

Effects of Spatial Realignment in Stereo PIV Calibration

Steven J. Beresh*

Sandia National Laboratories, Albuquerque, NM, 87185

Barton L. Smith†

Utah State University, Logan, UT, 84322

Simultaneous stereo PIV measurements of a round free jet were obtained from narrow and wide camera angles while a fifth camera viewed the laser sheet from 90 degrees to determine the two-component velocity field free of errors resulting from stereo calibration. Errors in mean velocities were small, but artificially reduced turbulent stresses were generated when self-calibration was not used, owing to a smearing effect that occurs when the two cameras are inadequately registered to each other. This difficulty worsened with increased laser sheet thickness. Spatial error in the stereo calibration process can artificially displace vector fields from the expected origin, which was detected through comparison to the simultaneous two-component measurement. Although this spatial offset typically is small with respect to statistical properties of a data set, it can be prominent when instantaneous snapshots of the velocity field are examined, particularly where the velocity gradient is momentarily large.

Introduction

As particle image velocimetry (PIV) has become a standard and widespread measurement technique, its use in practical environments has placed greater demands upon the ability to estimate its measurement uncertainty. Numerous potential sources of error exist and have motivated literally hundreds of conference papers and journal articles, which are well summarized in Raffel et al [1] and Adrian and Westerweel [2]. Stereoscopic PIV requires that two cameras are aligned and calibrated to a single imaging region defined by the laser sheet. Bias uncertainties arising from stereoscopic PIV calibration can be especially damaging. Many researchers since the initial development of stereoscopic PIV have examined calibration uncertainties, but these studies almost exclusively were concerned with the propagation of random error in the correlation algorithm (e.g., [3-5]). Unfortunately, bias errors in the stereoscopic alignment can easily dwarf the random errors in the PIV correlation algorithms, owing to the difficulty of precisely aligning both the calibration target and the laser sheet to a common imaging region in real space. This problem is known as image registration. A number of studies have expressed the difficulty of this registration process, in which misalignments of considerably less than 1 mm translation or 1 degree of rotation leads to errors exceeding 10% in velocity [5-9]. The establishment of image deformation as a routine component of PIV processing algorithms introduces further effects whose ultimate impact upon the data is hard to appreciate [10].

Fortunately, the uncertainty associated with misalignment between the laser sheet and the imaging plane may be reduced by employing the self-calibration technique, in which correlations between images acquired from each of the two cameras at the same moment in time are used to identify and correct any disparity in the registration between cameras [11, 12]. At its most successful, the calibration error can be reduced to levels below that commonly associated with random error in the PIV correlations [10-12]. However, this is misleading when considering the effect upon fully processed velocity data; instantaneous velocity fields may find the calibration error negligible with respect to random correlation error, but mean velocity fields and other quantities derived from entire

*Principal Member of the Technical Staff, Engineering Sciences Center, Associate Fellow AIAA, correspondence to: P.O. Box 5800, Mailstop 0825, (505) 844-4618, email: sjberes@sandia.gov

†Associate Professor

This paper is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

This work is supported by Sandia National Laboratories and the United States Department of Energy. Sandia is a multiprogram laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

data sets will reduce the random error and calibration error may become significant or even dominant. Moreover, self-calibration adjusts the intended stereo calibration in ways of which the user may not be fully cognizant.

Self-calibration is more effective for some experimental configurations than others. For thin laser sheets and small camera angles, each camera views similar particle images and they will correlate well; however, for thick laser sheets and/or sharper camera angles, perspective error will lead to different views of the same particles that will not correlate well. Particle image size and density are factors as well. As noted by Wieneke [11, 12], correlation summing over a set of images rather than a single image can reduce this difficulty, but it is not clear that this is universally effective. Also, self-calibration may be more problematic for configurations employing large magnification or short imaging distances, where the mapping functions exhibit greater variation over the measurement plane. In summary, a successful self-calibration clearly reduces error, though its effectiveness may be dependent upon specifics of the experimental configuration.

The question, then, is what error remains following stereo self-calibration? Even if the random error is negligible relative to correlation error, what bias errors are present due to stereo calibration? Are these errors introduced by the self-calibration, or do they originate in the initial target calibration? Does the self-calibration adjustment of the stereo calibration parameters affect the data accuracy? Is every self-calibration optimal, or might the error present in the stereo configuration only be partially corrected? In what way are these questions dependent upon the particulars of the PIV experimental configuration?

The present investigation uses multiple measurements of a round free jet to produce a data set ideal for uncertainty quantification. Simultaneous PIV measurements, both two component and stereoscopic, were obtained from different angles and based upon different camera calibrations. This approach provided redundant data acquired under varied experimental parameters, which may then be compared to determine what uncertainties arise and what dependencies they exhibit. This data set was gathered with two intended purposes: first, to test the effectiveness of self-calibration; and second, to facilitate a case study in PIV uncertainty quantification. The current paper addresses the former motivation.

Experimental Approach

Turbulent Jet

A turbulent free jet was chosen as the subject flow field for this experiment in PIV uncertainty quantification. Compressed air regulated at 137 kPa (20 psig) was supplied to a TSI six-jet atomizer, then the atomizer output was routed via a flexible hose to a stainless steel tube that served as the jet nozzle. The tube was 12.7 mm (0.500 inch) in outer diameter, 10.2 mm (0.402 inch) in inner diameter, and was carefully deburred at the exit. Its length was 360 mm (14 inch), or 35 L/D , which should be adequate to reach a fully developed state. The exit of the jet was raised 75 mm (3.0 inches) off a flat plate serving as an exit plane to isolate the jet from the flow control apparatus. To prevent seed particles from entering the laboratory and fouling equipment, the jet was placed underneath a fume hood. The first three feet of distance from the jet exit were surrounded by an acrylic enclosure 0.5 to 1 m (18 to 36 inches) on a side; this was less to prevent escape of seed particles and more to prevent random air currents in the room from affecting the jet propagation. The jet apparatus can be seen in the photograph of the experimental arrangement shown in Fig. 1. Two of the enclosure walls have been removed for this photograph, and the resulting asymmetric jet is due to exposure to room currents.

The six-jet atomizer produced droplets of mineral oil with a diameter of 1-2 μm . Considering the low velocities, particle tracking of the flow did not present a concern. Seeding density could be controlled by adjusting the number of atomizing jets in operation or a gas bypass around the atomizer, though in practice once an appropriate seeding level was found, it was not adjusted during the course of the experiment. No seeding was provided for the ambient air into which the jet exhausted. This may create a seeding bias at the edges of the jet as it mixes with the ambient, but the purpose of this experiment is not to make accurate jet measurements, it is to compare simultaneous measurements. By this perspective, any seeding biases will be identically present in every measurement.

The air flow rate through the jet was maintained at 50 slm, which yielded a bulk exit velocity of 13.6 m/s and a Reynolds number based on jet exit diameter of 7500. Though the flow rate was maintained precisely and the fume hood exhaust setting was not altered, the jet behavior was found to drift somewhat over time. However, all analysis in the present work was based on simultaneous measurements that therefore measure the equivalent jet variation.

PIV Systems

Three independent PIV systems measured the jet simultaneously as visible in the photograph of Fig. 1 and the sketch of Fig. 2. The three systems all viewed the same laser sheet, which was formed from the beams of a dual-cavity Nd:YAG laser (Continuum Minilite II) rated at 25 mJ/pulse and operating at 10 Hz. Prior to PIV data

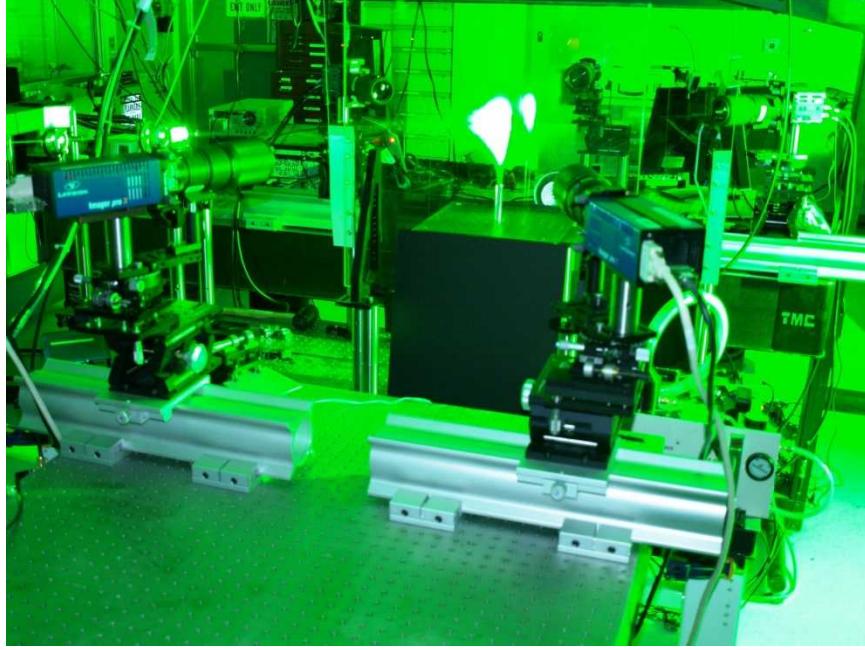


Fig. 1: Photograph of the experimental setup, including the free jet and five cameras used for three independent PIV measurements.

acquisition, the beams were confirmed to be round, Gaussian, and well overlapped. The time between pulses was always 30 μ s. Three laser sheet thicknesses were studied during the course of the experiments, estimated to be 0.8 mm, 1.3 mm, and 1.8 mm. These are denoted as “thin,” “middle,” and “fat,” respectively. Each laser sheet was aligned to the same imaging region, based on the laser sheet centerplane, which was centered about 5 inches downstream of the jet exit and aligned to its centerline. A series of transparent straightedges were aligned to the jet pipe and the flat plate in a repeatable fashion such that each of the three sheet thicknesses could be positioned as identically as possible. The imaging region nominally began at 109 mm (4.3 inch) downstream of the jet exit and covered about 43 mm (1.7 inch) in the streamwise (vertical) direction.

The three PIV systems included a single-camera two-component measurement and two stereoscopic measurements, each distinguished by the camera angle relative to the laser sheet. The two-component measurement was conducted using a Redlake ES-4/E camera with resolution 2048×2048 pixels and digitized at 8 bits. It was aligned normal to the laser sheet. It collected light using a 180-mm lens plus a 2x teleconverter and was set back from the laser sheet by 1.12 m (44 inch). The long lens and standoff distance was intended to minimize perspective distortion in the images as well as perspective bias in the velocity measurements. This approach should ensure that the x and y components measured are negligibly contaminated by the z component (out-of-plane). The lens was set to an f/stop of 5.6, though this effectively was 11.2 given the teleconverter.

On the opposite side of the laser sheet were two LaVision ProX cameras, each with a resolution of 2048×2048 pixels digitized at 14 bits. Each camera was equipped with a 200-mm lens placed on a Scheimpflug mount for an oblique focal plane. The lens angle with respect to the laser sheet normal was 30 deg for each camera. Lenses were set to an f/stop of 8. This stereo camera measurement is known as “narrow,” referring to its stereo lens angle.

A second stereoscopic measurement was made using a wider angle and was placed on the same side of the laser sheet as the two-component camera. This configuration used two LaVision sCMOS cameras, each with a resolution of 2560×2160 pixels digitized at 16 bits. Again, 200-mm lenses mounted on Scheimpflug adapters were used with an f/stop of 8. Here, the lens angle with respect to the laser sheet normal was 60 deg each, which earned this system a label of “wide.”

The cameras were synchronized with each other and with the laser sheet using multiple digital delay generators. Each of the stereo systems was timed using LaVision Programmable Timing Units, but these in turn were triggered by external delay generators (Stanford DG-645). A 5 Hz rate common to all five cameras was used to ensure that all images were matched in each measurement. Image sequences were triggered simultaneously by a simple button press. Careful attention to the contents of image sequences ensured that the five cameras stayed synchronized throughout the experiment.

A customized single-plane calibration plate was used to calibrate all three systems simultaneously. The target was an aluminum plate 3.05 mm (0.120 inch) thick carefully machined such that each surface was very flat and parallel, drilled with normal fiducial holes every 6.35 mm (0.25 inch). A stack of multiple translation, rotation, and tilt stages allowed an extravagant degree of freedom to be employed aligning the target. The target was carefully aligned to known coordinates along the jet centerline and the translation axis of the plate was aligned to the laser sheet normal. The target was translated through the volume of the laser sheet and seven calibration stations were imaged by each camera, then calibration was accomplished using a third-degree polynomial fit by the DaVis software. The two-component measurement was calibrated simply by determining a scale factor. Self-calibration was used to improve the calibrations for the two stereo configurations, and this is described in detail in the Results section. The image intensity detected by each camera in a stereo pair differed due to collection of the scattered laser light in forward scatter versus back scatter. This proved not be a significant difficulty for the narrow angle of the ProX cameras but it was substantial for the sCMOS cameras, necessitating that in the latter case images first were normalized in intensity for the purposes of self-calibration.

All data were calibrated and processed using DaVis 8.1. Typical image interrogation used three iterations; one pass using 64×64 pixel interrogation windows, followed by two passes of 32×32 pixel interrogation windows. A 50% overlap in the interrogation windows was used as well to oversample the velocity fields. The resulting vector fields were validated based upon signal-to-noise ratio, nearest-neighbor comparisons, and allowable velocity range.

Test Cases

The original motivation of this study was to test methods of uncertainty quantification for PIV. The multiple configurations of simultaneous measurements would permit a comparison not just of their data, but also of the effectiveness of their uncertainty estimates. The different camera angles and laser sheet thicknesses allow a study of the dependence of uncertainty on these parameters. The results of this aspect of the investigation will be described in a future publication.

The other component of this study, which matured as the data were examined in detail, was to attempt to quantify error due to stereo calibration and to scrutinize the efficacy of stereoscopic self-calibration under different configurations. The key variables were the two camera angles and the three laser sheet thicknesses. In particular, questions arose regarding the spatial positioning of the imaging regions by stereo calibration and then again following the self-calibration procedure. Additional experiences suggested velocity errors arising from the positional changes in the self-calibrated data. These form the subject of the current article.

The principal test parameters were the three independent camera configurations and three data sets corresponding to each of the three laser sheet thicknesses. Additionally, variations were tested in the PIV data processing and in particular the stereo calibration processes. Data sets were acquired in runs of 200 snapshots at a time. Each of the three datasets for the varying sheet thicknesses contained 3600-4200 individual snapshots, which provided for excellent convergence of turbulent statistics.

Jet conditions were not deliberately varied through the experiment, but some drift in the data was observed, common to all three measurements. Therefore, despite the temptation to compare results between the three laser sheet thicknesses, differences due to the measurement volume as seen by a single camera configuration cannot be distinguished from drift in the jet conditions. However, comparisons of the three measurements to each other systematically vary for the three sheet thicknesses, and this is independent of small jet variations. Possibly, the variation in jet conditions was not due to the jet itself, but the fume hood exhaust or ambient conditions.

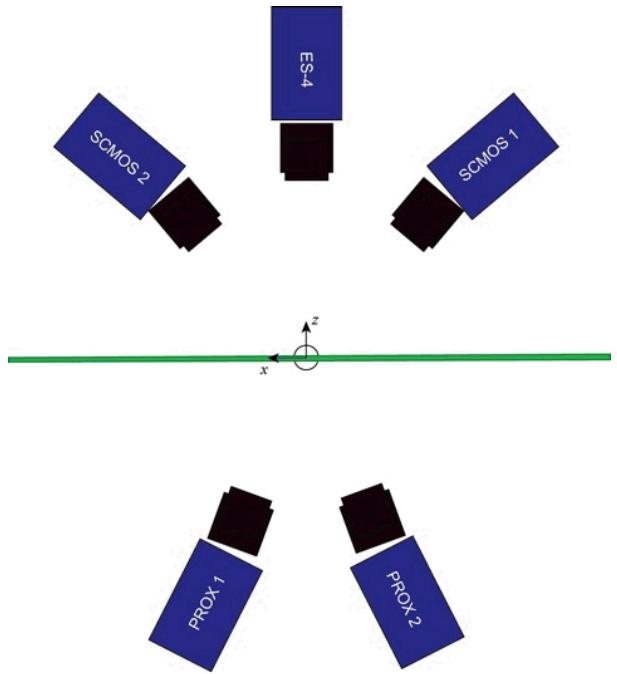


Fig. 2: Top-view schematic of the PIV setup, with the jet represented by the circle in the center. Camera angles and standoff distances are not to scale.

Results

Comparison of Camera Configurations

Velocity data for the free jet are given in Fig. 3, with Fig. 3a showing the mean velocity field and Fig. 3b the turbulent normal stress, both for the vertical velocity component (aligned with the jet axis). Vectors are subsampled by a factor of two horizontally and a factor of four vertically. The figure shows an unremarkable turbulent free jet. Some drift towards the $+x$ direction is evident in the mean vector field, which was observed visually during the experiments and is not a result of any measurement bias. Similarly, the turbulent normal stress displays unequal lobes on each side of the jet. Note that the asymmetry in the jet is actually much more subtle than that suggested by the photograph of Fig. 1, which was taken with a portion of the enclosure removed. These figures were produced from the narrow-angle stereo measurements using the thin sheet, but the other data are visually similar.

A quantitative comparison of the velocity data acquired by the three PIV configurations is facilitated by extracting cross-stream profiles through the plots of Fig. 3 and others like them. Horizontal profiles were extracted at $y=113.5$ mm (chosen because this corresponds to the position of the calibration origin dot). Three plots of the mean vertical velocity are given in Fig. 4, one for each laser sheet thickness. Each plot contains curves for the two component data, the narrow-angle stereo data, and the wide-angle stereo data. Additionally, the two stereo data sets have two curves each, one for the plate calibration and a second for the self-calibrated data. Here, a plate calibration refers to the calibration based entirely on the calibration target plate and without any self-calibration. The inset plots are zoomed views of the data regions near the maxima.

For all three sheet thicknesses, the same behavior of the different camera configurations may be seen. The position of the peak for the wide-angle configuration is in closer alignment with the two-component data than the narrow-angle, which probably is merely fortuitous. This is true for both the plate calibration and the self-calibration. The narrow data features a spatial offset with respect to the 2C data. Given the relative simplicity of establishing the spatial origin and scales in 2C PIV and the efforts to minimize perspective error by using a long lens and camera standoff distance, the approximately 0.5 mm offset is substantial. Note that this spatial offset has nothing to do with the actual alignment of the origin with respect to the jet; the offset is between a stereo measurement and a 2C measurement that use a common origin during calibration. The nature of this offset is the subject of greater scrutiny in the following section of this paper. As the sheet thickness is increased, self-calibration yields a small increase in the peak velocity magnitude for the wide-angle data, which brings it into closer agreement with the narrow data and the 2C data. This is indicative of the more precise image registration achieved by self-calibration, which reduces the real spatial area over which the vector is calculated. The effect only becomes apparent for the fat laser sheet where the increased volume through the thicker sheet exacerbates any misregistration between cameras. Differences in the peak velocity magnitude between sheet thicknesses are inconsequential because they can be attributed to variability in the jet behavior.

Figure 5 shows the equivalent plots for the turbulent normal stress in the vertical velocity component. Though second-order turbulence statistics are inherently noisier due to slower convergence, the same peak behavior can be seen, in which the narrow-angle data exhibit a small spatial offset from the wide data and the 2C data. Of greater interest is that self-calibration raises the magnitude of the turbulent stress for each stereo configuration relative to the plate calibration. This effect grows stronger with increasing sheet thickness, and for the fat sheet, it is evident in both the narrow and wide data. This is a consequence of the improved spatial overlap of the two cameras in a stereo configuration once self-calibration is employed. With only a plate calibration, some misregistration of interrogation windows occurs between the two cameras, leading to a matching of disparate vectors. This tends to act as a low-pass filter, smearing out the spatial region over which the vector is computed and hence reducing its magnitude. A successful self-calibration reduces this effect by more closely aligning the vectors from each camera and thus minimizing the measurement volume. Further analysis is found later in this document.

The plots of Figs. 4 and 5 suggest two effects of spatial realignment due to stereo calibration. First, the calibration can create a small but significant spatial offset from the origin of the data as established by the 2C origin, and this spatial offset occurs in the initial plate calibration and is not appreciably altered by the self-calibration. Second, without the use of self-calibration, the misalignment of the two cameras with respect to each other broadens the effective volume over which the vector is computed and tends to artificially diminish the mean velocity and, more dramatically, the turbulent stress. These matters are discussed in depth in the following two sections.

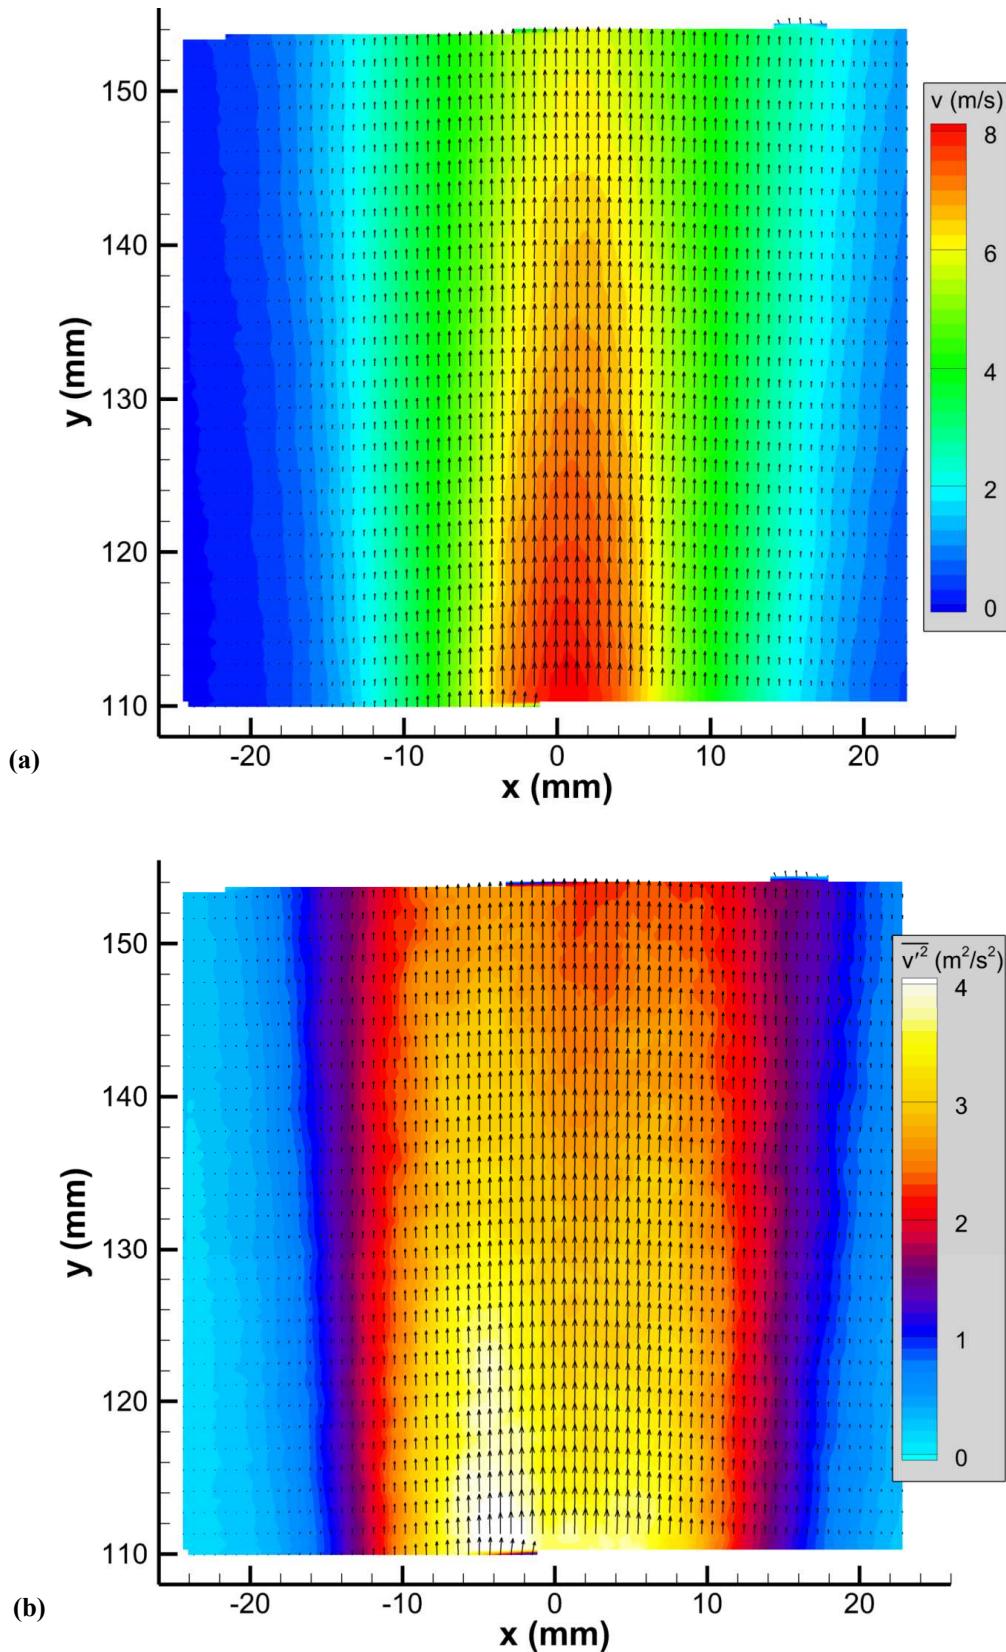


Fig. 3: Velocity data for the free jet taken from the narrow angle stereo configuration and the thin laser sheet. (a) mean vertical velocity field; (b) vertical turbulence normal stress field. Vectors are subsampled 2×4 .

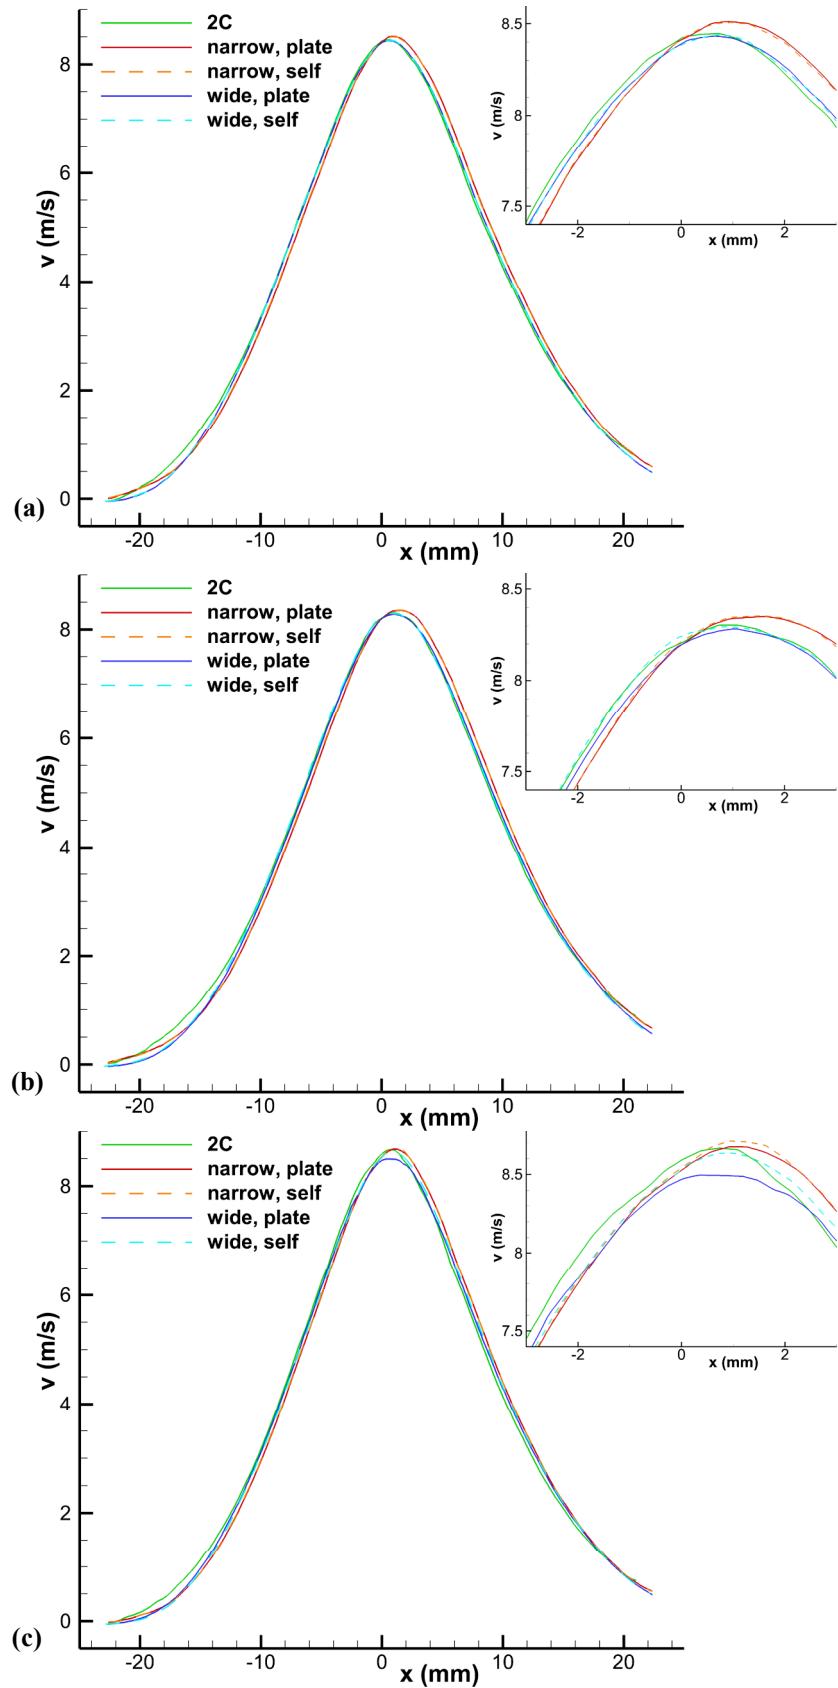


Fig. 4: Mean vertical velocity component at $y=113.5$ mm for all three camera configurations, with and without self-calibration (self and plate, respectively). (a) thin laser sheet; (b) medium laser sheet; (c) fat laser sheet.

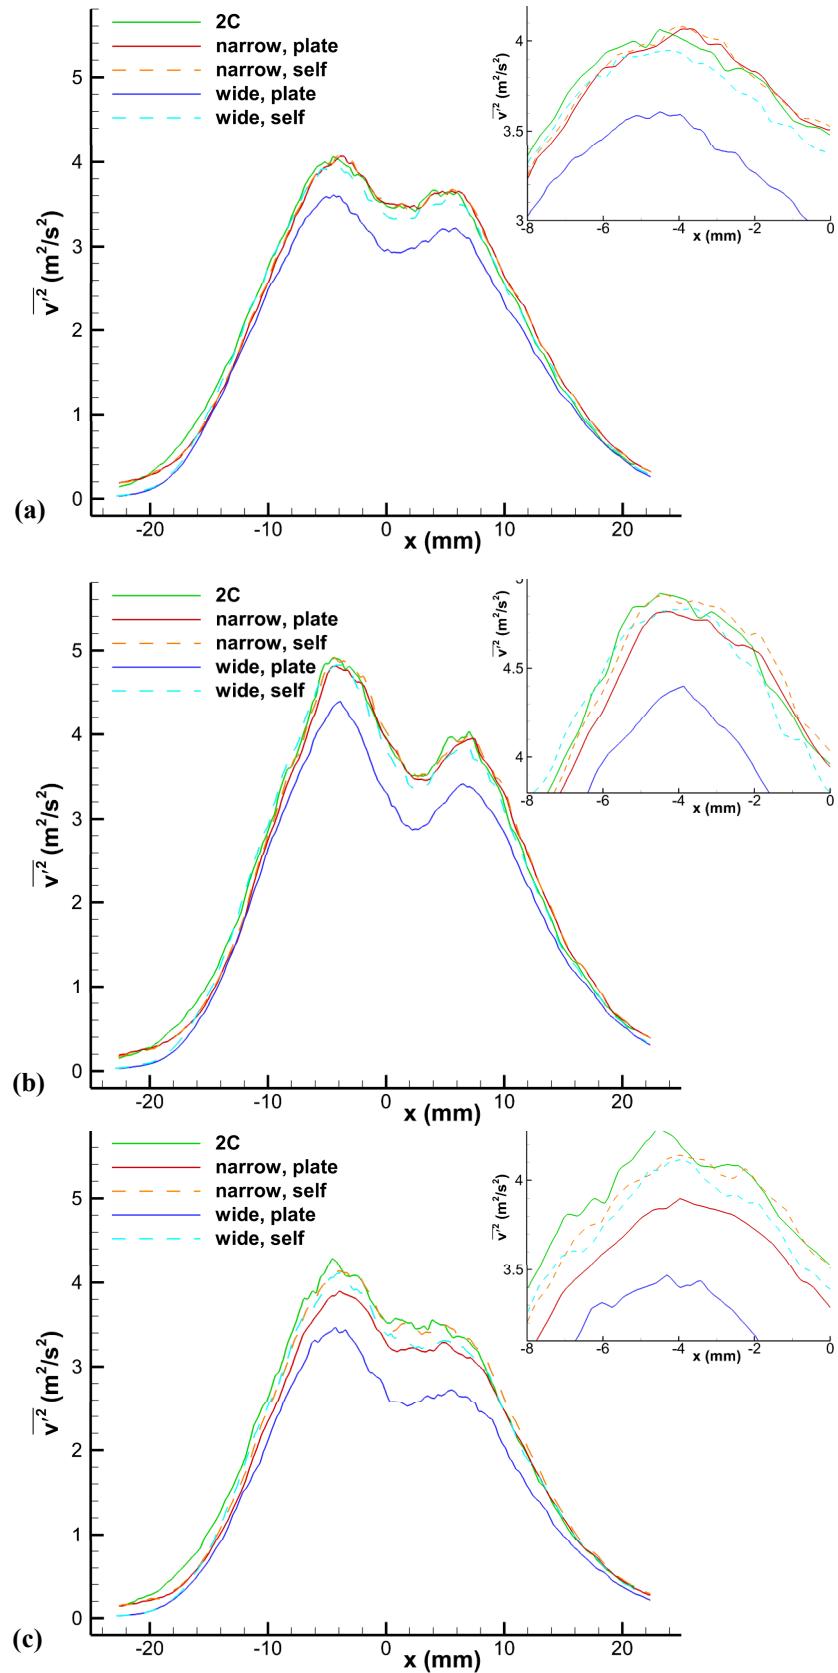


Fig. 5: Vertical turbulent normal stress at $y=113.5$ mm for all three camera configurations, with and without self-calibration (self and plate, respectively). (a) thin laser sheet; (b) medium laser sheet; (c) fat laser sheet.

Spatial Offset

To better appreciate the spatial shift created during the calibration process, an additional data processing algorithm is required to perform cross-correlation on the dewarped images that are used for the final image interrogation step. The algorithm employed by DaVis dewarps images based on the stereo calibration prior to performing a 2C cross-correlation on the image pair from each camera. While the dewarped images are not normally retained during processing, it is possible to export these images. However, the sCMOS cameras have an array size different than that of either the 2C camera or the ProX cameras. Furthermore, dewarping stretches the images uniquely to each configuration, resulting in different array sizes in each case. For this reason, it is necessary to resample the images prior to comparison. Since the correlation shape is unimportant and only the correlation peak location is required, resampling will not have significant adverse effects. Results are stated in mm rather than pixels since each camera has a unique resolution. By comparing the dewarped images from a stereo configuration to the unadulterated images from the 2C camera, it can be determined if the calibration process has resulted in any shift in the x - y origin. The results are presented in Fig. 6 as an error in x and y relative to the 2C camera. Again, the 2C measurements are regarded as being closest to the truth because the long lens and standoff distance minimize the perspective distortion and because the calibration is straightforward and requires no dewarping. It also bears reiterating that the three calibrations were performed simultaneously using a common target plate, which means that the spatial offsets are found relative to the origin dot as imaged by the 2C camera; the actual position of the origin dot relative to the jet exit is not relevant here.

For the results presented in Fig. 6, the interrogation windows were 10 mm square. The peak of each correlation was found and used to generate something similar to a disparity map from stereo self-calibration. In each case, the result from 50 images was averaged. This number was chosen based on a convergence study using the 0.8-mm laser sheet data. The result is the x - y mapping error as a function of position for a given camera. This method allows an examination of spatial variation in error. For the present case, most spatial error was uniform across the images, and large interrogation regions were used to trade resolution for precision.

Figure 6 indicates that significant spatial error is found in the narrow-angle data but less so in the wide data. The symbols show the spatial displacement in each individual camera whereas the lines show the displacement once the stereo measurements are assembled from the two views. The individual camera spatial errors indicate significant disparity along the x direction but not the y direction, which is entirely consistent with the stereo angle alignment to the x axis. The disparity is removed by self-calibration and the spatial error becomes an average of the disparities once the calibration is complete. No significant changes were found in these errors once self-calibration was performed, so they must arise in some fashion from the plate calibration. The x displacement of the narrow stereo configuration with respect to the 2C data is about 0.5 mm in magnitude. In the wide configuration, it is noticeable but only about 0.1 – 0.2 mm. This is consistent with observation of Figs. 4 and 5, in which the peaks of the narrow-angle data appear shifted from the 2C and wide data by approximately 0.5 mm in the $+x$ direction. The narrow data show a similar shift in the y direction, but of roughly half the magnitude; in the wide data it is nearly zero. The spatial offset in y seems particularly odd given that there is no stereo angle to correct or align along this axis.

Figure 6 also indicates that the disparity rises as the laser sheet thickness is increased. This is expected, as a

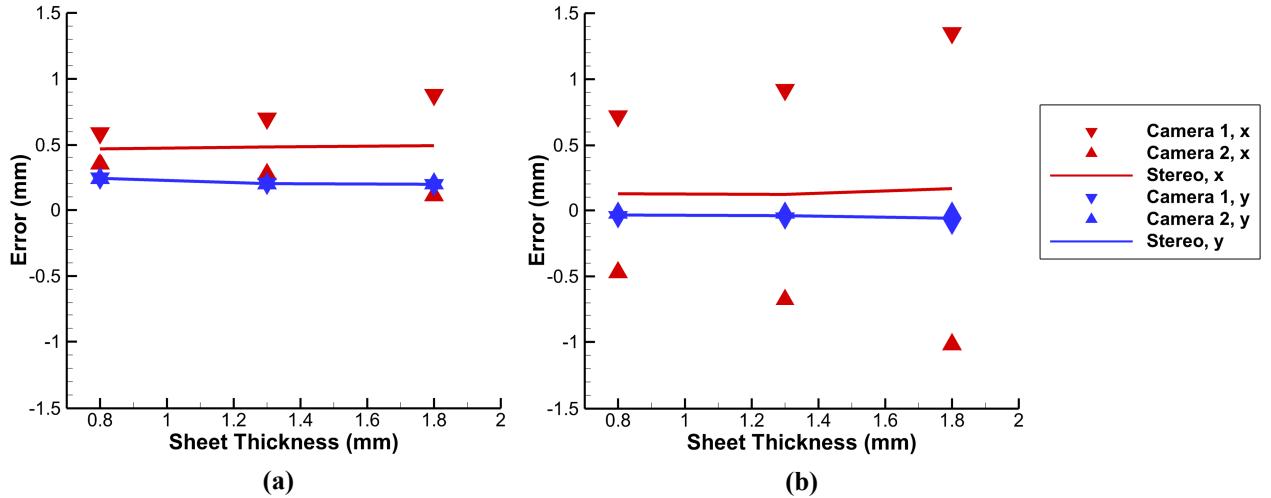


Fig. 6: Spatial offset from the two-component data following stereo calibration; (a) narrow-angle configuration; (b) wide-angle.

thicker laser sheet will exacerbate misregistration. It also is expected that the disparity should be greater for the wide-angle configuration, in which the cameras are more sensitive to misalignment in the z direction. However, self-calibration corrects the disparity and the spatial offset of the stereo combination remains reasonably constant with sheet thickness. The x and y error levels remained fairly stable throughout the experiment, even as different self-calibrations were conducted, indicating that the source of the spatial offsets lies in the original plate calibration and cannot be corrected by realignment due to self-calibration. The spatial error is dominantly a function of the initial camera orientation and little further drift occurs in the camera positions.

It is not clear why these spatial errors arise. The simple answer would be that the cameras shifted following the calibration, but there is ample evidence that this was not the case. Spatial offsets were studied run-by-run and found not to change significantly, demonstrating that the cameras did not move during the course of data acquisition. It seems unlikely that they would have moved immediately following calibration but never again. Furthermore, camera drift is more common along the y axis where gravity acts upon the camera and lens, but that is not consistent with Fig. 6. The y -axis camera drift also would have had to occur identically in each camera in the narrow-angle pair, which is implausible.

The spatial errors also do not appear to be due to a stretching of spatial scales during image warping since the cameras are stereoscopically separated only on the horizontal plane. Realignment in x may be a function of the stereo angle and any out-of-plane misalignment, but such an explanation does not hold for the origin change in y since this component is not sensitive to the stereo angle or out-of-plane translation. Furthermore, the error is greater for the narrow configuration, whereas it logically should be larger as the stereo angle is increased. Therefore, an additional error source appears to be present in establishing the origin of the calibrated axes.

It is tempting to ascribe the spatial error to some unidentified mistake during the calibration, but further investigation suggests that this is not the case. First, the original center-plane calibration images of the plate target were reloaded following the completion of the calibration. The calibration origin was found to exactly lie upon the desired origin dot. This indicates that the spatial offsets cannot be attributed to bias errors in locating the registration dots, which can occur if dots are misshapen due to a combination of sharp viewing angles and nonuniform illumination. This verification was reassuring, though accuracy of dot detection was carefully monitored during image acquisition and calibration. The spatial offsets were found to be as large as 0.5 mm compared to a dot diameter of 0.79 mm, which means a serious and obvious dot-finding error would be necessary to yield a spatial offset. It is possible that the spatial offsets only become apparent when data are acquired in a laser sheet of finite thickness, which creates volumetric effects and particle distributions that are unlikely to be perfectly symmetric along the z axis. A calibration check using a planar calibration image may mask an effect present in the real thickness of the laser sheet, even if the precise mechanism for this is unestablished.

Secondly, similar spatial offsets were detected in a subsequent experiment conducted entirely independently of the present one, using different equipment and calibration techniques (presently unpublished). The common element with this later experiment is that it too used a camera aligned normally to the laser sheet to supply an additional two-component measurement. This allowed the same correlation procedure to be applied between the 2C raw images and the stereo raw images, revealing similar spatial offsets. This suggests that the spatial offset phenomenon may be pervasive in stereo PIV and detectable only when the additional 2C measurement is available, though its source remains elusive.

The spatial resolution of a single vector in the 2C data is 0.35 mm and similar values are found for the two stereo configurations, with the caveat that the true spatial resolution of a PIV measurement is difficult to ascertain due to a complicated dependence upon the warping and filtering algorithms found in modern image interrogation algorithms [10]. Therefore, the detected error levels due to spatial offset are on the order of a single vector spacing. These levels of error may seem rather negligible in most circumstances, but they are significant in carefully conducted experiments. For statistically-averaged fields such as the mean velocities of Fig. 4 or the turbulence intensities of Fig. 5, the spatial offsets are a minor effect and not readily noticeable – though a close look does reveal their presence. However, the spatial offsets are quite apparent when examining instantaneous velocity fields.

The relevance of the spatial offset is most clearly appreciated by examining the effects on measurements of a single turbulent eddy. In Fig. 7, an instantaneous snapshot is shown over a smaller region of the larger field of view, taken from the thin sheet data set. This plot shows velocity fluctuation vectors to reveal the instantaneous turbulent structure, superposed on a contour plot of the total vertical velocity magnitude. This particular plot was created using the narrow-angle camera configuration, but the 2C and wide-angle data are much the same. Two distinct vortices are seen. The white lines indicate positions over which velocity profiles were extracted.

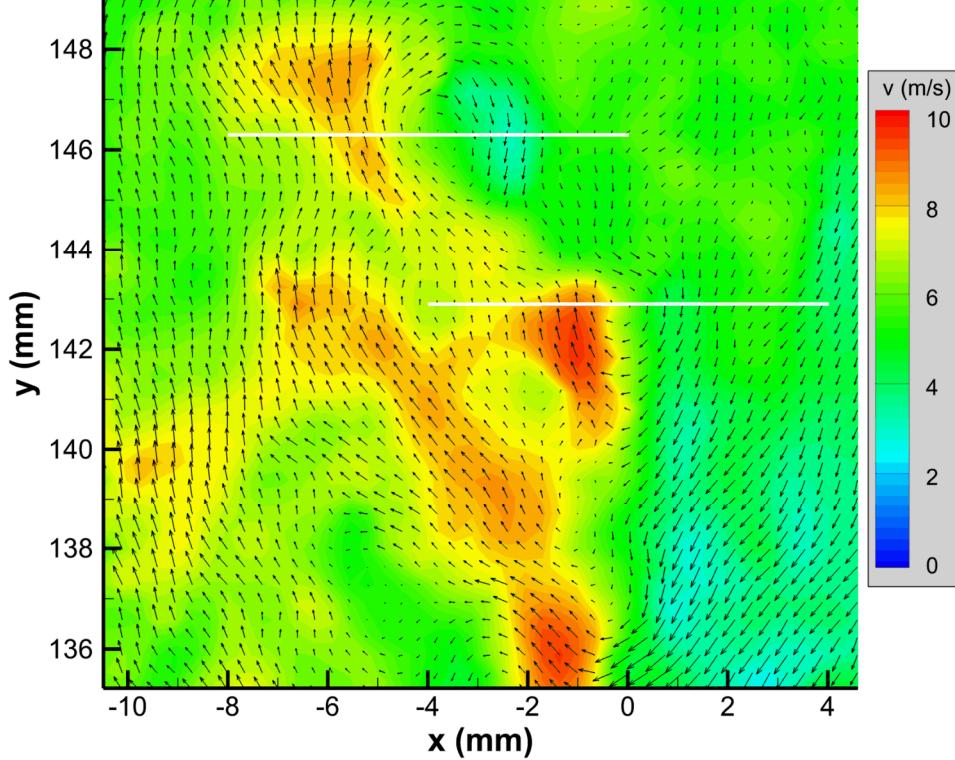


Fig. 7: A representative instantaneous snapshot from the thin laser sheet case. Vectors show velocity fluctuations and contours show total vertical velocity magnitude. The white lines indicate locations of velocity profiles extracted in Fig. 8.

The extracted vertical velocity profiles are shown in Fig. 8 at the two locations drawn in Fig. 7, displaying all three camera configurations. The solid lines denote the data as discussed thus far, including self-calibrations in the stereo cases. The narrow-angle data show a clear offset of about 0.5 mm in the $+x$ direction with respect to the 2C measurements, and the wide-angle data show a smaller offset in the same direction. This is consistent with the spatial offsets that were shown in Fig. 6. When corrections for the spatial offsets are employed on both x and y axes by incorporating the values found in Fig. 6 and extracting new profiles from Fig. 7, the dashed lines of Fig. 8 indicate that much better agreement with the 2C data is obtained. This is clearly evident in Fig. 8a at the $y = 146.2$ mm location and is mostly true in Fig. 8b at the $y = 142.8$ mm location; in the latter case, the wide-angle data is not as well corrected by the spatial offsets though the general trend is accurate. Velocity magnitudes all are essentially in agreement in Fig. 8a both before and after the spatial correction, but in Fig. 8b some reduction in the magnitude of the narrow-angle profile is found once spatial correction is applied. This occurs because the vertical offset moves the extracted velocity profile to a position where the velocity magnitude is weaker.

Higher velocity gradients create more dramatic impacts of the spatial offsets. Figure 9 shows an example of a profile extracted from a region where large gradients are found, particularly in the lateral component. This is seen along the marked line at about $x = 5$ mm, where vectors can be seen to turn from a generally leftward direction to a generally rightward direction over only about three vector spacings, all while slowing the magnitude of the downward velocity component. The extracted velocity profiles for both the u and v components are given in Fig. 10.

Figure 10a is particularly disturbing in that the curve for the narrow-angle configuration appears quite different once spatial offsets are corrected. The u component near $x = 6$ mm shows the narrow-angle data considerably positive of the 2C and wide data but then showing excellent agreement following spatial correction. The uncorrected curve yields an error of about 2 m/s, or roughly 20% of full-scale. This occurs primarily because of the vertical spatial offset as detailed in Fig. 6, which locates the velocity profile at an incorrect vertical position and hence extracts mismatched velocities. Similar but less dramatic observations are found in the v component in Fig. 10b.

Figures 7-10 demonstrate the degree of velocity error that can occur from what appears to be a relatively mild spatial offset that is generated by an unknown source within stereo calibration. Though this bias may have a

minimal effect on statistical properties of the data set, which tend to possess longer length scales, they can be quite prominent in instantaneous velocity fields of individual turbulent structures. The source of the error presently remains undetermined, but its presence in a carefully conducted experiment suggests it may be hidden in many other experiments as well, undetectable without the additional two-component measurement used as a dependable reference.

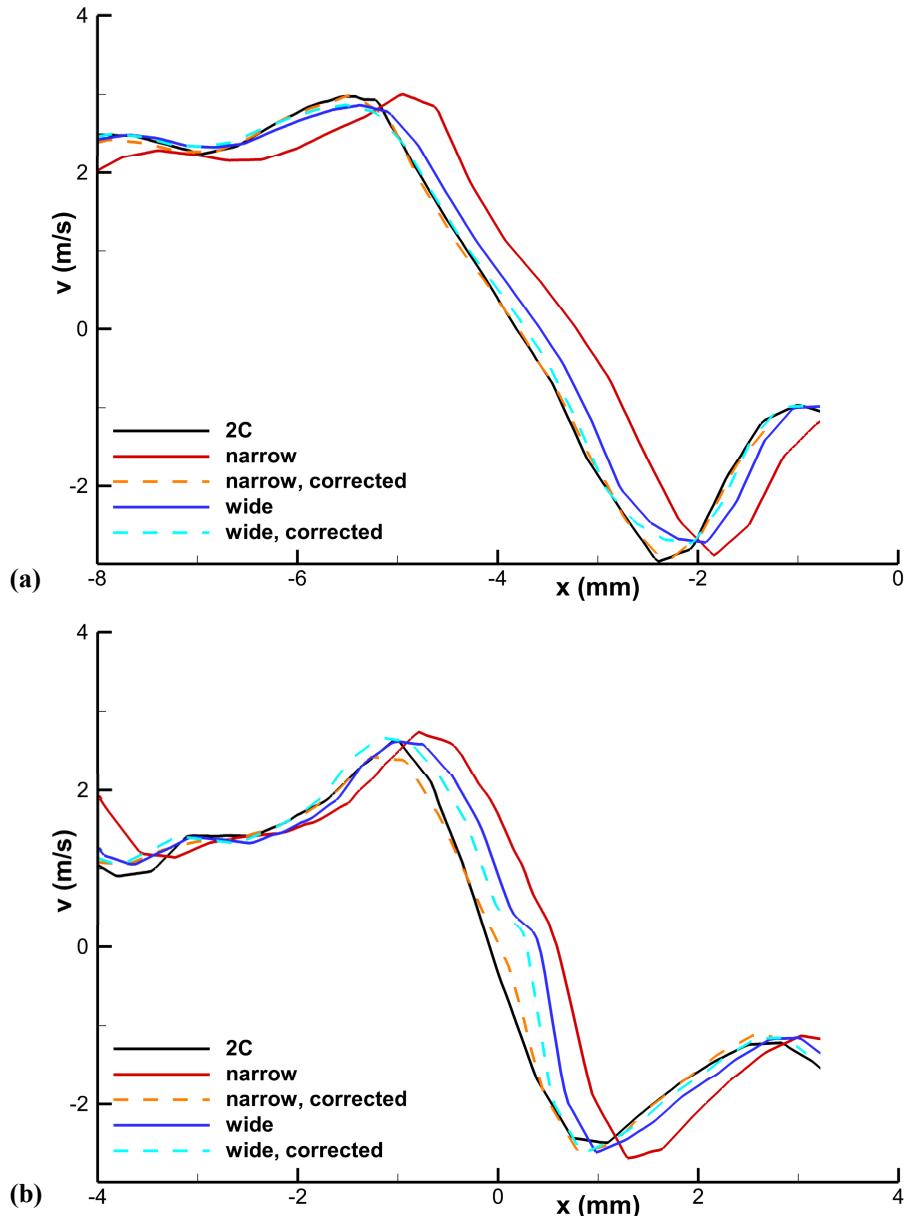


Fig. 8: Vertical velocity profiles extracted from the positions marked in Fig. 7. Solid lines are data following self-calibration and broken lines are data incorporating the spatial offsets of Fig. 6. (a) $y = 146.2$ mm; (b) $y = 142.8$ mm.

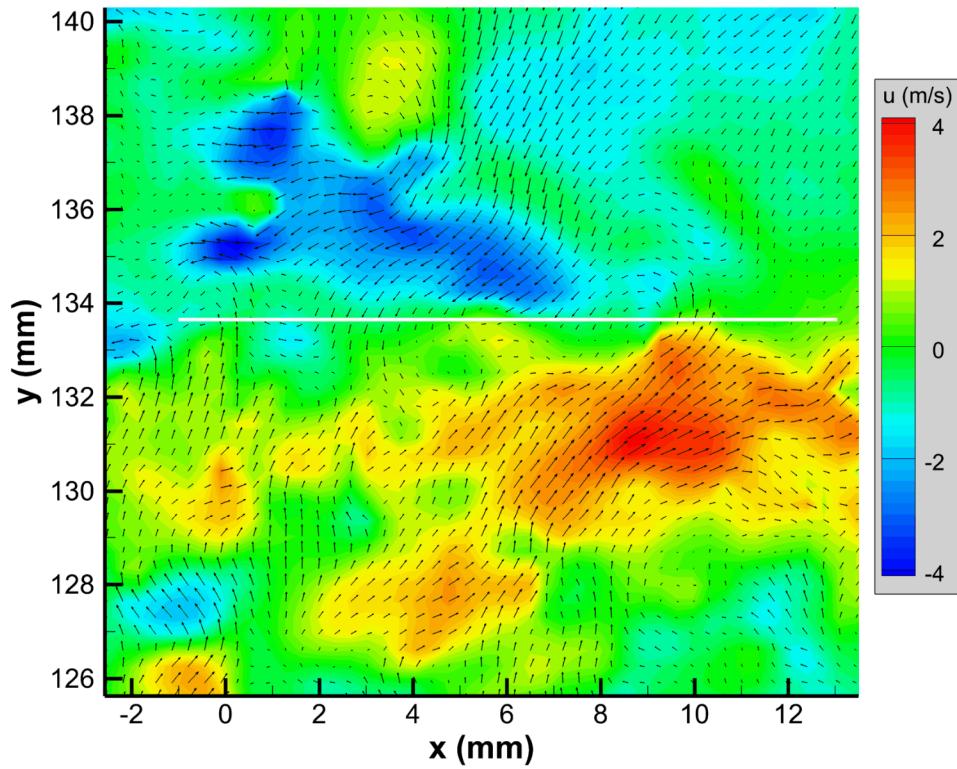


Fig. 9: An instantaneous snapshot from the thin laser sheet case containing a high-gradient region. Vectors show velocity fluctuations and contours show total lateral velocity magnitude. The white line indicates the location of velocity profiles extracted in Fig. 10.

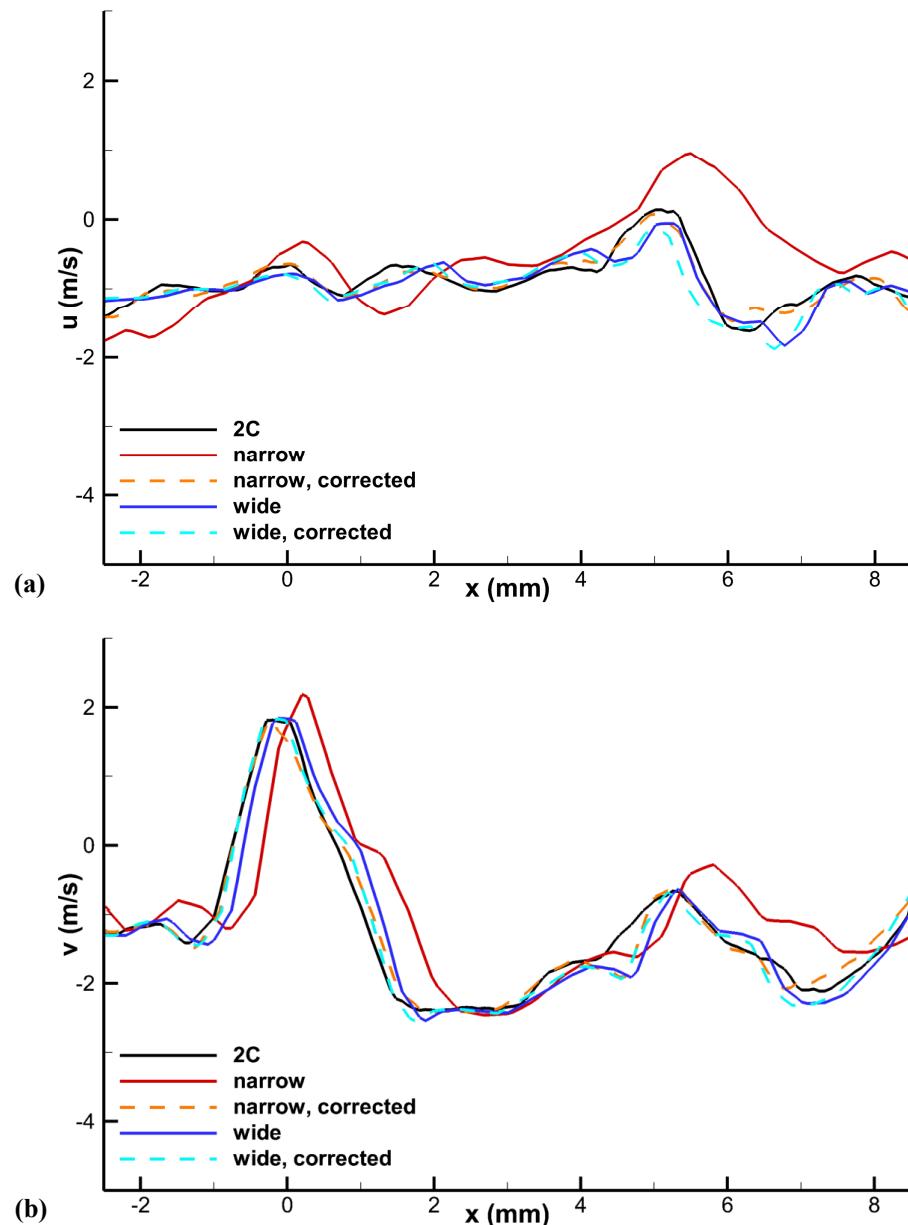


Fig. 10: Velocity profiles extracted from the position marked in Fig. 9, $y = 133.5$ mm. Solid lines are data following self-calibration and broken lines are data incorporating the spatial offsets of Fig. 6. (a) lateral component; (b) vertical component.

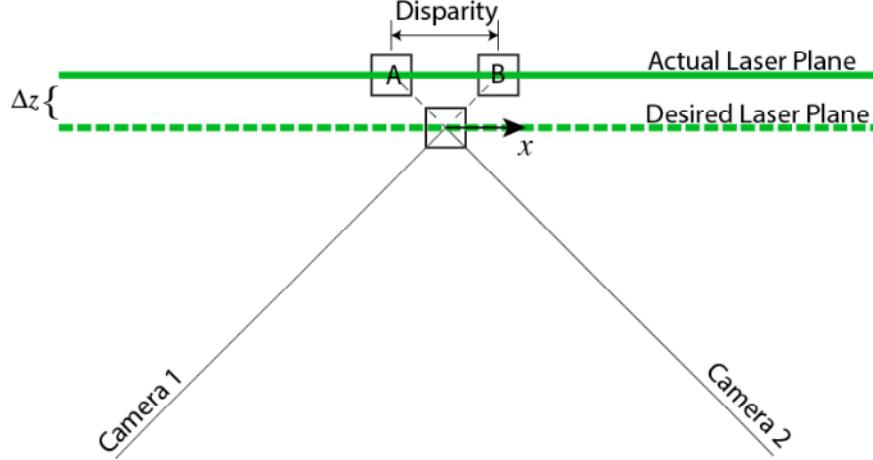


Fig. 11: Sketch depicting an SPIV setup for which the laser sheet lies on a plane different than the calibration plane.

Turbulent Normal Stress

Figure 5 demonstrated that misalignment in a stereo calibration that is not corrected by self-calibration leads to an artificial reduction in the turbulent normal stress. This occurs because the two cameras are not correctly registered to each other spatially. The magnitude of this bias, therefore, is expected to be a function of the local disparity vector that is generated by the self-calibration algorithm as it examines the correlation of the two cameras to each other.

Figure 11 depicts a situation in which the laser sheet does not lie in the plane for which the calibration was performed. In this case, each interrogation region for camera 1 is shifted to the right by an amount that is a function of the laser plane shift and the camera angle, while the interrogation region for camera 2 is shifted to the left. In general, the laser sheet also may be rotated with respect to the calibration plane. The distance between A and B is the disparity vector for this location. For the setup used here in which the rotations of the cameras around the x axis are nominally zero, each camera is equally sensitive to motion in the y direction, which is out of the page in Fig. 11. Therefore, the resultant instantaneous v value is a simple average of the v result from each camera, viewing locations denoted A and B.

$$v = \frac{1}{2}(v_A + v_B) \quad (1)$$

Assuming the time average of v is zero, the variance of the measured velocity fluctuation, or the turbulent stress in the vertical component, is

$$\overline{v'v'} = \frac{1}{N-1} \sum_{i=1}^N \left[\frac{1}{2} (v_{A,i} + v_{B,i}) \right]^2 \quad (2)$$

$$= \frac{1}{N-1} \sum_{i=1}^N \frac{1}{4} (v_{A,i}^2 + v_{B,i}^2 + 2v_{A,i}v_{B,i}) \quad (3)$$

$$= \frac{1}{4} (\sigma_{v_A}^2 + \sigma_{v_B}^2 + 2\sigma_{v_A}v_B) \quad (4)$$

The last term inside the parentheses can be related to the covariance coefficient of the velocity at points A and B, which has the definition

$$\rho_{AB} = \frac{\sigma_{v_A}v_B}{\sigma_{v_A}\sigma_{v_B}} \quad (5)$$

Assuming A is in the vicinity of B, it is likely that $\sigma_{v_A} \approx \sigma_{v_B}$. Making this assumption and using the definition of the covariance coefficient,

$$\overline{v'v'} = \frac{1}{4} (2\sigma_{v_A}^2 + 2\sigma_{v_B}^2) = \frac{\sigma_{v_A}^2}{2} (1 + \rho) \quad (6)$$

Therefore, for perfect correlation at the two locations ($\rho=1$), the reported turbulent normal stress would be the actual variance of the velocity at A or B. For the uncorrelated case, ($\rho=0$), the reported turbulent stress would be $\sigma_{v_A}^2/2$. Only for negative covariance coefficients (which should be rare) can the error become larger than 50%.

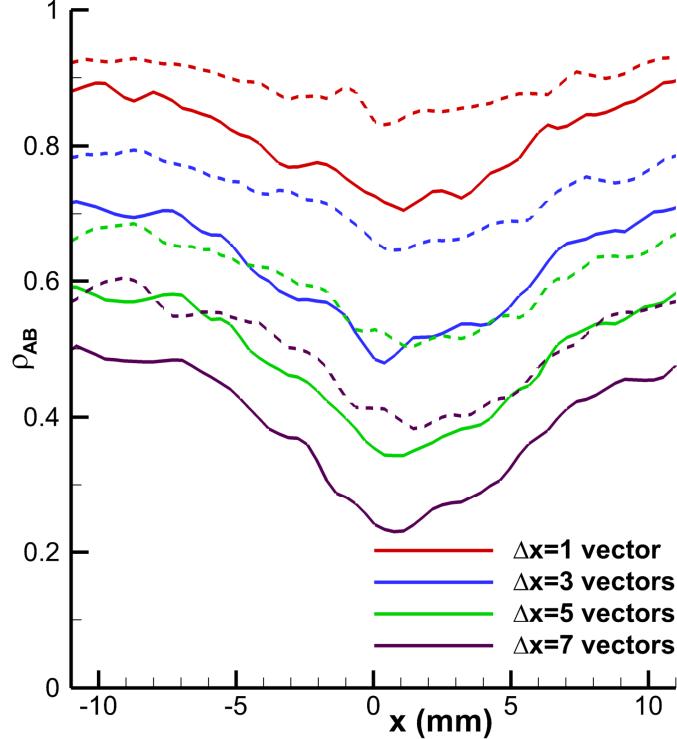


Fig. 12: Covariance coefficients computed from the fat-sheet 2C data for different disparity values. Solid lines are computed at $y = 113.5$ mm, dashed lines at $y = 143.5$ mm.

The covariance coefficient can be computed using the 2C data, since there is no risk of biases associated with stereo calibration. The covariance coefficients for a few different values of disparity, expressed in vector spacings, are shown in Fig. 12. These values were calculated using the 2C data for the fat laser sheet. Curves are shown for each disparity value along lines at $y = 113.5$ mm and $y = 143.5$ mm, which correspond to locations near the bottom and top of the imaging region. The covariance is a strong function of disparity and increases with downstream distance. All values are greater than zero.

Using equation 6, the error in the streamwise turbulent stress can be estimated as per $\epsilon_{v2}=0.5(1+\rho)-1$, which yields the estimated error as a fraction of the measurement. Errors are negative since disparity reduces the measured turbulent stress. Results are shown in Fig. 13 for four different points in the flow field as a function of the disparity in vector spacings. Errors vary in space because the covariance coefficient varies in space. Data at $x = 0$ mm lie near the jet centerline where the covariance coefficient is smallest; data at $x = -5$ mm lie near the peak of the turbulent stress, as evident in Fig. 5. This plot emphasizes that as the disparity increases, the error increases strongly, reaching values as high as 40%. Even for a small disparity of only a single vector spacing, the error ranges from 5% to 15% for the present experiment. Clearly, without self-calibration to minimize the disparity, a large negative bias error will occur in the turbulent stress. This was evident in Fig. 5 previously and noted to worsen with increased laser sheet thickness. The latter effect likely occurs because an oblique view of a thicker sheet effectively widens the volume of an interrogation window laterally as well, increasing the real spatial area over which the vector is calculated.

These data are applied to the worst-case scenario (see Fig. 5) of the fat laser sheet. The disparity before self-calibration was found to be 0.64 mm for the narrow-angle configuration and 2.35 mm for the wide angle, corresponding to spatial errors of roughly 2 and 7 vector spacings, respectively. Multiple iterations of self-calibration reduced the disparity to less than 0.1 mm, or about one-quarter of a vector spacing. This information, obtained during the self-calibration process, can be used to select a covariance coefficient from Fig. 12 or similar and equation 6 then predicts the turbulent stress based only upon the standard deviation of the velocity from a single camera.

Based on the covariances in Fig. 12 at $y = 113.5$ mm, the plate calibration for the wide-angle measurements is expected to report a turbulent stress about 40% lower than the 2C measurement, whereas self-calibration should reduce this error to nearly 0%. Lifting the pertinent values from Fig. 5c, the peak 2C turbulent stress is $4.3 \text{ m}^2/\text{s}^2$

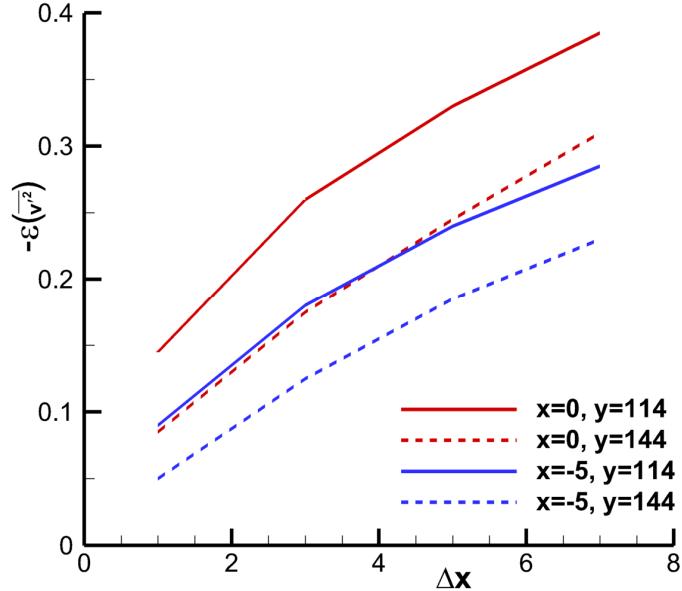


Fig. 13: Estimated error in the streamwise turbulent stress as a function of disparity at four points in the flow field. Error is given as a fraction of the measurement and Δx in vector spacings. Data are for the fat laser sheet.

and the wide-angle value is $3.5 \text{ m}^2/\text{s}^2$ prior to self-calibration and $4.1 \text{ m}^2/\text{s}^2$ following self-calibration. This equates to 19% and 5% reductions, respectively; the former value is about half that of the prediction and the latter indicates some residual error even after the self-calibration. Though not shown in a current plot, the same test was conducted at $y = 143.5 \text{ mm}$ on the same laser and camera configuration. At this location, the 2C turbulent stress is a maximum of $3.1 \text{ m}^2/\text{s}^2$ and the wide-angle value is $2.4 \text{ m}^2/\text{s}^2$ and $2.9 \text{ m}^2/\text{s}^2$ before and after self-calibration. These are reductions of 23% and 6%, respectively, compared to predictions of about 30% and 0%; this is somewhat better agreement than at $y = 113.5 \text{ mm}$.

In principal, the error estimates can be used to correct the measured turbulent stress. This is attempted in Fig. 14, which plots all three measurements of the turbulent stress profiles for the fat sheet and adds corrected profiles for both the narrow and wide configurations. Figure 14a repeats the data of Fig. 5c at $y = 113.5 \text{ mm}$ and Fig. 14b provides an analogous plot at $y = 143.5 \text{ mm}$, at which data were not previously shown. Clearly, the data are considerably overcorrected in Fig. 14a and mildly overcorrected in Fig. 14b. Therefore, the predictions of the bias error in the streamwise turbulent stress are estimates only and can be relied upon only to indicate the trend and order of magnitude of the error. These limitations likely are a function of the simplifying assumptions made in the analysis, in which the standard deviations were assumed identical for each camera and that self-calibration realignment to a new plane does not greatly change the velocity. These assumptions would be disturbed by substantial gradients in the instantaneous velocity field, either within the measurement plane or out of plane. Figures 7 and 9 show that such gradients are not necessarily small. For a realistic experiment, a successful self-calibration still is required to obtain an unbiased measurement.

Experiments that do not utilize self-calibration are likely to suffer a bias error in the measurement of turbulent stresses, the severity of which depends upon the precision of the stereo calibration as well as the correlation field. This would include stereoscopic PIV experiments found in the historical database. Although the present analysis has focused on only one component of the turbulent stress tensor, bias errors should be anticipated in other components as well. A successful self-calibration should reduce these errors nearly to zero but the quality of the self-calibration will determine whether some error remains. For configurations where a good self-calibration is difficult to obtain, such as a thick laser sheet and/or a wide stereo angle combined with high particle density, some residual error is likely.

Conclusions

The effectiveness of stereoscopic PIV calibration and its residual errors has been investigated by using redundant independent PIV measurements. A single-camera two-component measurement was conducted simultaneously with

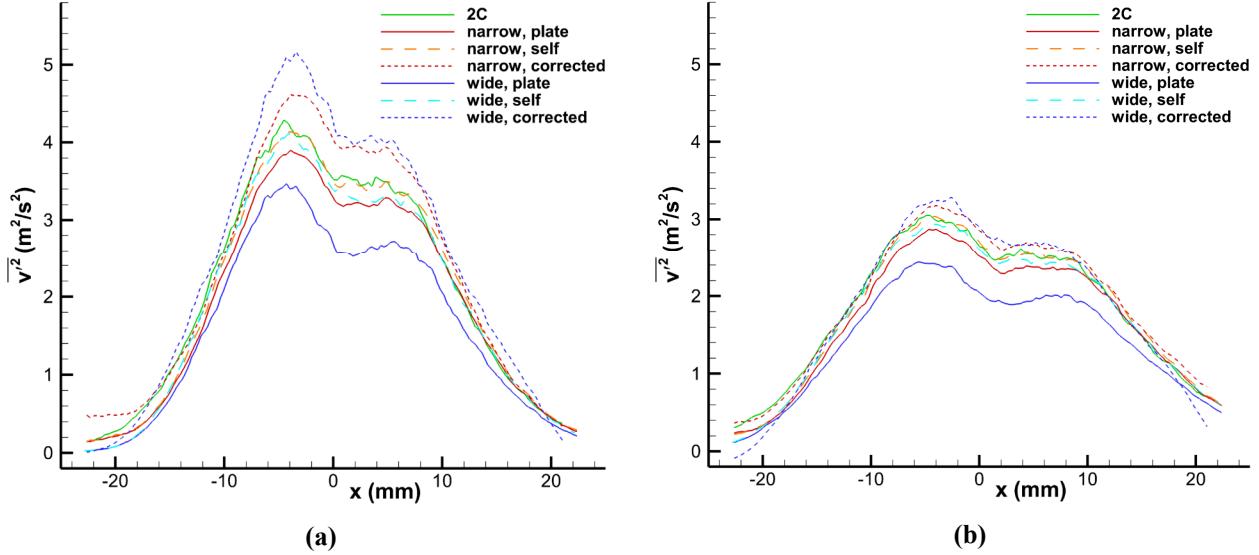


Fig. 14: Vertical turbulent normal stress for the fat sheet as Fig. 5, but additionally including corrected profiles using equation 6. (a) $y=113.5$ mm; (b) $y=143.5$ mm.

two separate stereo configurations at narrow and wide stereo angles. The two-component measurement was configured to minimize perspective distortion and is free of any error due to stereoscopic calibration or dewarping. A turbulent free jet has been used as a simple test flow. The laser sheet thickness also was varied as a test parameter.

Comparison of the stereo measurements with the two-component measurements reveals an erroneous spatial offset present in the stereo data sets, artificially displacing vector fields from the expected origin. This spatial error typically is small with respect to statistical properties of a data set, but can be prominent when instantaneous snapshots of the velocity field are examined, particularly in regions where the velocity gradient is momentarily large. Though the source of this spatial error remains elusive, it appears to arise during the target plate calibration and it is not altered by self-calibration. This spatial error was detectable only through comparison with the simultaneous two-component measurement and it may be hidden in many historical stereo experiments as well.

Self-calibration proved necessary to obtain high levels of accuracy in stereo measurements. Errors in mean velocities were small, but artificially reduced turbulent stresses were generated when self-calibration was not used, owing to a smearing effect that occurs when the two cameras are inadequately registered to each other and therefore effectively broaden the volume over which the vector is computed. This difficulty worsened with increased laser sheet thickness. Bias errors in the turbulent normal stress using only a target calibration were found to be on the order of 20% of the measurement found using self-calibration and the two-component data. This observation suggests that historical experiments conducting stereo PIV measurements of turbulent stresses that did not employ self-calibration may retain significant bias errors.

References

- [1] Raffel, M., Willert, C. E., Wereley, S. T., and Kompenhans, J., *Particle Image Velocimetry: A Practical Guide*, 2nd Ed., Springer, 2007.
- [2] Adrian, R. J., and Westerweel, J., *Particle Image Velocimetry*, Cambridge University Press, 2011.
- [3] Lawson, N. J., and Wu, J., "Three-Dimensional Particle Image Velocimetry: Error Analysis of Stereoscopic Techniques," *Measurement Science and Technology*, Vol. 8, No. 8, 1997, pp. 894-900.
- [4] Zang, W., and Prasad, A. K., "Performance Evaluation of a Scheimpflug Stereocamera for Particle Image Velocimetry," *Applied Optics*, Vol. 36, No. 33, 1997, pp. 8738-8744.
- [5] Prasad, A. K., "Stereoscopic Particle Image Velocimetry," *Experiments in Fluids*, Vol. 29, No. 2, 2000, pp. 103-116.
- [6] Prasad, A. K., and Adrian, R. J., "Stereoscopic Particle Image Velocimetry Applied to Liquid Flows," *Experiments in Fluids*, Vol. 15, No. 1, 1993, pp. 49-60.
- [7] Willert, C., "Stereoscopic Digital Particle Image Velocimetry for Application in Wind Tunnel Flows," *Measurement Science and Technology*, Vol. 8, No. 12, 1997, pp. 1465-1479.
- [8] Coudert, S. J. M., and Schon, J.-P., "Back-Projection Algorithm with Misalignment Corrections for 2D3C Stereoscopic PIV," *Measurement Science and Technology*, Vol. 12, No. 9, 2001, pp. 1371-1381.

- [9] Van Doorne, C. W. H., Westerweel, J., and Nieuwstadt, F. T. M., "Measurement Uncertainty of Stereoscopic PIV for Flow with Large Out-of-Plane Motion," *Particle Image Velocimetry: Recent Improvements*, Proceedings of the EUROPIV2 Workshop, 2004, pp 213–227.
- [10] Stanislas, M., Okamoto, K., Kahler, C. J., Westerweel, J., and Scarano, F., "Main Results of the Third International PIV Challenge," *Experiments in Fluids*, Vol. 45, No. 1, 2008, pp. 27-71.
- [11] Wieneke, B., "Stereo-PIV Using Self-Calibration on Particle Images," 5th International Symposium on Particle Image Velocimetry, Paper 3106, Busan, Korea, September 2003.
- [12] Wieneke, B., "Stereo-PIV Using Self-Calibration on Particle Images," *Experiments in Fluids*, Vol. 39, No. 2, 2005, pp. 267-280.