Intelligent Structural Health Monitoring With Guided Ultrasonic Waves

A thesis submitted in partial fulfillment of the requirement
for the degree of Bachelor of Science
Physics from the College of William and Mary in Virginia,

by

Corey A. Miller

Williamsburg, Virginia
May 2008

Abstract

The purpose of this project is to evaluate a promising technique for the structural health monitoring of aluminum airplane stringers. By sending guided ultrasonic waves throughout the length of a stringer, we are able to examine the structural health of the aluminum. We have recorded waveforms produced with two different transducers in both an incremental milling and an accelerated corrosion sequence, and have thoroughly analyzed the milling waveforms. In order to interpret the signals received from the transducers we used a dynamic wavelet fingerprinting technique which provides us with a simple visual way to view areas of interest within the signal. We found that we are able to accurately extract mode arrivals from the fingerprint images using both a nanopulser np3 and a toneburst source along with shear transducers in SV mode at 1MHz. Using these mode arrivals, we were able to produce group velocity values that agreed with values determined from the measured thickness values to within 4% accuracy.

Acknowledgements

I would first like to thank Dr. Hasso Weiland of the Alcoa Technical Center for providing the stringer test samples, and to Mr. Don Kennamer of Oceana Sensor Technologies for providing the nanopulser np3. I would also like to thank Cara Campbell for the images of the corroded T-stringer that she scanned in the US tank. I would like to give many personal thanks to Jill Bingham and Dr. Mark Hinders for all of their support and guidance over the year.

Contents

1. Introduction

Suppliers such as Alcoa, who provide aluminum parts for airliners, are looking for ways to continue to provide their materials for airplanes in a market that is switching to composite materials. Alcoa's goal is to make their aluminum both lighter and cheaper, a goal which is accomplished by reducing the factor of safety and hence using less aluminum. Unfortunately, using less aluminum in their pieces means that any deterioration of the pieces will cause significantly higher percent damage. Using sensor systems which provide low-cost, real-time inspection of the materials, damage can be detected before failure occurs. Ultrasonic guided waves, Lamb waves, allow for large areas of plate-like structures such as airframes to be inspected with fewer sensors than conventional point by point measurements as they monitor the entire region between sensors [2][3]. However due to the complicated nature of the propagation of these guided waves, algorithms need to be developed that will analyze the waveforms and process the important information in a way that doesn't require mechanics to have extensive knowledge of the wave physics. Because the propagation of the Lamb wave modes depends on both the vibrational frequency and thickness of the material, the waveforms can be processed to determine flaws that can be associated with effective changes in thickness, such as corrosion or cracks.

This project deals with exploiting our understanding of guided wave propagation in airplane stringers. Using accelerated corrosion as well as incremental milling to represent material loss, we monitor the resulting waveforms using a dynamic wavelet fingerprinting technique (DWFT) that we use to extract mode arrivals [6].

2. Lamb Waves

Lamb waves (ultrasonic guided waves) are through thickness vibrations frequently used in nondestructive evaluation, a form of testing that does not damage the test object [8]. They propagate relatively long distances allowing them to monitor extended regions between sensors, making them much more practical for monitoring objects like airframes and industrial piping than conventional point by point measurements. However, since Lamb waves are not straightforward to interpret, a method needs to be developed to automatically analyze the waveforms and process the necessary information.

Guided waves occur in the form of symmetric and anti-symmetric modes. Symmetric modes are signified by a breathing motion, in which the wave expands and contracts thickness as it travels through the plate. Antisymmetric modes, on the other hand, travel in a sinusoidal flexural motion from one boundary to the other as the wave moves outward through the plate [7].

Dispersion curves are used to describe and illustrate the relationship between group velocity of a mode and thickness of the material. A dispersion curve for aluminum is shown in Figure 1.The red curves are the symmetric lamb wave modes while the blue curves are the antisymmetric modes. As can be seen, using a low frequency-thickness value will stay below the frequency-thickness regime for the higher modes, preventing the signal from becoming too complex. Since the aluminum stringers used here had flanges that were 1.6mm thick, we used transducers that were excited at a frequency of 1MHz, giving us a frequency-thickness of 1.6MHz-mm. At this value, we stay below all frequencies of the higher nodes; we only encounter the S0 and A0 nodes. The group

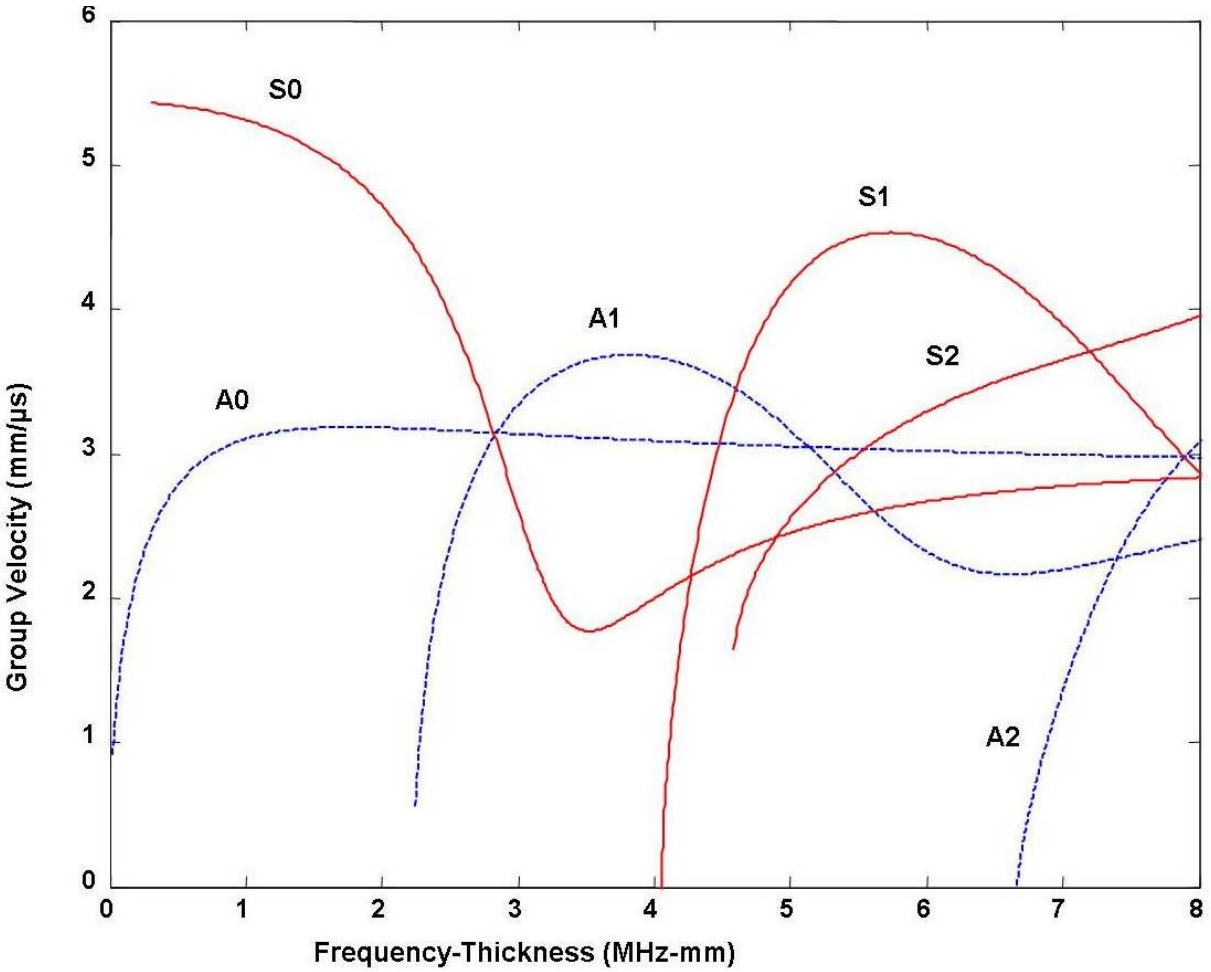

Figure 1: An aluminum dispersion curve. The red curves are the symmetric Lamb wave modes while the blue curves are the antisymmetric modes. We were using a frequency of 1MHz, and an initial thickness of 1.6 mm, giving us a frequency-thickness of 1.6MHz-mm. At this frequency-thickness, the group velocity of the lowest symmetric mode S0 is roughly 5 mm/μs and the group velocity of the lowest antisymmetric mode is roughly 3 mm/μs. Notice that at 1.6 mm/μs, the S0 mode depends on the frequency-thickness, making it sensitive to thickness loss at a constant frequency. The A0 mode is relatively flat, so it is non-dispersive. Once there is a loss in thickness, the S0 mode flattens out, while the A0 mode becomes dispersive.

velocity of the lowest symmetric mode S0 is roughly 5 mm/μs and the group velocity of the lowest antisymmetric mode is roughly 3 mm/μs. Notice that starting from 1.6 mm/μs and decreasing the thickness (moving to the left on the dispersion curve), the S0 mode's group velocity will increase, however the rate at which it increases is steadily decreasing. That is, the slope of the S0 curve levels off as the thickness decreases. The A0 mode's group velocity, however, does just the opposite; it begins with essentially zero change as the thickness initially begins to decrease, but around 1.0mm/μs the group velocity will increase more and more as the thickness gets lower and lower.

These A0 and S0 mode group velocity trends can be used to relate the effect of corrosion and material loss in airplane stringers to our measurements. Starting with a 1.6mm thick aluminum stringer, any thickness loss will be represented as an increase in the group velocity of the first arriving, S0 mode, while the second A0 mode will have a relative small change in arrival over the same change in thickness. Once the stringer thickness approaches 1.0mm, the first arriving S0 mode will show a relatively small change in arrival as the second A0 mode will show a decrease in group velocity.

3. Dynamic Wavelet Fingerprinting Technique (DWFT)

Due to the need to extract multiple mode arrivals from the complicated time-series data received by the transducers, joint time-frequency methods that have the ability to isolate individual spectral components while retaining their absolute relationships in time domain are commonly used. Wavelet transforms have been found to be optimal because mother wavelets can be chosen that best represent a signal feature, i.e. a mode, and can better isolate the signal features of interest in a joint time-scale representation.

An additional benefit of this approach is that mathematical operations transform a 1D time-series signal into a 2D image. This enables 2D image processing operations, but is also advantageous for the human eye's natural ability to detect patterns among a messy 2D image as opposed to a complex 1D signal. Our recently developed dynamic wavelet fingerprint technique (DWFT) renders wavelet time-scale signal patterns as binary contour plots of the wavelet transform coefficients [4]. The resulting fingerprint-like image is amenable to computerized interpretation for automated extraction of signal features of interest.

Figure 2 illustrates how the wavelet fingerprints are formed. An isolated ultrasonic pulse is shown in (a) along with a standard surface plot (b) of the corresponding continuous wavelet coefficients. In (b) we also see the coefficients displayed as a traditional contour plot underneath the surface plot. To form a wavelet fingerprint, we use the "thick contours" shown in (c). The wavelet coefficients are first normalized into the range of [0, 1], and then equal slices of the 3D-coefficient surface are projected back onto the time-scale plane. This results in the black and white, 2D representation shown in (d) which we refer to as the wavelet fingerprint [5].

We have implemented the wavelet fingerprint technique into a simple MATLAB tool that allows us to easily import the recorded waveforms and view them under different wavelet basis functions and parameter settings, identifying the best combinations that distinguish the waveform features of interest. In the resulting fingerprints, the peaks (positive values) are shown in white while the valleys (negative values) are shown in gray.

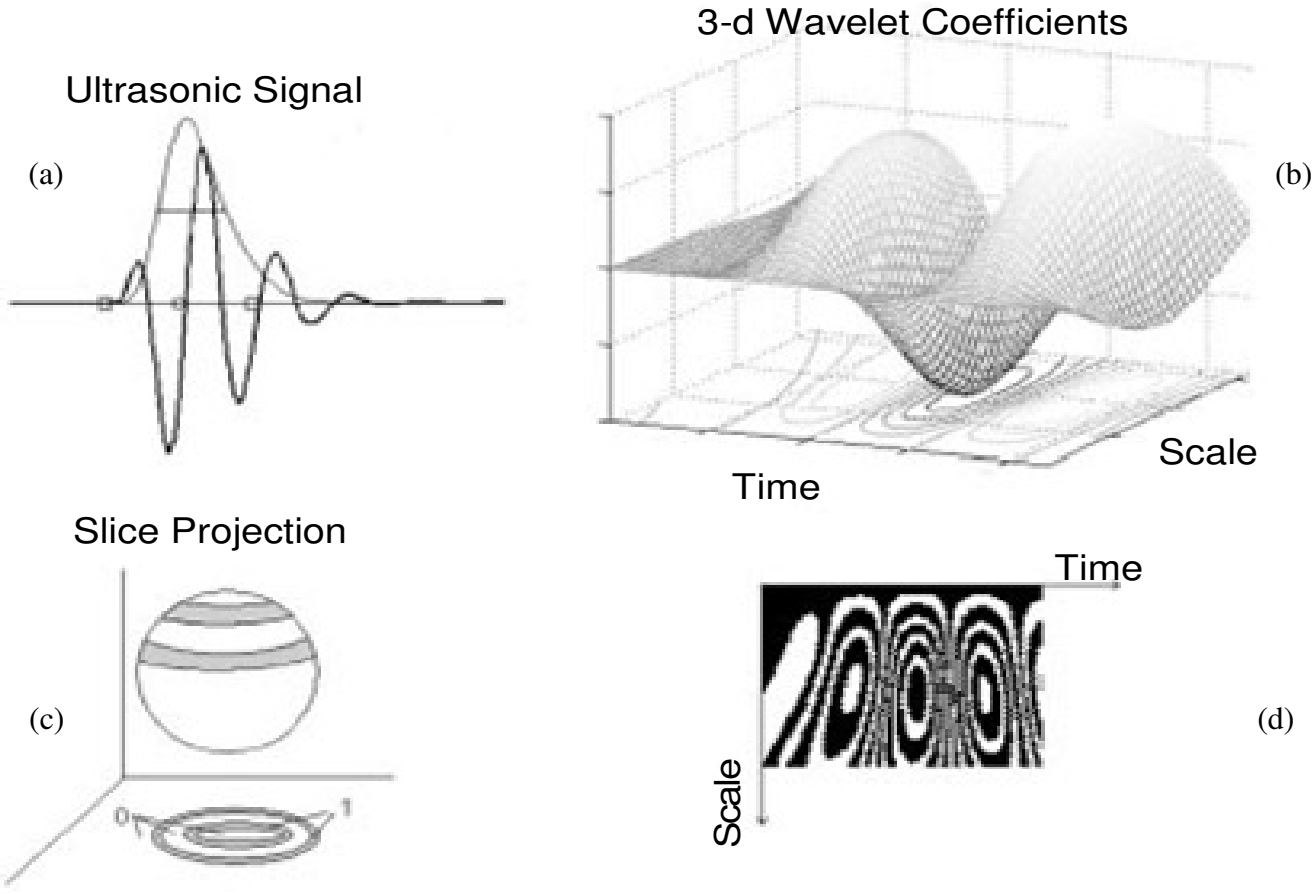

Figure 2. Wavelet fingerprints generation. An isolated ultrasonic pulse is shown in (a) along with a standard surface plot (b) of the corresponding continuous wavelet coefficients. In (b) we also see the coefficients displayed as a traditional contour plot underneath the surface plot. To form a wavelet fingerprint, we use the "thick contours" shown in (c). The wavelet coefficients are first normalized into the range of [0, 1], and then equal slices of the 3D-coefficient surface are projected back onto the time-scale plane. This results in the black and white, 2D representation shown in (d) which we refer to as the wavelet fingerprint.

4. Experimental Details

4.1 Equipment Setup

Our lab setup can be seen in Figure 3. We used a desktop computer that contains a Matec TB1000 toneburst board along with a Gage A/D board. Alternatively, the toneburst card is capable of pulse excitation, cycling the voltage at a desired frequency for a number of cycles. The nanopulser np3 has its own A/D card in it and is connected to the computer through a USB cord. The np3 uses a spike excitation for a desired number of nanoseconds, causing the signal to contain many frequencies.

The main transducers we used were a standard shear (SV) transducer, as can be seen in Figure 4 as well as a custom transducer we were testing for a local company Oceana Sensors (O), as can be seen in Figure 5. Both transducers are shear contact transducers, however their configurations are different. In both cases, when the voltage hits the transducer the PZT crystal vibrates with the voltage cycle or spike. In the SV transducer the crystal expands and contracts in line with the connection cable, which we lined up parallel to the surface of the stringer. The Oceana transducer is a donut disc shape that expands out in all directions within the plane of the contact surface.

We also began our recordings using a pair of longitudinal (L) transducers in addition to the SV and O transducers, however after several separate waveform recordings using the L transducers we discontinued using them. It was apparent from the poor quality of the waveforms we were seeing that they were not likely to be of much use to us. The reason we were seeing such poor waveforms is most likely due to the fact that the transverse-wave transducers are more directional, whereas the longitudinal transducers are omni-directional. This causes the transverse-wave transducers to launch

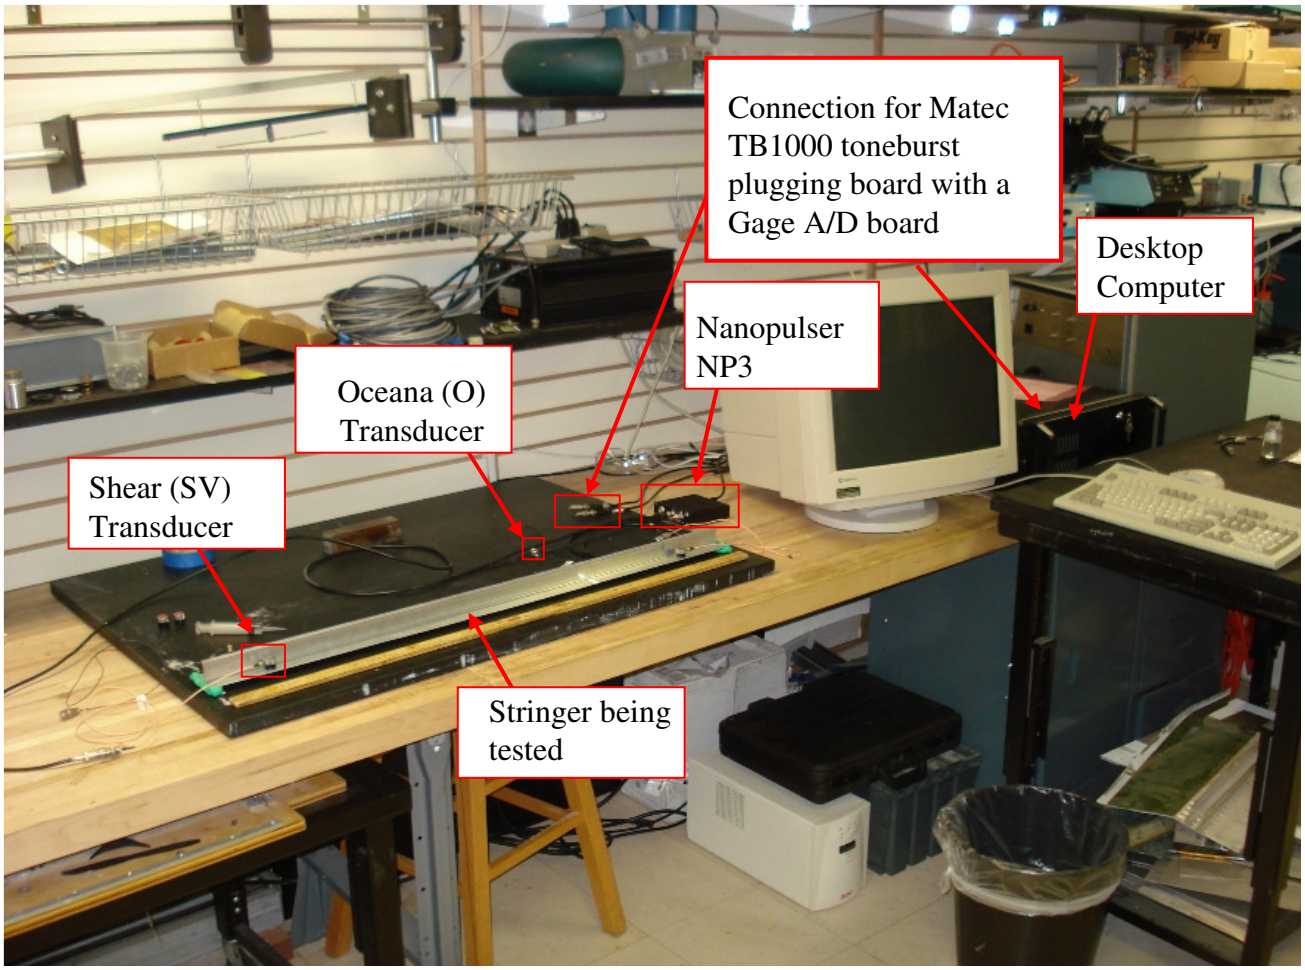

Figure 3: We used a desktop computer that contains a Matec TB1000 toneburst board with a Gage A/D board. The toneburst card is capable of pulse excitation, cycling the voltage at a desired frequency for a number of cycles. The nanopulser np3 has its own A/D card in it and is connected to the computer through a USB cord. The np3 uses a spike excitation for a desired number of nanoseconds, causing the signal to contain many frequencies. The main transducers we used were a standard shear (SV) transducer as well as a transducer we were testing for a local company Oceana Sensors (O). Both transducers are shear contact transducers, however their configurations are different. In both cases, when the voltage hits the transducer the PZT crystal vibrates with the voltage cycle or spike. In the SV transducer the crystal expands and contracts in line with the connection cable, which we lined up parallel to the surface of the stringer. The Oceana transducer is a donut disc shape that expands out in all directions within the plane of the contact surface. The aluminum stringers were raised off the table on modeling clay.

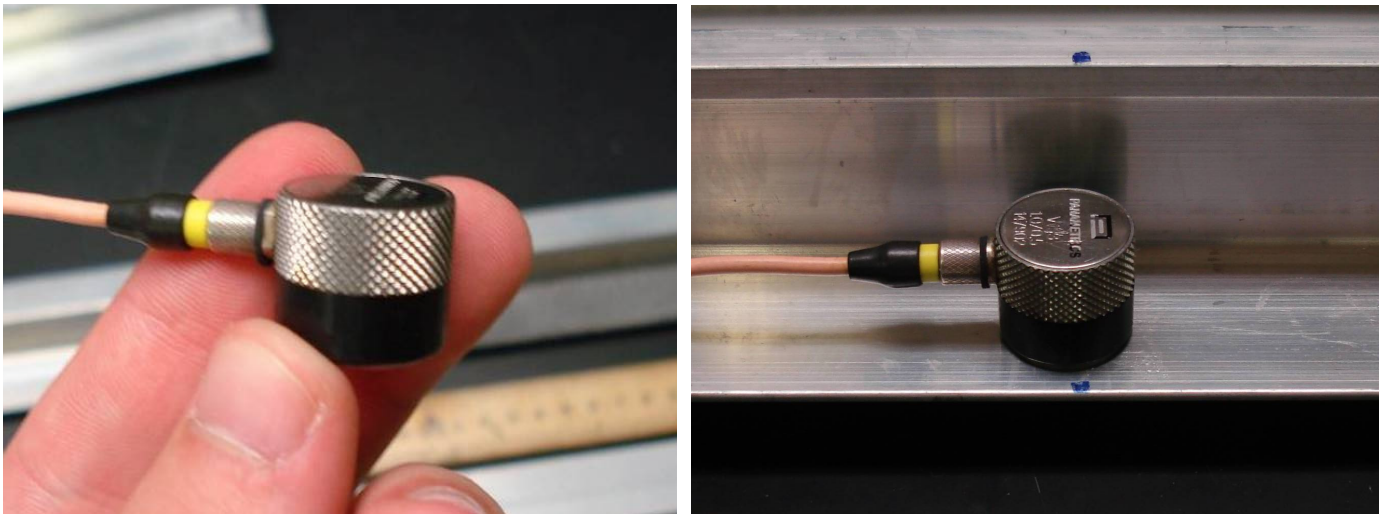

Figure 4: Relative size of the shear vertical (SV) transducer (left) and placement of the SV transducer on aluminum stringer (right)

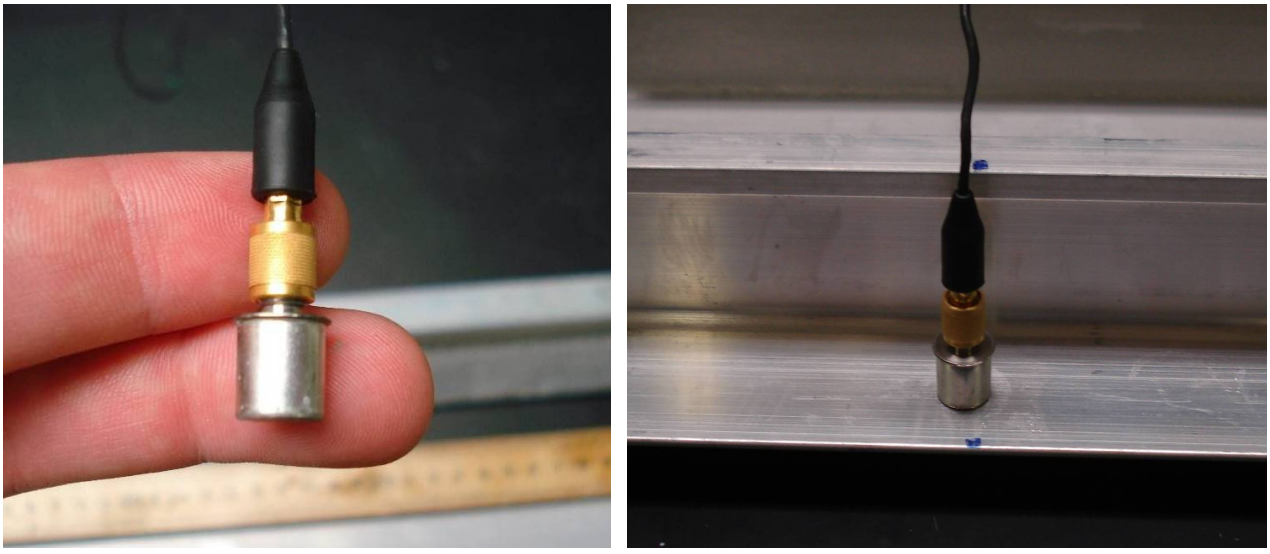

Figure 5: Relative size of the Oceana (O) transducer (left) and placement of the O transducer on aluminum stringer (right)

the waves down the stringer flange as opposed to longitudinal transducers which make other waves that end up bouncing around the stringer in other ways. This bouncing tends to complicate the signal, which is what we ended up seeing. Because of this, we discontinued using the longitudinal transducers for the remainder of the project.

4.2 Accelerated Corrosion Test

In order to study how real-world corrosion would generate material-loss on the stringers, we performed a controlled, accelerated corrosion test. For this, we used three types of aluminum airframe stringers. As can be seen in Figure 7, we used "Z" cross-section stringers that are 1.0m in length, "T" cross-section stringers that are 1.0m in length, and longer "T" cross-section stringers that were 2.4m in length. These will be referred to as Z, T, and L stringers, respectively.

We started by creating a PVC dunking vessel which was to be used for cleaning the corrosion products off the samples. This was a piece of PVC piping that was cut to 1.0m in length, with one end sealed and the other with a removable top. This dunking vessel was then filled with nitric acid (0.1M HNO_3).

The solution used in the accelerated corrosion followed an ASTM standard test method known as the EXCO test [1]. The EXCO solution consisted of sodium chloride (NaCL), potassium nitrate (KNO_3), and nitric acid (HNO_3), which converts to 234g NaCL, 50g KNO_3, and 6.3mL HNO_3 which were diluted to 1L with de-ionized water in our actual solution. To achieve a faster corrosion rate, we also included 10mL of hydrogen peroxide on the recommendation of Alcoa.

In an idealized experiment we would be using infinitely long stringers so the signal of interest would continuously move past our recording transducer, thereby recording a complete waveform. In the real world, however, there are boundaries. Assuming right angle boundaries, guided waves will reflect back towards their origin upon reaching an endpoint. We took a few precautions in order to reduce the amount of unwanted wave reflection off the ends of the stringers in our recorded signals. One method was to place the transducers 10cm from each end of the stringers. This gives the transducer time to record all modes of interest before the fastest modes have the opportunity to bounce back off the end of the stringers towards the transducer again. Another method of prevention was to place the ends of the stringers in modeling clay. Doing so allowed some of the signal to be absorbed by the clay upon reaching the end, thereby effectively damping the reflecting signal. The clay also keeps the stringer from resting directly on the benchtop.

Because the T and Z stringers were only 1.0m in length, the placement of the transducers was limited to being 10cm from each end of each stringer, giving an 80cm propagation distance between them. For the 2.4m L stringer, one transducer was placed 10cm from one end of the stringer. The second transducer was then placed 90cm from the same end, giving an 80cm propagation distance between the two. It was then moved 20cm further away from the original transducer, followed by another 60cm further away, followed by another 60cm further away, making the wave propagation distance 100cm, 160cm, and 220cm respectively. This can be seen in Figure 6. Note that the final transducer location with a propagation distance of 220cm allows the reflection distance of 10cm on either side as stated above.

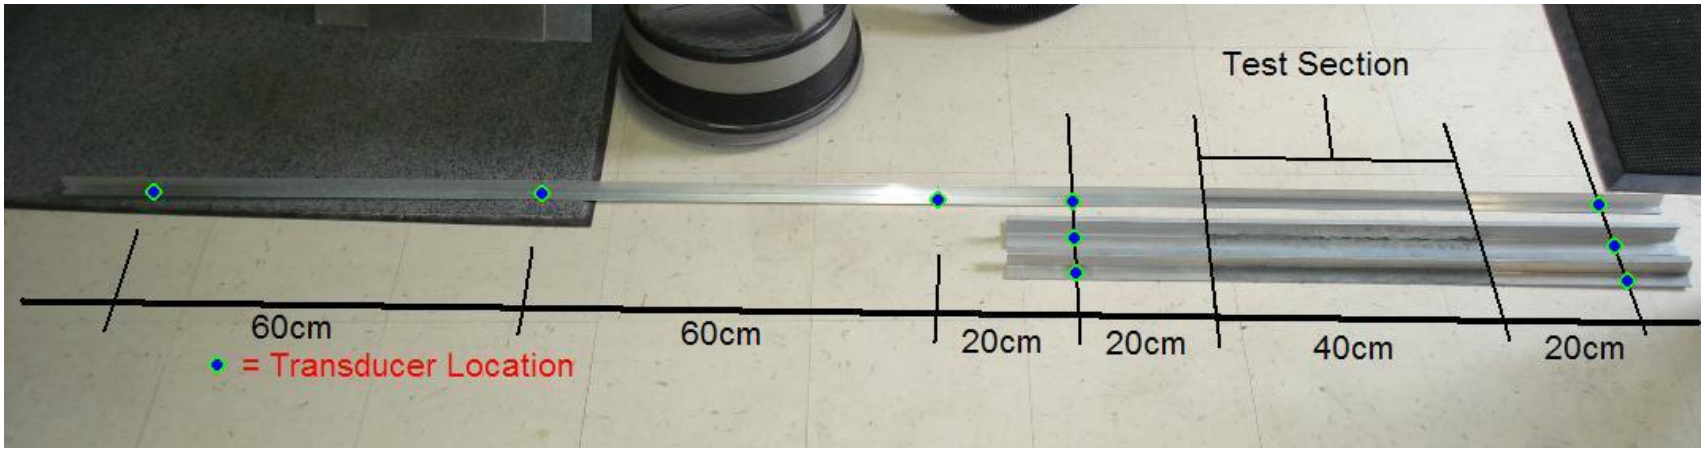

Figure 6: Setup of transducers on T, Z, and L stringers. On the T and Z stringers, the transducers are placed 10cm from each end of the stringer while the test section is located in the middle 40cm of each stringer. On the L stringer, the transducers are initially placed 20cm from either side of the test section (same position relative to the test section as the T and Z stringers), then one is moved 20cm (for a total of 100cm between them), then it is moved 60cm further (making 160cm between them), and finally another 60cm (for a total of 220cm between the two transducers).

Figure 7: Cross section of the 2.4m "L" cross-section stringer (left), the 1.0m "T" cross section stringer (center), and a 1.0m "Z" cross section stringer (right).

We took baseline measurements on the clean stringers before applying any corrosion solution. Our measurements included four sampling combinations of both a nanopulser NP3 and a Toneburst (TB) source, each with Oceana (O) and shear vertical (SV) transducers. We also took 5 thickness measurements of each stringer in the area to be corroded using a micrometer. Using a pitch-catch setup for all of our measurements, we used 80cm propagation sections on the Z and T stringers and propagation sections of 80cm, 100cm, 160cm, and 220cm on the L stringer.

After taking the baseline measurements, we masked the stringers with painters tape 30cm from each of the ends, leaving 40cm of exposed flange in between which was to corrode, in order to create "walls" for the corrosion solution as can be seen in Figure 8 and Figure 9. This was done to all the stringers, with the L stringer having this test section placed 30cm from one of the ends (making it similar to the other T stringer but with one end extending out to the 2.4m length). We then placed the EXCO solution onto the exposed 40cm flange sections on each stringer, covering the entire exposed surface. The EXCO solution was left on the stringers for an exposure period of roughly 12 hours. At the end of the exposure period, the solution was left on while all necessary waveforms were recorded once again.

Next, the stringers were rinsed off with water and the T and Z stringers were soaked in nitric acid located in our dunking vessel for about 10 seconds each (the L stringer was too long for our vessel, so its primary method of rinsing was with water in the sink). Once the stringers were cleaned, we recorded all necessary waveforms with the clean surfaces before re-taping the "walls" and applying the corrosion solution again.

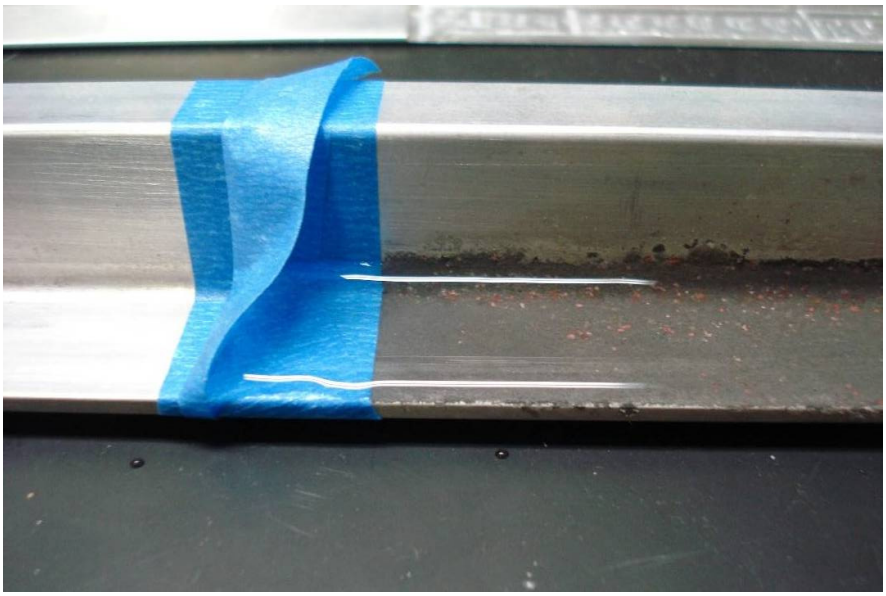

Figure 8: Close-up of blue painters tape barrier. A piece of painters tape was placed firmly on the surface of the stringer. A second piece of painters tape was put on top of the first at an angle that when folded over created a wall. This wall prevented any of the solution from leaving the test section.

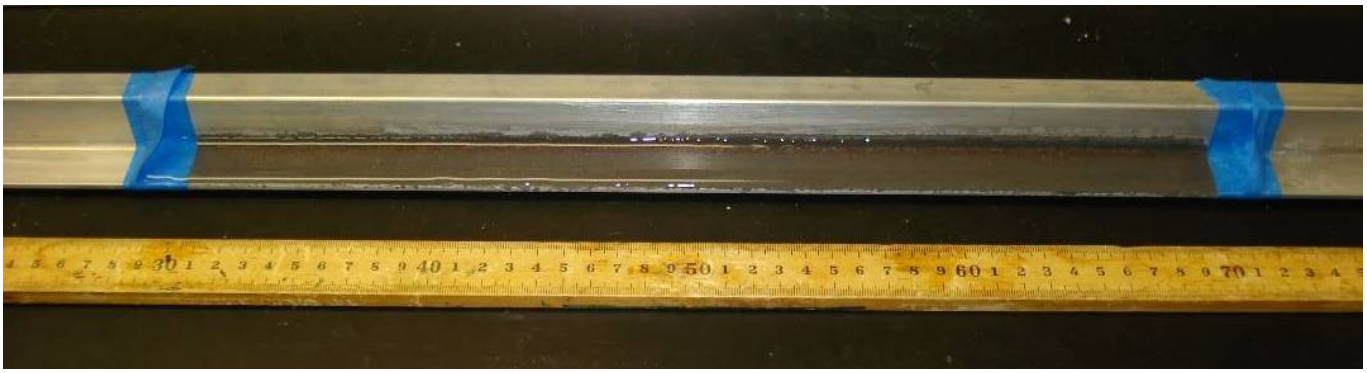

Figure 9: Setup of accelerated corrosion test section. Blue painters tape was placed on either end of 40cm test section creating a barrier for the corrosion solution. The solution was then left on the stringer for periods of 12 hours.

This process of taking measurements with the solution on, rinsing the stringers, taking measurements on the clean surfaces, and re-applying the solution was repeated until noticeable corrosion had occurred; this ended up taking 14 cycles.

The accelerated corrosion testing produced an appearance of flaking on the T stringer as can be seen in Figure 10. This flaking effect is known as exfoliation corrosion, where wedges of corrosion product build up along the grain boundaries of the metal and exert pressure that forces layers of metal upwards, giving rise to a layered appearance. This layered product made it difficult to accurately measure the thickness of the flange with a micrometer, as there could very well be empty space below the top layer.

The Z and L stringers, however, were made up of different alloys than the T stringer and experienced a different type of corrosion. Upon inspection, there appeared to be multiple holes covering the surface of the flange as can be seen in Figure 11. The spotted effect is known as pitting corrosion, where undermining pitting gives the appearance of incipient exfoliation.

4.3 Incremental Milling Test

While the accelerated corrosion test provided us with stringers that were close to what would be found in real-world situations, we also wanted to study the effect of material loss in a more controlled situation. As opposed to the accelerated corrosion test, where it was difficult to determine the thickness of each flange mid-test, we were able to precisely know how much material we are removing by milling the flange in incremental steps. The milling machine we used can be seen in Figure 12. We used a single “T”

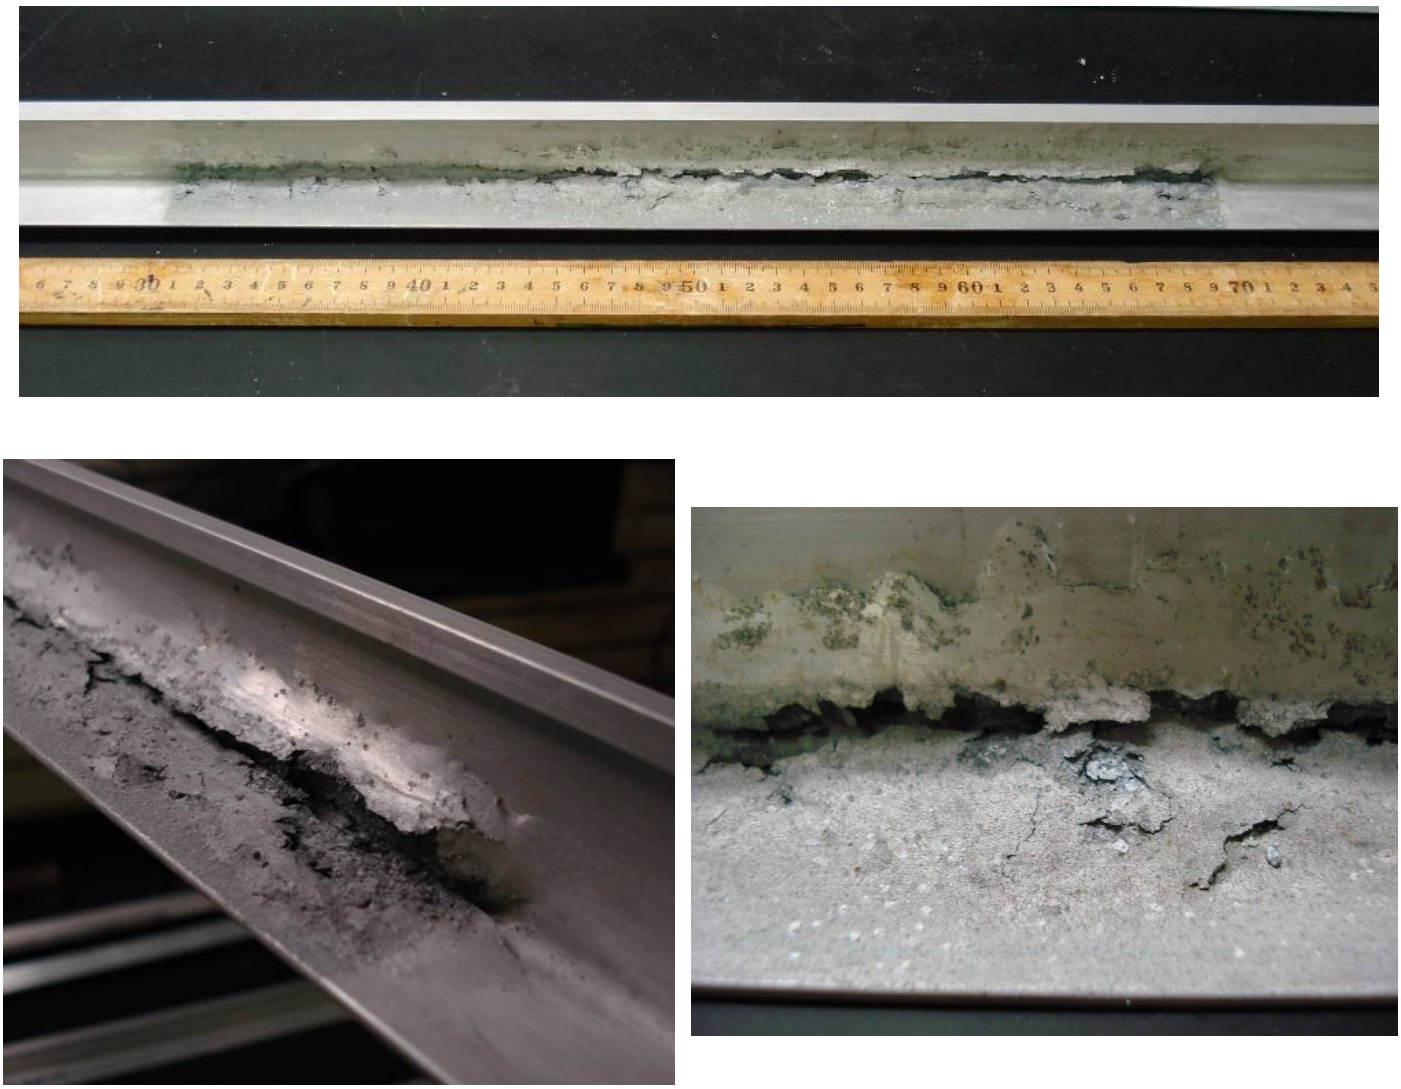

Figure 10: The accelerated corrosion test section on a 1.0m “T” stringer at the end of the accelerated corrosion test (top). Taking a close-up look at the test section, the flakes of aluminum can easily be seen (bottom left, bottom right). The flaking effect is known as exfoliation corrosion, where wedges of corrosion product force layers of metal upwards, giving rise to a layered appearance.

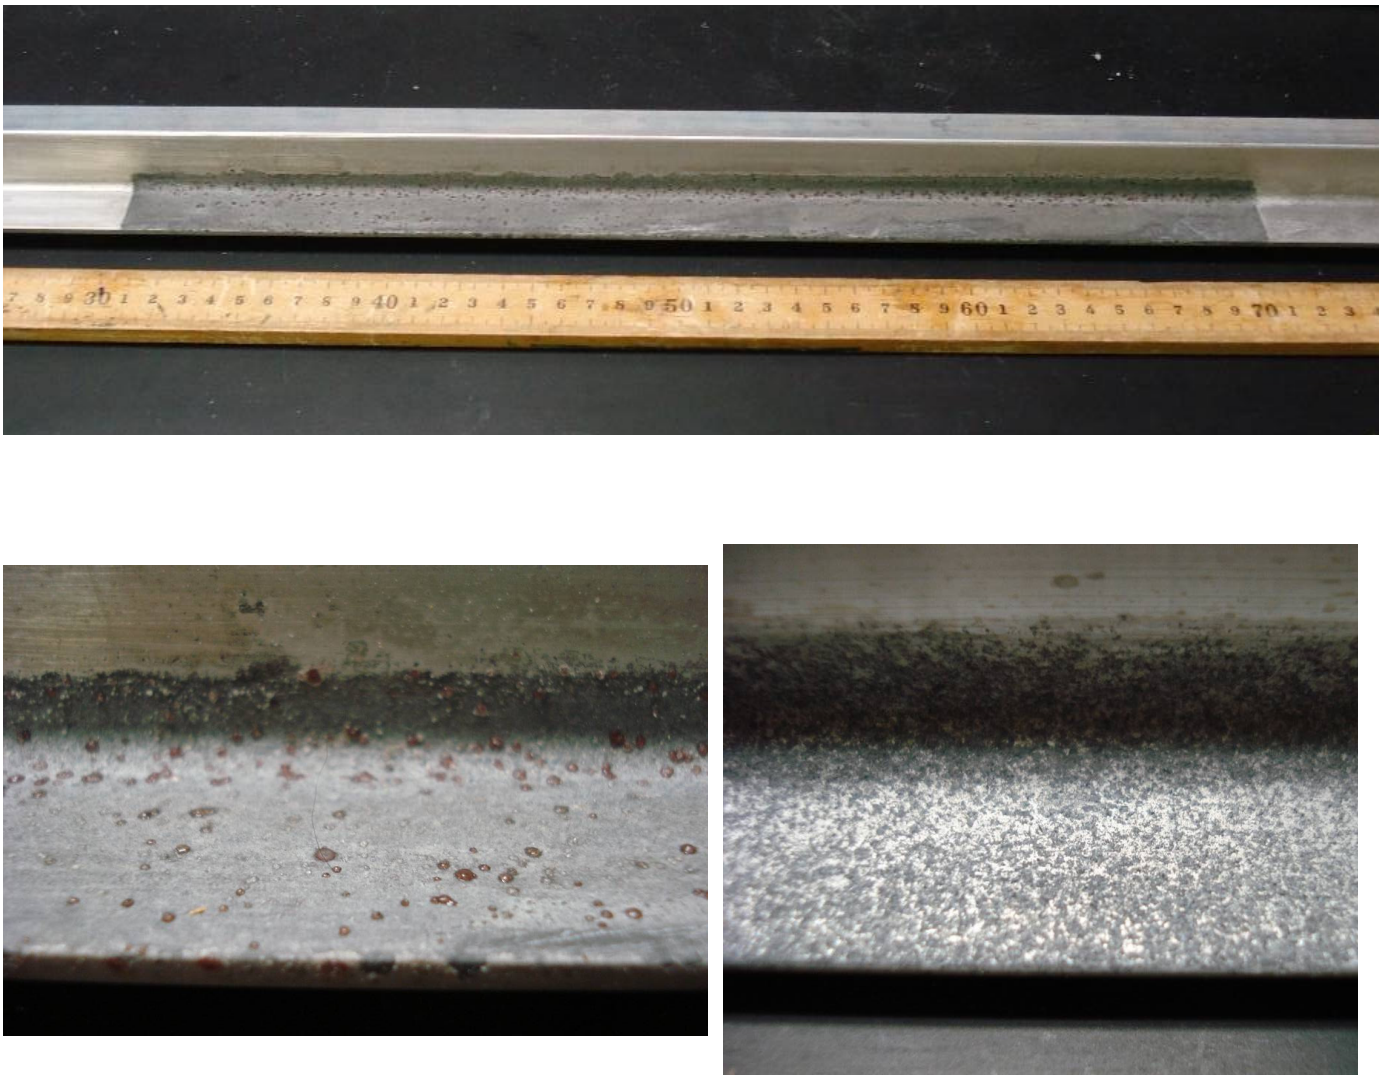

Figure 11: The accelerated corrosion test section on a 1.0m “Z” stringer at the end of the accelerated corrosion test (top). Taking a close-up look at the test section, small holes in the aluminum can be seen (bottom left for Z stringer, bottom right for L stringer). The hole effect is known as pitting corrosion, where undermining pitting gives the appearance of incipient exfoliation. The different alloys in the Z and L stringers was the cause of the slightly different pitting corrosion results.

stringer 1.0m in length for the milling portion of the project, setting the stringer up in exactly the same fashion as before; the transducers were placed 10cm from each end of the stringer, and the milled test-section was located 30cm from each end of the stringer (giving a total milled length of 40cm). This can be seen in Figure 13.

Along with measuring the initial thickness of the flange in 5 places, we recorded initial waveforms using the O and SV transducers, each paired with the NP3 and TB sources on the stringer before performing any milling. We then removed a ~.10mm layer off the middle 40cm test section of the flange. The stringer was then wiped clean of all shavings. We recorded the waveforms and flange thickness using the same placement and method as before. This process of milling off .05-.15mm layers and recording the new waveforms was repeated for total of 10 increments (which removed about 1.1mm).

5. Waveform Analysis

5.1 File editing

Once we had all of our waveforms recorded, we first needed to "clean" the files. Every time we recorded a waveform throughout our experiment, we would actually record 10 waveforms in a row. These 10 waveforms were stored as binary files, with the waveforms saved back-to-back within the file. The editing process of these binary files was to average them together into one single waveform which we would then use in our fingerprint analysis. The purpose of saving 10 waveforms at a time as opposed to a single waveform was to ensure we recorded at least one good waveform for each measurement; during the actual recording of the waveforms, we were relying on human means to physically place the transducers on the stringers and press record on the computer. While

Figure 12: Milling machine setup in the NDE Lab at the College of William and Mary. The 40cm test section of a 1.0m "T" stringer was milled off in ~0.05mm increments, with measurements taken after each milling step.

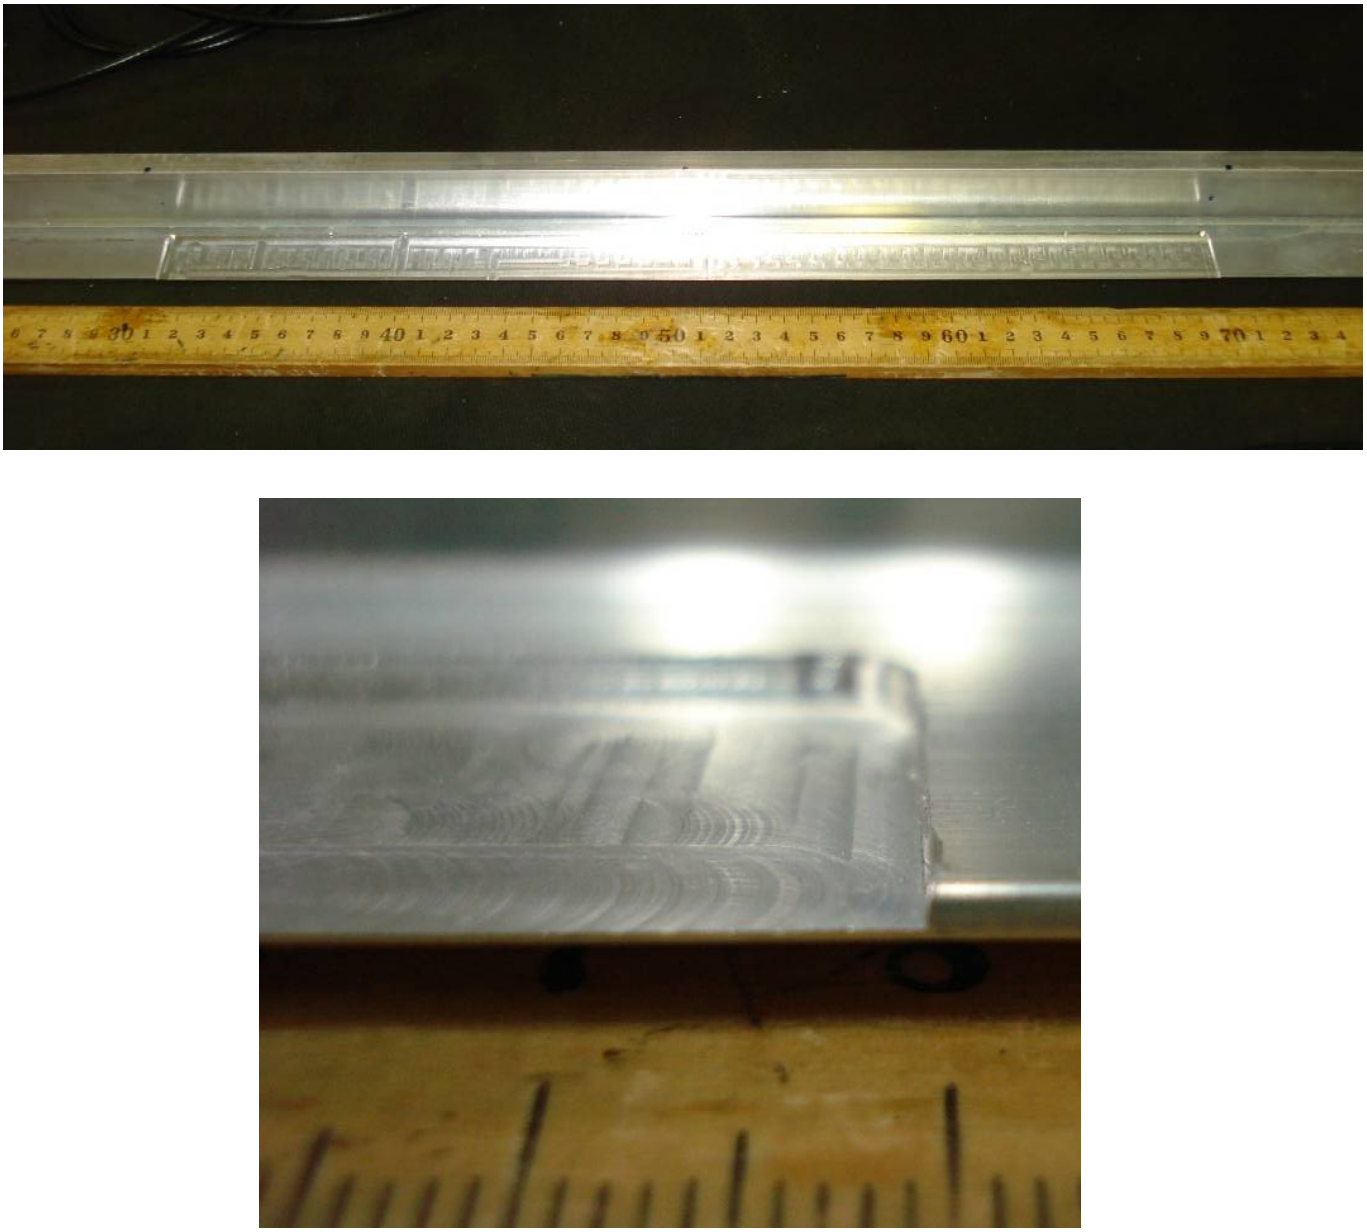

Figure 13: The milled test section on a 1.0m "T" stringer at the end of the incremental milling test (top). Taking a close-up look at the test section, the groves produced by the milling machine can be seen. The initial thickness of the stringer was 1.6mm, which was then reduced to 0.5mm by the end of the test.

this seems an easy task, performing all the measurements alone can quickly become daunting. For whatever reason, from the inability to simultaneously hold the transducers while starting the waveform recording to simply having poor contact between the transducer and the stringer, we recorded multiple “bad” waveforms, usually at the start of the 10 sequential waveforms taken during each measurement. These bad waveforms, therefore, needed to be removed from our files before we averaged the remaining “good” waveforms together, as can be seen in Figure 14.

In a real application, the transducers would not have their large external casing and would be cemented in place on the material being monitored. This secure placement would ensure consistently clean waveforms being recorded, and would remove the need to clean the bad waveforms from within the files. We would also develop algorithms which automatically asses each waveform and discard any ‘bad” ones, but that is beyond the scope of this project.

5.2 Incremental Milling Analysis

We began our analysis by looking at the milling data we recorded. Because milling is a controlled process in which we know how much material is being removed each step, it will produce waveforms that accurately follow expected mode arrival trends based on thickness measurements. We manually inspected all 10 waveforms in each of the 44 files we saved. When bad waveforms were found within a file, we would manually cut out the bad portion of the file and re-save the file. Once we had removed all the bad waveforms, we then averaged the remaining good waveforms within each file together, producing a single, clean waveform that we would use in our fingerprint analysis.

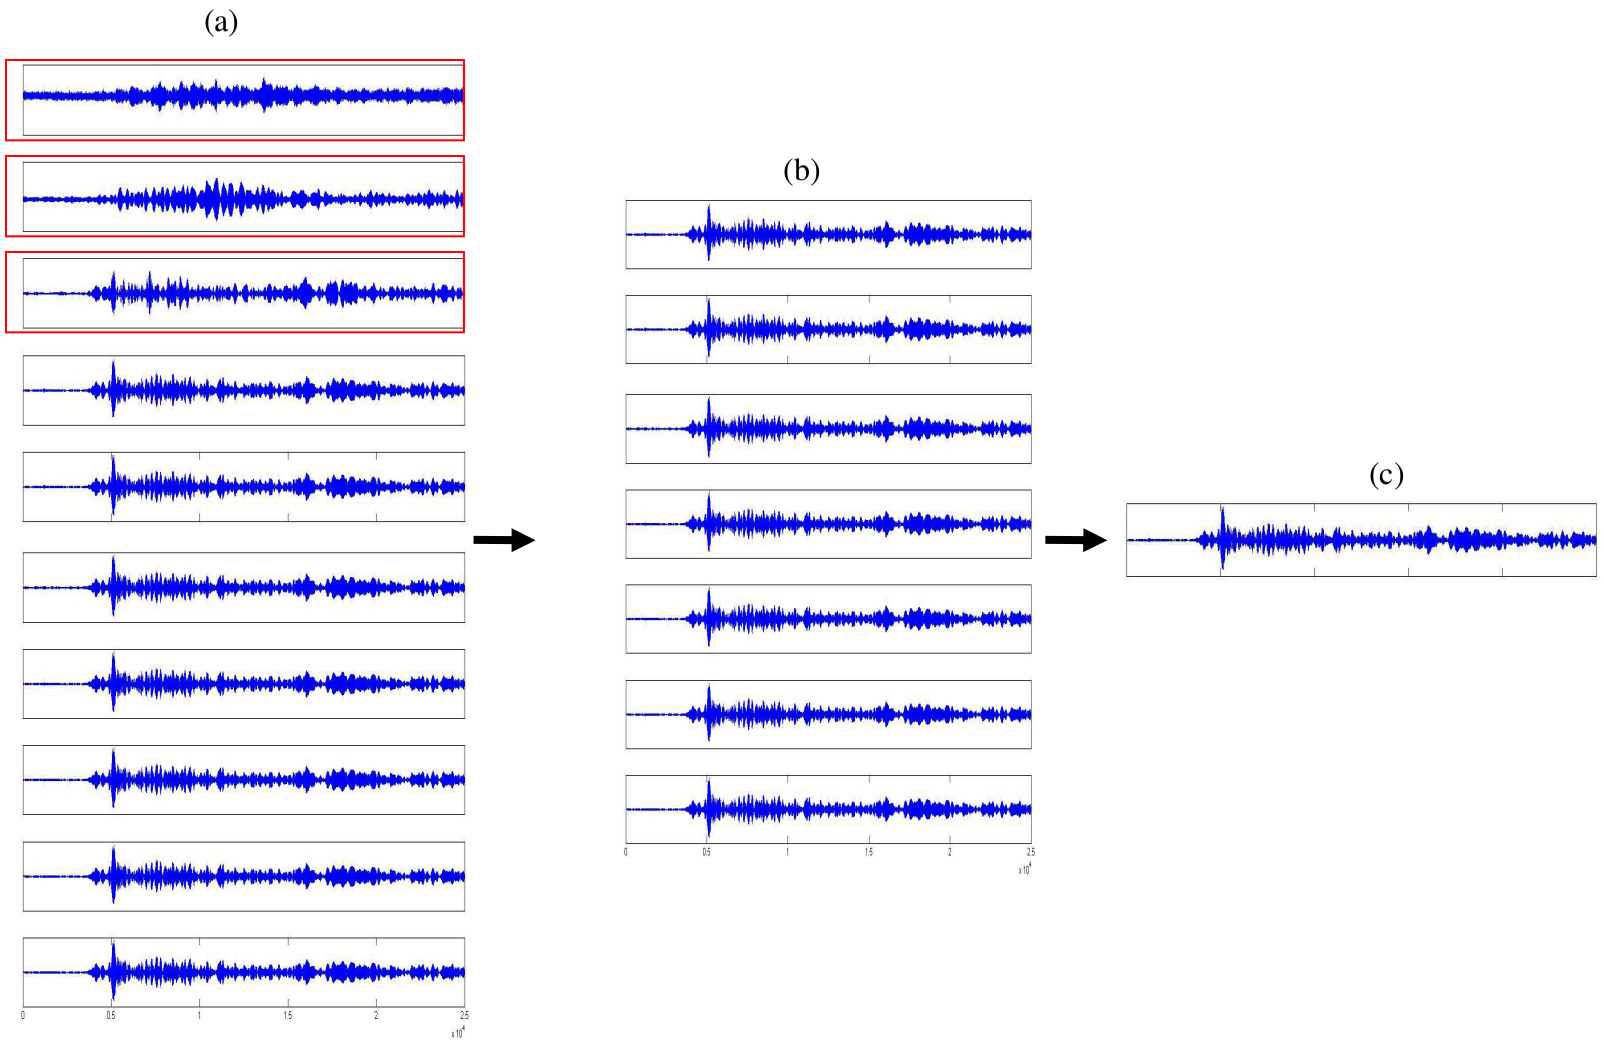

Figure 14: Editing process of recorded data sets. Each file we recorded contained 10 individual signals, recorded in succession. Occasionally, bad individual signals (outlined in red) were recorded due to poor transducer contact, incorrect alignment, etc (a). Each file was manually inspected and the bad data sections were removed from the file using MATLAB (b). The remaining clean data was then averaged together (c). This process produced a single, clean, averaged waveform per file based on the initial 10 recorded waveforms.

We then wanted to know where to look for the mode arrivals within the waveforms. Using an aluminum dispersion curve, we determined the approximate S0 and A0 mode group velocities based on our measured thicknesses. We plotted these mode group velocity values against the milling step, and used these plots to determine trends in the expected arrival times as the milling step increased. We found that the S0 mode sped up from roughly 158 μs to 153 μs, a difference of 5 μs, while the A0 mode slowed down from roughly 252 μs to 270 μs, a difference of 18 μs. This means that the overall change in A0 mode arrival time would be more significant than the overall change in S0 mode arrival time. We also calculated the overall expected arrival times of the S0 and A0 modes based on the sampling rate for both the NP3 and TB sources and the calculated group velocities. Figure 15 shows these results. Using these expected arrival times, we were able to narrow down where to look for the mode arrivals.

Once we knew where to look, and what trends to look for, we were able to begin an initial comparison of the waveforms using the wavelet fingerprinting GUI run out of MATLAB. We chose the data from the NP3 source using the SV transducers to begin with and imported data from every other milling step. The reason we didn't look at all of the milling steps initially is because importing the files and viewing each fingerprint is a time-consuming process, so by initially looking at just a few we were able to get an accurate sense of what a further analysis would produce. We then used the ability to easily change wavelet pre-filters and fingerprint filters to search for the desired mode arrivals within the initial un-milled stringer recordings. Once we determined which filters proved to be optimal for producing easily readable fingerprints for the initial measurements, we looked at the rest of the NP3 SV waveforms as the milling step

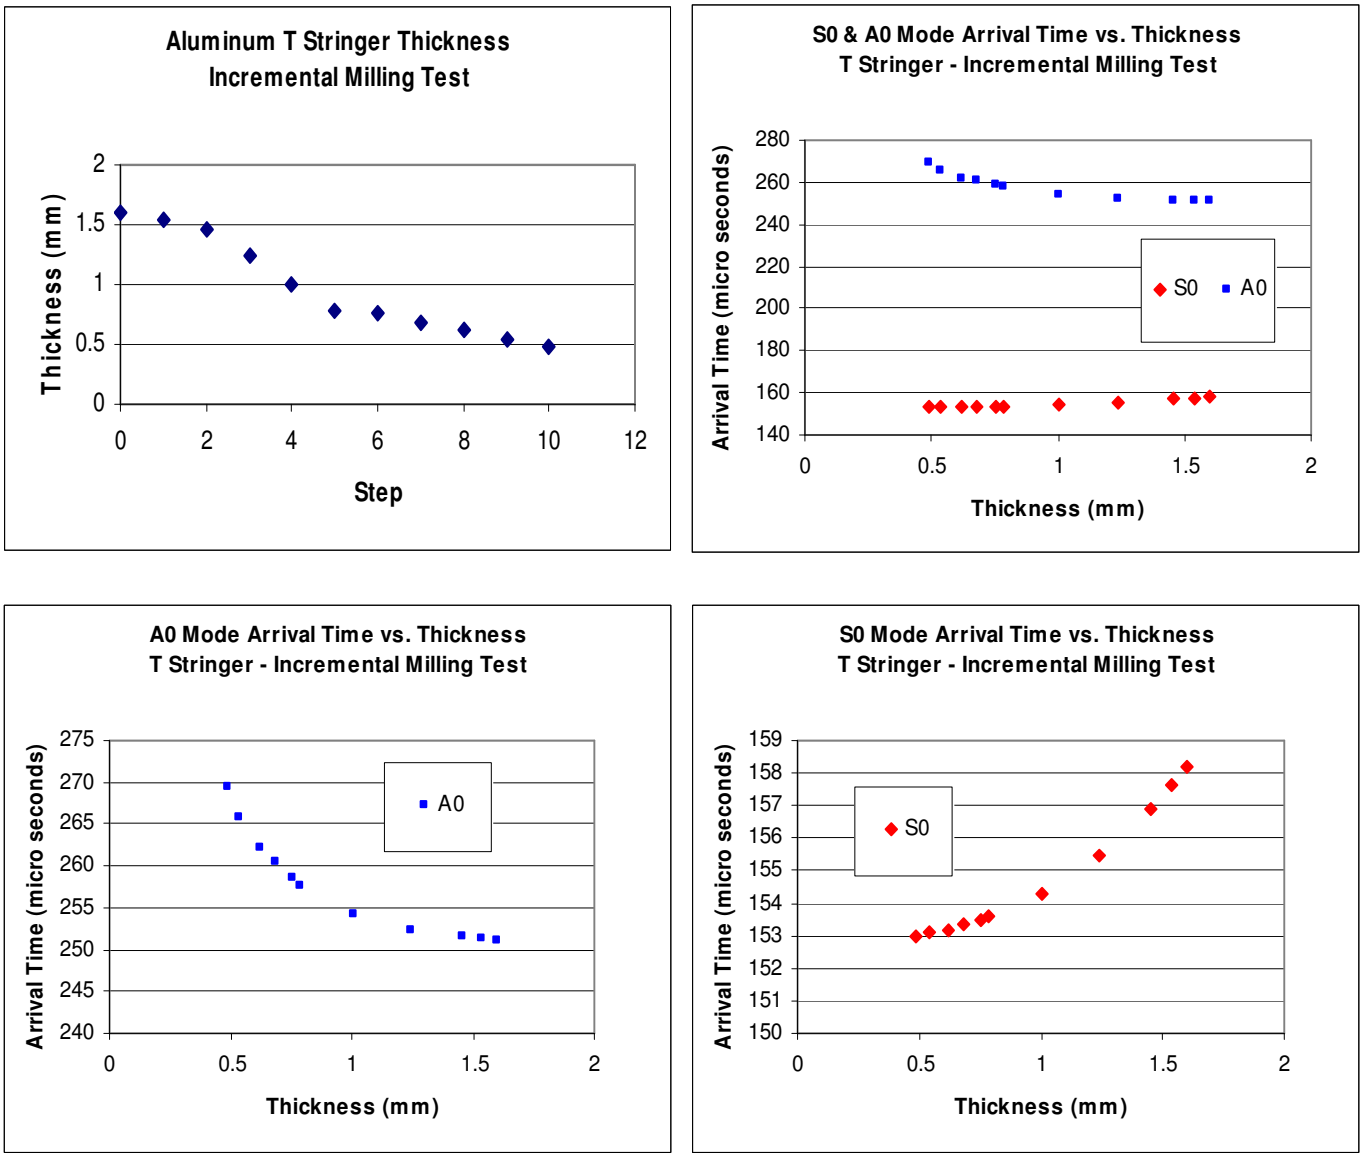

Figure 15: Thickness measurements taken on aluminum T stringer during incremental milling test (top left). Notice the decrease in thickness as the milling removes clear layers of aluminum. Resulting expected S0 and A0 arrival times (top right). Close up of expected S0 (bottom right) and A0 (bottom left) arrival times. Notice that the S0 mode speeds up from roughly 158 μs to 153 μs, a difference of 5 μs, while the A0 mode slows down from roughly 252 μs to 270 μs, a difference of 18 μs.

increased. Using the “compare thumbprints” ability of the GUI, we looked at each file’s respective fingerprints directly next to each other. This ability to compare the fingerprints directly made picking out the desired features and trends of that feature remarkably easier than looking at the raw waveforms themselves.

As can be seen in Figure 16, we were able to pick out features in the wavelet fingerprints that represent the arrival of the S0 mode when comparing steps 00, 02, 04, 06, 08, and 09 of the milling test with the NP3 SV setup with a ‘coif3’ discrete stationary wavelet transform using the ‘coif3’ mother wavelet. We know it is the S0 mode because of the location in the respective raw waveform. The red dots were added to show where the mode is expected to arrive in each fingerprint based on our earlier calculations. These dots were placed with respect to the 00 step. The red boxes show what we determined to be the first sign of the mode’s arrival in the fingerprint. As can be seen, the expected arrival position coincides with the actual arrival for the most part, illustrating our ability to use the fingerprints to track the material loss in the aluminum stringers. We confirmed this ability on the TB SV data as well, as can be seen in Figure 17.

The TB O data proved very difficult to analyze. The filter properties we determined for the un-milled stringer were not fitting well for the milled stringers, and vice-versa. We found that the signals were not constant enough for us to apply a filter and have consistent results. Because of this, we were not able to produce any fingerprint trends that could be associated with a thickness loss. We therefore did not analyze the milling data with the TB source and O transducer any further.

The NP3 O data was also hard to interpret. The raw waveforms did not filter as well as the NP3 SV waveforms, producing much more complicated fingerprint images.

We were eventually able to extract trends that followed the expected trends of the S0 mode arrival however they were again not as simple as the SV data. As can be seen in Figure 18, the trend we found was determining where the fingerprint began having more rings again. Notice that we said “more rings” without being more specific. This is not as concrete as determining where the fingerprint began having a specific number of rings, which later deterred us from further analysis using the NP3 source and O transducers. On the other hand, we were able to locate the A0 mode arrival with about the same certainty as the S0 mode arrival. While the fingerprints turned out to be messy, we did find a feature that shifted as we expected it to within the fingerprint. This can be seen in Figure 19. This is the only time we were able to find the arrival of the A0 mode with any of the NP3/TB and SV/O combinations. However, as the A0 mode arrival was not easy to locate, we did not perform any further analysis looking for the A0 mode.

Once we finished our initial examination of the waveforms and determined that the NP3 SV as well as the TB SV source/transducer combinations provided the simplest waveform fingerprints, we turned our attention to performing a more complete analysis on those sets of data. If we were to solely use the fingerprint GUI for this analysis, which can be seen in Figure 20, loading in each file every time would have taken more time than the analysis itself. In order to speed things up, we used MATLAB to write code that automated the fingerprint comparison output. The code depends on input values similar to those chosen within the wavelet fingerprint GUI, and performs the comparison without having to click a button each time to upload the files and view the waveforms. Using this automated process, we were able to compare waveforms and their respective fingerprints in a much timelier manner.

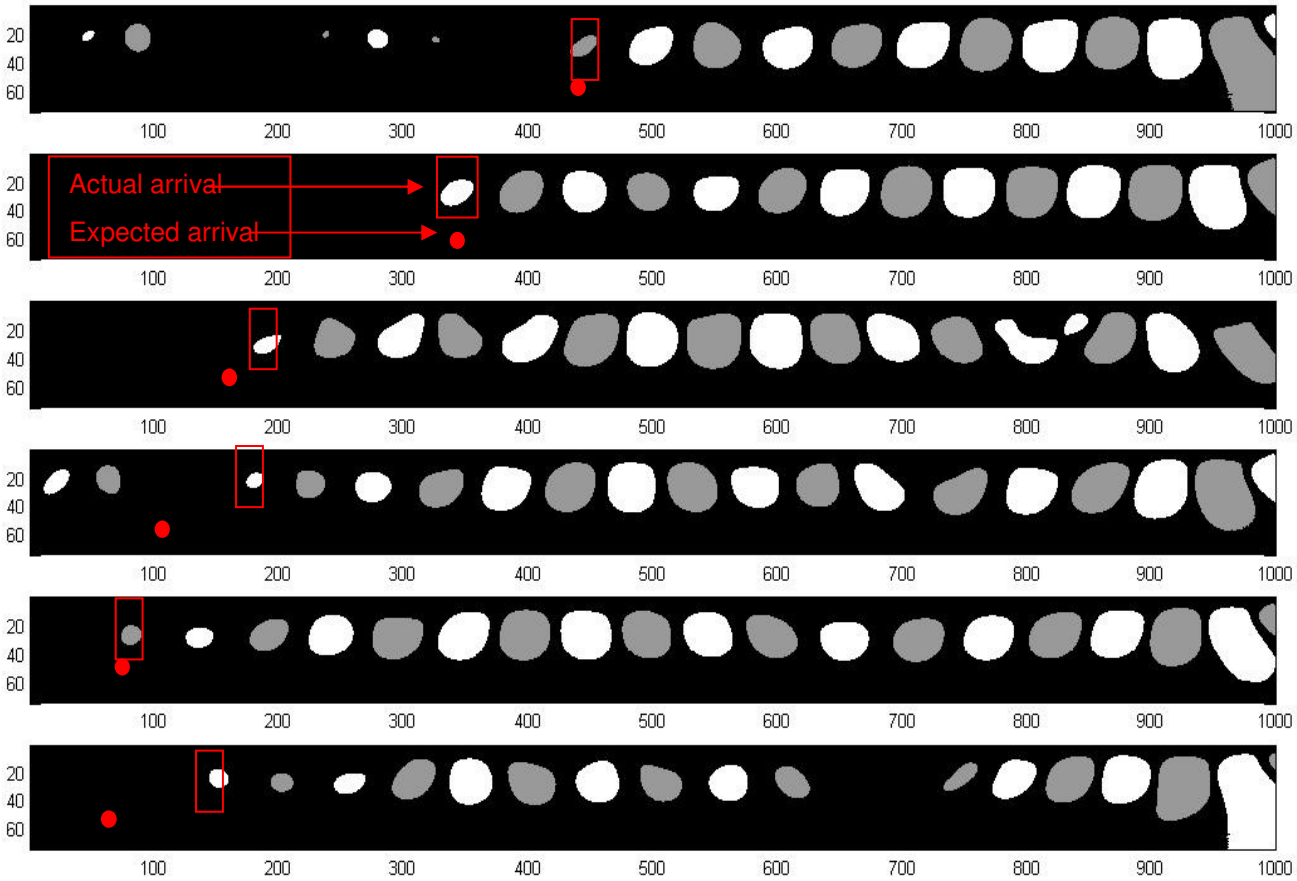

Figure 16: Showing runs 00,02,04,06,08, and 09, this was from our initial analysis on the fingerprints of the NP3 SV data. This shows the S0 mode arrival, as well as the expected shift in arrival time due to thickness loss. The red boxes show were we are considering the mode to have arrived. We have added red dots to show where the mode should (roughly) arrive based on the 00 run, and it can be seen that the two agree for the most part.

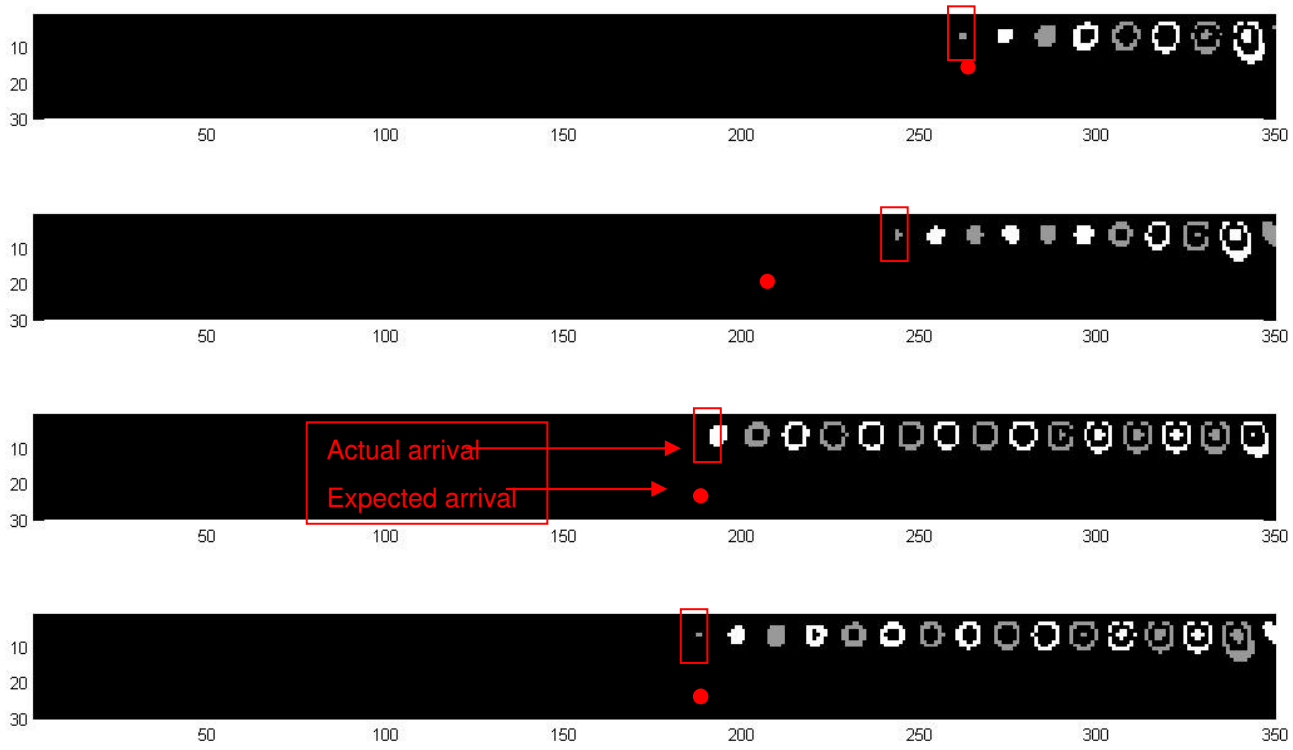

Figure 17: Showing runs 00,04,07, and 10, this was from our initial analysis on the fingerprints of the TB SV data. This shows the S0 arrival, as well as the expected shift in arrival time due to thickness loss. The red boxes show were we are considering the mode to have arrived. We have added red dots to show where the mode should (roughly) arrive based on the 00 run, and it can be seen that the two agree for the most part.

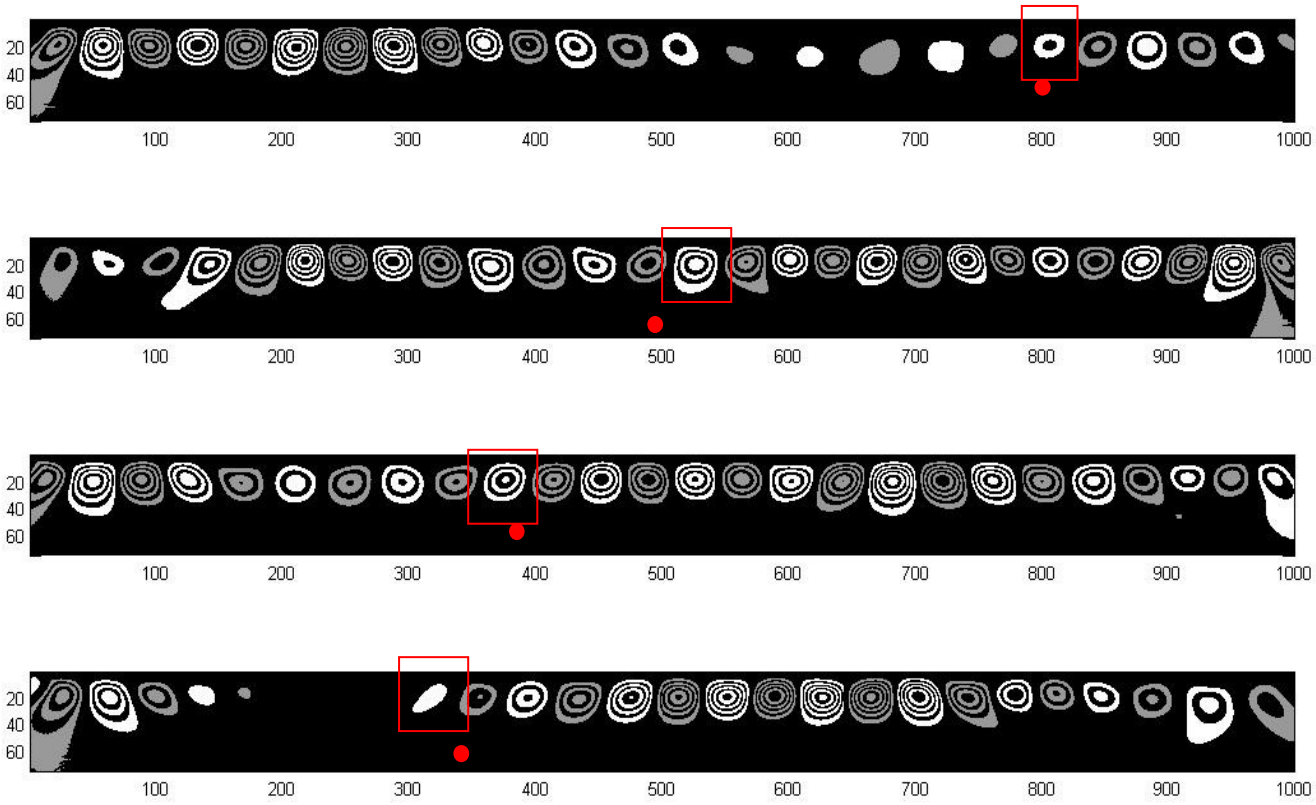

Figure 18: Showing runs 00,04,07, and 10, this was from our initial analysis on the fingerprints of the NP3 O data. This shows the S0 arrival, as well as the expected shift in arrival time due to thickness loss. The arrival on these is where the fingerprint starts to have more rings again. The red boxes show where we consider the mode to have arrived. We added red dots to show were the mode should roughly arrive based on the 00 run. It can be seen that the two agree for the most part.

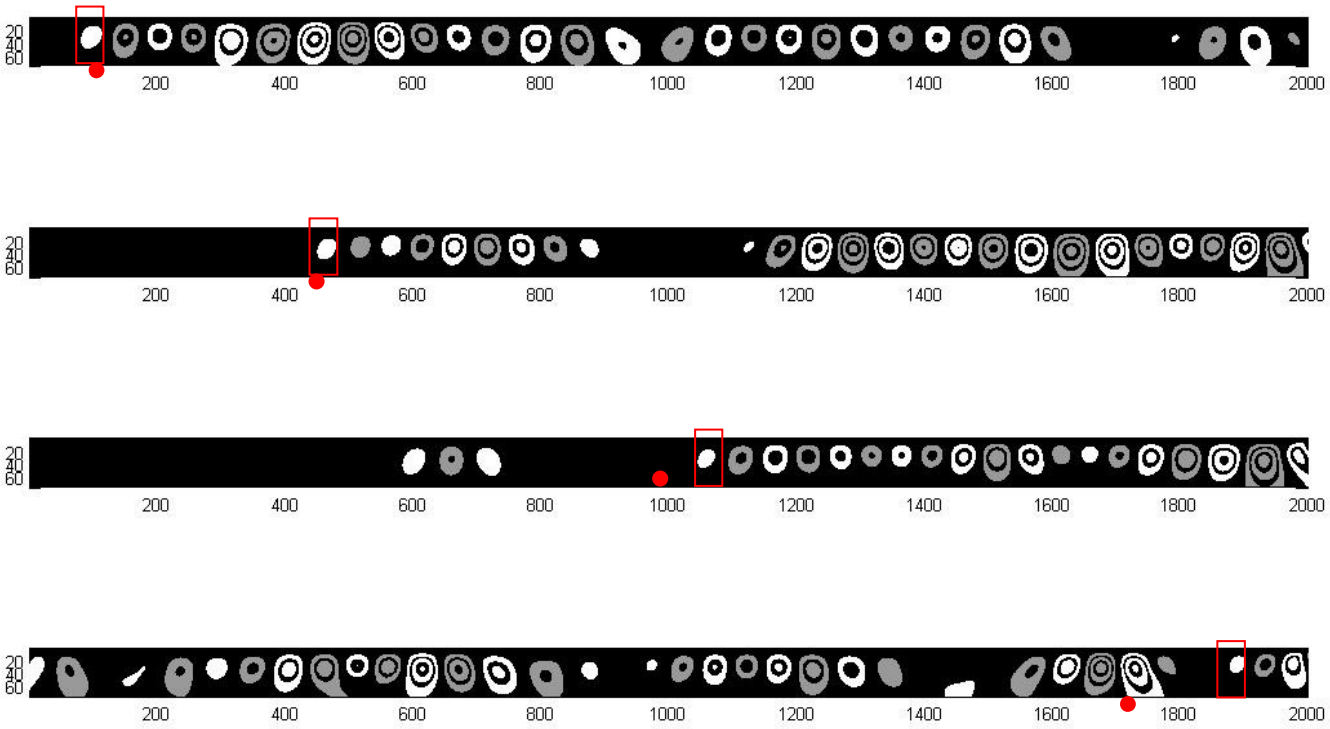

Figure 19: Showing runs 00,04,07, and 10, this was from our initial analysis on the fingerprints of the NP3 O data. This shows the A0 arrival, as well as the expected shift in arrival time due to thickness loss. The red boxes show where we consider the mode to have arrived. We added red dots to show were the mode should roughly arrive based on the 00 run. As can be seen above, the fingerprints are messy this far into the signal and it is much more difficult to pick out the A0 mode arrival.

We then looked at all of the milling steps using the NP3 SV source/transducer combination. We started with the same filtering parameters that we found optimal in our initial analysis and saw that the trend we found earlier applied to the entire set of data. Figure 21 shows the raw waveforms along with the resulting fingerprints from the filtered signal for the S0 mode. The pre-filter used here is a discrete stationary wavelet transform using the 'coif3' mother wavelet which removes several layers of detail to de-noise the signal. There is a clear shift in the arrival of the S0 mode with the material loss due to the milling increments.

We did the same in-depth look at the TB SV source/transducer combination as well. Again, we went back and used the same parameters that we used in our initial analysis of the TB SV data and applied them to the entire set. We saw the same trend apply to the entire set of data that we did initially. This can be seen in Figure 22, which compares the complicated raw waveforms with the resulting fingerprints from the filtered signal at the location of the S0 mode. We again used a discrete stationary wavelet transform using the 'coif3' mother wavelet as our pre-filter. There is again a clear shift in the arrival of the S0 mode with the material loss due to the milling increments.

Once we had verified that we could successfully see the shift in mode arrival due to material loss in the incremental milling test, we wanted to develop a method in which we could process the data, extract the arrival time of the S0 mode, and use that arrival time to determine the thickness of the stringer. For example, if we had a set of fingerprints that contained only one white point that shifted with respect to thickness, we could then write a program to go in and find the location of the one white point. As the fingerprints are binary images and can be processed digitally, one could scan a column at

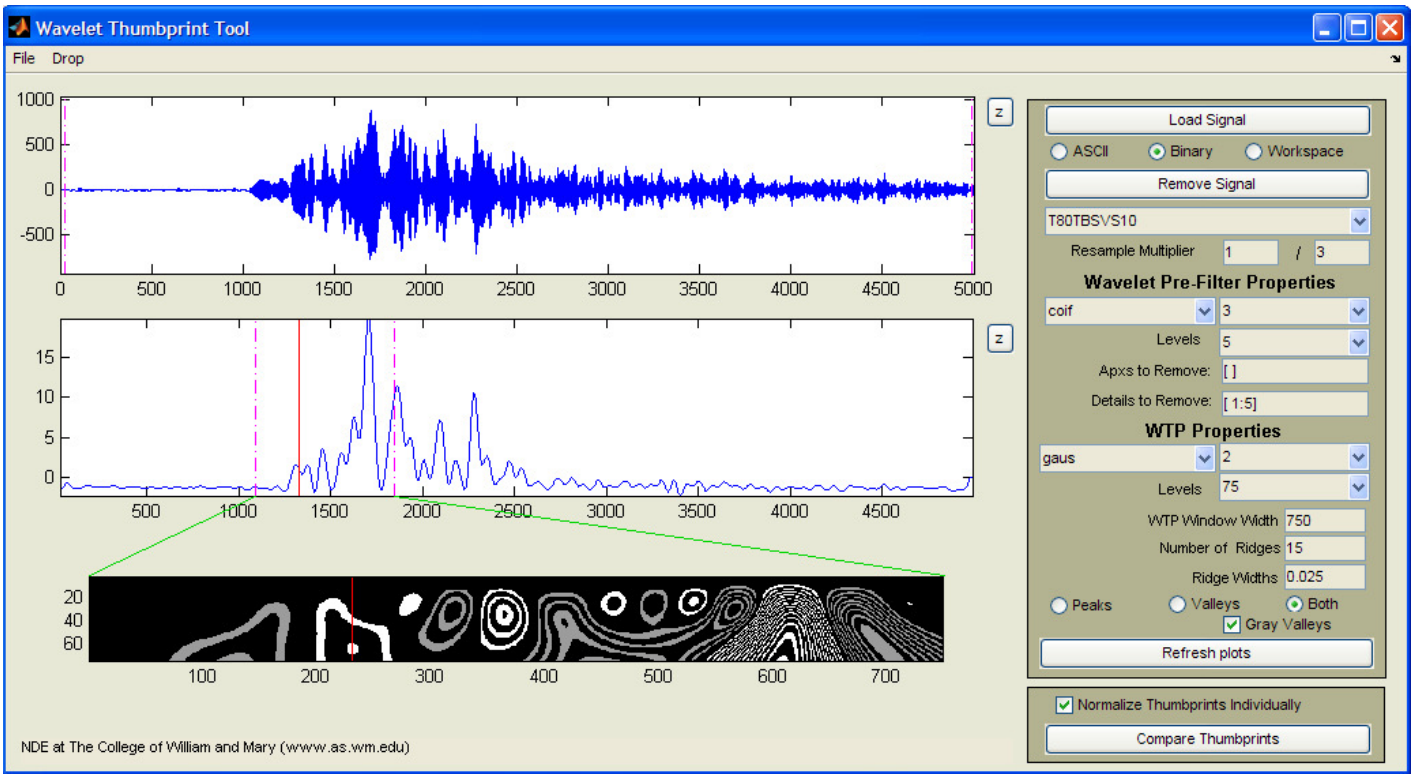

Figure 20. Wavelet fingerprinting GUI run out of MATLAB. Waveform shown is from the 1.0m "T" stringer after 10 steps of the incremental milling test. The GUI allows the user to import multiple signals at once and view them individually, picking which wavelet filters to use for analysis. This allows the user to quickly identify which wavelet basis functions and parameter settings best identify and distinguish the waveform features of interest.

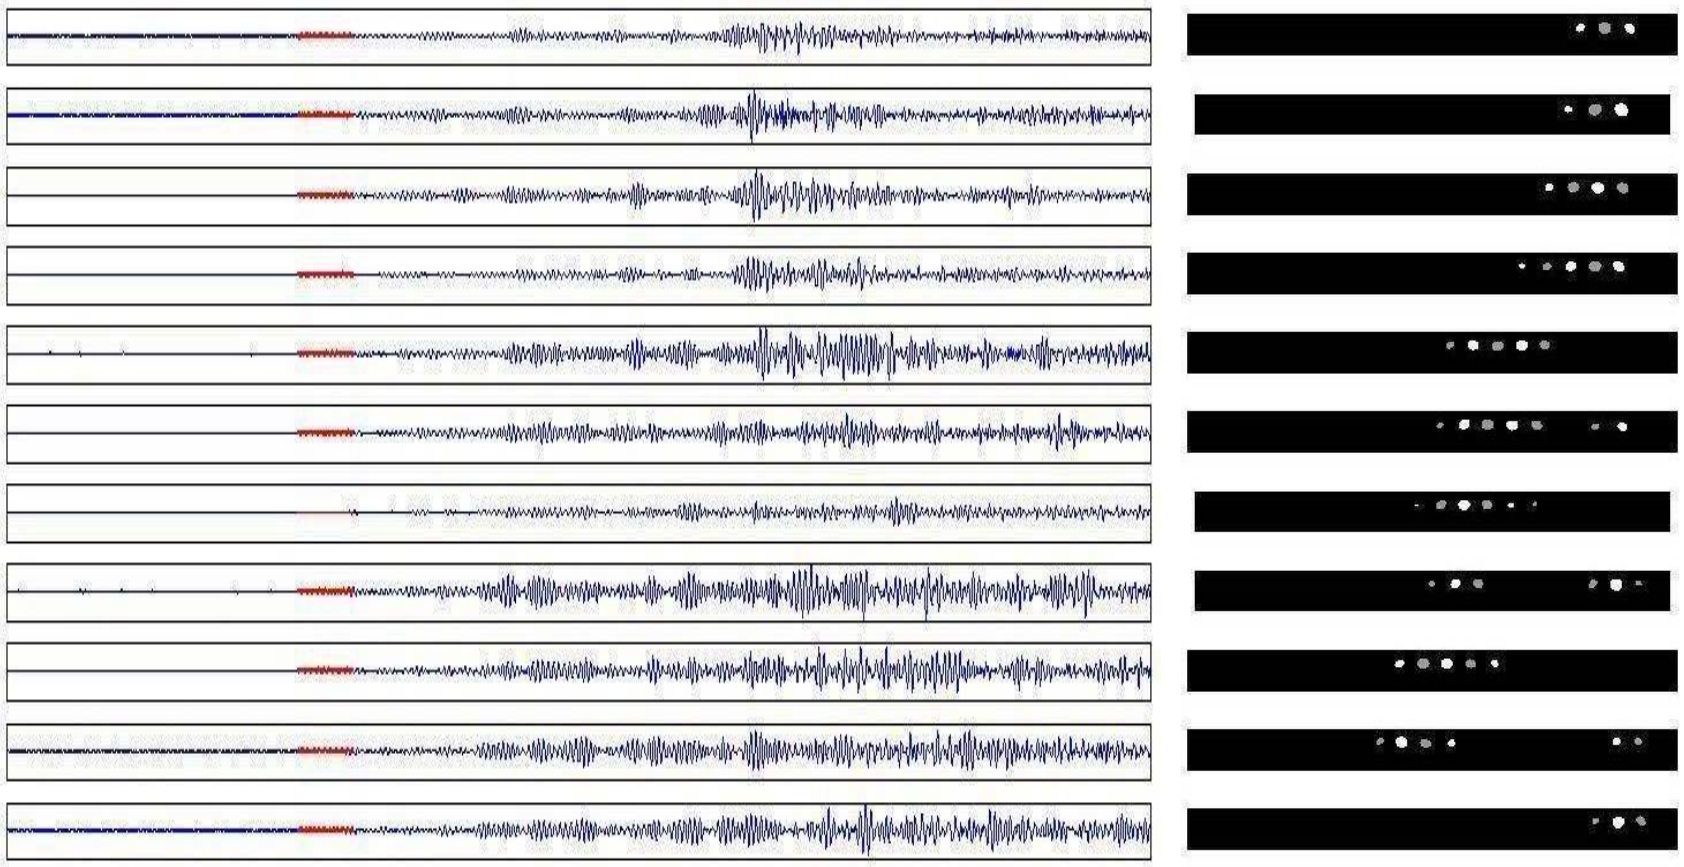

Figure 21: The raw waveforms from the NP3 SV milling T stringer (above). The red section of the waveform corresponds to the windowed filtered region that was passed to the wavelet fingerprinting algorithm. This was found using the expected arrival time of the S0 mode for the clean sample. The resulting fingerprints from the filtered window for the S0 mode (below). The pre-filter used here is a discrete stationary wavelet transform using the 'coif3' mother wavelet and it removes several layers of detail to de-noise the signal.

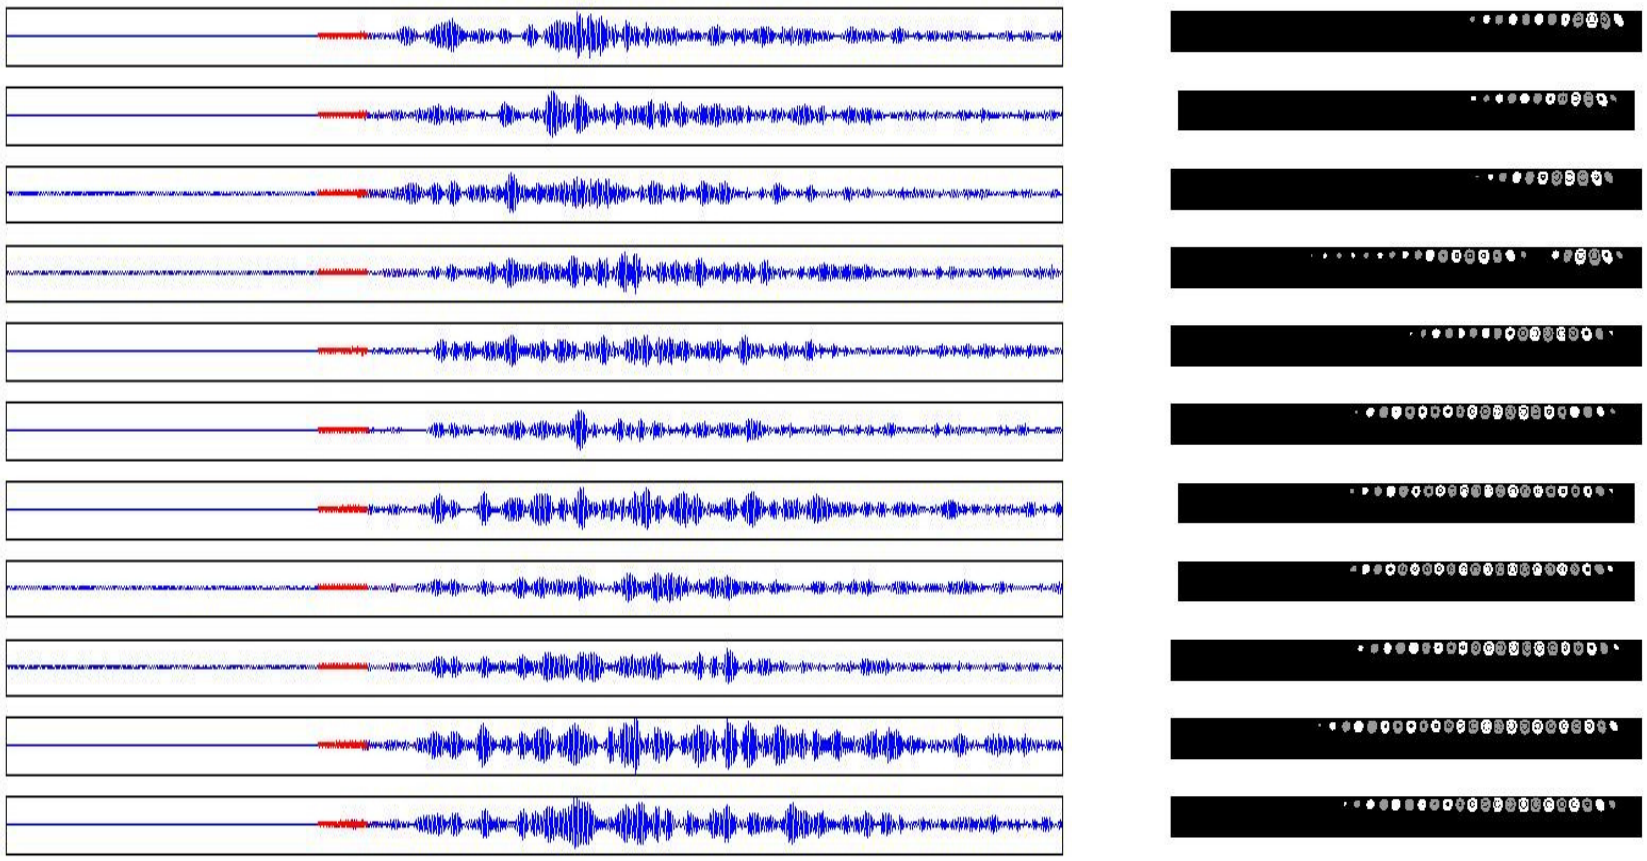

Figure 22: The raw waveforms from the TB SV milling T stringer (above). The red section of the waveform corresponds to the windowed filtered region that was passed to the wavelet fingerprinting algorithm. This was found using the expected arrival time of the S0 mode for the clean sample. The resulting fingerprints from the filtered window for the S0 mode (below). The pre-filter used here is a discrete stationary wavelet transform using the 'coif3' mother wavelet and it removes several layers of detail to de-noise the signal.

a time in order to determine where the first non-black pixel is located, as can be seen in Figure 23. It would be very straightforward to determine the arrival time of the mode using this location. This program could then be run for a series of fingerprints, and we would then have automated our analysis process.

Naturally, the next step in our analysis was to work backwards in the sense that we wanted to look at the fingerprint images, determine the resulting thickness and compare that to the measured thickness at the same step. We took the image in Figure 21 and determined where within the fingerprint each S0 mode arrival occurred. In other words, we found the relative distance into the fingerprint that the first point, gray or white, was located. Since we knew where the fingerprint was located with respect to time on the raw waveform, we then knew the overall time into the complete recorded signal of the S0 mode arrival. Using this overall time of S0 mode arrival, we then calculated the S0 mode's group velocity. Going back to the aluminum dispersion curve in Figure 1, we used the group velocities to determine corresponding frequency-thickness values for each step. As we used a frequency of 1MHz, the thickness values followed directly. Table 1 shows the steps taken in these calculations for the NP3 SV data.

We compared these calculated thickness values to the actual micrometer-measured thickness values in Figure 24. While a downward trend in our calculated thickness values can be seen, it does not fit our measured thickness values very well. The difference in thickness values increased from 0.410mm to over 1.0mm throughout the experiment. This is not what we expected to find.

We decided to run the same calculations on our TB SV data to see if we could find better results. Our calculated values can be found in Table 2. We experienced similar

results, as can be seen in Figure 25. The calculated thickness values still exhibited a downward trend however it was not consistent with our measured thickness values. The difference in thickness values increased from 0.287mm to over 1.0mm again.

To asses why we were not finding accurate thickness values, we went through our equations again and again trying to find a fault. We noticed that our expected group velocity values which were based on the measured thicknesses had a span from 5.057 mm/μs to 5.414 mm/μs, a difference of 0.357 mm/μs, while our calculated group velocities based on the fingerprint images had a span from 4.725 mm/μs to 5.027 mm/μs, a difference of .302 mm/μs. These difference values are similar however our calculated group velocity values were just shifted downward, as can be seen in Figure 26 (b). We also noticed that the S0 dispersion curve is non-linear, meaning if we had two group velocities that were a set distance apart, the range of the resulting thickness values would depend on where on the curve we were looking. This means that because our calculated group velocity values were lower than our expected group velocity values yet the two shared a similar range, we would see a higher set of thickness values with a narrower range than the measured thickness values we were looking to match.

With this, we decided to see what our results would look like if we shifted our group velocity values up to have a same initial step 00 value as the measured thickness values. A graph of this can be seen in Figure 27 (a). It is clear from this graph that both the NP3 and TB SV fingerprints generated an essentially correct range of group velocity values, with the NP3 SV fingerprints being slightly more accurate overall. A graph of the percent error of the calculated-shifted group velocity values and the expected group velocity values is shown in Figure 27 (b). The TB SV line shows an average percent error

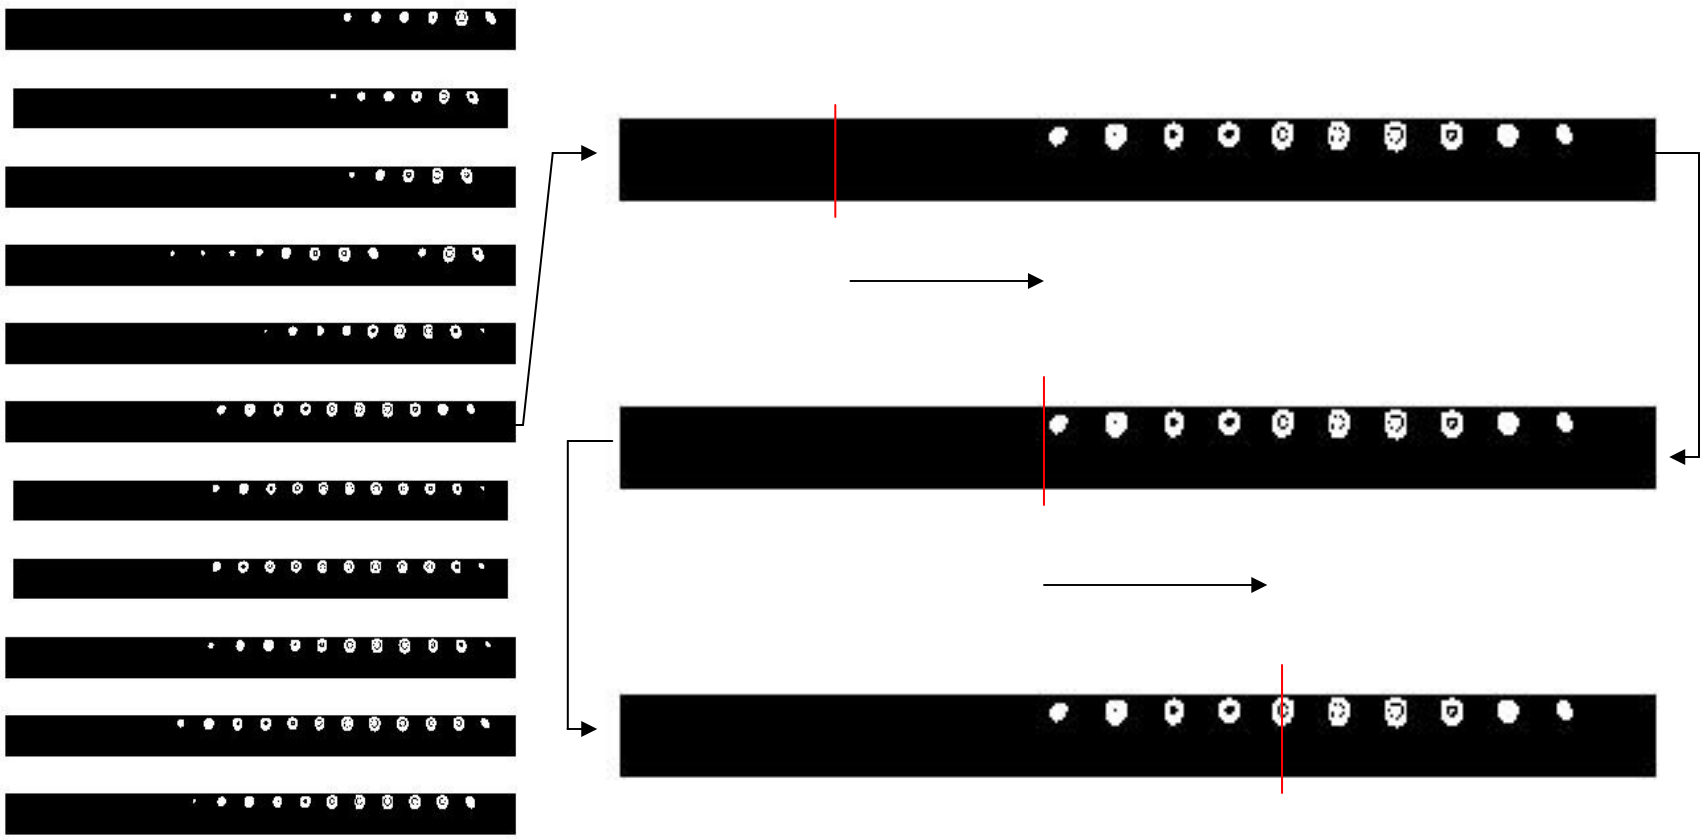

Figure 23: Possible automation of the mode extraction process. Starting with a set of fingerprints (in this case, the TB SV fingerprints of the S0 mode arrival with peaks only), the process scans each fingerprint left to right. In this scan, it is easy to determine when it reaches the first white point, i.e. the first dot or ring, based on the color of the pixels in the column of the image being scanned. The location of this point is recorded, and used to determine the thickness of stringer.

Step	Location in signal	Arrival time (μs)	Group Velocity (mm/μs)	Thickness (mm)	Measured Thickness (mm)	Difference (mm)
00	13100.33	163.75	4.725	2.008	1.598	0.410
01	13085.74	163.57	4.735	1.992	1.539	0.453
02	13038.16	162.98	4.769	1.954	1.454	0.500
03	12983.75	162.30	4.808	1.900	1.240	0.660
04	12836.08	160.45	4.917	1.777	1.005	0.772
05	12815.36	160.19	4.933	1.754	0.786	0.968
06	12770.24	159.63	4.967	1.731	0.755	0.976
07	12799.65	160.00	4.945	1.746	0.683	1.063
08	12729.87	159.12	4.998	1.677	0.621	1.056
09	12693.60	158.67	5.027	1.629	0.538	1.091
10	13134.01	164.18	4.702	2.031	0.488	1.543

Table 1: Fingerprints from the NP3 SV data (see Figure 21) were analyzed in order to derive a thickness measurement from the location of the S0 mode arrival (top). This was compared to the actual measured thickness and the difference was calculated

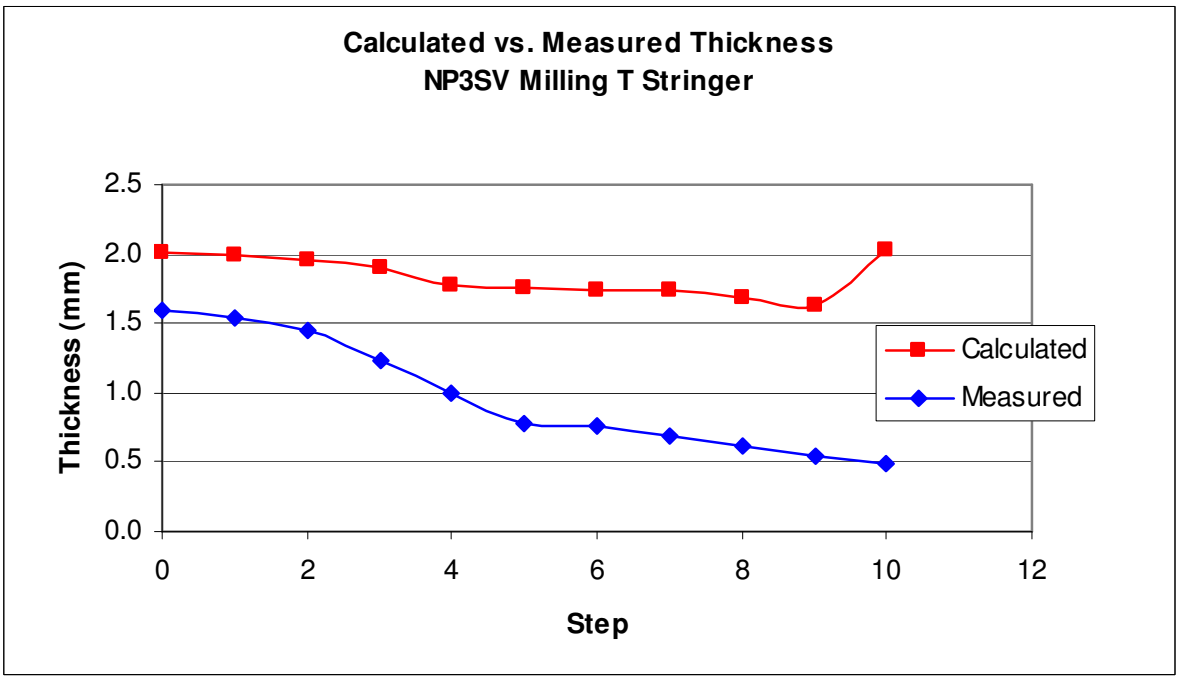

Figure 24: Thickness measured versus thickness calculated from fingerprints generated by the NP3 SV data on the T stringer of the incremental milling test at the S0 mode arrival location. The calculated thickness values do not fit the measured thickness values, and have an increasing incremental difference between calculated and measured values.

Step	Location in signal	Arrival time (μs)	Group Velocity (mm/μs)	Thickness (mm)	Measured Thickness (mm)	Difference (mm)
00	12942.31	161.78	4.838	1.885	1.598	0.287
01	12949.38	161.87	4.833	1.892	1.539	0.353
02	12952.66	161.91	4.830	1.892	1.454	0.438
03	12605.52	157.57	5.097	1.508	1.240	0.268
04	12815.36	160.19	4.932	1.762	1.005	0.757
05	12698.78	158.73	5.023	1.654	0.786	0.868
06	12684.68	158.56	5.034	1.623	0.755	0.868
07	12684.68	158.56	5.034	1.623	0.683	0.940
08	12703.96	158.80	5.020	1.639	0.621	1.018
09	12621.06	157.76	5.085	1.546	0.538	1.008
10	12675.47	158.44	5.041	1.615	0.488	1.127

Table 2: Fingerprints from the TB SV data (see Figure 22) were analyzed in order to derive a thickness measurement from the location of the S0 mode arrival (top). This was compared to the actual measured thickness and the difference was calculated

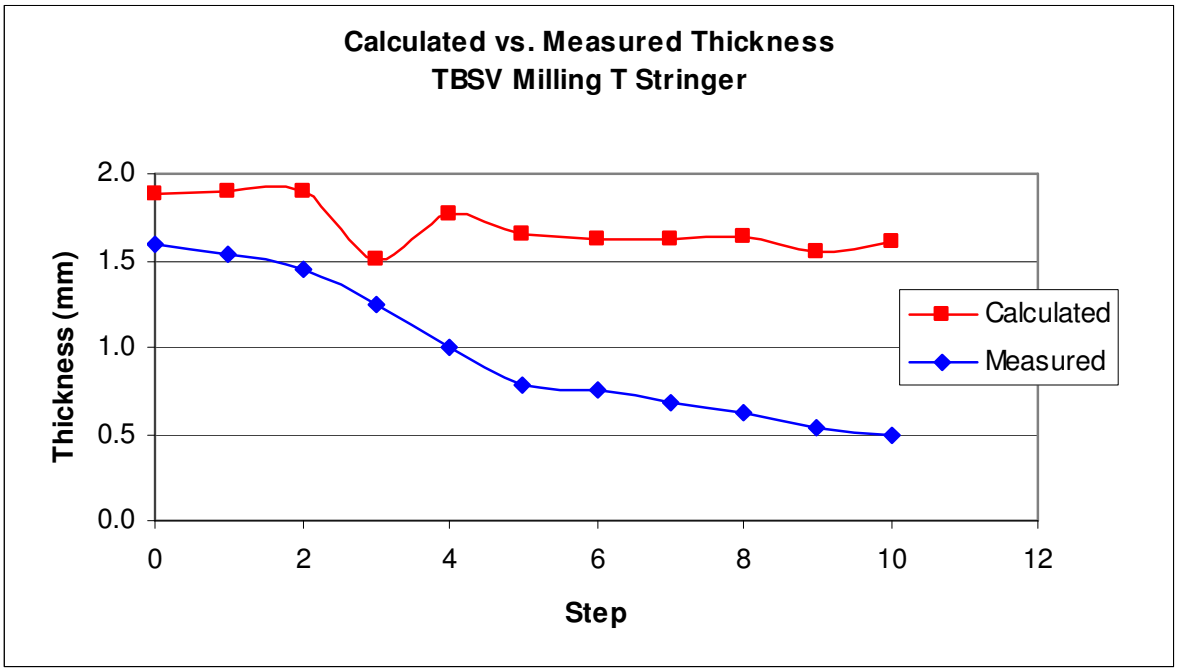

Figure 25: Thickness measured versus thickness calculated from fingerprints generated by the TB SV data on the T stringer of the incremental milling test at the S0 mode arrival location. The calculated thickness values do not fit the measured thickness values, and have an increasing incremental difference between calculated and measured values.

(a) (b)

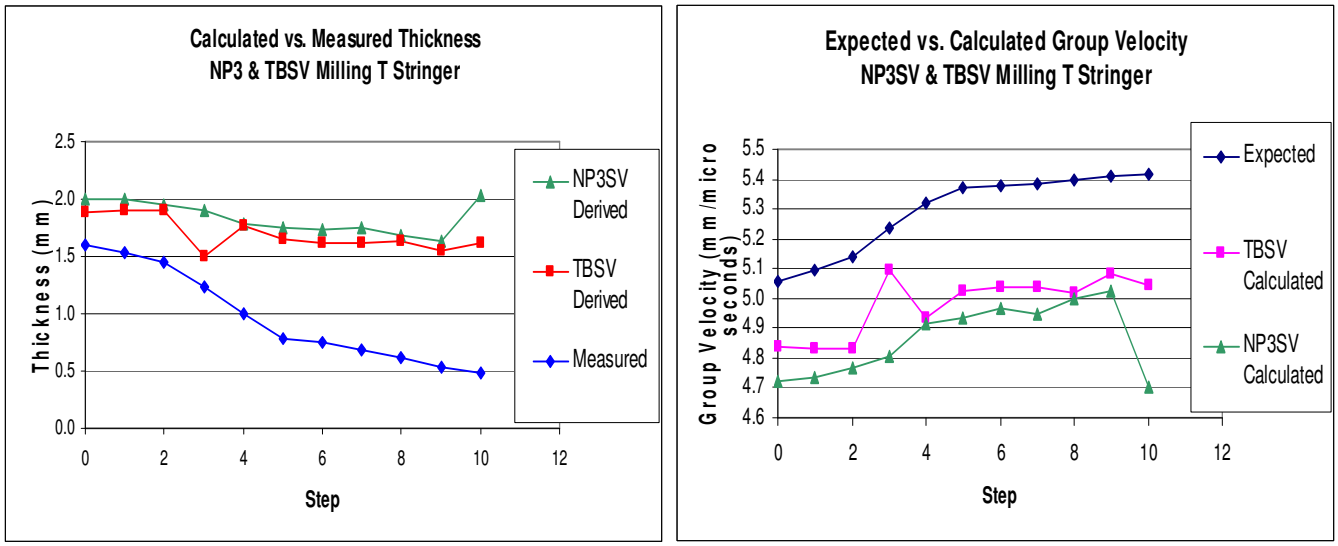

Figure 26: (a) shows the thickness calculations based on fingerprint S0 mode arrival for both the NP3 SV and TB SV data on the T stringer for the incremental milling test, as well as the actual measured thickness. (b) shows the corresponding group velocities that were used to derive the thickness measurements in graph 5. As you can see, while they are shifted downards, they both follow the trend set by expected group velocities.

(a) (b)

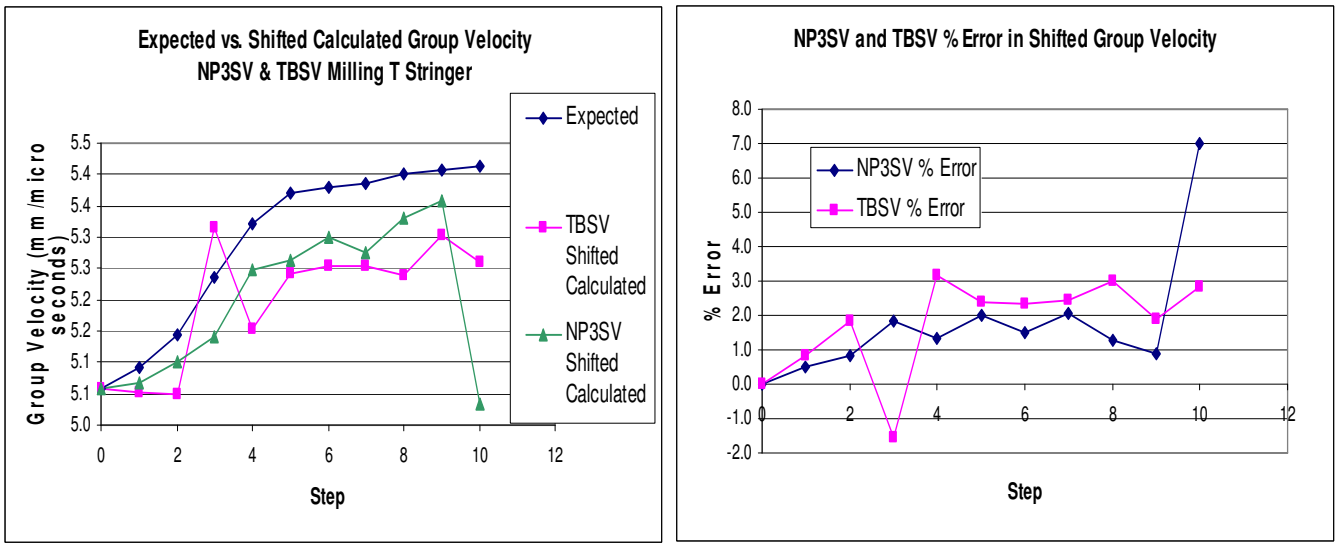

Figure 27: (a) shows the group velocities found in graph 6 shifted so they all have the same initial 00 step value. It can be seen that both the NP3 SV and the TB SV data follows the same trend set by the expected group velocities. However, the NP3 SV values tend to be closer to the expected group velocities than the TB SV values do. (b) shows the % error for the shifted group velocities for each step. Again, it can be seen that the NP3 SV data is better fitting to the expected group velocity than the TB SV data, with the NP3 SV data having most values within a w% error of the expected values.

of about 3% while the NP3 SV line shows an average percent error of about 2%. The shift that we needed to apply to the numbers could be due to a variety of things, the most likely being some sort of delay in the recorded waveform that we weren't aware of. One possible source of this could be the delay lines in the transducers themselves. The point of this is that we can assume the delay was constant throughout the experiment and counter it in the initial clean stringer values. Doing so, we find correct values of group velocity throughout the incremental milling steps which then give us accurate calculated thickness values as well.

5.3 Accelerated Corrosion Analysis

The point of performing the incremental milling test was to determine if our method of monitoring material loss in a controlled setting would provide us with accurate results. The accelerated corrosion test was performed in order to apply our results to a real-world situation. We "cleaned" the NP3 Oceana transducer files from the corrosion test in the same fashion as before, that being the manual inspection of the actual files to remove any poorly recorded waveforms.

We again wanted to know where to look for the mode arrivals within the recorded waveforms. We determined approximate S0 and A0 mode group velocities based on our measured thickness values using the aluminum dispersion curve. These results can be found in Figure 28, Figure 29, and Figure 30. The T stringer experienced an increase in external thickness due to the flaking effect produced by the exfoliation corrosion that occurred. Because corrosion removes material rather than making it thicker, there is underlying material missing that we cannot measure using a micrometer, so our expected

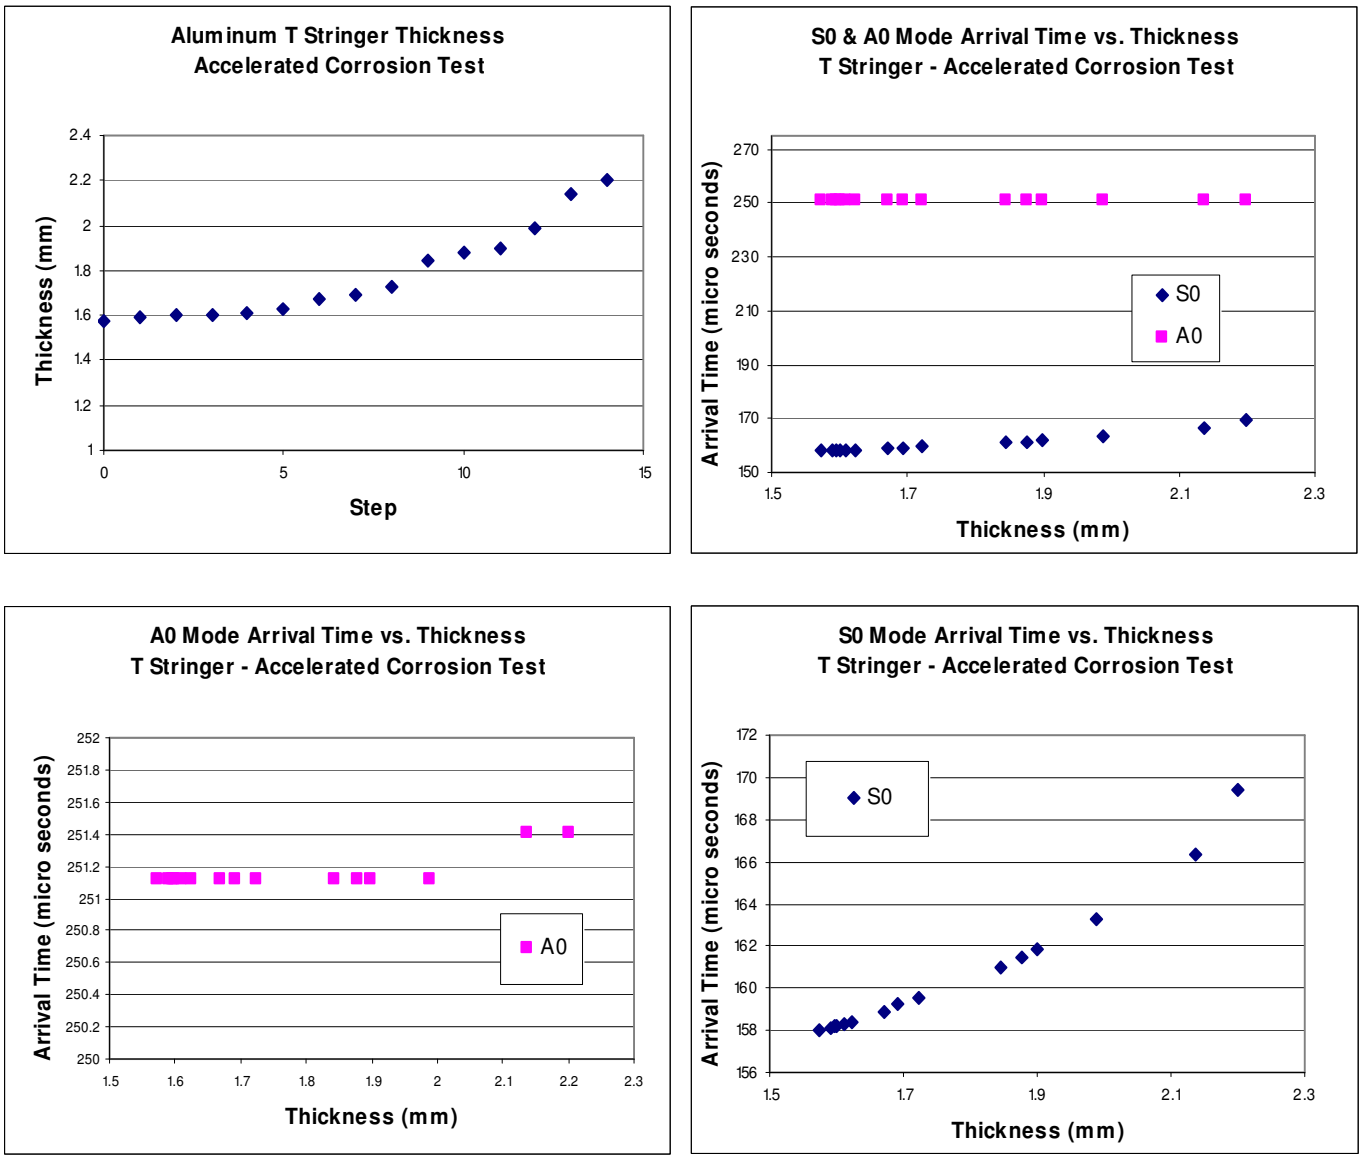

Figure 28: Thickness measurements taken on aluminum T stringer during accelerated corrosion test (top left). Notice the increase in thickness as the exfoliation corrosion forces metal layers upwards. Resulting S0 and A0 arrival times (top right). Close up of S0 (bottom right) and A0 (bottom left) arrival times. Notice that the S0 mode slows down from roughly 158 μs to 170 μs, a difference of 12 μs, while the A0 mode is relatively constant at 251 μs.

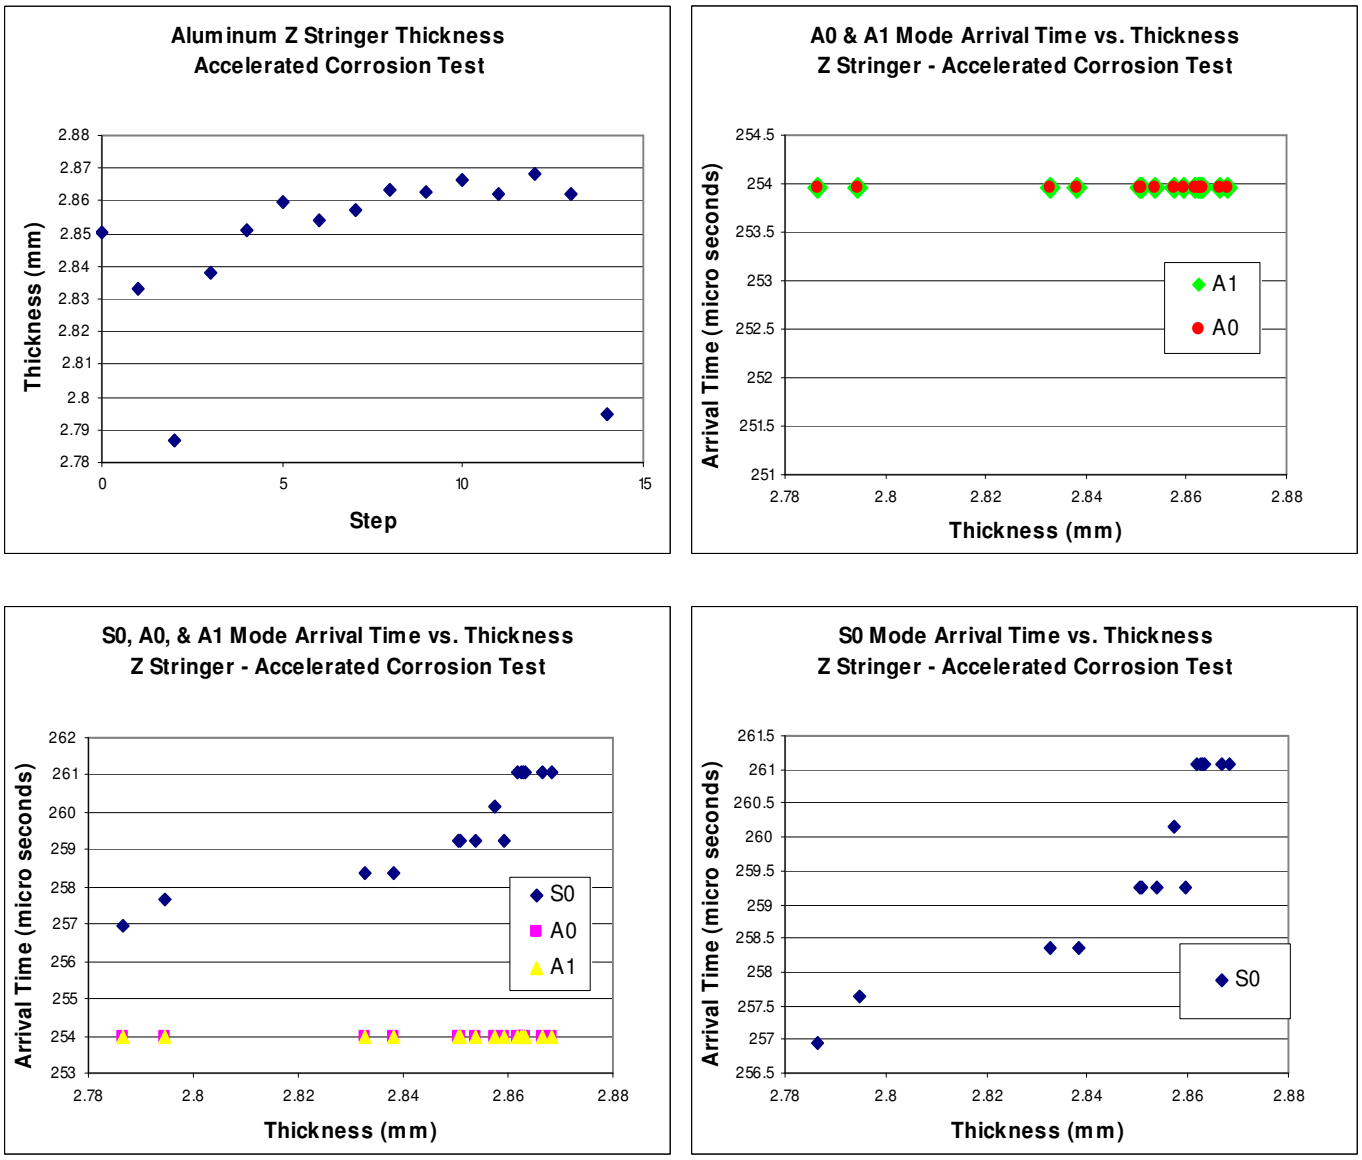

Figure 29: Thickness measurements taken on aluminum Z stringer during accelerated corrosion test (top left). There is a only slight increase in thickness which is due to the fact that pitting corrosion forms internal damage as opposed to external damage, causing our micrometer measurements to remain relatively constant. Resulting S0, A0, and A1 arrival times (top right). Close up of S0 (bottom right) and A0 & A1 (bottom left) arrival times. Notice that the S0 mode slows down from roughly 257 μs to 261 μs, a difference of 4 μs, while the A0 & A1 modes stay constant at 254 μs.

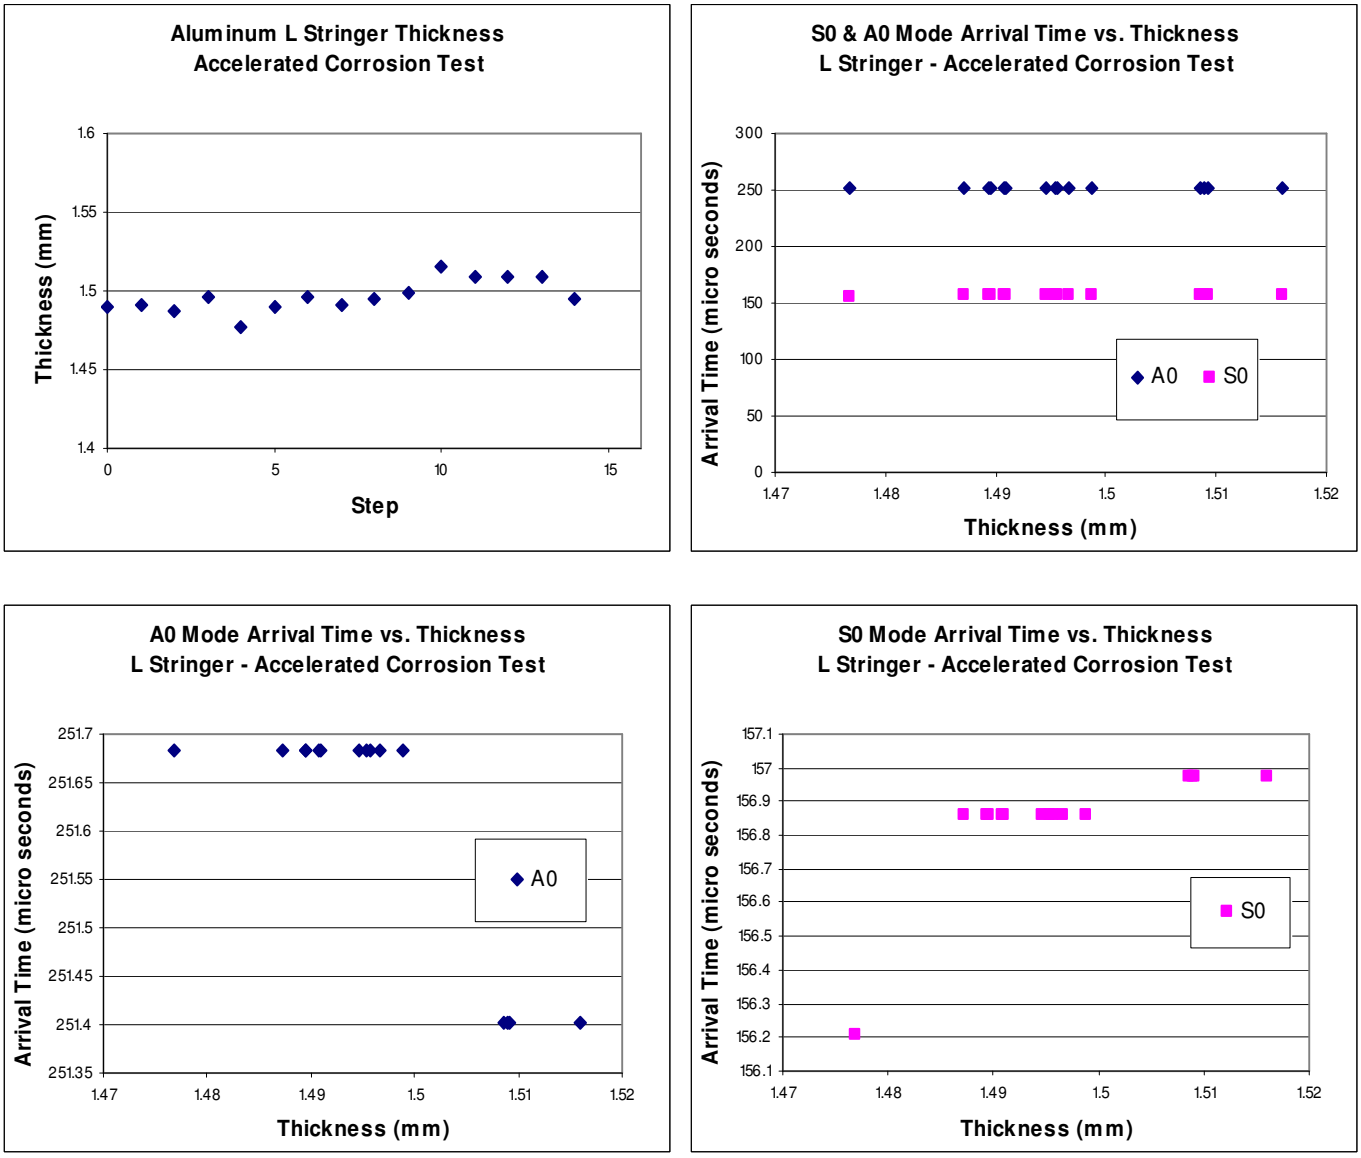

Figure 30: Thickness measurements taken on aluminum L stringer during accelerated corrosion test (top left). There is a only slight increase in thickness which is due to the fact that pitting corrosion forms internal damage as opposed to external damage, causing our micrometer measurements to remain relatively constant. Resulting S0 and A0 arrival times (top right). Close up of S0 (bottom right) and A0 (bottom left) arrival times. Notice that the S0 mode slows down only slightly from 156 μs to 157 μs, a difference of 1 μs, while the A0 mode stay constant at 251 μs.

mode arrival times based on these external thickness measurements are incorrect. Also, the pitting corrosion that took place on both the Z and L stringers is an internal form of material loss, causing the external thickness values to seem essentially constant. This inconsistence between external thickness and material loss will again skew our expected mode arrival times for both of these stringers. One thing to note here is that while we were computing the expected arrival times, we found that the S0, A0 and A1 mode should show up at the same time in the clean Z stringer's waveforms. This is due to the fact that the Z string had a flange that was 2.8mm thick, which on the dispersion curve is right at the point where the S0, A0, and A1 curve all intersect.

One way to overcome this discrepancy between external thickness measurements and actual material loss is to use another method of measuring thickness. We scanned the corroded T stringer in an ultrasound immersion tank, as can be seen in Figure 31. By doing so, ultrasound passes through the back side of the metal and reflects from the bottom of the corroded surface as opposed to the top of the corroded surface which we were using our micrometer to measure. A computer scans point-by-point over the corroded area and records the ultrasonic waveforms. MATLAB is then used to extract the time-delay of the reflection from the corrosion, which is mapped out in Figure 32. This technique gives us a much more accurate measurement of the actual material loss due to the corrosion.

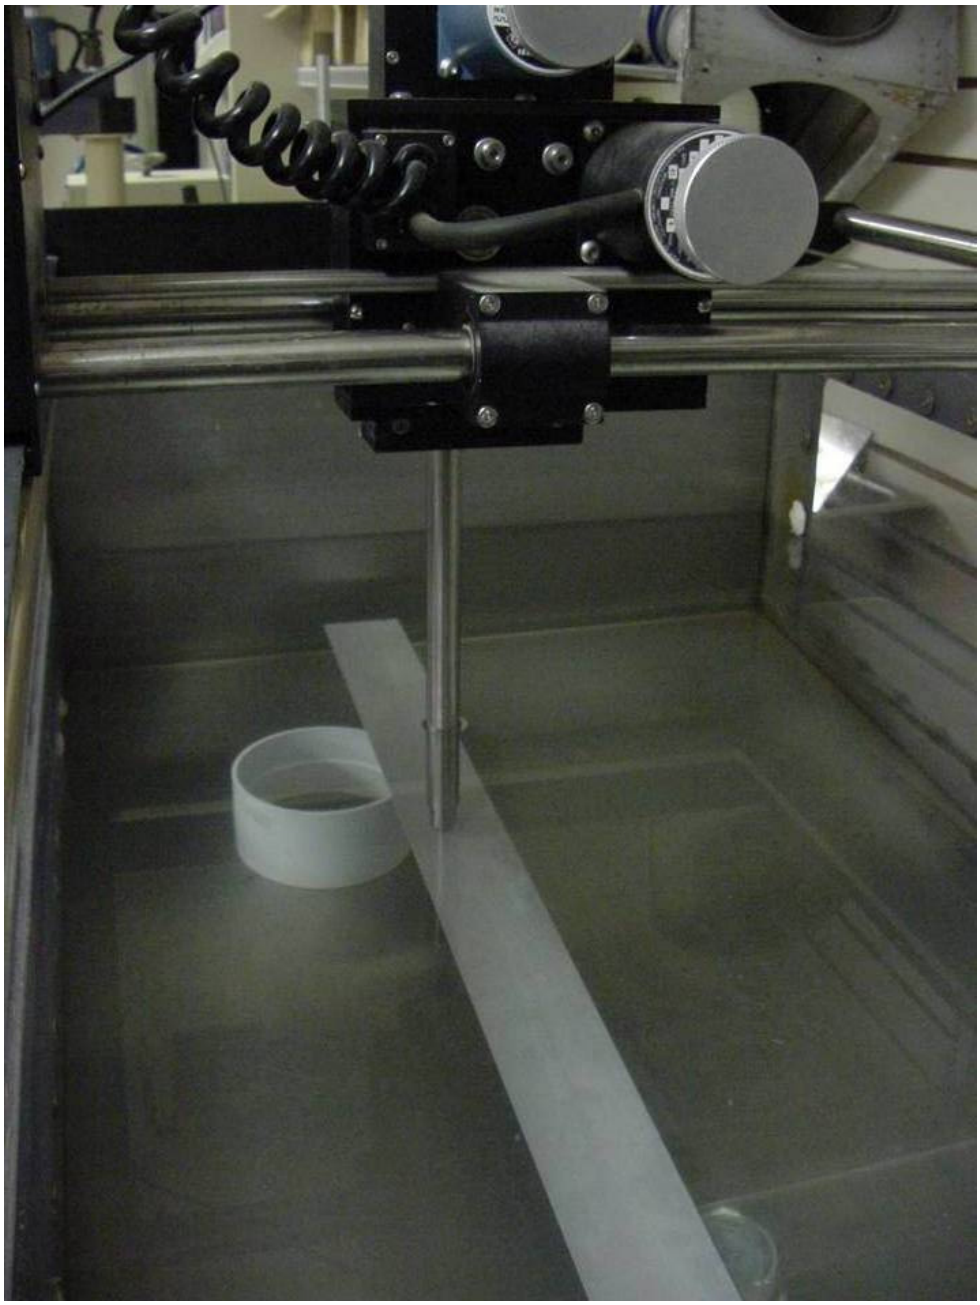

Figure 31: Corroded stringer shown being scanned in an ultrasound immersion tank. Ultrasound passes through the back side of the metal and reflects from the bottom of the corroded surface. A computer scans point-by-point over the area and records the ultrasonic waveforms. MATLAB is used to extract the time-delay of the reflection from the corrosion.

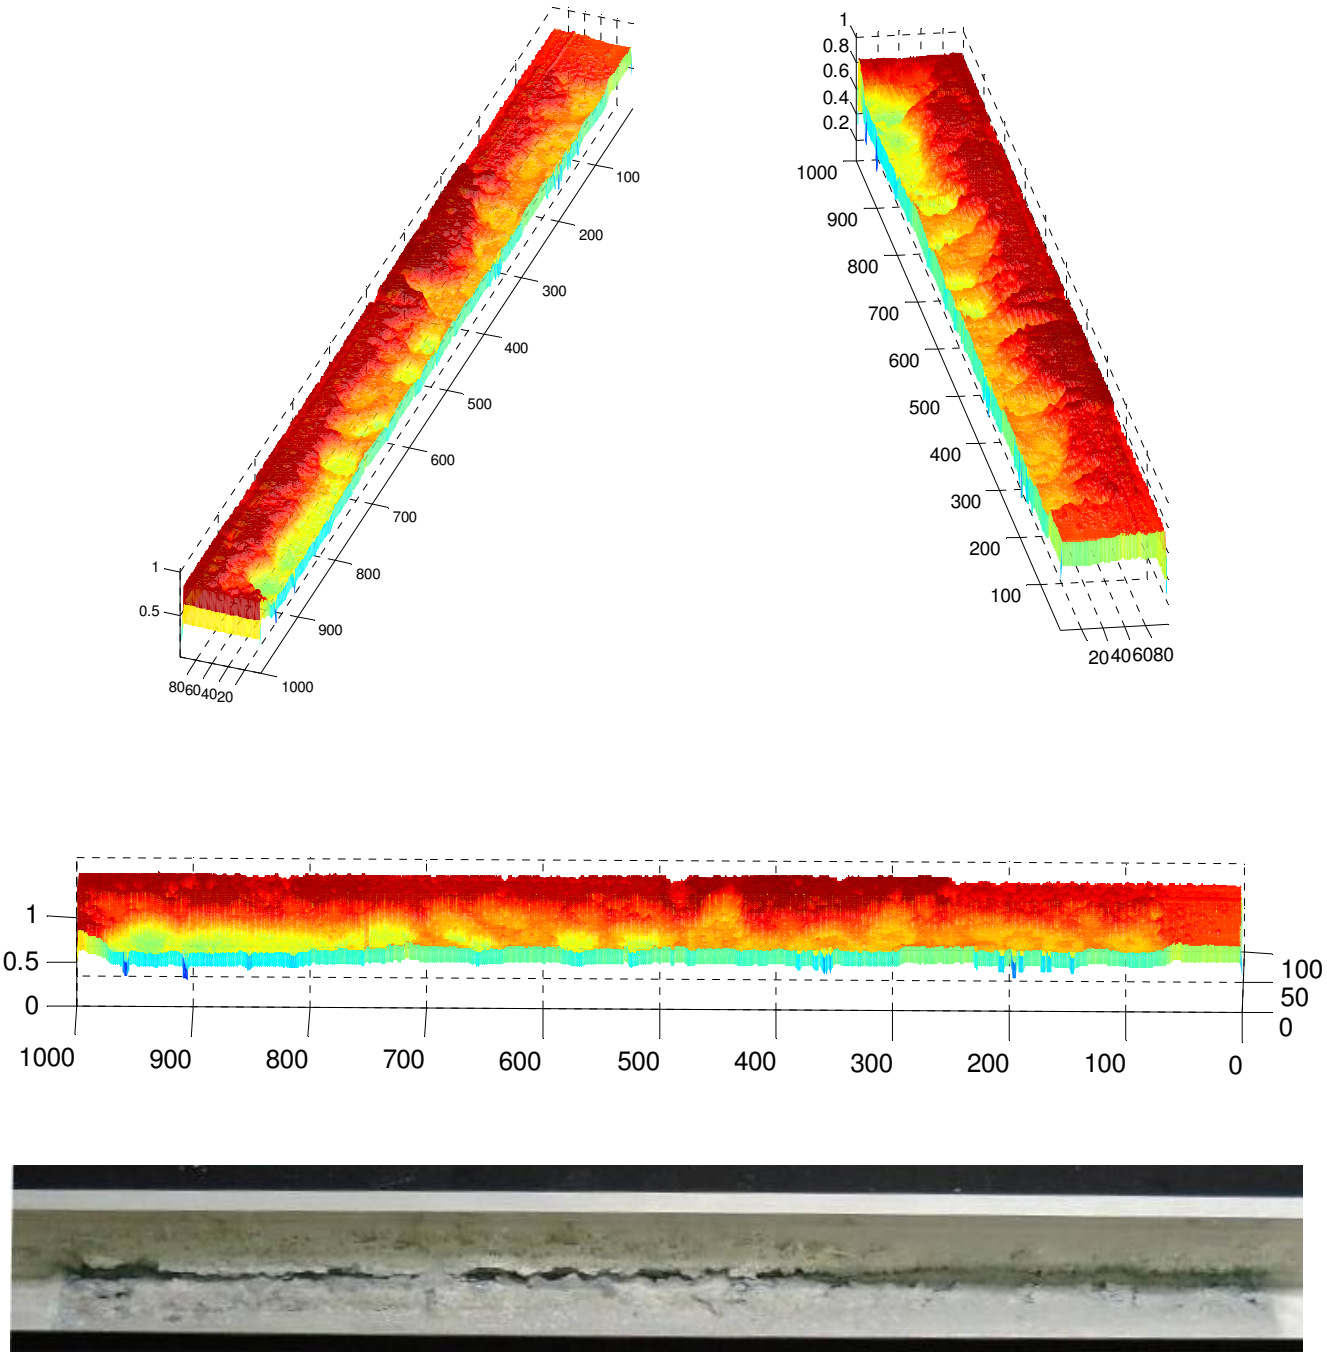

Figure 32: Images produced by MATLAB of the corroded T stringer (top left, top right, and middle). The stringer was scanned in an ultrasound immersion tank, and MATLAB used the point-by-point measurements to construct these images. As can been seen in the comparison between the scanned image and an actual photo of the stringer (middle and bottom), using this ultrasound scanning technique gives us a much more accurate picture of the actual material loss due to corrosion than visually inspecting the stringer.

6. Future Work

The next step in this project is to perform the manual analysis of the corroded T, Z, and L stringers as was performed on the milled T stringer. The Z and L stringers from the accelerated corrosion test also need to be scanned in the ultrasound immersion tank to provide us with accurate models of the material loss within the stringers. Using the scanned stringer images, accurate thickness measurements can be extracted that will provide for better expected arrival modes than the micrometer measurements we recorded. Using these expected arrival modes, the corresponding fingerprint images need to be viewed to determine if features can be found relating to the respective mode arrivals. We expect that similar trends in the fingerprint images should be found, although they will most likely be messier than the controlled milling fingerprints. The two different types of corrosion we experienced in our experiment due to the different alloys our stringers were made of will be a factor to keep in mind during that analysis.

We are currently finishing development on computer code which will provide us with an alternate form of analysis. Using large parallel computers, we can simulate the elastic wave propagation with the finite integration technique. This technique using computational models will be faster and easier to setup than their experimental counterparts. There is also potential to use the scanned images from the ultrasound immersion tank to replicate the extent of stringer corrosion in a 3D computer model. If successful, we should be able to replicate the results we found in our manual analysis of the experimentally corroded stringers. This technology would allow us to be able to systematically examine a large number of scenarios with computer simulations in order to develop a robust corrosion detection technique.

Ultimately, we hope to develop algorithms capable of automatically detecting trends in the waveform fingerprints that correspond to flaws in the aluminum stringers. Using such algorithms, a system of sensors could be developed which monitor the structural health of a material in real-time. This system would then analyze and interpret the complex and ambiguous signals, making the need for trained technicians obsolete.

7. Conclusions

The work done throughout this project shows the ability of the dynamic wavelet fingerprinting technique in determining the propagation and arrivals of multiple guided wave modes. Testing a variety of source/transducer combinations, we found that the shear (SV) transducer produced the best waveforms when paired with the nanopulser np3 as well as the toneburst (TB) sources. The NP3 SV and TB SV fingerprints produced obvious shifts in S0 mode arrival representing a loss in material, and were able to produce S0 group velocity values within 2% and 4% of the measured values, respectively. Since the wavelet fingerprints are binary images, they are easy to store and straightforward to process digitally. By developing computer algorithms which use fingerprint recognition to automatically identify features within the fingerprint images, this technology can be applied in the inspection and structural health monitoring of aluminum aircraft pieces.

References

[1] *Standard Test Method for Exfoliation Corrosion Susceptibility in 2XXX and 7XXX Series Aluminum Allows (EXCO Test),* ASTM International, Designation: G 34 – 01

[2] *Guided Wave Interpretation for Integrated Vehicle Health Management Sensors* (J. Bingham,) Nondestructive Evaluation Laboratory, Department of Applied Science, The College of William and Mary

[3] *Lamb Wave Tomography for Monitoring Aircraft Structural Integrity* (M. Hinders, E. Malyarenko and J. Heyman), in USAF Aircraft Structural Integrity Program Conference Proceedings, 2001.

[4] *Dynamic Wavelet Fingerprint Identification of Ultrasound Signals* (M. Hinders, J. Hou) Materials Evaluation 60, #9 1089-1093, (2002).

[5] *Ultrasonic Lamb Wave Tomography* (M. Hinders, K Leonard and E. Malyarenko) Inverse Problems Special NDE Issue, 18, #6, 1795-1808 (2002).

[6] *Automatic Multi-mode Lamb Wave Arrival Time Extraction for Improved Tomographic Reconstruction* (M. Hinders, J. Hou, Kevin R. Leonard) Inverse Problems 20, 1873-1888 (2004).

[7] *Multi-mode Lamb Wave Tomography with Arrival Time Sorting* (M. Hinders, Kevin R. Leonard) JASA, Vol. 117, #4, 2028-2038 (2005).

[8] Rose, Joseph L. "Dispersion Curves in Guided Wave Testing." Back to Basics. Jan 2003. The American Society for Nondestructive Testing. 10 Apr 2008. <http://www.asnt.org/publications/materialseval/basics/jan03basics/jan03basics.htm>