ABSTRACT

Title of Thesis: EXCITATION AND PHOTOGRAMMETRY ANALYSIS OF FLUID-STRUCTURAL VIBRATIONS

Matthew Vincent Killian
Master of Science, 2023

Dissertation Directed by: Professor Stuart Laurence
Department of Aerospace Engineering

The field of fluid-structural interactions (FSI) requires specifically designed measurement systems that can be used to interpret the results of experiments without interfering with any aspects of the tests. In our hypersonic experiments, the chosen method of measurement is photogrammetry, the processing of 2-dimensional images to obtain 3-dimensional positional information about a structure exposed to a flow. To examine vibrations, the test specimen is painted with evenly spaced markers which are then photographed by a stereo digital image correlation setup of two high-speed cameras. There exists an effective algorithm for processing these images to obtain displacement and deflection data that can in turn be analyzed using spectral proper orthogonal decomposition (SPOD) to find vibrational modes. However, the current method for locating markers within an image is computationally expensive and slow, so a new algorithm was adapted to perform the same task. This adapted method differs from the old method by not being iterative, allowing it to run more quickly as it detects the markers and tracks them between images. We verified the efficacy of this new algorithm with two calibration tests,

one with artificial marker images and one with real images of a painted plate translated at known displacements. After characterizing the errors of the method, it was tested on FSI experimental data collected at the NASA Langley Facility in their Mach 10 wind tunnel. The results of these tests showed that the algorithm can be quick and accurate, but it is not robust with regards to non-ideal image conditions. The images obtained in FSI photogrammetry are often not ideal, so this method must be developed further. A mechanism for test specimen excitation was also explored. We evaluated a solenoid-based prototype by performing a modal test on a compliant panel with a vibrometer. The results of this test show that the prototype is effective in producing strong and reliable vibrations in the test panel, and as such this should be developed further for use in hypersonic wind tunnel tests.

EXCITATION AND PHOTOGRAMMETRY ANALYSIS OF FLUID-STRUCTURAL VIBRATIONS

by

Matthew Vincent Killian

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Master of Science
2023

Advisory Committee:
Professor Stuart Laurence, Chair/Advisor
Professor Christopher Cadou
Professor James Baeder

Acknowledgments

Firstly, I would like to thank Dr. Stuart Laurence for giving me the chance to work on this project. I am incredibly thankful for the opportunity to work as a graduate researcher and I have certainly learned a great deal throughout my time.

I also want to thank Dr. Christopher Cadou and Dr. James Baeder for agreeing to be part of my thesis committee. I have learned a lot from both in my coursework, so it is special to have them be a part of this aspect as well.

For teaching me how to be a graduate student, I thank Dr. Shaun Skinner. He helped me settle into labwork, even though it was initially part of a different project.

For all his help teaching me about fluid-structural interactions, sharing data, and explaining code, I thank Antonio Schöneich, who also collected the experimental data for this project.

For keeping me sane as I wrote this thesis, I thank the television programs *Columbo* and *Seinfeld*, which were consistently playing in the background.

Of course I must also thank my family who kept me going as I worked in the lab as an undergraduate and a graduate student.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

DAQ	Data acquisition device
DC	Direct current
DIC	Digital image correlation
FOV	Field of view
FSI	Fluid structural interactions
PSD	Power spectral density
SPOD	Spectral proper orthogonal decomposition
STFT	Short-time Fourier transform

Chapter 1: Introduction

1.1 Overview

High-speed flow is often analyzed on its own in experiments as we try to understand and characterize compressibility, real gas effects, heating, and more. However, in many high-speed applications we must also be aware of how the flow interacts with whatever structures exposed to it. This field of fluid-structural interactions (FSI) is important but can be difficult to study. When designing experiments for these interactions we must consider methods of excitation and measurement that do not affect the natural properties of the test specimens. This non-invasive requirement forces us to develop different ways to carry out these experiments.

There are several methods of measurement for FSI tests, including temperature and pressure sensitive paint, but we will focus on a fully non-invasive method of displacement measurement. We use a multi-camera setup during hypersonic wind tunnel experiments to take high-speed photographs of a test specimen from outside of the flow. This method can determine displacements and deflections of the specimen through a process of marker tracking on the specimen in tandem with camera calibration that can solve for a 3-dimensional estimate based on 2-dimensional images. The algorithm that processes these images can be quite computationally intensive, so there is a considerable effort to devise a method that can be quick while still being effective.

Current tests generally focus on the interactions only between the structure and the flow, but further experiments are to involve active excitation of the structure as it interacts with the flow. This can help test how a structure will react if it has induced vibrations, which are common in aircraft applications due to engines, control surfaces, and more. Once again, there is a non-invasive requirement when designing such an excitation mechanism. There cannot be any direct attachments to the test specimen or else there would be changes in the natural properties. As such, there must be a device that can impact the test specimen without affecting it outside of the excitation.

1.2 Past FSI Work

The present work directly follows that done by Schöneich et al.[1] All experiments are further steps in the same campaign that aims to analyze the vibrational modes of a compliant panel in a hypersonic flow. This is motivated by the design of hypersonic aircraft, namely their slender, lightweight components.[2] There is significant research to be done in this field, specifically involving the interactions of shock-waves and boundary layers.[3][4][5] In the past, researchers such as Willems et al.[6] used mounted sensors like accelerometers and strain gauges to measure structural vibrations, but these results are affected by the invasive nature of the sensors. Experiments relating specifically to hypersonic FSI are unfortunately limited. Currao et al.[7] has done work involving a cantilevered plate beneath a Mach 6 Prandtl-Meyer expansion measured with schlieren imaging. Research into panel vibrations has also been done by Casper et al.[8] using a 7° half-cone at Mach numbers of 5,6 and 8. These vibrations were measured with accelerometers. While the results of hypersonic FSI experiments are in and of themselves important, the data collected by Schöneich et al. are being used here as tests for a new photogrammetry algorithm. We cannot simply use low-speed experiments for these tests because the algorithm is specifically being developed for hypersonic applications. There have been limited programs dedicated to studying these flows, and as such there is not much existing work to inform the data processing aspect of the experiments. The design of the experiments is something of an extension of work done at supersonic speeds involving modal analysis[9] and digital image correlation (DIC)[10]. The efficacy of these methods must be proven at the hypersonic conditions of the current FSI experiments which are performed at Mach 6 and, for this study, Mach 10. The results of the new adapted algorithm are compared to the method described by Whalen et al.[11], which will be further explained in Ch. 2.

1.3 Objectives

The goal of this work is to further the development of better experimental methods in the field of FSI. Chapters 2,3, and 4 are focused on the adaptation and verification of a new processing technique for photogrammetry. These chapters aim to explain the differences between the current method and this new method while also evaluating the performance of the new algorithm. The evaluations include benchtop tests with known data to measure errors in performance and application to real FSI experiments to see how the algorithm handles the true conditions. Chapter 5 discusses the design and performance of an excitation mechanism prototype. This is the first step in a line of work that will culminate in the addition of a similar mechanism in the FSI campaign. We seek only to prove that the proposed design is adequate for exciting the natural modes of the test specimen that is used in the FSI experiments.

Chapter 2: Explanation of Photogrammetry Algorithms

2.1 Overview

For examining the fluid-structural interactions in high-speed flows, a method must be devised that can extract displacement information from each test. The most significant concerns with any method are how invasive it is and how accurate it is. Photogrammetry, the processing of 2-dimensional images to obtain 3-dimensional positional information about a structure exposed to a flow, is ideal since it is non-invasive and it utilizes equipment already used in hypersonic facilities, namely high-speed cameras. To address the other concern, any image processing for these photos must accurately track the displacement of any object of interest. The algorithm chosen must be able to deal with noise and non-ideal camera placement while retaining precision. For the specific experiments with which we are concerned, the most important aspects of the structure's displacement are its vibrational modes. These are determined by tracking small, painted markers on the test specimens which are naturally more difficult to track due to limitations in image resolution. Our lab is currently considering two different methods to track these markers, each with its own benefits and drawbacks. The first of these has been used for analysis in the past, but the second method is adapted from a different hypersonic application. This adapted method uses a simpler algorithm, so it may be much quicker. However, this may come at the cost of sufficient accuracy. We will describe both methods and investigate the performance of the new method to determine whether it will be more effective.

2.2 Current Processing Method

The currently implemented photogrammetry routine uses an iterative method to find and track markers in each frame of high-speed video. Starting with a 12-bit black-and-white image, any image processing algorithm must first determine where markers may be based on the intensity in each pixel. The raw data will show which pixels are dark and which are bright, but anything beyond that requires some processing. The current method calls for user input to begin the process. Since the markers are arranged in a 2-dimensional grid of known size, the user need only input a corner point, a point in the x-direction, and a point in the y-direction. This triad, along with the predefined dimensions, allow the algorithm to trace along the grid to where each point should be as an initial guess. To begin finding the true location of each marker, the method creates a subimage around the initial guess. Within the subimage, it performs a ”center-of-mass” calculation, essentially assigning weights to each pixel based on its intensity and placing a guess for the marker center at the centroid of the subimage. After completing these initial calculations, a hyperbolic tangent function centered at the centroid is used to created an intensity curve that guesses the edges of the marker. As these guesses are traced out, forming an ellipse, they are compared to the actual points that define the edge of the shape. The iteration begins as the algorithm uses the Nelder-Mead numerical method[12] to minimize the error between the modeled shape and the actual subimage by varying several parameters of the ellipse. Once the parameters have been altered such that the generated shape matches sufficiently, the position of the marker is judged to have been determined and the information is saved for use in the displacement tracking.

This method is very accurate and especially good for the low-resolution images produced in high-speed photography. Accuracy is improved with subpixel detection which samples 100 points within each pixel to compute average density for the detection parameters.[11] The drawbacks come from the sheer number of calculations that must be performed for each test. In the FSI experiments, the panel is painted with 625 markers, each of which should be tracked to accurately characterize the vibrational modes. Additionally, since the images are sampled at such a high rate (30-31 kHz), each test produces over 60,000 images per camera. Every one of these markers must be tracked for the best results, and with iteration for each one the time and computational costs grow immensely. When the code for this method was written, it was hard-coded to use 8 cores of the CPU, which can be draining even on high-powered computers. That being said, since the algorithm was written specifically for the FSI application, it is optimized for these conditions and does not require much tweaking to run for each test. All of these aspects must be considered when comparing the current, effective method to any proposed replacement.

2.3 Adapted Method

A single-step method, at first glance, would be preferable to an iterated approach like the one described in Sec. 2.2. The simplicity would certainly ease computational costs, but the method needs to be comparably accurate. Luckily, there exists a single-step edge-finding method that is used for a different field of hypersonic research. Previously used for free flight experiments, this algorithm is optimized for large objects and displacements. For example, one experiment involved the free flight of a sphere in a hypersonic flow recorded with high-speed cameras. The algorithm was used to find the edge of the sphere in each frame of video and then determine its center. This positional information was then used to determine the sphere's displacement, velocity, and the forces acting on it. This method performs very well with these higher resolution bodies, whereas our current application requires precision at low resolution.

Starting again with a 12-bit black-and-white image, this method then uses the intensities of each pixel differently. Instead of only generating subimages, it generates a filtered version of the image with predicted edges. This is done with a Gaussian filter with a user-defined width. The standard deviation of the filter can be changed to detect edges with varying discretion and to remove some sharpness in the filtered image. Subimages can then be created from this predicted edge plot. This algorithm also includes subpixel detection with added correction for diagonal edges, allowing for a more accurate elliptical fitting of each marker after the edges have been tracked. Similar to the iterative method, this method requires user input to give initial guesses for the markers. The user input is identical, a triad of points in the corner to give a starting estimate for spacing and direction for the x- and y-dimensions of the marker layout. However, this is where the similarities end. To keep the computation limited, the algorithm then finds the nearest edge on the filtered image to the inputted point. After finding this starting point, it then searches in a user-defined direction (north, northeast, east, southeast, south, southwest, west, or northwest) for the next connected edge point. Then it will continue this step-by-step search for neighbors until it returns to the starting point. This requires no iteration, and quickly traces each marker's outline. In order to find the center of each markers, the algorithm then uses the edge point information to generate an ellipse that matches the edge using a least-squares fit. This is done with numerically stable method developed by Halíř and Flusser.[13] The position of the ellipse's center is then saved as the position of that marker.

At the surface, this adapted method seems quicker and much easier to implement. However, it must be tested in the conditions of the FSI tests, which we know are not ideal for this algorithm. The least-squares fit is certainly simple, but this can lessen the precision of the generated ellipses. Additionally, the filter that detects the edges in the first place is global across the image, which requires each marker's edge to satisfy the same sensitivity conditions. In Ch. 3, the efficacy of this method will be tested in two types of calibrations, followed by an applied use with FSI data in Ch. 4.

2.4 Comparison of Methods

Both methods clearly function adequately for detecting edges in grayscale images, and their respective methods for determining the centers of these edges are effective in their own applications. Before looking into the calibration test, let us examine these two methods qualitatively.

Comparing each algorithm chronologically as it performs its tasks, we will start by looking at how they handle edge detection. By using subimages, the iterative method only worries about the intensity of and around each marker, treating each one as independent. The adapted method, partially for speed, detects edges across the entire image. This means that the parameters defining its filter remain unchanged for each marker. Even when subimages are used for this method, the Gaussian filter would need to be altered between markers to have the same effect as the center-of-mass method from the iterative algorithm. This decreases the benefits of decreased time and computation, which is why the adapted method is being considered in the first place.

The primary computational cost for each method is its center-finding procedure. This procedure is also what makes each method accurate and useful for obtaining 3-dimensional displacements through photogrammetry. By iterating to minimize error with the Nelder-Mead method, the current algorithm can be much more accurate in matching the actual edges in the image. This optimization between accuracy and computational cost is exactly what we are concerned with when implementing the new method. We will attempt to weigh the advantages and disadvantages to make an overall judgment.

Chapter 3: Verification and Adaptation of the New Method

3.1 Overview

As previously stated, the single-step algorithm was originally written for large, high resolution free flight images. Luckily, the edge finding and tracking processes do not need to be changed to detect bodies of any size, allowing the change to small marker tracking to involve minimal code alterations. The largest change was adding the loop that made the algorithm automatically cycle through each marker in the image, but this was easy to borrow from the iterative code. As the code was tested, some runtime bugs were encountered such as detecting the wrong edges and looping across the image incorrectly. These were fixed as each set of images was processed.

There were two methods of verification for the adapted algorithm. The first, discussed further in Sec. 3.2, was entirely artificial, using generated images of pixelated ellipses of varying size, orientation, position, and noise level. Since these images were designed and generated, the exact parameters are known, so it is easy to compare the outputs of the algorithm with the actual marker properties. This clearly showed different sources of error that could affect the algorithm's performance. The second calibration setup, discussed further in Sec. 3.3, used real photographs of a plate with evenly spaced markers, similar to that used in the FSI tests. The plate was translated a specific amount between each frame, once again giving a known parameter to measure error for the algorithm.

Since each verification method was designed to mimic different aspects of the FSI test, any errors found will be able to characterize the algorithm's performance.

3.2 Artificial Ellipse Verification

3.2.1 Setup

The first step of evaluating the adapted algorithm was to generate a test set of ellipses. To accurately characterize the algorithm's performance a wide variety of properties were considered, detailed in Table 3.1. The size of the image, assumed to be square, was purposely kept small to reflect the low resolution of the small plate markings in FSI tests. The semi-major axis of each generated ellipse was constrained to be 1/3 of the image dimension, directly correllating full image size with ellipse size. To account for oblique viewing angles from experimental images, the shape and angle of each ellipse was also varied. For ease of generation and computing, the shape was modeled by varying the ratio between the minor and major axes. Since the major axis is constrained to the image size, this property constrains the minor axis. The angle was modeled as a simple angular rotation about the center of the ellipse. Since each marker could easily be displaced within a captured image, the generated images varied the center position of the ellipse. This was done by running through a range of x- and y-directional offsets defined as a percentage of the total image size. As with all experimental data, noise is expected in test images. This aspect was modeled at several percentages of total intensity and used a uniform random number function to affect the base value of each pixel within the noise percentage. Lastly, when generating each ellipse, the minimum and maximum pixel intensity were limited to be between 10% and 90% intensity to mimic the abilities of a real camera photographing painted markers, especially in non-ideal conditions with poor lighting and marker intensity.

Table 3.1: Overview of Ellipse Parameters.

Parameter	Range of values
Size	8, 10, 12, 14, 16 pixels
Axis ratio	0.75, 0.80, 0.85, 0.90, 0.95, 1.00
Angle of rotation from horizontal	0, $\pi/8$, $\pi/4$, $3\pi/8$, $\pi/2$, $5\pi/8$, $3\pi/4$, $7\pi/8$ radians
Center offset in x-direction	-10%, -5%, 0%, 5%, 10% of image size
Center offset in y-direction	-10%, -5%, 0%, 5%, 10% of image size
Image noise	0.5%, 1.0%, 2.0%, 5.0%

By generating ellipses through each of these parameters, a full factorial experiment of 24,000 images was produced. Example ellipses can be seen in Figs. 3.1, 3.2, and 3.3. The algorithm then estimated edge points and center coordinates for each of the ellipses, the latter of which being important for FSI applications. This method is described in Ch. 2, Sec. 2.3. Each generated ellipse acted similarly to a subimage of a marker from an FSI test. This made determining a starting guess much easier as there was only one edge that needed to be detected. For quick computation, any user input was eliminated by choosing a starting point directly from the detected edges, knowing there would not be any extra markers to interfere. Since the images were artificially generated, the center coordinates are known and can be compared with the values found by the algorithm. The distance between the actual and calculated coordinates was recorded as the error for each image. This was calculated in raw pixel values and as a percentage of the average semi-axis length of each ellipse.

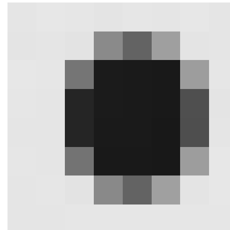

Figure 3.1: Generated ellipse of size 8 pixels, axis ratio of 0.9, angle of π/2, x-offset of 5%, and noise of 0.5%.

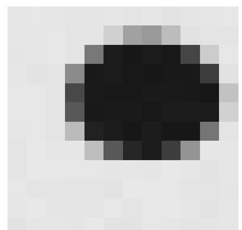

Figure 3.2: Generated ellipse of size 12 pixels, axis ratio of 0.8, angle of 0, x-offset of 10%, y-offset of -10%, and noise of 1.0%.

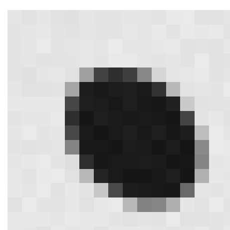

Figure 3.3: Generated ellipse of size 16 pixels, axis ratio of 0.75, angle of 3π/4, x-offset of 5%, y-offset of 5%, and noise of 2.0%.

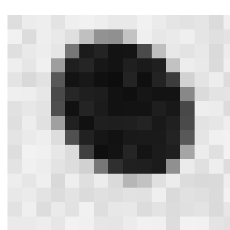

Figure 3.4: Generated ellipse of size 16 pixels, axis ratio of 0.8, angle of 3π/4, x-offset of 0%, y-offset of -10%, and noise of 5.0%.

3.2.2 Results

To more easily see what the results were for each image, the errors were plotted in several different ways. The data were separated into a few groups separated by size and noise level as well as a full grouping. To show the distribution of error, all groups were plotted on histograms in Figs. 3.5, 3.6, and 3.7.

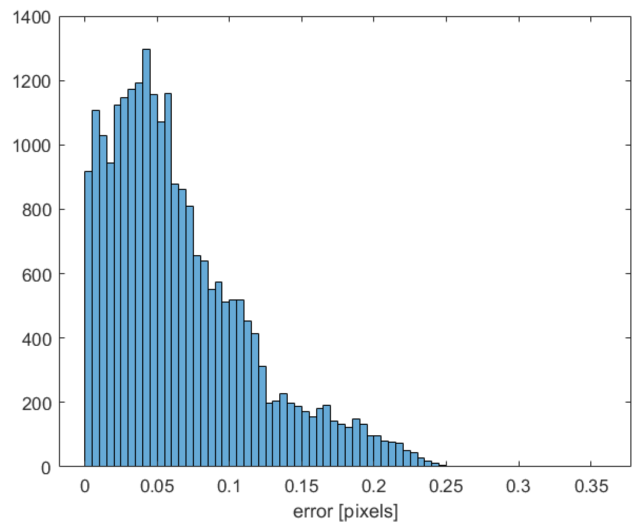

(a) Raw error.

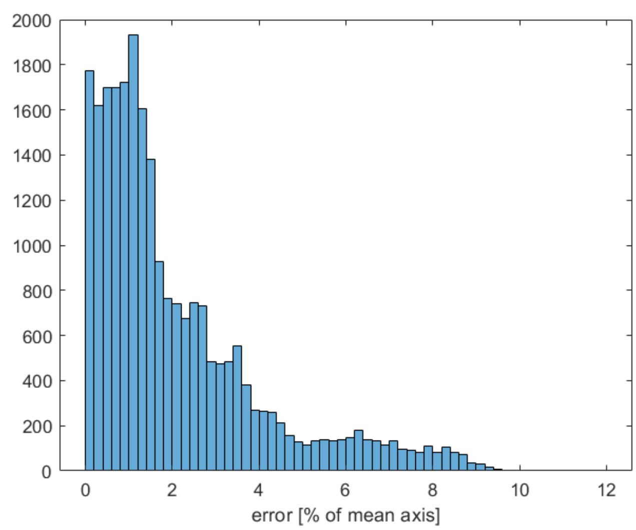

(b) Percent error.

Figure 3.5: Errors of full data set plotted on histograms.

Examining the error of the full data set, the first thing we can do is find a maximum error. We see in Fig. 3.5a that the majority of the generated images produced an error less than 0.25 pixels, with only a very small number reaching an error of 0.35 pixels. This is an encouraging result that shows us that even at low resolution we should not expect more than 1/4 pixel of deviation from the true center of a marker. Even considering this, we must also look over the percent error. The maximum value on Fig. 3.5b is 10%, much less encouraging than the raw error implied. This shows that for some conditions the error relative to the size of the marker is much larger than ideal for precise measurements. That being said, the majority of the data is grouped far below this 10% mark. The largest histogram bars lie down near the 1% error mark, with most of the data falling below the 5% mark. In order to determine which cases produce the higher errors, we will group the data by image size and noise percentage.

Since the percent error is calculated based on the mean axis length, the image size can directly affect this number. We must also see if the image size affects the raw error. In the histograms in Fig. 3.6 we can see how the error distribution shifts smaller as the image size grows larger. While this may be intuitive, as resolution improves so does performance, it is still significant to note how much the error decreases. Once the images reach a size of 12 pixels, the maximum error is only slightly above 0.1 pixels. Note that, based on how the ellipses are generated, the major axis of the ellipse for this case is only 8 pixels. As we increase to an ellipse with a major axis 10-2/3 pixels across (image size of 16 pixels), the raw error is decidedly below 0.1 pixels. We further see in Fig. 3.6b that the only percent errors that near 10% are for an image size of 8, which corresponds to a major axis of only 5-1/3 pixels. For the largest tested images the error decreases to below 2%, which shows good performance by the algorithm.

Before testing the adapted algorithm, image noise was a significant concern. Based on

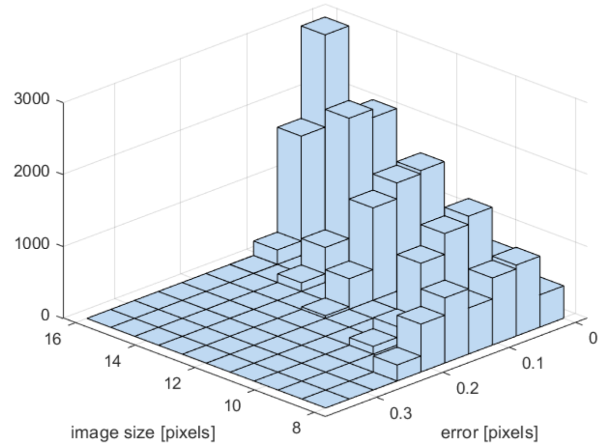

(a) Raw error.

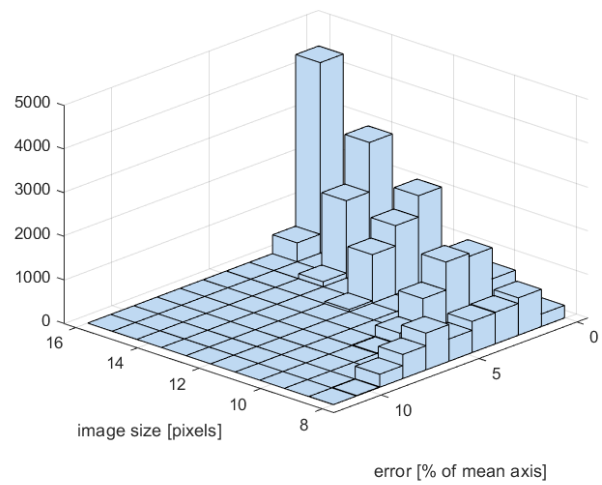

(b) Percent error.

Figure 3.6: Errors of data set grouped by size plotted on histograms.

previously performed experiments, we expect that the maximum noise in the images would be 2%, but to better characterize the performance we also tested smaller percentages of 1% and 0.5% as well as a larger percentage of 5%. The unexpected outcome of this, as seen in Fig. 3.7, is that the noise level, if sufficiently low, does not seem to affect the performance of the algorithm. The histograms for each noise level is nearly identical for both raw and percent errors, each one reflecting the shape of the full data set's histograms. The majority of the data has relatively low

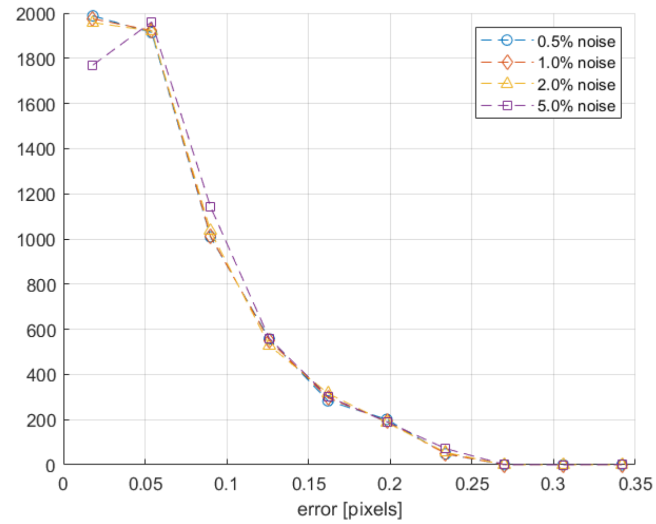

(a) Raw error.

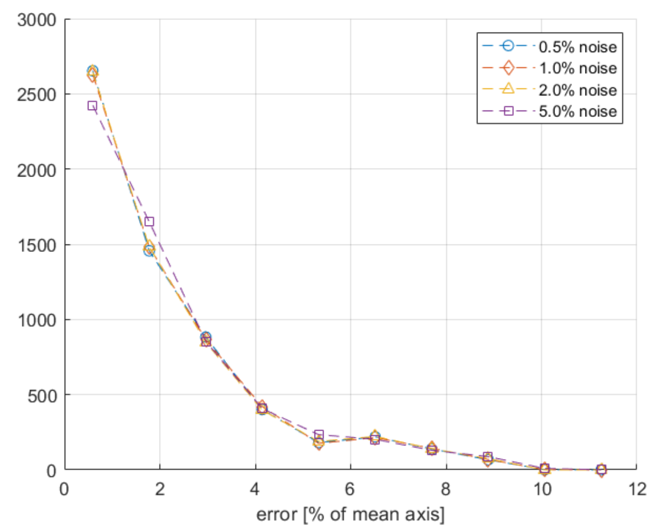

(b) Percent error.

Figure 3.7: Errors of data set grouped by noise level plotted on histograms.

error with the percent error being mostly below 5%. We see in Fig. 3.3 how even with a low noise of 2%, it is still noticeable in the image as it darkens the white background and lightens the black marker. This is even more apparent in Fig. 3.4, which has 5% noise. The consistency of the algorithm even with varying noise in the image is certainly encouraging when evaluating its performance.

3.3 Real Photograph Verification

3.3.1 Setup

To mimic the FSI test conditions, a calibration plate was painted with the same stencil used for the experiments. On a white background, a 15 by 9 grid of black markers spread by 6.35mm was painted. This small plate of markers was then photographed as it was translated. The photographs were taken with a Phantom high-speed camera with an Irix 150 f/2.5 macro lens. This gave appropriate focus while still being able to see the plate as it was moved between frames. For the photographs, the camera was setup at a resolution of 1024x976 with a 5000 microsecond exposure. The full setup is shown in Fig. 3.8, which includes the marker plate mounted to a linear stage across from the camera. The linear stage had a vernier micrometer readout with 0.01mm gradations. This introduced some uncertainty in the measurement of each translation, but this was a known issue for a benchtop test. Mounted on the stage was a small metal plate with stenciled markers matching the size and spacing of those used in the FSI experiments. Instead of altering the ellipse shapes between images, each frame was at a known displacement from the previous frame. This was in order to have each calibration test characterize different aspects of the algorithm, in this case its accuracy in tracking motion. Three different camera orientations were used: head-on, oblique angle towards the right of approximately 9.5°, and oblique angle towards the left of approximately 13.5°. The plate was translated at 10 micron increments across a range of 200 microns for the head-on arrangement and 300 microns for each oblique arrangement. This gave a better idea as to how the algorithm would perform in an FSI test where camera angle will likely not be ideal. Since the effects of shape variation were already observed with the artificial images they can be accurately considered in the benchtop test.

Figure 3.8: Photographic calibration setup with camera, linear stage, and marker plate.

Examples of the grayscale images can be seen in Fig. 3.9. Once these photos were taken, they were analyzed the same way the FSI images would be. User input was taken to give the location and spacing of the markers, then the algorithm ran to determine the actual centers and spacing of the markers. These center positions were compared between frames to see how much displacement there was from frame to frame. Using the known marker spacing of 6.35mm along with a measured spacing from the images, the calculated pixel displacement was able to be converted to physical displacement in microns.

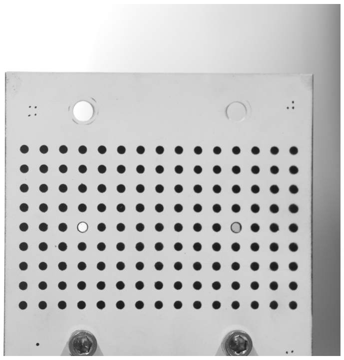

(a) 9.5° Oblique angle towards right.

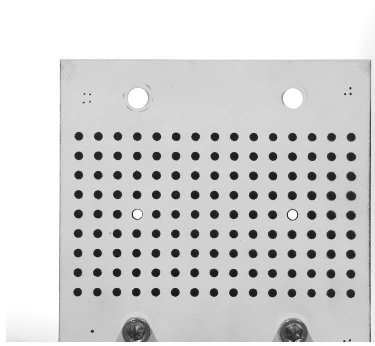

(b) Head-on.

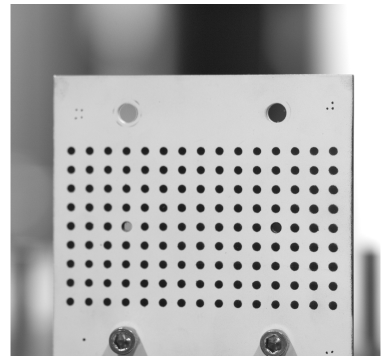

(c) 13.5° oblique angle towards left.

Figure 3.9: Sample photographs from the verification test.

3.3.2 Results

The adapted algorithm outputs a displacement for each marker in each image, meaning there are 135 displacements calculated between each image. In order to decrease the raw data acquired for these tests, the 135 markers were condensed into two data points for each image: the mean and median of all displacements in a frame. The mean is the average displacement across the marker plate, which should be near the 10 micron translation, while the median can help mitigate effects of outliers that may come from errors in the algorithm's output. In addition to plotting these means and medians, measurement uncertainty was plotted to compare with the data. The uncertainty was very roughly assumed to be 1/3 of the smallest measurable distance on the micrometer, which would be 3-1/3 microns. While this is a fairly significant amount compared to the translation of 10 microns, it is a decent estimation with which to compare our outputs.

Starting with the simplest camera orientation, head-on, we see fairly acceptable results. As shown in Fig. 3.10, the displacements stay between the uncertainty bounds. Additionally, there is minimal difference between the mean and median plots, which can imply that there are not any

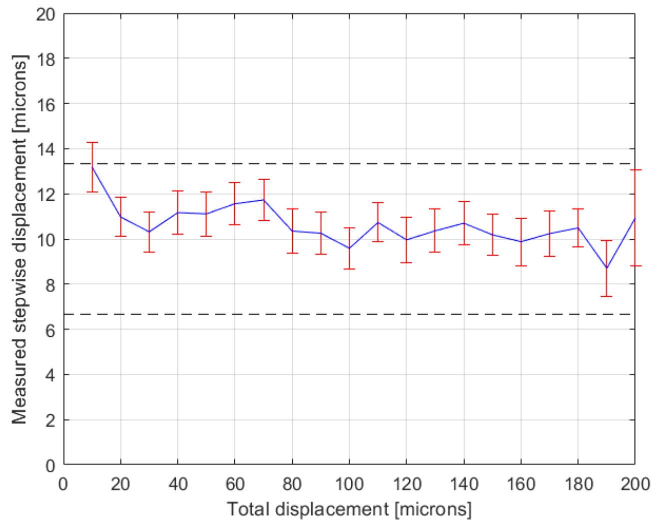

(a) Mean of marker displacement with error bars representing standard deviation.

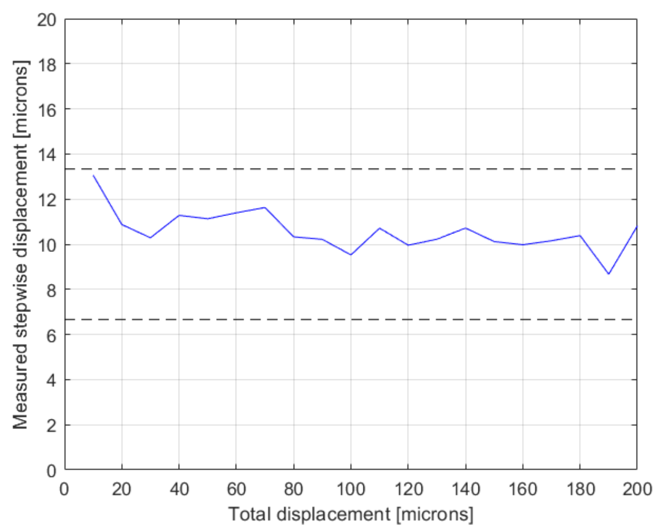

(b) Median of marker displacement.

Figure 3.10: Head-on orientation: measured displacement between images. Dotted lines represent limits of uncertainty.

major errors in the algorithm's output. This is confirmed by the small standard deviations seen in the error bars, most of which being less than 1 pixel. We can also see that most of the measured displacements are within 2 microns of the ideal displacement of 10 microns. While this seems small, that error can be significant considering the minuscule vibrations in FSI tests. However, since all but one of the means are still within the assumed bounds of uncertainty, we can only say that the algorithm may perform sufficiently well with a head-on camera angle.

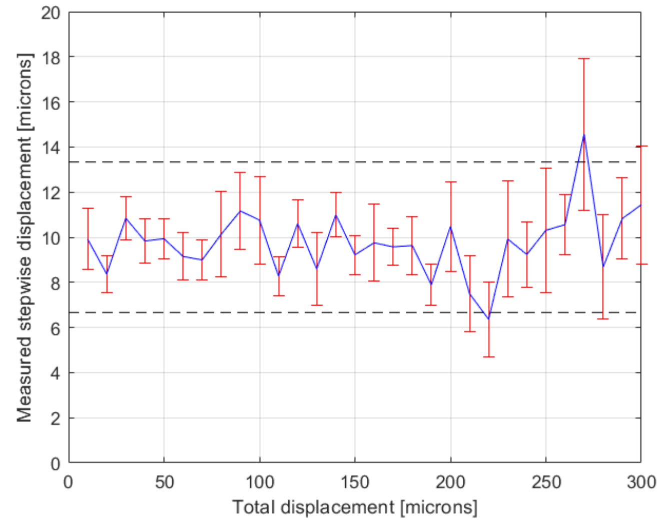

(a) Mean of marker displacement with error bars representing standard deviation.

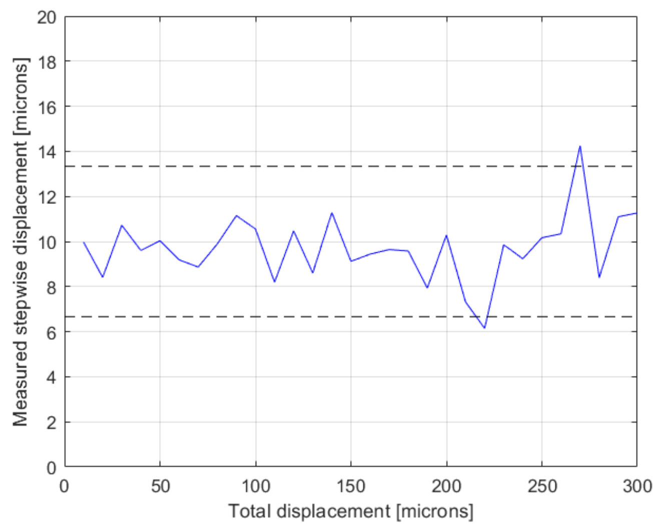

(b) Median of marker displacement.

Figure 3.11: 9.5° rightward orientation: measured displacement between images. Dotted lines represent limits of uncertainty

As we see in Fig. 3.9a, the camera angle is not severely different from the head-on case, but it is still slightly off center. This is reflected in the results plotted in Fig. 3.11. Once again, there is little difference between the means and medians, but the displacement range is wider. Immediately, we can notice that the displacements no longer stay within the assumed uncertainty bounds. Additionally, the standard deviations exceed 1 pixel, in some cases nearing 3 to 4 pixels. This shows that even with a slight angle, the algorithm may be losing accuracy. Most of the data is still within 2 microns of the 10 micron mark, but the plot is definitely more erratic than for the head-on case in Fig. 3.10. Once again, this does not conclude much about the performance, but it does imply that any off conditions may affect the outputs of the algorithm.

The last, most intense camera angle can be seen in Fig. 3.9c where the edge of the plate is clearly visible. Based on the performance of the algorithm on the past two cases, we may expect to see a further drop in performance for this case. This is unfortunately confirmed in Fig. 3.12, where we see an even more erratic plot with a wider spread of displacements. The standard deviations are similar to those in the rightward case, being around 2 pixels, but there are two means that have extreme outliers near total displacement of 120 μm evidenced by their large standard deviations of over 40 pixels. There is also significant difference between the mean and median plots as the median seems to catch a few outliers that that cause the mean plot to be far off the actual displacement. This is most apparent for the same points that have the largest standard deviation seen in Fig. 3.12a. Even with the correction for outliers, there is a great deal of data near or past the uncertainty boundaries, and the range about the 10 micron mark is much larger than for the rightward and head-on cases. For both the mean and median plots we see large spike exceeding 30 microns near total displacement of 220 μm. This is clearly an error in the algorithm since it is so far outside of the assumed uncertainty. While this leftward case is not immensely oblique, only being 13.5°, it seems to be off-center enough to affect the results significantly. There is clear evidence of poor performance for this case.

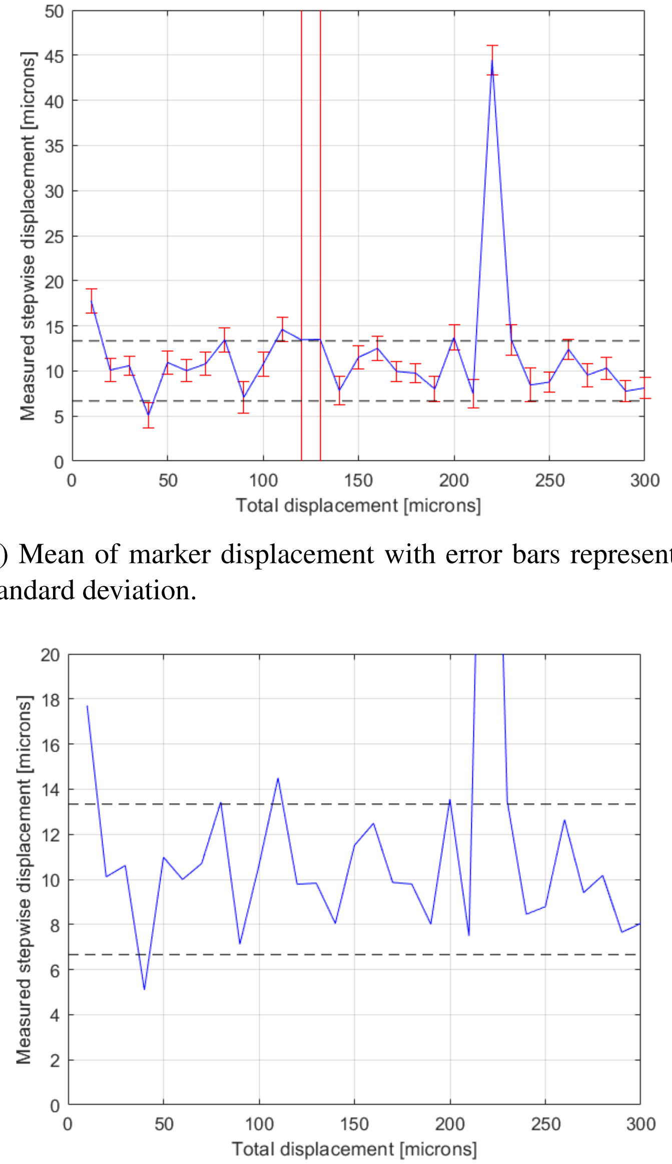

(a) Mean of marker displacement with error bars representing standard deviation.

(b) Median of marker displacement.

Figure 3.12: 13.5° leftward orientation: measured displacement between images. Dotted lines represent limits of uncertainty.

3.4 Conclusions

Combining the results from both verification tests, we can endeavor to understand how the adapted algorithm will perform in the real FSI experiments. Based on the artificial ellipse testing we were able to clearly identify any issues with edge detection for a variety of marker shapes. The results showed that the biggest cause of error was the size and resolution of the marker, and for an ellipse exceeding the size of 8 pixels the error is quite limited. As we shifted to photographs of a marker grid for the second verification test, extra sources of error seemed to arise. At oblique angles of 9.5° and 13.5°, errors far greater than those encountered in the artificial test were recorded. As seen in Fig. 3.9, there are some deviations from the elliptical shape of each marker. While the generated images accounted for noise and viewing angle, they all assumed the basic shape of an ellipse. This means that even though the edges were detected for the real images, the fitting that is used to determine the marker's center may be affected by the deviated shape of the markers. We also saw in the artificial test that noise did not significantly affect the algorithm's accuracy, so there can only be a limited number of causes. When looking at the images in Fig. 3.9 there is definitely some blur in the oblique cases that may be caused by certain parts of the marker plate being out of focus. This problem, combined with the shape variations of the markers are likely to be sources of error for these cases in the photographic verification test. For a standard FSI test, we can expect a variety of deflections across the plate, so the displacements for each marker may be significantly different. This means that we require confidence in the measurements from every marker. If we can truly attribute some errors to the blurriness of the photograph then we can say that the issues may have come from the setup of the test instead of the algorithm itself. This idea is somewhat corroborated by the performance of the algorithm with the generated images, but cannot be confirmed. Since the setup was nearly identical for all three cases in the photographic test, the point can also be made that the only error must be from the processing algorithm. The best and most thorough conclusion that can be drawn is that the algorithm does not work well for off-condition cases and that errors should be considered for non-ideal camera angles.

Chapter 4: Application of Adapted Algorithm for FSI Experiments

4.1 Description of Experiments

The data from the FSI experiments was acquired at the NASA Langley facility using their 20-inch Mach 6 Hypersonic Wind Tunnel and 31-inch Mach 10 Hypersonic Wind Tunnel[14]. The tunnels are blow-down arrangements using dried, heated air. Nominal reservoir conditions for the Mach 6 test were a stagnation pressure of 480 psia and a stagnation temperature of 530 °F. Nominal reservoir conditions for the Mach 10 test were a stagnation pressure of 1350 psia and a stagnation temperature of 1280 °F. For the Mach 6 test, two Vision Research Phantom cameras, one v2512 and one v2640, were used. Each camera was equipped with Nikon Nikkor 60-mm lenses; the v2512 imaged at a resolution of 768x768 pixels whereas the v2640 imaged at 768x864 pixels. For the Mach 10 test, two Photron SA-Z cameras was used for stereo DIC photography with one mounted on the left of the test section and the other on the right. The FOV for each camera was 768 x 888 using an Edmund Optics #89-772 filter and a Nikon 105mm f/2 lens.

The Mach 10 test specimen was a steel compliant panel painted with a 25 x 25 grid of markers spaced 6.35 mm apart, while the Mach 6 panel had two 25 x 10 grids spaced 5.35 mm apart. Both panel fixtures also included a plenum that could provide back-pressure during each test. The panel was mounted in the tunnel at different angles for each test which resulted in different deflections in the metal panel. Upstream of the panel was a mounting point for a trip mechanism designed to trigger transition to turbulence in the flow. Two runs were analyzed to test the new algorithm, one of which at Mach 6 and one at Mach 10. Neither run for this analysis used the trip mechanism. A 3D CAD model of the experimental setup is shown in Fig. 4.1. The run conditions are shown in Table 4.1.

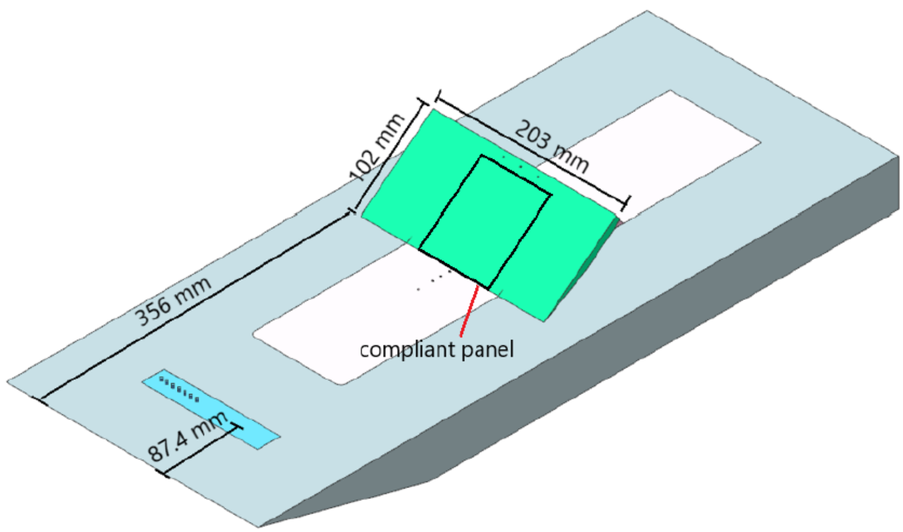

Figure 4.1: 3D CAD model of the experimental setup.

Table 4.1: Test Conditions.

Conditions	Run A	Run B
Mach number	6	10
Ramp angle	34°	10°
Stagnation pressure	476 psia	1354 psia
Stagnation temperature	475 °F	1294 °F
Sample rate	31 kHz	30 kHz

Images of the panel from each DIC camera can be seen in Figs. 4.3 and 4.2. Note that the Run A angle is visibly steeper, as we expect from Table 4.1. Close visual inspection of the markers themselves reveals evidence of deflection as the dots no longer fit in straight rows and columns. The two different tests were chosen specifically to give a wide range of image conditions for the algorithm to process. While the standard analysis of these FSI tests is for the determination of vibrational modes and boundary layer interactions, these images were chosen to evaluate the adapted algorithm. Additionally, the images from Run A were already analyzed by Schöneich et al.[1], so we can compare the results to determine the accuracy of the adapted algorithm.

(a) Left camera (designated 19148).

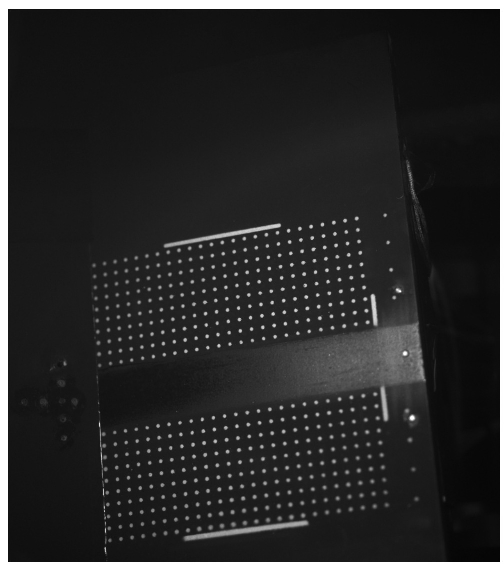

(b) Right camera (designated 26184).

Figure 4.2: Run A camera views.

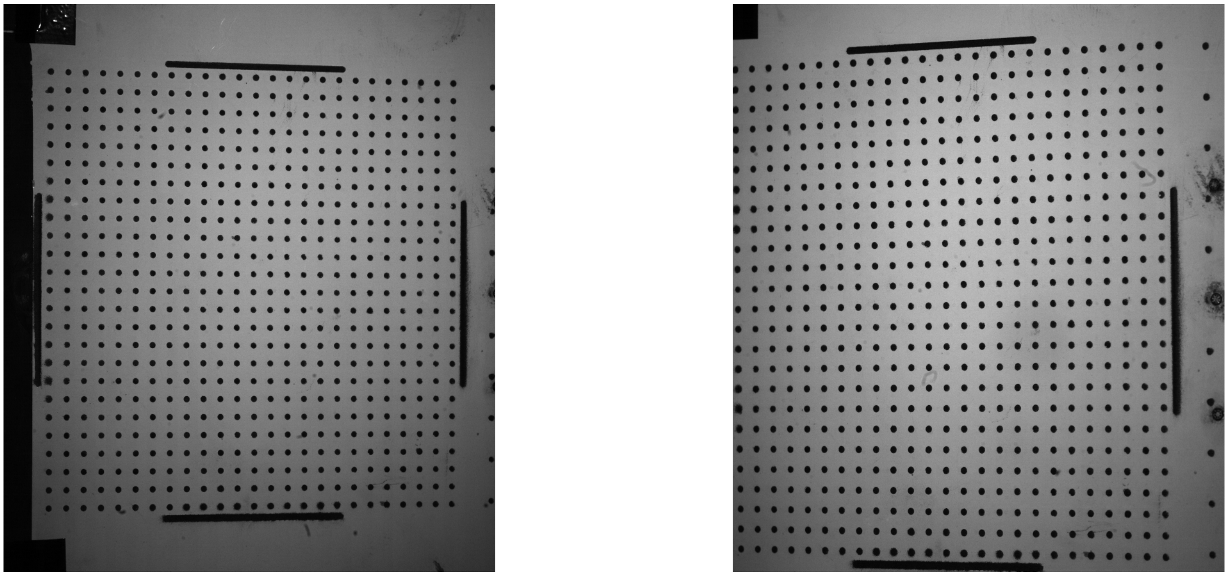

(a) Left camera.

(b) Right camera.

Figure 4.3: Run B camera views.

4.2 Data Analysis

4.2.1 Overview

The analysis process for the FSI tests must convert photographs into vibrational frequency data. This takes several steps, starting with tracking the x- and y-coordinates of each marker in each frame which is done with the adapted algorithm as described in Ch. 3. In order to obtain any deflection information, the motion in and out of the plane must also be determined. This is the purpose of the stereo DIC camera setup. First, a LaVision white QR-code calibration plate with 8.5 mm spacing is photographed at the same camera angles as for the FSI tests but at various positions on the plate. These images are then fed into a calibration algorithm that uses the coordinate data to solve where the cameras are in 3D space relative to the plate. This calibration data can then be used with the coordinates from the test data to determine the 3D location of each marker in each frame. Once these 3D points are found, the frequency analysis can begin. This is done by feeding the 3D points into an SPOD algorithm provided by Towne et al.[15] that will show the power spectra of each mode. The time series of data are separated into overlapping bins chosen to give a reasonable number of modes as well as sufficient frequency resolution. A few individual points were also analyzed using Welch's method of PSD averaging in order to determine the strong frequencies. This is close to a one-dimensional analog to SPOD analysis. For Run A 6000 frames were analyzed, and for Run B 2000 frames were analyzed.

4.2.2 SPOD Analysis

Since the time series length and sample rate for each run was different, different SPOD parameters were used, seen in Table 4.2. The spectral resolution for under these conditions was 1201 discrete frequencies for Run A and 401 discrete frequencies for Run B. The number of modes displayed in the SPOD analysis matches the number of bins. The examination of the power spectra will be mostly qualitative.

Table 4.2: SPOD Parameters.

Parameter	Run A	Run B
Spectrum type	One-sided (real-valued)	One-sided (real-valued)
Bin size	2400 frames	800 frames
Overlap	1800 frames	600 frames
Number of bins	7	7
Window type	Hamming	Hamming
Weighting	Uniform	Uniform

First looking at Run A, we know that this test is at a high ramp angle and as such we expect more panel deflection. Based on the results of Ch. 3, this may also mean that the performance of the adapted algorithm is decreased. The spectra are plotted on a semi-logarithmic scale and the frequency scale is limited to the Nyquist frequency of 15.5 kHz. We first examine the 7 SPOD modes in Fig. 4.4. The highest energy mode, plotted on its own in Fig. 4.5, seems to have minor peaks, but there is nothing clearly seen in any of the spectra. We also see that the highest intensity of any of these peaks is less than 10^0 μm^2/Hz, which can be compared to the peaks near 10^2 μm^2/Hz found by Schöneich et al. shown in Fig. 4.6. In addition to the low intensity of the modal energies, we were unable to use the SPOD analysis to model any of the expected mode shapes seen by Schöneich et al. While the SPOD spectra do not show any clear modes on their own, we can look near the known modal frequencies for any energy peaks. The highest peaks are near 1 kHz and 2 kHz, corresponding to the spanwise-symmetric (1,2) mode at 1.07 kHz and (1,3) mode at 1.83 kHz. Other modes seen in Fig. 4.6 include the (1,1) mode at 0.759 kHz and the spanwise-antisymmetric (2,2) mode at 1.36 kHz. Additional peaks in Fig. 4.5 appear between 2.5 and 3 kHz and near 4 kHz. These do not seem to coincide with any modes found by Schöneich et al. We will look at single-point analyses to determine whether other high-energy modes are noticeable using the adapted algorithm.

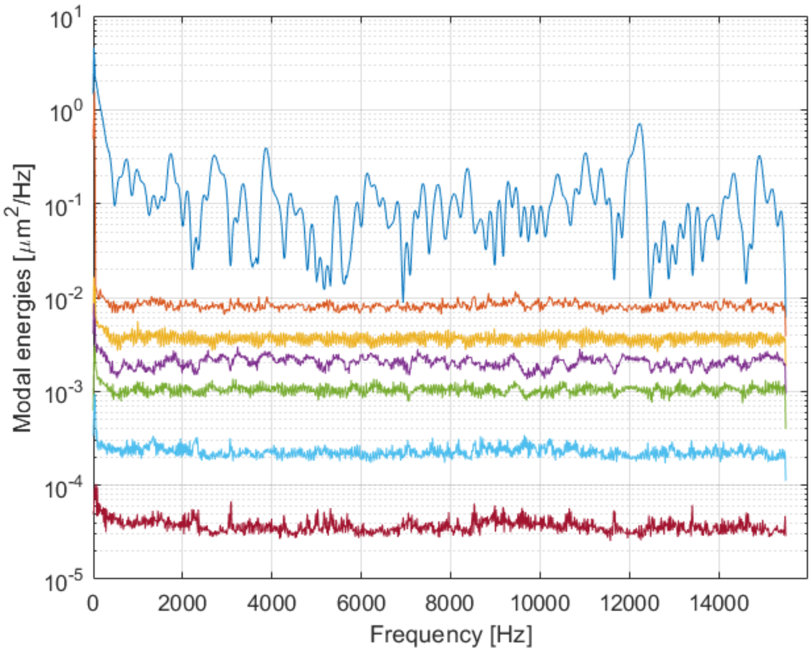

Figure 4.4: Power spectra for each SPOD mode in Run A.

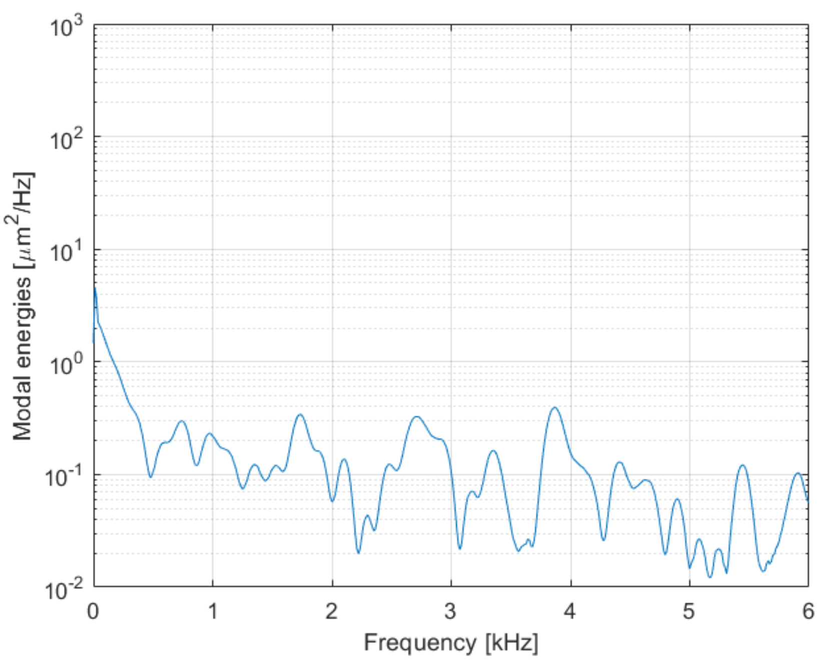

Figure 4.5: Highest energy SPOD mode in Run A.

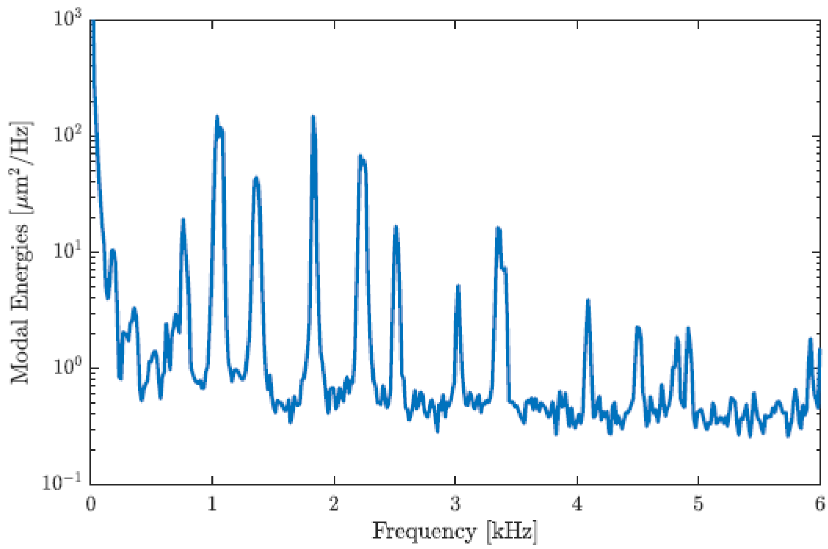

Figure 4.6: SPOD spectra from Schöneich et al.[1]

For Run B, the ramp angle is decreased and the deflection in the panel follows suit. The angle of the camera is expected to affect the performance of the algorithm based on the results in Ch. 3. The SPOD plot is in the same form as for Run A, and we see in Fig. 4.7 a significant difference from Run A. First we notice lower energies for the SPOD modes, all of which being below 10^{-1} μm^2/Hz. Even when we look at Fig. 4.8, which only displays the first SPOD mode, we do not see any obvious energy peaks. There seems to be a slight increase near 2000 Hz, which may correspond to a physical vibration such as the (1,3) mode at 1860 Hz found in Run A. Unfortunately, there is a lack of clarity for this SPOD analysis which may be based on the errors shown in Ch. 3 Sec. 3.3. Similar to the procedure for Run A, we will perform single-point analysis to try and see any energy peaks.

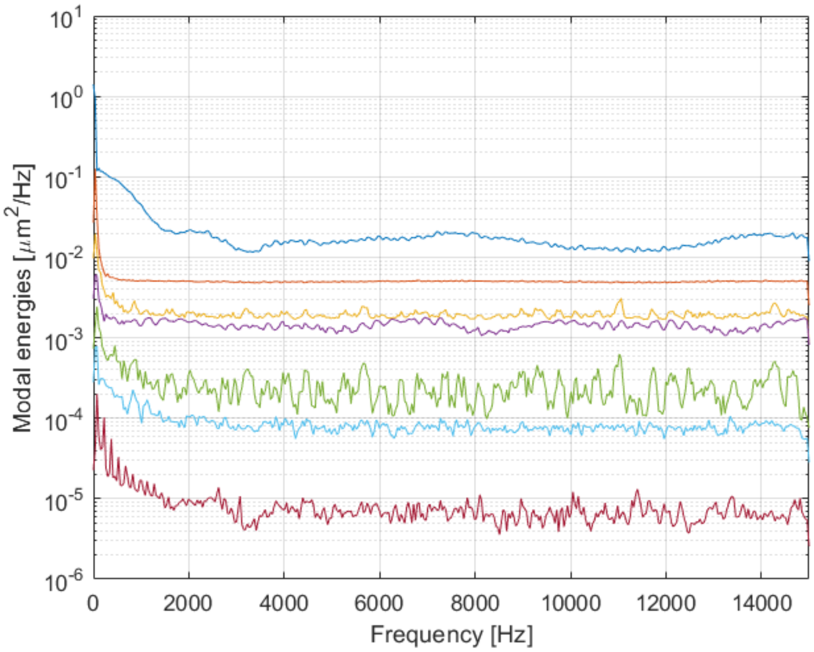

Figure 4.7: Power spectra for each mode in Run B.

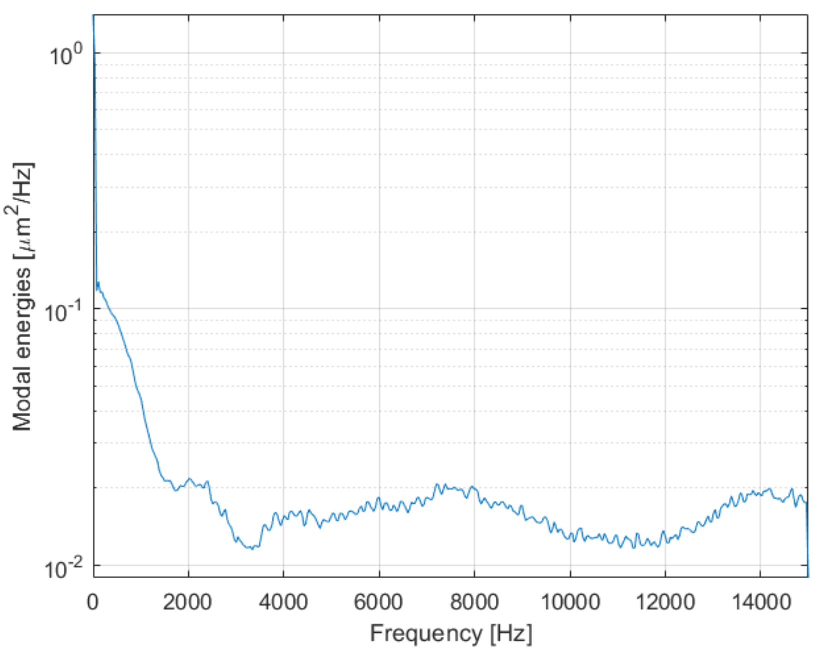

Figure 4.8: Highest energy SPOD mode in Run B.

4.2.3 Single-Point Analysis

Several points for individual analysis were chosen from each run. They were specifically picked to be in different areas of the panel, all of which being near antinodes of the panel for high-energy vibrational modes. Each point was meant to be sufficiently far from the edge of the panel but also far enough apart from the other points. For Run A, the choice of markers was limited by the blank band in the center of the panel, but there were still points present near the antinodes. The markers chosen for Run A were in the 10th and 13th streamwise columns and the 6th and 18th rows, shown in Fig. 4.9. These markers were chosen to be near the antinodes of the (1,2), (2,2), and (1,3) modes, which are at points 1/6 and 1/4 the length of the panel. The positional data for each point have some error spikes, but by smoothing with a moving average, these errors can be mitigated.

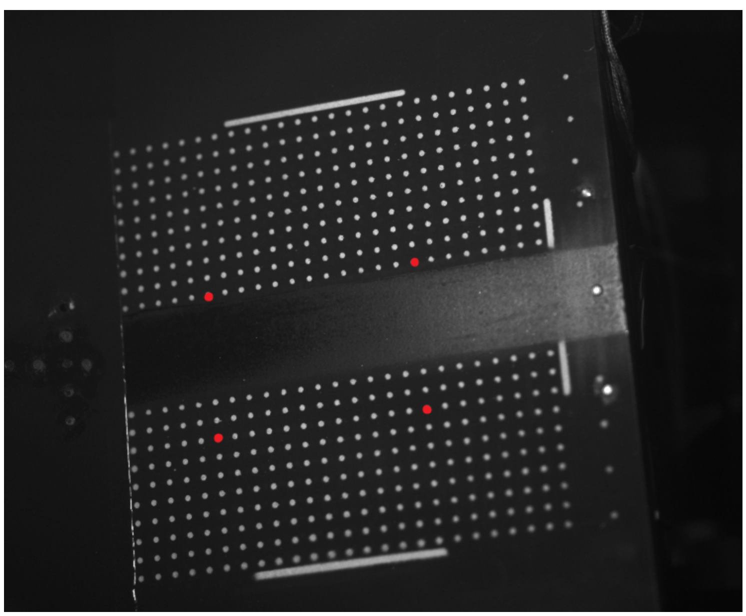

Figure 4.9: Run A single points chosen for analysis highlighted in red.

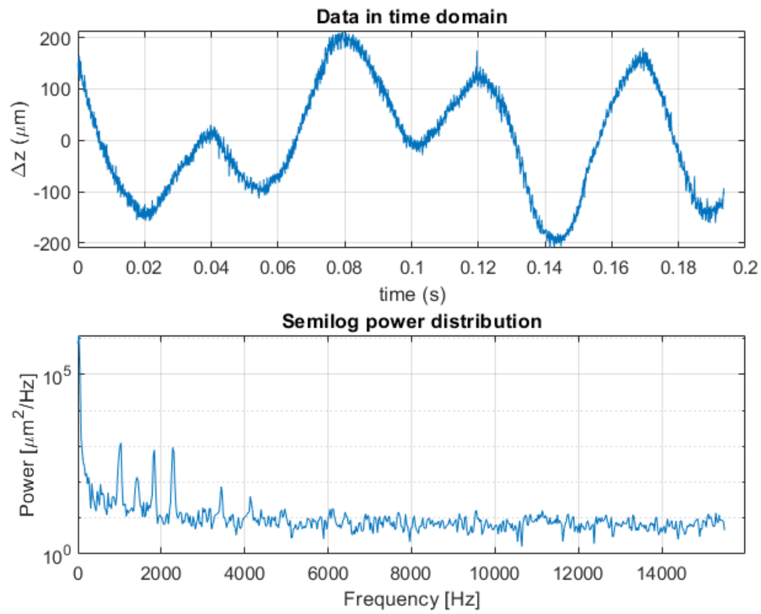

(a) 10th column, 6th row.

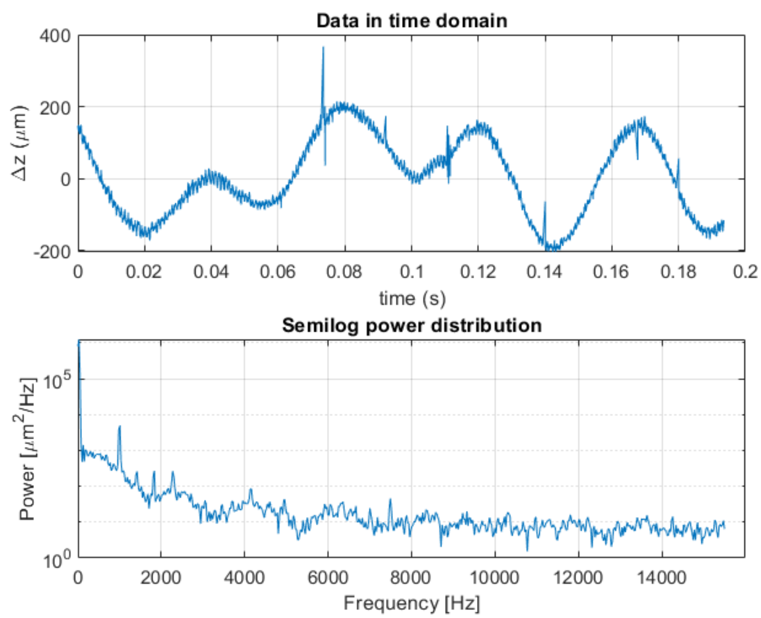

(b) 10th column, 18th row.

Figure 4.10: Run A single-point Welch's method plots.

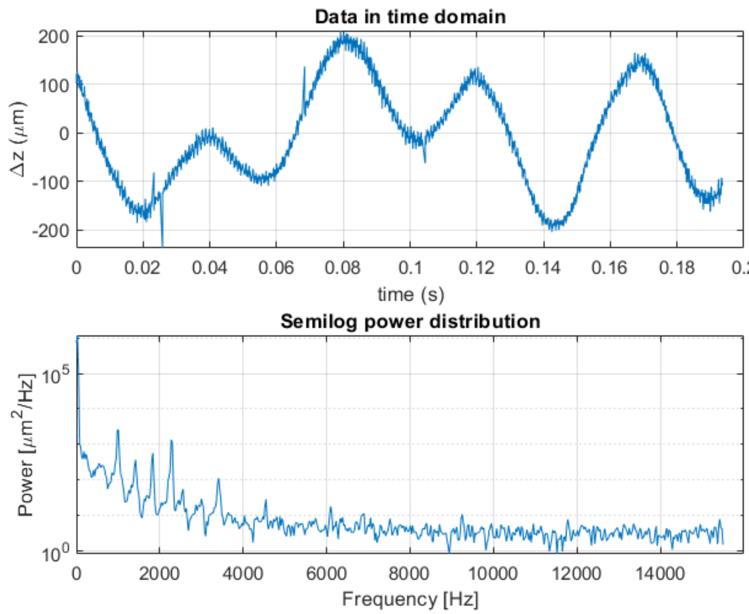

(c) 13th column, 6th row.

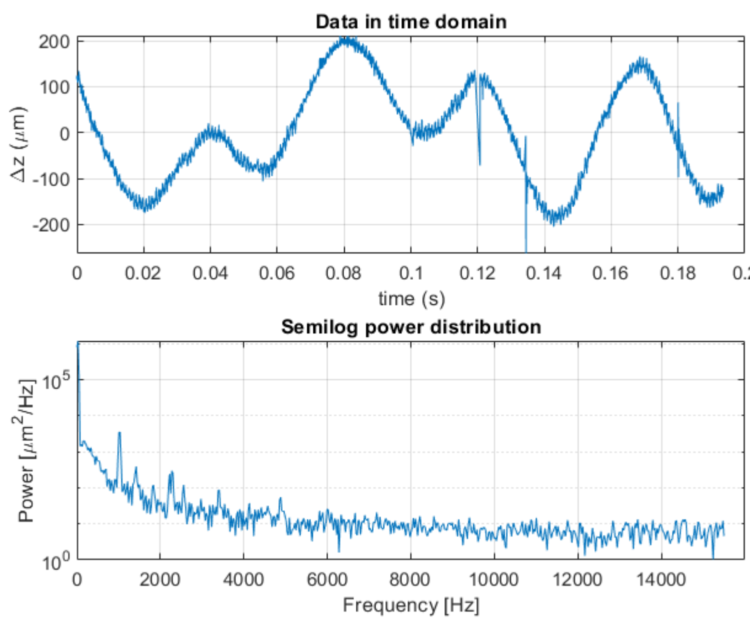

(d) 13th column, 18th row.

Figure 4.10: Run A single-point Welch's method plots. (cont.)

The plots in Fig. 4.10 are all very similar, save for what appears to be some noise in the data. The frequency peaks in the power spectra are also very similar, with the largest peaks reaching an intensity near 10^3 μm^2/Hz. These peaks are near 1000 Hz and 2000 Hz, with an additional smaller peak at around 1500 Hz seen in Figs. 4.10a & 4.10c. Comparing to the SPOD spectra in Fig. 4.6, we see similar peaks near the same frequencies. These correspond to the (1,2) and (1,3) modes, which is expected based on the locations of the chosen points. Even though the SPOD anaylsis in Figs. 4.4 & 4.5 did not show any particularly clear vibrational mode, the use of Welch's method shows that the adapted algorithm can achieve similar results to the older method. To further analyze the adapted algorithm, the raw positional data from each camera was compared between each method. The difference between each x-y point was calculated as a distance in the x-y plane and plotted for each camera and each point in Fig. 4.11. This gives us a raw difference between the two methods. These plots show that the difference between each method is less than 1.5 pixels for all the points, barring some large error spikes. However, since the data were smoothed using a moving average, these errors are not as significant in the analysis.

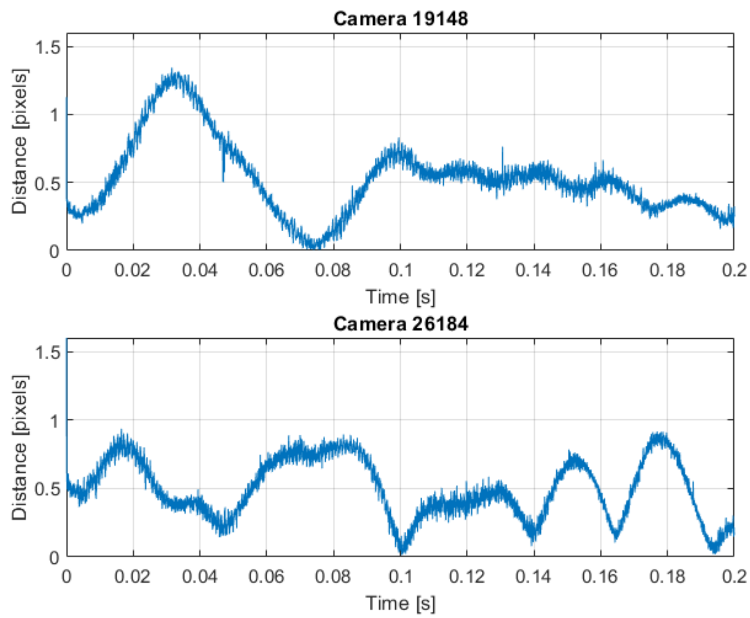

(a) 10th column, 6th row.

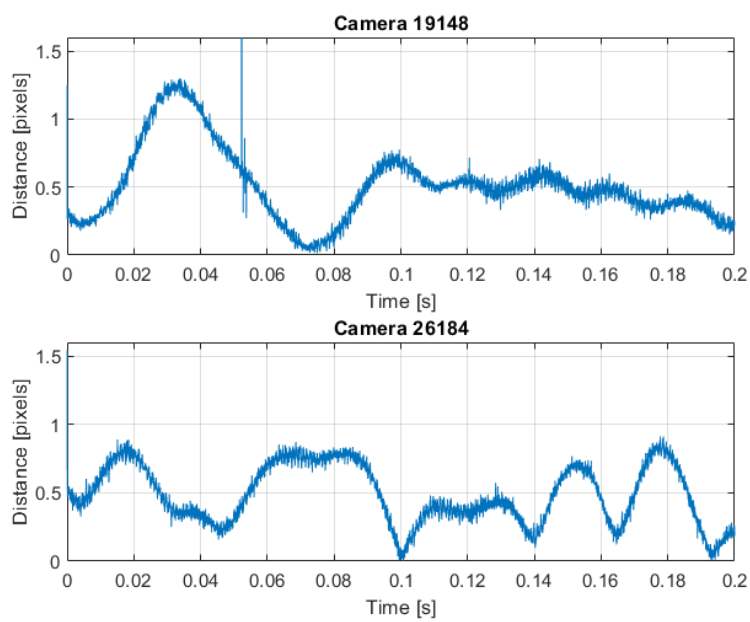

(b) 13th column, 18th row.

Figure 4.11: Run A single-point position differences.

When analyzing Run B, we also want to focus on antinodes of the physical modes. We once again expect to see the antinodes for the (1,2) and (1,3) modes near the 1/6 and 1/4 length marks of the plate, so the points chosen were in the 9th and 19th streamwise columns and the 6th and 18th rows, shown in Fig. 4.12. Compared to the points chosen for Run A, these are near the same locations on the plate.

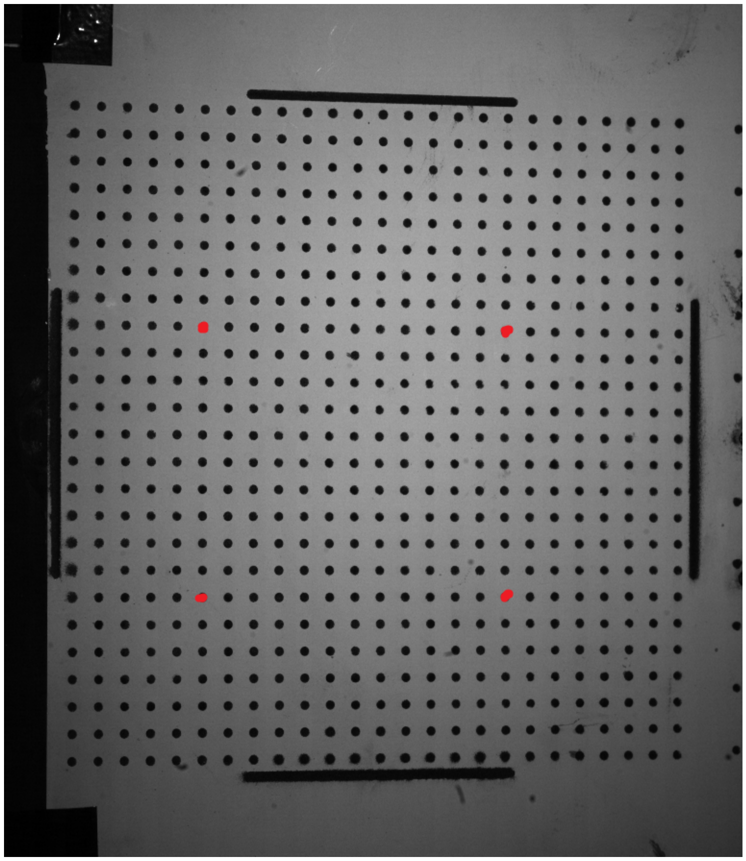

Figure 4.12: Run B single points chosen for analysis highlighted in red.

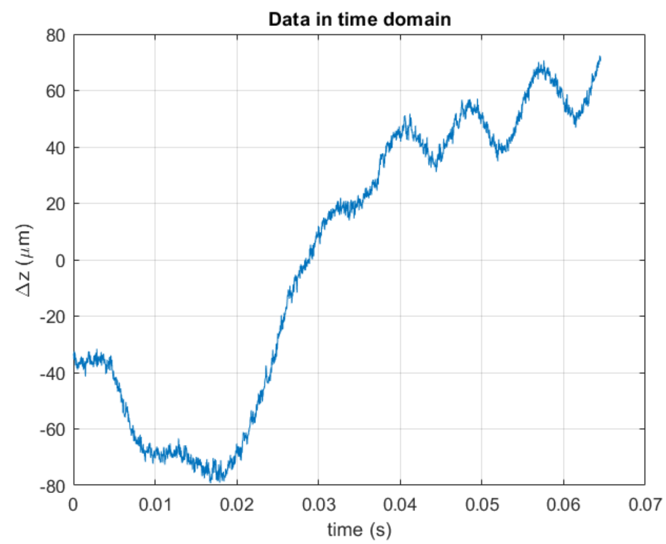

Figure 4.13: Run B single point time series.

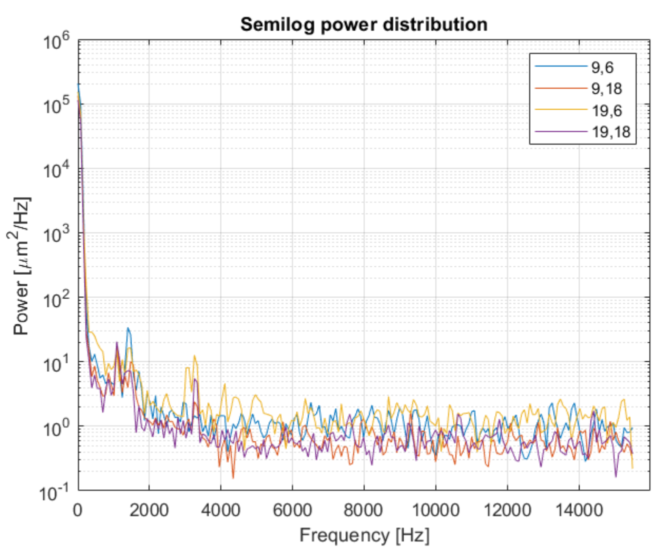

Figure 4.14: Run B single point Welch's method spectra, all four points.

The plots in Fig. 4.14 are also very similar, save for some small errors in the data. We know that the SPOD analysis of Run B is significantly different from that of Run A, but the single-point analysis implies more consistency than we saw. This may indicate issues with applying the adapted marker tracking algorithm globally. However, even with these issues we can see some evidence of periodicity in the data seen in Fig. 4.13. The power spectra also have some slight spikes between 1000 Hz and 2000 Hz as well as near 3000 Hz. The intensities of these peaks are less than those seen in Run A, but there is some evidence of vibrational modes at these frequencies. If we consider the known modal frequencies for the (1,2) and (1,3) modes, we once again see a correlation. Since the points were chosen to be near the antinodes of these modes, it is encouraging to see energy peaks near the expected frequencies. The other major peak near 3000 Hz also seems to appear in Fig. 4.6, but other high-energy modes are not noticeably present in Fig. 4.14. As discussed in Ch. 2, the old algorithm iterates to tailor the edge detection to each marker. This is more computationally intensive, but it seems to be quite effective. We may consider utilizing a similar method for the single-step adapted algorithm to increase precision which may result in better analysis and performance. This idea and other further work will be described in Ch. 6 Sec. 6.3.

Chapter 5: Panel Excitation Mechanism

5.1 Overview

The next step of the FSI experimental campaign is specifically concerned with exciting vibrations in the compliant panel while it is within a hypersonic flow. This must be done without attaching anything to the plate so as to not affect the natural vibrational modes present. When testing the modes of the plate on the benchtop, an impact hammer is used that simply pings the metal plate. This approximates an impulse function on the compliant panel. For the experiments, this effect must be achieved with an automated apparatus able to fit within the test facility and be operated from the outside. The chosen apparatus consists of a push-pull solenoid with a spring. For the desired effect, the solenoid must be normally retracted, so the spring is mounted to return it to this position. The solenoid selected for a prototype was a 12V DC component that extended its impactor when given a signal. The proposed usage would involve mounting the solenoid below the desired point, likely a node or antinode, on the compliant panel, and sending an impulse signal of 12V to actuate the impactor which in turn would ping the panel. In order to confirm the efficacy of this idea, the solenoid was used to perform a standard modal test on the compliant test panel.

5.2 Setup

With an apparatus to excite the compliant panel, there must also be some way to measure those vibrations. For this test, a Polytec OFV 056 vibrometer was used that included a large scanner head equipped with a laser and a desktop control station. The output of the vibrometer was sent to a DAQ that allowed any data to be saved and later processed. To ensure that only relevant data was recorded, a trigger was devised that would only activate the DAQ when the panel was pinged. This methodology is based off of the procedure when an impact hammer is used that is described by Schöneich et al.[16] The hammer contains a piezo-electric crystal in its head that, when impacted, would send a signal to the DAQ to start data collection. Since the solenoid is actuated with a 12V signal, this signal was split to also act as a trigger for the DAQ. A DC plug with 12V was connected to a button that, when pressed, would send a signal to the solenoid and the DAQ simultaneously. Once triggered, the DAQ would record 1 second of data at a 100kHz sample rate. The solenoid was held above the panel such that when the impactor extended it would hit the panel for only a moment.

The panel itself was marked up with a grid of lines. These lines formed 16 boxes and 9 vertices. Each box was assigned a number and each vertex was assigned a letter to make keeping track of each test easier. These designations can be seen in Fig. 5.1. 90 total tests were carried out on the panel, each with a different measurement and excitation location. In order to primarily measure at nodes, the vibrometer was always pointed at one of the numbered boxes. The solenoid pinged at each of the vertices. For ease, some points were not measured, but every vertex was pinged for the test. Specifically, the boxes marked 1, 2, 3, 4, 6, 7, 10, 11, 13, and 16 were measured. For each box, all nine vertices were pinged. Ideally, every ping should produce the same response from the panel, so comparing this can tell us whether or not the proposed excitation mechanism is effective.

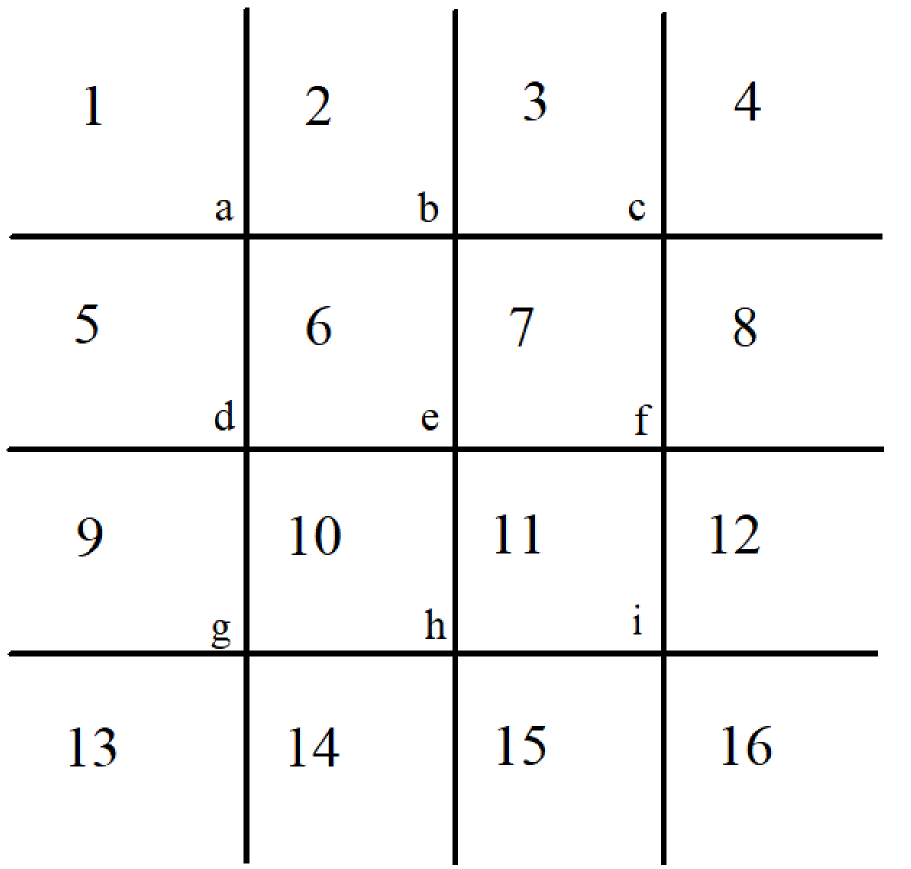

Figure 5.1: Diagram of test panel with marked boxes and vertices.

5.3 Results

The vibrometer data only gives us a time series of displacement in and out of the plane. We must then process this positional data to obtain spectral data. This is done by first filtering the data with a 6th order, low-pass Butterworth filter then running it through two processes. The simpler process is a basic log-decrement analysis of the entire time series. This does not determine any modes within the data, but it can show an overall damping for the ping which can tell us if the excitation mechanism has sufficient power. For log-decrement, the frequency and damping ratio are determined as follows, where n is the number of periods between measurements, T is the period in seconds, $u(t)$ is the displacement at time t, ζ is the damping ratio, ω_d is the damped frequency in radians per second, and ω_n is the natural frequency:

$$\delta = \frac{1}{n} \ln \frac{u(t)}{u(t+nT)} \tag{5.1}$$

$$\zeta = \frac{\delta}{\sqrt{4\pi^2 + \delta^2}} \tag{5.2}$$

$$\omega_d = \frac{2\pi}{T} \tag{5.3}$$

$$\omega_n = \frac{\omega_d}{\sqrt{1-\zeta^2}} \tag{5.4}$$

The second process is more intensive, starting with a STFT and then analyzing the power spectra to then find the strongest frequencies. To account for a delay between the trigger and the impact of the solenoid, the peaks in the power spectra were found at a time 1/4 of the way through the recording window. This also lets any transient vibrations to dampen out quickly, allowing us to easily determine the true modal frequencies. Once these frequencies are determined, the time series of amplitudes at each frequency are fit with an exponential curve of the form ae^{-bt} to try and determine an individual damping ratio. The damping ratio is defined as follows, where b is defined from the exponential and ω is from the STFT and converted to radians per second:

$$\zeta = \frac{b}{\sqrt{b^2 + \omega^2}} \tag{5.5}$$

The comparison of frequencies between pings should tell us whether or not the mechanism is consistent, and the damping ratios will tell us whether or not the pings are sufficiently strong.

5.3.1 Test Set 1

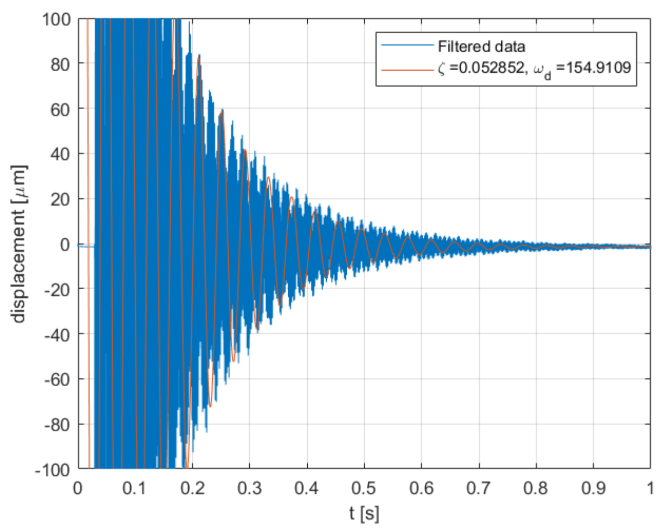

Figure 5.2: Plot of filtered data for ping measured at (6), excited at (a).

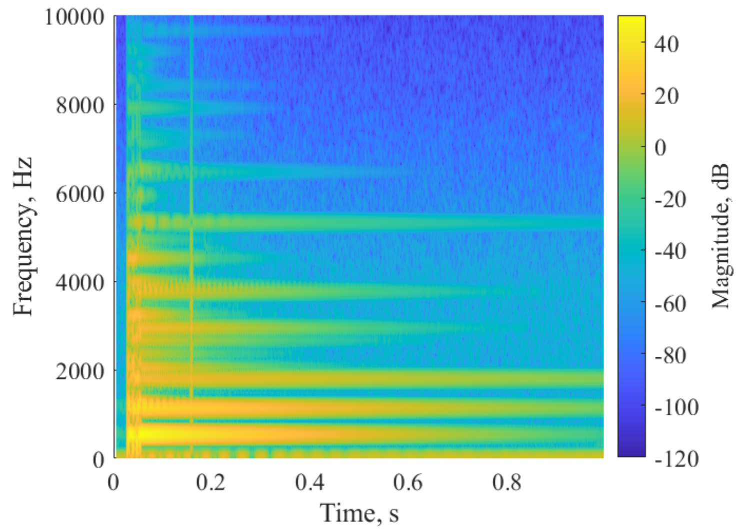

Figure 5.3: STFT of filtered data for ping measured at (6), excited at (a).

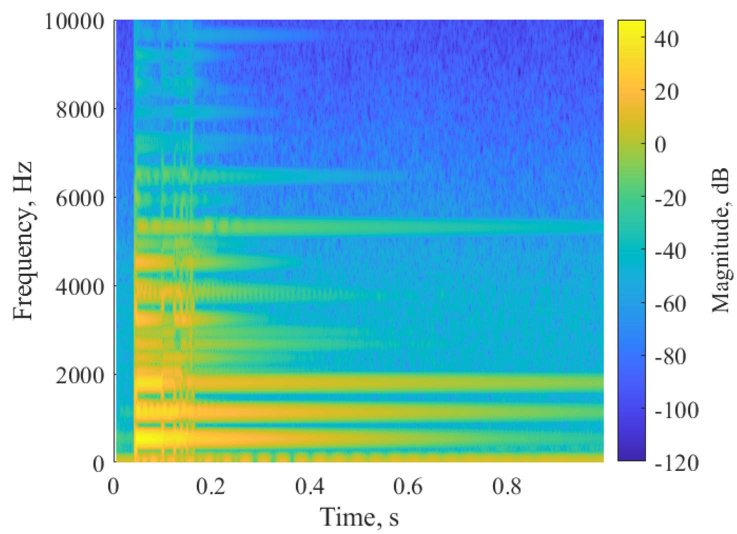

Figure 5.4: STFT of filtered data for ping measured at (6), excited at (c).

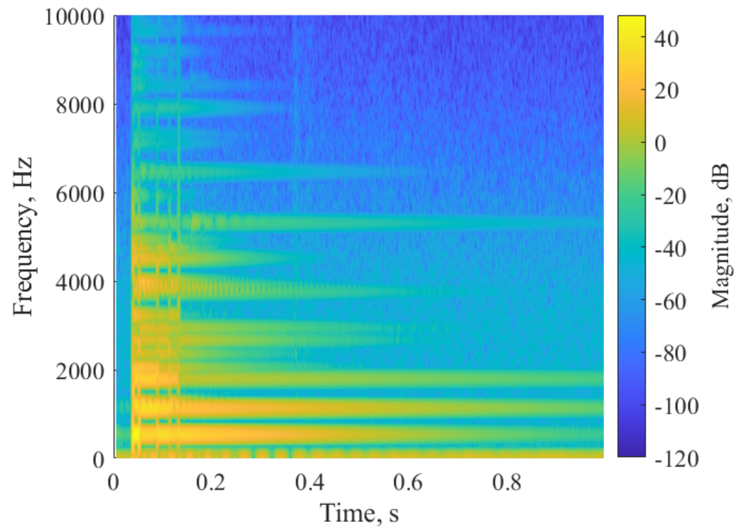

Figure 5.5: STFT of filtered data for ping measured at (6), excited at (g).

A total of 90 tests were done on the compliant panel, but we will only review a subset. The first set of tests were measured in box (6) and excited at points (a,c,e) (see Fig. 5.1). Looking at Fig. 5.2 we see a standard impulse response with fairly good overall damping as the vibrations take over 0.25 seconds to die below 10 microns. For a better understanding, we can examine Figs. 5.3, 5.4, and 5.5 where we see how these displacements sit on a frequency scale. The brighter areas indicate a higher magnitude, which allows us to find rough estimates of the peak frequencies in the natural vibrations of the compliant panel. There is a clear, strong vibration near 2000 Hz with some weaker vibrations near 1000 Hz and 500 Hz. Using the STFT procedure, we sifted through the magnitude data and found the exact magnitude peaks in the frequency domain and examined time series at each frequency. These frequencies and their estimated damping ratios are shown in Tables 5.1, 5.2, and 5.3. We see similar frequencies for all 3 excitation points, also noticeable in the STFT plots. Interestingly, the damping ratios decrease as the frequencies increase. This is different from what we would expect physically, which is for high frequencies to damp out more quickly. This may indicate an error in the analysis of the damping, but the frequencies are definitely present in the STFT analysis. The frequencies seen appear to coincide with the symmetric modes of the panel, especially the (1,3) mode near 2000 Hz.

Table 5.1: Frequencies and damping ratios measured in box (6) and excited at (a).

Frequency	Damping ratio
537.1 Hz	0.0024
1806.6 Hz	0.0003
2929.7 Hz	0.0005
3759.8 Hz	0.0004
5322.3 Hz	0.0002

Table 5.2: Frequencies and damping ratios measured in box (6) and excited at (c).

Frequency	Damping ratio
537.1 Hz	0.0025
1806.6 Hz	0.0024
2661.1 Hz	0.0012
3857.4 Hz	0.0006
5322.3 Hz	0.0002

Table 5.3: Frequencies and damping ratios measured in box (6) and excited at (g).

Frequency	Damping ratio
537.1 Hz	0.0021
1806.6 Hz	0.0020
2661.1 Hz	0.0007
3808.6 Hz	0.0006
5322.3 Hz	0.0002

5.3.2 Test Set 2

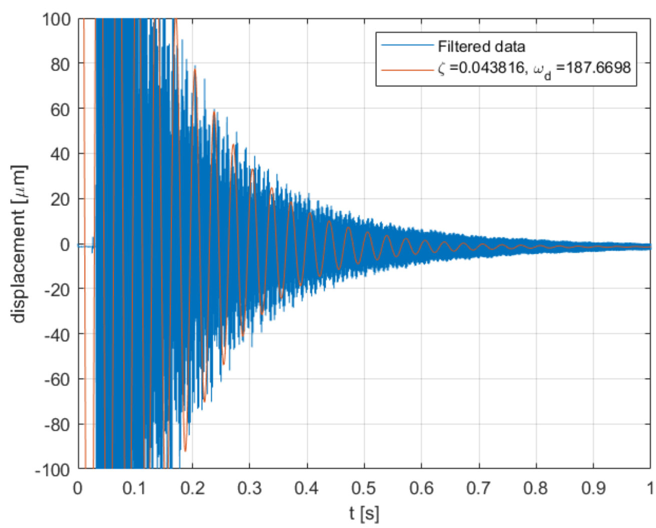

Figure 5.6: Plot of filtered data for ping measured at (11), excited at (b).

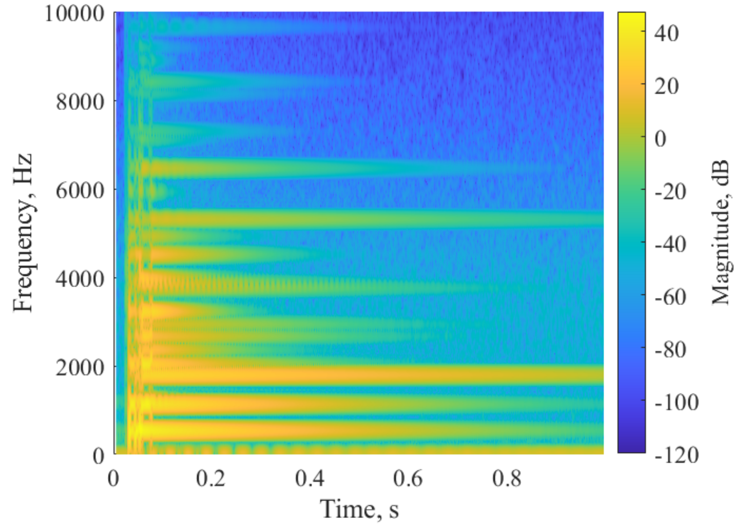

Figure 5.7: STFT of filtered data for ping measured at (11), excited at (b).

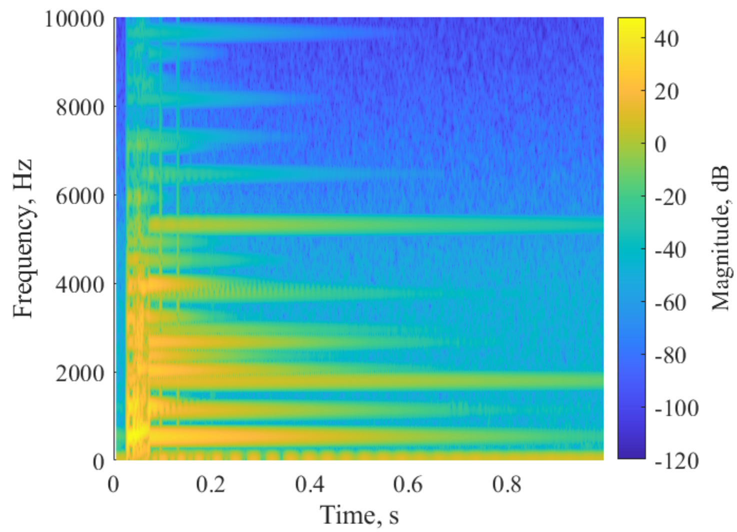

Figure 5.8: STFT of filtered data for ping measured at (11), excited at (d).

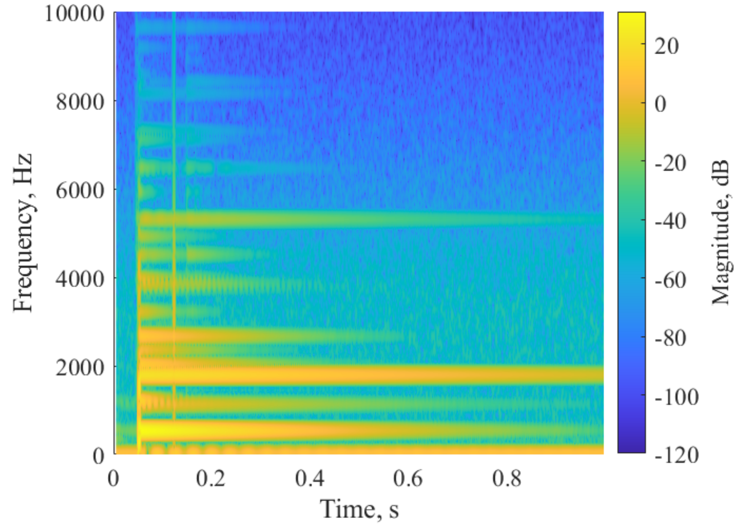

Figure 5.9: STFT of filtered data for ping measured at (11), excited at (f).

The second set of tests were measured in box (11) and excited at points (b,d,f) (see Fig. 5.1). Looking at Fig. 5.6 we see the same kind of response seen for the first set of tests. Examining Figs. 5.7, 5.8, and 5.9 we see how these displacements sit on a frequency scale. There is again a clear, strong vibration near 2000 Hz with some weaker vibrations near 500 Hz. Using the same procedure as before, we sifted through the magnitude data and found the exact magnitude peaks in the frequency domain and examined time series at each frequency. These frequencies and their estimated damping ratios are shown in Tables 5.4, 5.5, and 5.6. We once again see similar frequencies for all 3 excitation points, also noticeable in the STFT plots. The frequencies are also similar to those seen in the first set of tests, which indicates consistency across the panel. The damping ratios seem to have the same problem as previously seen, but the frequencies still coincide with the symmetric modes of the panel, especially the (1,3) mode near 2000 Hz. This display of consistency when the measurement and excitation are at vastly different locations on the plate is very encouraging. Combined with the relatively high strength of the vibrations measured, we have excellent evidence of the efficacy of the solenoid mechanism.

Table 5.4: Frequencies and damping ratios measured in box (11) and excited at (b).

Frequency	Damping ratio
537.1 Hz	0.0021
1806.6 Hz	0.0003
2685.5 Hz	0.0007
3735.4 Hz	0.0004
5297.9 Hz	0.0002

Table 5.5: Frequencies and damping ratios measured in box (11) and excited at (d).

Frequency	Damping ratio
537.1 Hz	0.0026
1782.2 Hz	0.0003
2661.1 Hz	0.0007
3808.6 Hz	0.0012
5397.9 Hz	0.0001

Table 5.6: Frequencies and damping ratios measured in box (11) and excited at (f).

Frequency	Damping ratio
537.1 Hz	0.0020
1806.6 Hz	0.0003
2661.1 Hz	0.0006
5297.9 Hz	0.0002

Chapter 6: Conclusion

6.1 Algorithm Adaptation

When evaluating the adapted marker tracking algorithm, we must consider not only the results but also the methodology. We know that it is computationally simpler than the current method, as stated in Ch. 2, but the exact aspects that are simpler are important. Since it is a single-step algorithm, it does not have the ability to adjust its performance without user input. This is evident in the verification methods in Ch. 3 and the FSI tests in Ch. 4. where we see significantly lower quality results when the algorithm must operate in suboptimal conditions. This is clearest in Ch. 3 Sec. 3.3 where the oblique camera angle greatly affects the displacement calculations, in some cases beyond what is considered acceptable. Since the FSI tests often require severe angles for photogrammetry, this is an issue that requires resolving.

Further examining the inner workings of the algorithm, it must be noted that the edge detection parameters required near-constant alteration when processing the FSI photographs. Ideally, user input would only be required on the first frame and the algorithm would carry out detection across the entire data set. Since the x- and y-displacements should be limited, this ideal case is somewhat expected. Unfortunately, there were several cases when the image smoothing and sensitivity threshold parameters needed to be changed to allow the algorithm to detect all of the markers in the image. Since one of the reasons for adapting this algorithm was saving time, there must be alterations to the method for it to accomplish this goal.

On the positive side, for the on-condition tests, the algorithm performed well. In the first FSI run, the SPOD results were encouraging, showing clear modes as would be required when carrying out a full analysis of the experiment. Additionally, the results of the artificial ellipse verification were good, especially considering that the marker size in the FSI images is generally close to the maximum size in the verification test. The results of the head on photography verification test were also quite good, staying within the uncertainty bounds and showing that the algorithm is capable of accurately measuring the displacements of individual markers.

6.2 Excitation Mechanism

While the testing of the solenoid prototype was limited, it showed excellent results. The mechanism produced consistent and strong vibrations in the compliant panel. The main concerns for applying any excitation to an FSI experiment are that of amplitude and reliability, and the solenoid appears to satisfy both criteria. The ease of actuating the solenoid is also quite beneficial, as it will not require much additional equipment for a test. Using only an impulse signal, the circuitry would be far from complex and the only issue would be finding a way to keep the test at pressure while running wires from outside the tunnel to send the signal. However, this is done for other applications such as pressure sensors, so we do not see it being a difficult problem to solve.

6.3 Next Steps

The new algorithm is only at the beginning of its adaptation. Verification tests have shown it can be effective in the proper setting, and we know that it works well for the larger image detection applications. Based on the findings of the FSI tests and the structure of the old method, there are several immediate considerations for improving the adapted algorithm. Firstly, after implementing subimages, performance was improved as it allowed the algorithm to focus on detecting edges only in a small region. To further improve results, a small amount of iteration may be necessary, specifically with regards to the edge detection parameters. As previously stated, there were many cases in which the image smoothing and sensitivity thresholds needed alteration. If the algorithm were able to iterate through several values on its own for optimal results, it would likely run more accurately and as a whole faster due to not requiring as much tinkering. Currently, it may prove better to continue using the old iterative method for FSI analysis due to its accuracy, but there should be continued work in editing the single-step method such that it can work precisely but retain its speed and computational cost.

For the excitation mechanism, the next steps are to improve the prototype. There only testing so far is a benchtop modal scan with a vibrometer. While encouraging, this does not prove the efficacy of the mechanism under more realistic conditions. Devising a mounting method for the apparatus is imperative so we can determine whether there are any unforeseen effects on the vibrations of the panel. Also, properly wiring the solenoid is important for using the mechanism in a wind tunnel. The mounting of the apparatus will also require slight alterations to the panel which may include redesigning any flow diverters to make room for the solenoid underneath the panel.

Much of this work is still in the early stages, so there is much to be done with regards to refining and verifying both the adapted algorithm and the excitation mechanism. All of this is in service of improving the results of hypersonic FSI experiments, and we believe that both ideas have merit and should be explored further.

Bibliography

[1] A. Schöneich, S. Laurence, F. Dettenrieder, D. Bodony, G. Buck, J. Weisberger, B. Bathel. ”Hypersonic Fluid-Thermal-Structural Interactions on a Compression Ramp with an Embedded Compliant Panel,” AIAA Aviation Forum (2023).

[2] J. McNamara and P. Friedmann, ”Aeroelastic and Aerothermoelastic Analysis in Hypersonic Flow: Past, Present, and Future,” AIAA Journal, Vol. 49, No. 6, 1089-1122 (2011).

[3] D. Dolling, “Fifty Years of Shock-Wave/Boundary-Layer Interaction Research: What Next?” AIAA Journal, Vol. 39, No. 8 (2001).

[4] B. Zuchowski, ”Predictive Capability for Hypersonic Structural Response and Life Prediction, Phase 1—Identification of Knowledge Gaps,” Air Force Research Laboratory, Tech. Rept. AFRL-RB-WP-TR-2010-3069, Wright-Patterson Air Force Base, OH (2010).

[5] B. Zuchowski, ”Predictive Capability for Hypersonic Structural Response and Life Prediction: Phase II—Detailed Design of Hypersonic Cruise Vehicle Hot-Structure,” Air Force Research Laboratory, Tech. Rept. AFRL-RQ-WP-TR-2012-0280, Wright-Patterson Air Force Base, OH (2012).

[6] S. Willems, A. Gülhan, and B. Esser, “Shock Induced Fluid-Structure Interaction on a Flexible Wall in Supersonic Turbulent Flow,” Progress in Flight Physics, Vol. 5, 285–308 (2013).

[7] G. Currao, A. Neely, C. Kennell, S. Gai, and D. Buttsworth, “Hypersonic Fluid-Structure Interaction on a Cantilevered Plate with Shock Impingement,” AIAA Journal, Vol. 57, No. 11, 4819–4834 (2019).

[8] K. Casper, S. Beresh, J. Henfling, R. Spillers, P. Hunter, and S. Spitzer, “Hypersonic Fluid-Structure Interactions due to Intermittent Turbulent Spots on a Slender Cone,” AIAA Journal, Vol. 57, No. 2, 749–759 (2019).

[9] S. Spottswood, T. Beberniss, and T. Eason, “Full-Field, Dynamic Pressure and Displacement Measurements of a Panel Excited by Shock Boundary-Layer Interaction,” AIAA Aeroacoustics Conference, AIAA Paper 2013-2016 (2013).

[10] T. Beberniss, S. Spottswood, D. Ehrhardt, and R. Perez, “Dynamic Response of a Thin Panel Subjected to a ShockWave Impingement and Thermal Buckling,” AIAA Aviation Forum, AIAA Paper 2017-3555 (2017).

[11] T. Whalen, A. Schöneich, S. Laurence, B. Sullivan, D. Bodony, M. Freydin, E. Dowell, G. Buck, ”Hypersonic Fluid–Structure Interactions in Compression Corner Shock-Wave/Boundary-Layer Interaction,” AIAA Journal, Vol. 58, No. 9 (2020).

[12] J. Nelder and R. Mead, “A Simplex Method for Function Minimization,” The Computer Journal, Vol. 7, Iss. 4, 308-313 (1965).

[13] R. Halíř and J. Flusser, ”Numerically Stable Direct Least Squares Fitting of Ellipse,” Proc. 6th International Conference in Central Europe on Computer Graphics and Visualization, Vol. 98, 125-132 (1998).

[14] K. T. Berger, K.E. Hollingsworth, S.A. Wright, and S.J. Rufer, “NASA Langley Aerothermodynamics Laboratory: hypersonic testing capabilities,” 53rd AIAA Aerospace Sciences Meeting (2015).

[15] A. Towne, O. Schmidt, and T. Colonius, “Spectral Proper Orthogonal Decomposition and Its Relationship to Dynamic Mode Decomposition and Resolvent Analysis,” Annual Review of Fluid Mechanics, Vol. 847, 821–867 (2018).

[16] A. Schöneich, T. Whalen, S. Laurence, B. Sullivan, D. Bodony, M. Freydin, E. Dowell, L. Stacey, G. Buck, ”Fluid-Thermal-Structural Interactions in Ramp-Induced Shock-Wave/Boundary-Layer Interactions at Mach 6,” AIAA SciTech Forum (2021).