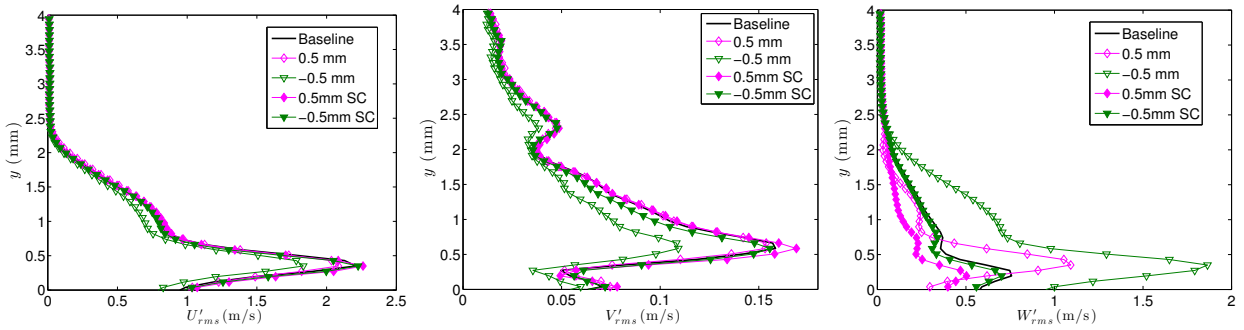
(c) V-profile

(d) W-profile

Figure 10: Spanwise-averaged profiles for the two shifted calibrations and their self-calibrations downstream of the 1.7 mm FFS.

The results for the stationary crossflow profiles are shown in Figs. 11a-11c. There is not much difference between the two +0.5 mm calibrations, though there is some change in the W-perturbation profile once the self-calibration is applied. Recall from Fig. 8 that the -0.5 mm calibration was the worst performer in terms of the stationary crossflow profiles and amplitudes. Surprisingly, once the self-calibration is applied, these results improve significantly for every velocity component. It is not immediately clear how or why the

self-calibration would result in improvement in the spanwise flow modulation results while causing larger discrepancies in the spanwise-averaged results. This result illustrates the importance of evaluating the effect of the error sources on the final quantity of interest.



(a) U-perturbation profile **(b) V-perturbation profile** **(c) W-perturbation profile**

Figure 11: Mean perturbation profiles for the two shifted calibrations and their self-calibrations downstream of the 1.7 mm FFS.

E. Approach for Estimating Uncertainty Due to Registration Errors

The angular uncertainties that were estimated using the approach of Bhattacharya et al.[8] for a 0.5 mm planar offset were on the order of 0.2 to 0.35 m/s near the edge of the boundary layer, but much lower near the wall. The errors that were produced by performing a 0.5 mm offset, in a spanwise-averaged sense (such as those results shown in Fig. 10), tend to fall within this estimate, although errors near the wall tend to be larger than those estimated by the angular uncertainties. However, this level of estimated uncertainty definitely does not account for the errors that are found in the stationary crossflow profiles due to the offset, such as those shown in Fig. 11. The reason for this is that the primary source of error due to misalignment of the calibration target and laser plane is the combination of velocities from two different locations (as illustrated in Fig. 12). This effect is not accounted for in the approach by Bhattacharya et al.[8] In a situation where there is no significant velocity gradient, this may not produce significant errors. However, in the current situation, there are significant velocity gradients present, both in the spanwise and wall-normal directions. Both of these gradients can result in errors due to the misalignment, however, since the cameras are at a nearly 0° angle with respect to the x-z plane, the wall-normal gradient was found to have no significant effect and is neglected in the following analysis.

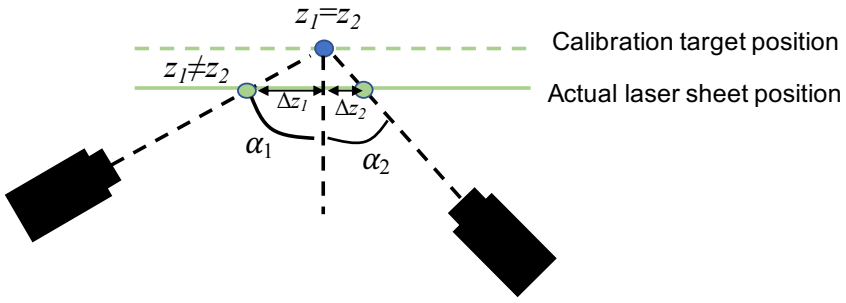
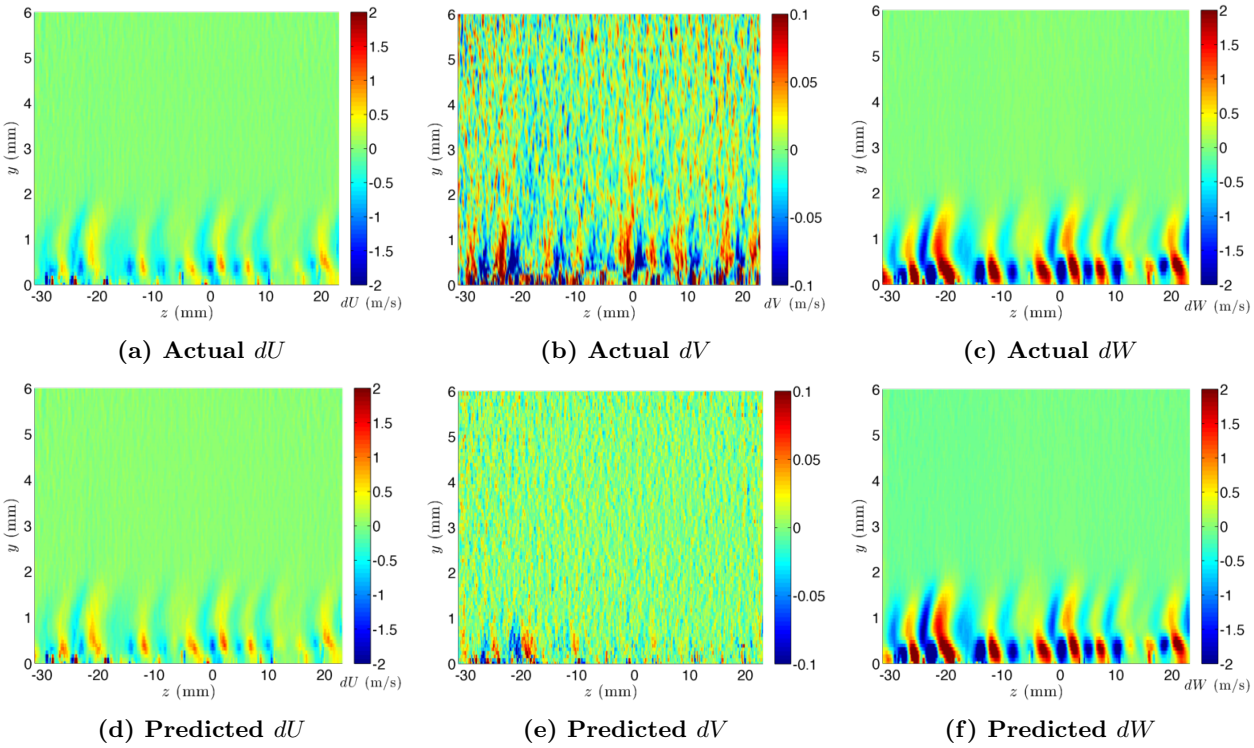


Figure 12: Sketch of registration error effect.

A simple approach was undertaken to estimate the bias errors given a constant planar offset or rotation of the laser sheet from the calibration target. First, we compute the stereo angles for each camera from the calibration mapping function gradients, following Giordano et al.[25] As mentioned earlier, for the current analysis we only consider the in-plane (α) angles and neglect the out-of-plane (β) angles since the latter are very close to zero. Once the angles are computed, the Δz values for each camera can be computed given a known shift in the plane location. Next, we use the velocity field that was computed from our

baseline calibration to reconstruct the velocity field with a known shift. To do this, we first deconstruct the velocity components (u, v, w) back into their U and V components for each camera (U_1, V_1, U_2, V_2). This is done following the velocity reconstruction equations of Giordano et al.[25] Next, we find the deconstructed velocities (U_1, U_2, V_1, V_2) at the $z + \Delta z$ locations for every point in the plane for each camera, and then recombine them to obtain the new u, v, and w components. The difference in the velocity from the baseline result (dU, dV, dW) is shown in Fig. 13 for a positive 0.5 mm shift. Results are shown for the method described in Sec. IIIC in Figs. 13a-13c and are referred to as the actual results. The predicted results, using the current technique, are shown in Figs. 13d-13f. As illustrated in these figures, a very small offset of 0.5 mm can result in very large local velocity errors on the order of 1-3 m/s for the U and W components, which is an order of magnitude larger than the errors predicted using any of the other methods. It is also evident from these figures that the technique for estimating the errors caused by the offset is very successful for the U and W components. Unfortunately, the method does not accurately predict the error in the V component, which is an order of magnitude smaller than the other velocity components. This discrepancy could be due to the noisy nature of the V-velocity results, which would hinder an accurate recombination of the velocities at two different locations.



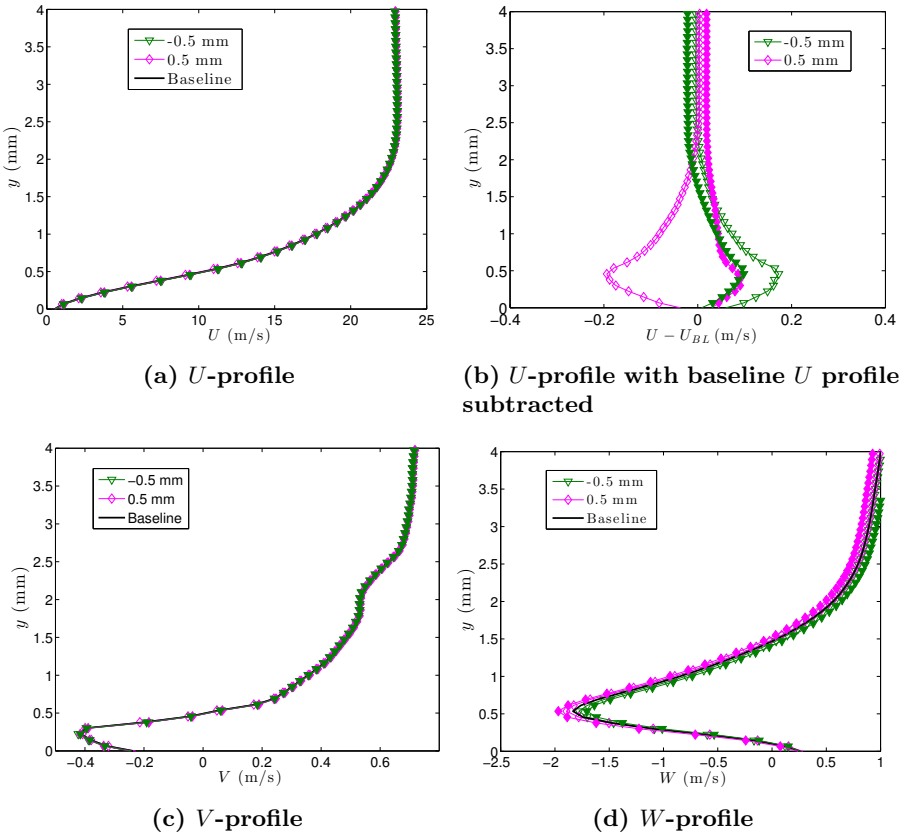
(a) Actual dU **(b) Actual dV** **(c) Actual dW**

(d) Predicted dU **(e) Predicted dV** **(f) Predicted dW**

Figure 13: Differences between the baseline and 0.5 mm shifted calibration results downstream of the 1.7 mm FFS.

These results are investigated further in Figs. 14 and 15 for the 1.7 mm FFS case. Fig. 14 shows the spanwise-averaged results for each velocity component. The actual results are shown as the empty symbols, and the predicted results are the filled symbols. There are some discrepancies between the mean profiles, particularly evident in Fig. 14b, which shows the difference between the shifted U-profiles and the baseline U-profile. While the prediction technique does predict some difference in the mean U profile, it does not match the actual results. There is also some discrepancy in the freestream. The prediction approach also somewhat overpredicts the difference in the W-profiles (Fig. 14d), though it does follow the same trend as the actual results. The discrepancies near the wall for this approach may be attributable to the slight shift in wall location that occurs when the calibration plane (or laser plane) is shifted. As shown in Sec. IIIB, a small vertical shift can have a large effect on the results. This effect was not accounted for in the current approach, and may not have been perfectly corrected for when the shifted calibrations were applied. Further analysis is required to determine if this is the true source of the discrepancy.

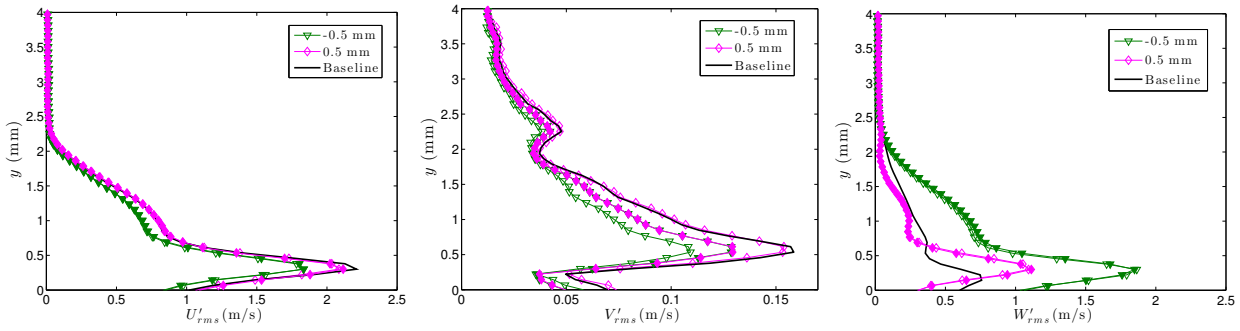
Results for the mean-perturbation profiles are shown in Fig. 15. The predicted results agree very well with the actual results for the U and W components. There is, however, some discrepancy between the predicted and actual V-perturbation profiles, as is also evident from Figs. 13b and 13e.



(a) U-profile **(b) U-profile with baseline U profile subtracted**

(c) V-profile **(d) W-profile**

Figure 14: Spanwise-averaged profiles for the two shifted calibrations (empty symbols) and the estimated results (filled symbols) downstream of the 1.7 mm FFS.

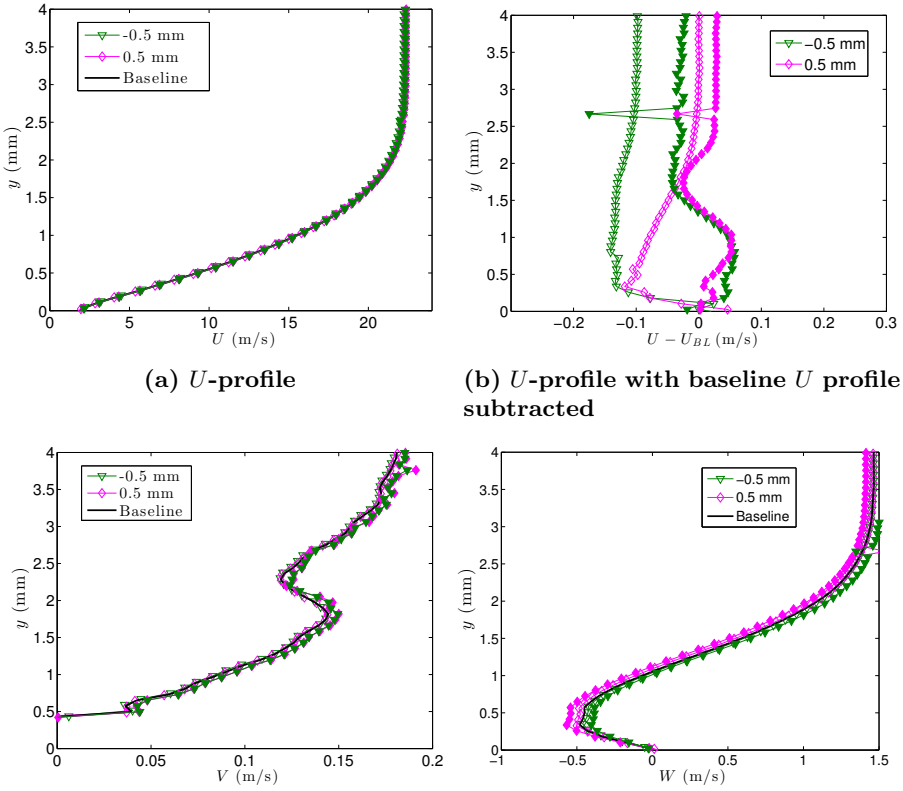


(a) U-perturbation profile **(b) V-perturbation profile** **(c) W-perturbation profile**

Figure 15: Mean perturbation profiles for the two shifted calibrations (empty symbols) and the estimated results (filled symbols) downstream of the 1.7 mm FFS.

The same analysis was applied to the measurements from the no-step case, and similar results were obtained, as shown in Figs. 16 and 17. Once again, the spanwise-averaged results show some discrepancies in all velocity components, but the perturbation profiles match very well. The V-perturbation profile is not shown since the stationary crossflow amplitude was much smaller for this case, resulting in a very noisy and

unrealistic V-perturbation profile. Overall, the results show that this technique can be useful for estimating the velocity error caused by a misaligned calibration target or laser plane. Given that the self-calibration may not completely correct for these misalignment errors, it is useful to have an approach that can be used to bound the uncertainty by prescribing a maximum possible offset or angle.



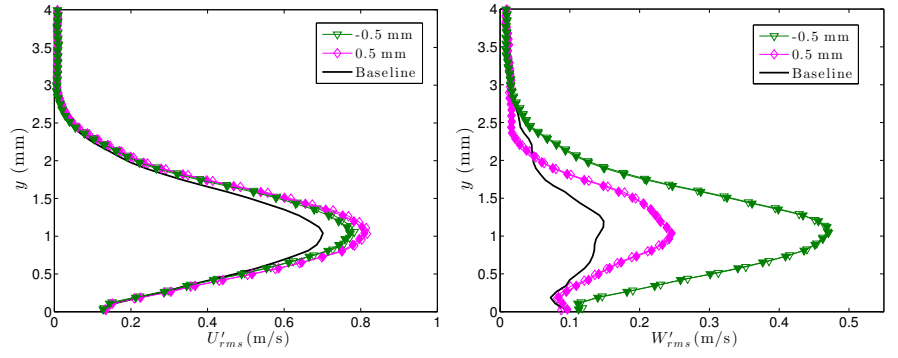
(a) U-profile

(b) U-profile with baseline U profile subtracted

(c) V-profile

(d) W-profile

Figure 16: Spanwise-averaged profiles for the two shifted calibrations (empty symbols) and the estimated results (filled symbols) for the no-step case.



(a) U-perturbation profile

(b) W-perturbation profile

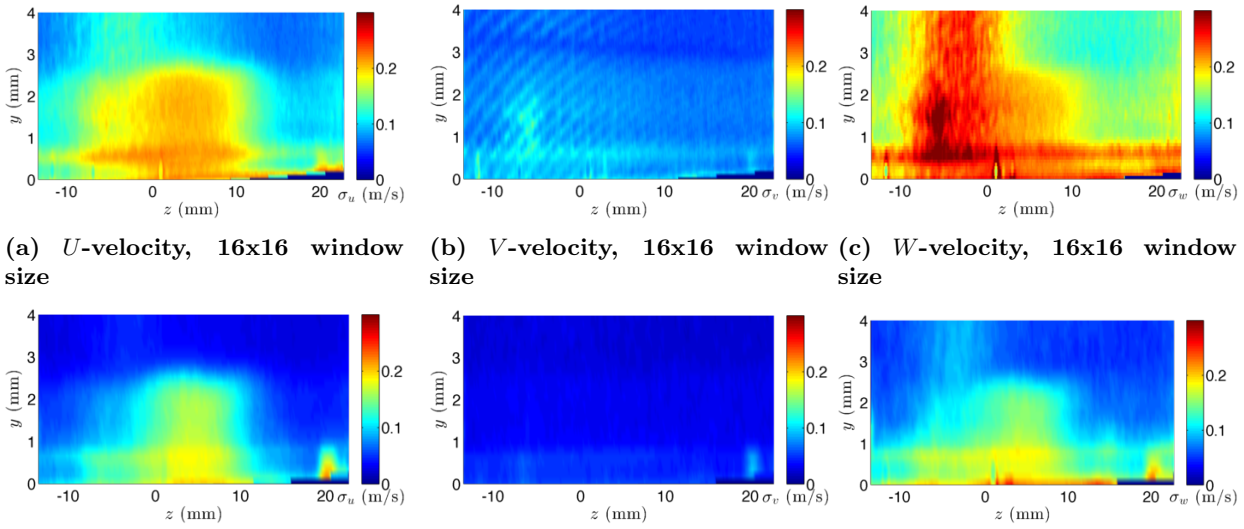
Figure 17: Mean perturbation profiles for the two shifted calibrations (empty symbols) and the estimated results (filled symbols) for the no-step case.

IV. Comparison between PIV and HW Results

In this section, we will compare results for the TRPIV and HW measurements performed downstream of a 1.18 mm backward-facing step (BFS). The reason the BFS case is used for this comparison is that this case resulted in a broad spectrum of unsteady disturbances. Thus, this case makes for a good comparison between both the mean and time-resolved PIV and HW results. Note, however, that this experiment was not designed or carried out with the intent of doing a quantitative comparison between two different measurement techniques. There are several factors that make an exact comparison impossible. First, measurements were likely not performed in exactly the same locations either in the streamwise or spanwise directions for the two techniques. The data that are compared with the two techniques are estimated to be within ± 5 mm of each other in the streamwise direction. The spanwise offset could be as much as 25 to 50 mm, and could result in differences in the measurements due to small amounts of spanwise nonuniformity. The second factor that could affect the comparison is the fact that we used a different batch of DREs for the PIV measurements, which were thicker than the original batch and were shown to result in approximately a 25% difference in stationary crossflow initial amplitude.[26] Despite these difficulties, it is still advantageous to show comparisons between the two techniques. Since PIV is only recently and very rarely being used for transition research, it is worthwhile to compare the results to a well-trusted and accepted technique (such as HW anemometry) to ensure that the introduction of particles into the flow does not significantly affect the transition process, and to ensure that the measurement technique is able to resolve the instabilities.

A. Uncertainty of Time-Resolved Data

The planar uncertainties were calculated for each image using the correlation statistics approach[4] that is built into DaVis. The results shown in Fig. 18 are acquired by taking the standard deviation of the uncertainty from each of the 10,000 images. Results are shown for both window sizes used during the analysis: the 16 x 16 results are on the top, and the 32 x 32 results are on the bottom. The maximum U and W uncertainties tend to be localized near the center of the figure and reach values near 0.3 m/s. The V uncertainty is much lower, typically less than 0.1 m/s. The uncertainty is lower when the window size is increased, as the maximum values for the 32 x 32 cases are typically less than 0.2 m/s.



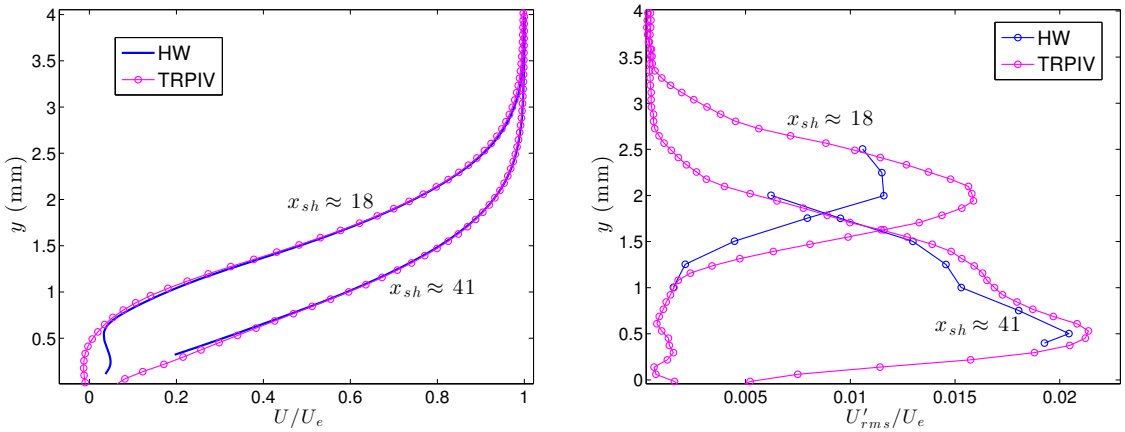
(a) U-velocity, 16x16 window size **(b) V-velocity, 16x16 window size** **(c) W-velocity, 16x16 window size**

(d) U-velocity, 32x32 window size **(e) V-velocity, 32x32 window size** **(f) W-velocity, 32x32 window size**

Figure 18: Planar uncertainties for each velocity component, 40 step heights downstream of the 1.18 mm BFS.

B. Comparison of Mean Flow Results

Some comparisons of the mean flowfield results for the HW and TRPIV measurements are shown in Fig. 19. These include spanwise-averaged U-profiles (Fig. 19a) and stationary crossflow profiles (Fig. 19b) for two different locations downstream of the step. The quantity x_{sh} is the streamwise distance downstream of the step normalized by the step height. The mean U-profiles for the HW data were computed from $y-z$ (or planar) HW scans, while the stationary crossflow results were computed from multiple z-scans (spanwise scans at a constant y-location) that were performed at the same streamwise location, since the z-scans allowed a larger spanwise extent to be covered in a shorter time. The mean U-profiles compare very well between the two techniques, though there are some slight discrepancies near the wall. Note that the HW data near the wall for the x_{sh}=18 location (where U begins to increase again as the wall is approached) is not expected to be correct due to the separated flow at this location, and also due to our inability to resolve the direction of the flow with the HW technique. The stationary crossflow profiles (Fig. 19b) compare fairly well qualitatively, but the amplitude of the TRPIV result at the upstream station (x_{sh}=18) is significantly larger than the HW amplitude. This difference is expected due to the different batch of DREs that was used for the TRPIV measurements, as discussed earlier. The amplitudes at the downstream station are much closer and exhibit the same distorted mode shape. The better agreement in amplitude at this location is likely due to saturation of the stationary crossflow mode.



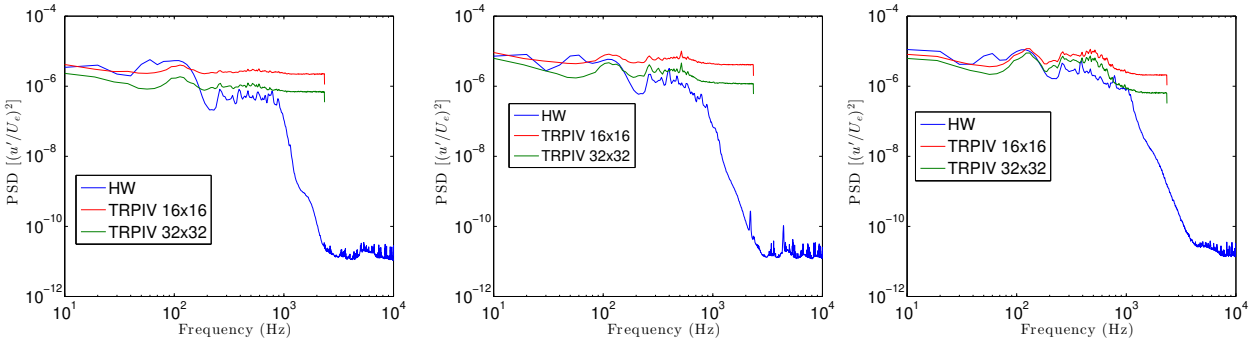
(a) Spanwise-averaged U-profiles **(b) U-perturbation profile**

Figure 19: Spanwise-averaged mean and stationary crossflow profile comparison for two locations downstream of the 1.18 mm BFS.

C. Comparison of Unsteady Results

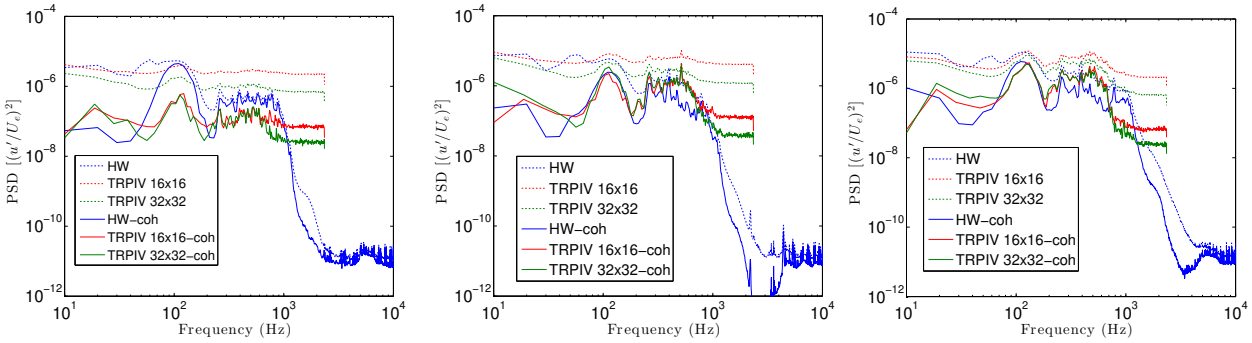
Streamwise velocity (u') spectra from the TRPIV and HW results are compared in Fig. 20 for three different streamwise locations. These spectra were computed at a constant height of $y \approx 0.75$ mm and were averaged across the span of approximately 60 mm. The spectra were computed using 50 averages for the HW data, and 20 averages for the TRPIV data, which resulted in similar frequency resolutions (Δf = 10 and 9.4 Hz, respectively). The spectra were then multipled by this Δf to obtain a more direct comparison. The lines in these figures represent the spectra computed without any additional post-processing performed. Results are shown for both window sizes tested for the TRPIV post-processing: 16 x 16 and 32 x 32 pixels. The instabilities of interest lie in three different frequency ranges, as found from the previous HW measurements.[22] The frequency ranges are 80 to 200 Hz, 200 to 800 Hz, and 800 to 1500 Hz, corresponding to traveling crossflow, Tollmien-Schlichting (TS), and shear-layer instabilities, respectively. These features are evident in the HW spectra, particularly the two lower frequency ranges. The TRPIV data is significantly noisier than the HW data. In fact, the noise floor is about 6 orders of magnitude higher. Notice that the noise floor is much higher for the 16 x 16 TRPIV data compared to the 32 x 32. In general, the results for the 32 x 32 TRPIV data agree better with the HW data. For the farthest downstream location (x_{sh}=50 in Fig. 21c), the features corresponding to the traveling crossflow and TS instabilities (80 to 200 Hz and

200 to 800 Hz) are visible in the total spectra for the TRPIV data (red and green lines). However, these features become less and less evident as we move upstream due to the decreasing amplitude relative to the noise floor. One approach to resolve the noise issue is to plot the coherent spectra, since it is expected that the instabilities should be coherent over a significant spatial distance. These results are shown in Fig. 21, in which the dotted lines are the total spectra (as shown in Fig. 20), and the solid lines are the coherent spectra. The HW scans that are presented here were performed with a second wire at a fixed location inside the boundary layer. Thus, we can compute the coherence between the two HW signals, and we can do the same with the TRPIV results by choosing a fixed location within the plane to act as the reference point. Then, the coherence at each frequency is multiplied by the spectra to obtain the coherent part of the spectra. The coherent spectra agree much better between the two techniques, but some of the higher frequency ($f \approx$ 1000 Hz) TRPIV energy is lost. For the most part, the coherent spectra for the two TRPIV results agree well. The only noticeable difference is at the higher frequencies (around 1000 Hz) at x_{sh}=31 (Fig. 21a). The 32 x 32 TRPIV coherent spectra has a lower amplitude than the 16 x 16 TRPIV coherent spectra at these higher frequencies.



(a) $x_{sh} \approx$ 31 **(b) $x_{sh} \approx$ 41** **(c) $x_{sh} \approx$ 50**

Figure 20: HW and TRPIV total spectra at three different streamwise locations, y=0.75 mm, averaged across the span.

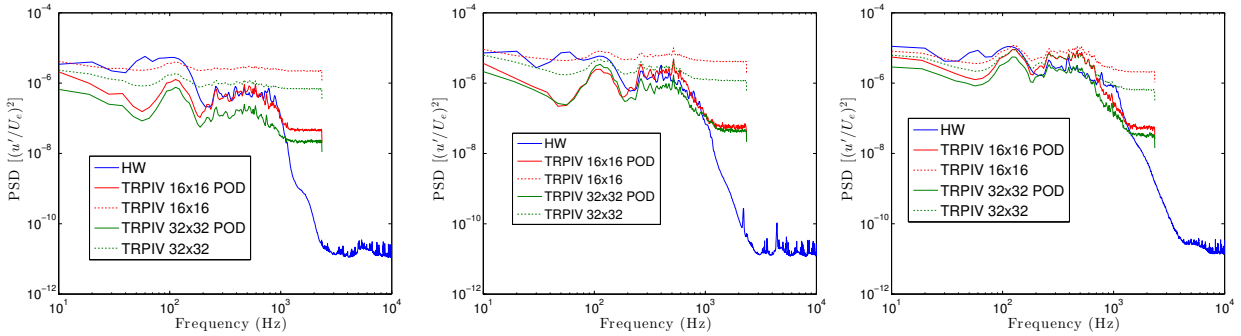


(a) $x_{sh} \approx$ 31 **(b) $x_{sh} \approx$ 41** **(c) $x_{sh} \approx$ 50**

Figure 21: HW and TRPIV total and coherent spectra at three different streamwise locations, y=0.75 mm, averaged across the span.

Another approach that can be taken is to perform proper orthogonal decomposition (POD) on the TRPIV data and reconstruct the data using a subset of the modes. This should help reduce the noise while maintaining the primary instability structures and frequencies. The results are shown in Fig. 22 for the same three locations. The TRPIV data were reconstructed using the first 50 POD modes. The original HW and TRPIV spectra are also included in the figures for comparison. The POD reconstruction is successful at reducing the noise, and we recover similar features to what is seen in the HW spectra. In fact, the overall magnitudes of the spectra are very similar, particularly in the 200 to 800 Hz range. There tends to

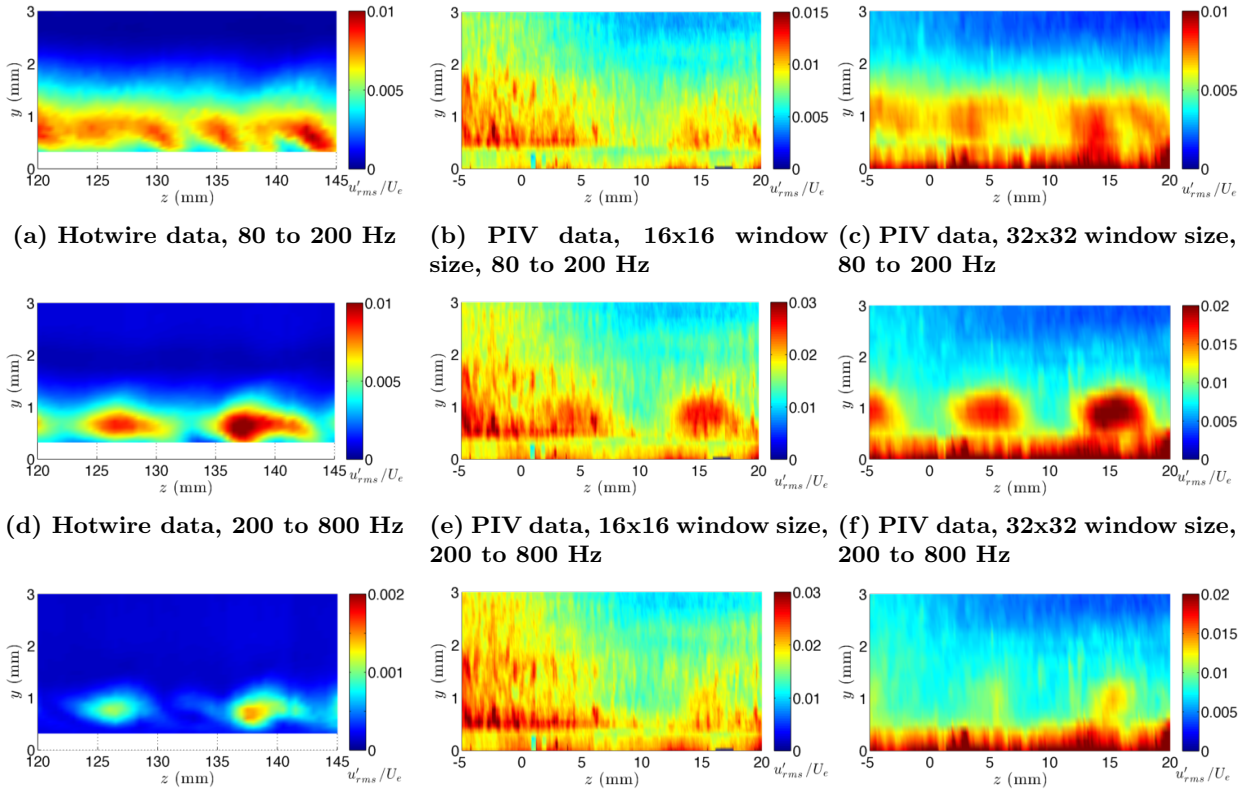
be some discrepancy in the lower (<200 Hz) and higher (>800 Hz) frequency ranges. Even with the POD reconstruction, the shear-layer instabilities are hard for the TRPIV to resolve because they are so small in amplitude compared to the noise floor.



(a) $x_{sh} \approx$ **31** **(b)** $x_{sh} \approx$ **41** **(c)** $x_{sh} \approx$ **50**

Figure 22: HW and TRPIV spectra, including the POD-reconstructed TRPIV spectra, at three different streamwise locations, y=0.75 mm, averaged across the span.

The unsteady instabilities are further examined by plotting u'_{rms} contours for each of the frequency bands of interest. Figure 23 shows the u'_{rms} values across the whole measurement plane, integrated across each of the three frequency bands. Results are shown for the HW and both of the TRPIV results. Note the different color scales on these figures.



(a) Hotwire data, 80 to 200 Hz **(b) PIV data, 16x16 window size, 80 to 200 Hz** **(c) PIV data, 32x32 window size, 80 to 200 Hz**

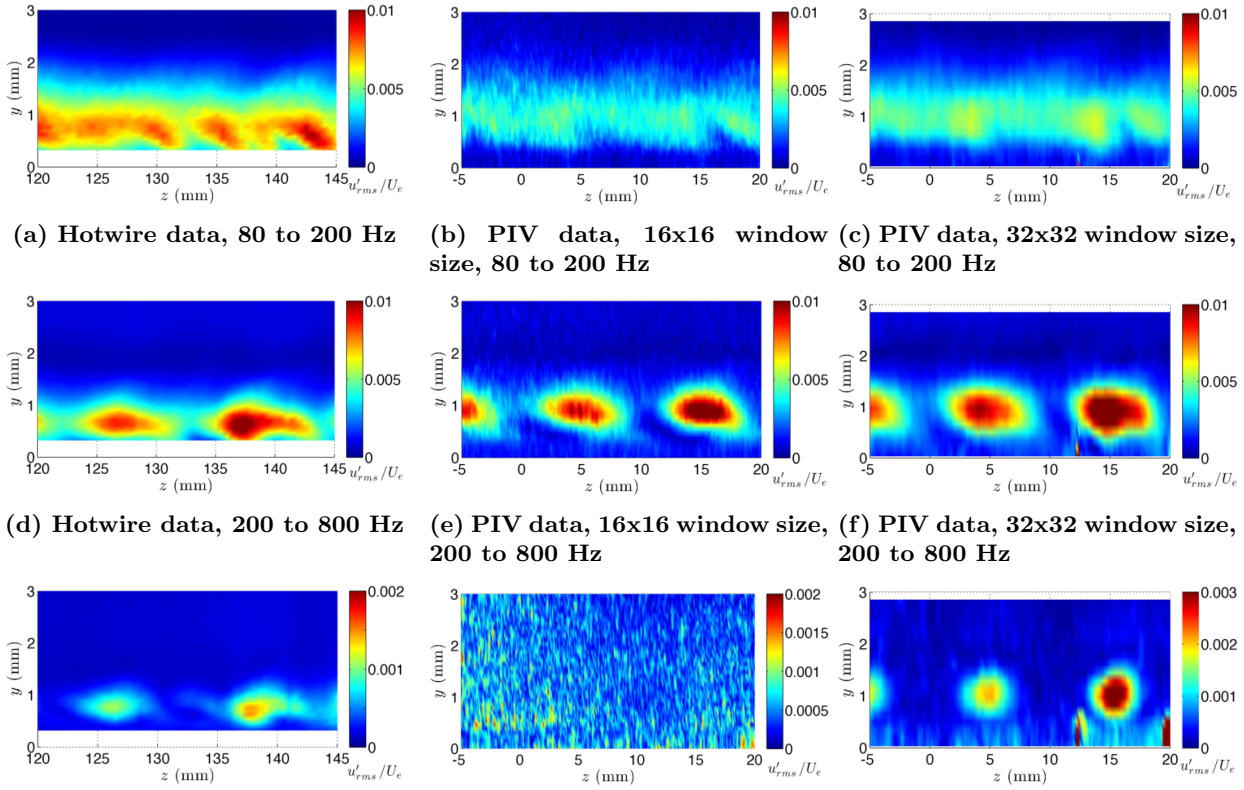
(d) Hotwire data, 200 to 800 Hz **(e) PIV data, 16x16 window size, 200 to 800 Hz** **(f) PIV data, 32x32 window size, 200 to 800 Hz**

(g) Hotwire data, 800 to 1500 Hz **(h) PIV data, 16x16 window size, 800 to 1500 Hz** **(i) PIV data, 32x32 window size, 800 to 1500 Hz**

Figure 23: Comparison of hotwire and PIV u'_{rms} contours for different frequency bands and different PIV interrogation window sizes, $x_{sh} \approx$ 40.

Due to the interaction of the different instabilities with the stationary crossflow vortices, the unsteady instabilities become modulated in the spanwise direction. The results are very noisy for the 16 x 16 TRPIV window size, but the 32 x 32 TRPIV results show much better agreement with the hotwire data, particularly for the two lower frequency bands. For the 800 to 1500 Hz frequency range (Fig. 23i), there are some structures present, but the results are too noisy to really make out what is happening.

The next approach we take to study the instabilities is to bandpass filter the time-resolved data within the three frequency ranges, and then perform the POD analysis on each bandpass filtered set, and finally reconstruct each data set from the first ten POD modes. We then take the spectra of the reconstructed data, and again plot the u'_{rms} contours. These results are shown in Fig. 24. The POD analysis helps immensely with the 16 x 16 TRPIV data, however, there are still no structures evident for the highest frequency range (Fig. 24h). The POD analysis also helps clean up the results for the 32 x 32 data, especially for the highest frequency band (Fig. 24i), in which clear structures are now evident.



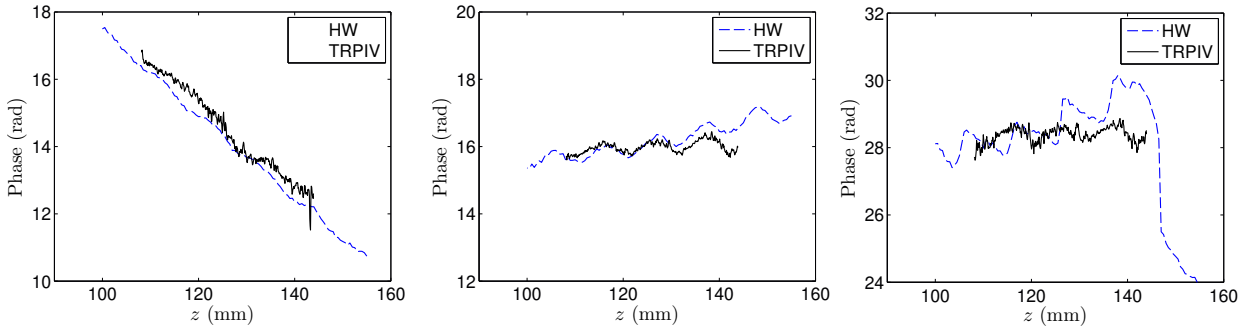
(a) Hotwire data, 80 to 200 Hz **(b) PIV data, 16x16 window size, 80 to 200 Hz** **(c) PIV data, 32x32 window size, 80 to 200 Hz**

(d) Hotwire data, 200 to 800 Hz **(e) PIV data, 16x16 window size, 200 to 800 Hz** **(f) PIV data, 32x32 window size, 200 to 800 Hz**

(g) Hotwire data, 800 to 1500 Hz **(h) PIV data, 16x16 window size, 800 to 1500 Hz** **(i) PIV data, 32x32 window size, 800 to 1500 Hz**

Figure 24: Comparison of hotwire and PIV u'_{rms} contours for different frequency bands and different PIV interrogation window sizes, with PIV data reconstructed from first 10 POD modes, $x_{sh} \approx 40$.

Another way to compare the behavior of the measured instabilities is to plot the phase of the instability as it moves across the plane. This gives a measurement of the in-plane phase speed of the instability. These results are shown at the most upstream measurement location in Fig. 25 for a chosen frequency within each of the three frequency bands of interest. Similar to the coherence results, the HW results were obtained using data from two simultaneous hotwire signals, with one at a fixed location and the other traversing in the spanwise direction. The TRPIV results were obtained by choosing a reference point in the middle of the plane and computing the phase at each spanwise location relative to the reference point. The slope of the phase vs. z plot of the 100 Hz mode (Fig. 25a) indicates that the traveling crossflow instability is moving through the plane (from inboard to outboard) with a phase speed of 5 m/s, and a spanwise wavelength of 50 mm. The results agree very well between the HW and TRPIV results.



(a) $x_{sh} \approx$ **31** **(b)** $x_{sh} \approx$ **41** **(c)** $x_{sh} \approx$ **50**

Figure 25: HW and TRPIV phase of three modes at $x_{sh} \approx$ 31, y=0.75 mm.

The results for the 300 Hz mode (Fig. 25b) indicate that this mode, which corresponds to the TS instability, is primarily traveling out of plane. This is evident from the nearly zero in-plane slope of the phase plots. Again, the results for the HW and TRPIV techniques agree fairly well. The 1000 Hz mode (Fig. 25c), which corresponds to the shear-layer instability, also exhibits a slope that is close to zero, indicating that the mode is primarily traveling out of plane. Overall, these phase results illustrate that the TRPIV and HW techniques are measuring similar in-plane phase speeds and directions of the instabilities.

In every case, the HW and PIV results agree well qualitatively, showing the same types of instabilities, the same behavior of the instabilities, and the same type of breakdown to transition. As seen in the results presented here, the mean velocity profiles agree well, and the stationary crossflow profiles agree well qualitatively, though there are some quantitative differences. The velocity spectra, once the noise is reduced, agree well qualitatively and fairly well quantitatively. A larger window size is found to be beneficial in terms of reducing the noise and helps improve the measurement of the unsteady instabilities. The TRPIV technique is able to resolve the lower frequency instabilities (80 to 200 Hz and 200 to 800 Hz ranges) well. In some cases, if the amplitude is large enough, the TRPIV technique can also resolve the higher-frequency instabilities in the range of 800 to 1500 Hz. Overall, the results illustrate that stereo TRPIV is a viable measurement technique for boundary-layer transition research, though some care must be taken in minimizing the stereo registration error and in dealing with the extra noise introduced by the TRPIV system.

V. Summary and Concluding Remarks

Results were presented from the uncertainty analysis of time-resolved PIV results from a low-speed transition experiment. The uncertainty was estimated for several cases using several different approaches, and the primary potential error source was found to be from stereo registration errors. The registration errors can lead to very large errors in the measured stationary crossflow amplitude (on the order of 20 to 50%) due to the large spanwise velocity gradients, which are not accounted for in any existing stereo-PIV uncertainty analysis approach. The effectiveness of the self-calibrations was assessed and it was found that the self-calibrations do not always sufficiently correct for the misalignment. Interestingly, the self-calibrations often resulted in somewhat poorer performance when comparing the spanwise-averaged results, but worked very well in terms of correcting the stationary crossflow mode shapes and amplitudes.

A simple approach was demonstrated for predicting the error caused by stereo registration errors. First, the in-plane displacement for each camera is estimated given an assumed offset of the laser plane from the calibration target. Then, the processed velocity field is utilized to predict the velocities that would be measured by each camera using the in-plane shifts, and these velocities are recombined to obtain the predicted velocity. This approach was found to work very well for both cases studied. Thus, this approach could be useful for quickly and easily predicting the uncertainty for a given setup if a maximum offset can be estimated. This can provide more confidence in the results since the self-calibration cannot always be trusted to completely remove the errors due to the misalignment.

An in-depth comparison was performed of the time-resolved PIV and HW results that were acquired downstream of a backward-facing step. In addition, the window-size for the TRPIV processing was varied from 16 x 16 to 32 x 32 to determine the effect of the window size on the uncertainty and noise of the

measurement. The TRPIV and HW results agreed very well qualitatively: the same types of instabilities were measured and they exhibited similar behavior leading up to breakdown to turbulence. Quantitatively, there were some minor differences in the instability amplitudes, but the results actually agree quite well considering the possible differences in amplitude due to different measurement locations and different initial stationary crossflow amplitudes. The most encouraging result is that the TRPIV was able to resolve the unsteady instabilities, even in the higher frequencies, at very low amplitudes (<1% U_e). There was clearly quite a bit of noise present in some of the frequency bands, particularly the highest frequency band that was studied, but the structures were still present and clearly discernible, particularly after performing proper orthogonal decomposition. The window size had a noticeable effect on the results. The larger window size (32 x 32) helped to significantly reduce the noise and uncertainty that was present in the time-resolved measurements, and, in fact, made the higher frequency structures visible when they had been washed out by noise in the 16 x 16 results. It is possible that the improvement in noise at the larger window size is simply due to the increased signal (i.e., number of particles) within that larger window size. Thus, one may benefit from increased seeding to improve the noise of the measurements at the smaller window size, without the need to increase the window size (and thus reduce the spatial resolution) to obtain better results. However, this was not directly tested. The good agreement that was found between the HW and time-resolved PIV results allows us to conclude that the time-resolved PIV technique is well-suited for boundary-layer transition measurements since the transition process was not significantly affected by the particles that were introduced into the flow. For time-resolved measurements, HW anemometry is probably still necessary if accurate measurements of low-amplitude instabilities are desired. However, the TRPIV technique is extremely useful for understanding the flow structures and interactions, while also allowing measurements of multiple velocity components, and all of this in a very timely fashion.

References

[1]Lazar, E., DeBlauw, B., Glumac, N., Dutton, C., and Elliott, G., "A Practical Approach to PIV Uncertainty Analysis," *27th AIAA Aerodynamic Measurement Technology and Ground Testing Conference*, 2010, p. 4355.

[2]Westerweel, J., "Theoretical Analysis of the Measurement Precision in Particle Image Velocimetry," *Experiments in Fluids*, Vol. 29, No. 1, 2000, pp. S003–S012.

[3]Raffel, M., Willert, C. E., Wereley, S. T., and Kompenhans, J., *Particle Image Velocimetry: A Practical Guide*, Springer, 1998.

[4]Wieneke, B., "Generic a-posteriori Uncertainty Quantification for PIV Vector Fields by Correlation Statistics," *17th International Symposium on Applications of Laser Techniques to Fluid Mechanics, Lisbon, Portugal*, July 2014.

[5]Sciacchitano, A., Wieneke, B., and Scarano, F., "PIV Uncertainty Quantification by Image Matching," *Measurement Science and Technology*, Vol. 24, No. 4, 2013.

[6]Timmins, B. H., Wilson, B. W., Smith, B. L., and Vlachos, P. P., "A Method for Automatic Estimation of Instantaneous Local Uncertainty in Particle Image Velocimetry Measurements," *Experiments in Fluids*, Vol. 53, No. 4, 2012, pp. 1133–1147.

[7]Beresh, S. J., Wagner, J. L., and Smith, B. L., "Self-Calibration Performance in Stereoscopic PIV Acquired in a Transonic Wind Tunnel," *Experiments in Fluids*, Vol. 57, No. 4, 2016, pp. 48.

[8]Bhattacharya, S., Charonko, J. J., and Vlachos, P. P., "Stereo-Particle Image Velocimetry Uncertainty Quantification," *Measurement Science and Technology*, Vol. 28, No. 1, 2016.

[9]Foucaut, J.-M., Carlier, J., and Stanislas, M., "PIV Optimization for the Study of Turbulent Flow Using Spectral Analysis," *Measurement Science and Technology*, Vol. 15, No. 6, 2004, pp. 1046–1058.

[10]Gericke, T. and Scholz, P., "PIV Compared to Classical Measurement Techniques Using the Example of Turbulent Flow in the Wake of an Air Outlet," *17th International Symposium on Applications of Laser Techniques to Fluid Mechanics*, July 2014.

[11]Stanislas, M., Foucaut, J., and Dupont, P., "Experimental Study of a High Reynolds Number Turbulent Boundary Layer Using DPIV," *International Symponsium on Applications of Laser Techniques to Fluid Mechanics*, 1998.

[12]Li, C.-T., Chang, K.-C., and Wang, M.-R., "PIV Measurements of Turbulent Flow in Planar Mixing Layer," *Experimental Thermal and Fluid Science*, Vol. 33, No. 3, 2009, pp. 527–537.

[13]Lin, J., Foucaut, J.-M., Laval, J.-P., Pérenne, N., and Stanislas, M., "Assessment of Different SPIV Processing Methods for an Application to Near-Wall Turbulence," *Particle Image Velocimetry*, Springer, 2007, pp. 191–221.

[14]Wernet, M. P., "Time-Resolved PIV for Space-Time Correlations in Hot Jets," *Instrumentation in Aerospace Simulation Facilities, 2007. ICIASF 2007. 22nd International Congress on*, IEEE, 2007, pp. 1–11.

[15]Dróżdż, A. and Uruba, V., "Comparison of PIV and Hot-Wire Statistics of Turbulent Boundary Layer," *Journal of Physics: Conference Series*, Vol. 530, IOP Publishing, 2014.

[16]Eppink, J. L. and Yao, C., "Time Resolved Stereo Particle Image Velocimetry Measurements of the Instabilities Downstream of a Backward-Facing Step in a Swept-Wing Boundary Layer," AIAA Paper 2017-4415, 2017.

[17]Serpieri, J. and Kotsonis, M., "Time-Resolved PIV Investigation of the Secondary Instability of Cross-flow Vortices," International Symposium on Turbulence and Shear Flow Phenomena, TSFP10, 2017.

[18]Serpieri, J. and Kotsonis, M., "Three-Dimensional Organisation of Primary and Secondary Crossflow Instability," *Journal of Fluid Mechanics*, Vol. 799, 2016, pp. 200–245.

[19]Eppink, J. and Yao, C.-S., "Stereo Particle Image Velocimetry Measurements of Transition Downstream of a Backward-Facing Step in a Swept-Wing Boundary Layer," AIAA Paper 2017-0305, 2017.

[20]Saric, W. S. and Reshotko, E., "Review of Flow Quality Issues in Wind Tunnel Testing," AIAA Paper 1998-2613, 1998.

[21]Eppink, J., *The Interaction of Crossflow Instabilities and a Backward-Facing Step in Swept Boundary Layer Transition*, Ph.D. thesis, Tufts University, 2014.

[22]Eppink, J. L., Wlezien, R. W., King, R. A., and Choudhari, M., "Interaction of a Backward-Facing Step and Crossflow Instabilities in Boundary-Layer Transition," *AIAA Journal*, 2017, pp. 1–12.

[23]Mei, R., "Velocity Fidelity of Flow Tracer Particles," *Experiments in Fluids*, Vol. 22, No. 1, 1996, pp. 1–13.

[24]Sciacchitano, A. and Wieneke, B., "PIV Uncertainty Propagation," *Measurement Science and Technology*, Vol. 27, No. 8, 2016, pp. 084006.

[25]Giordano, R. and Astarita, T., "Spatial Resolution of the Stereo PIV Technique," *Experiments in fluids*, Vol. 46, No. 4, 2009, pp. 643.

[26]Eppink, J. L., Wlezien, R. W., King, R. A., and Choudhari, M., "The Effect of Backward-facing Step Height on Instability Growth and Breakdown in Swept Wing Boundary-Layer Transition," AIAA Paper 2015-2776, 2015.