

**The Autonomous Collection of Highly Nonlinear Solitary Waves for Corrosion Detection**

by

**Ritesh Misra**

B.S. in Computer Engineering, University of Pittsburgh, 2018

Submitted to the Graduate Faculty of the  
Swanson School of Engineering in partial fulfillment  
of the requirements for the degree of  
Master of Science in Electrical and Computer Engineering

University of Pittsburgh

2021

UNIVERSITY OF PITTSBURGH

SWANSON SCHOOL OF ENGINEERING

This thesis was presented

by

**Ritesh Misra**

It was defended on

July 22, 2021

and approved by

Jingtong Hu, PhD, Associate Professor, Electrical and Computer Engineering

Zhi-Hong Mao, PhD, Professor, Electrical and Computer Engineering

Thesis Advisor: Samuel Dickerson, PhD, Associate Professor, Electrical and Computer Engineering

Thesis Co-Advisor: Piervincenzo Rizzo, PhD, Professor, Civil Engineering

Copyright © by Ritesh Misra

2021

# **Autonomous Collection of Highly Nonlinear Solitary Waves for Corrosion Detection**

Ritesh Misra, M.S.

University of Pittsburgh, 2021

The aging of existing infrastructure calls for the capability to monitor structures to find weakening components and prevent life-threatening failures before they happen. The fields of non-destructive evaluation (NDE) and structural health monitoring (SHM) are dedicated to investigating sensor technologies and corresponding solutions that can satisfy this need. The use of highly nonlinear solitary waves (HNSW) has emerged within the fields of NDE/SHM as a cost-effective technique to monitor a variety of structures and materials. HNSWs are mechanical waves that can form and travel in highly nonlinear systems, such as granular particles in Hertzian contact. While this technique has frequently been investigated to inspect and characterize materials, the use of this kind of sensor in field applications where portability is required has yet to be explored. Apart from conventional civil engineering inspection scenarios, practitioners within the field of robotics have explored the construction of robots specifically meant to inspect pipes through the usage of specialized sensors for NDE, such as ultrasonic transducers.

This thesis proposes the design of portable, low-cost, low-power wireless HNSW devices with the ability operate autonomously for long periods without human intervention and classify materials with the same level of accuracy as expensive, specialized laboratory equipment systems. This thesis also explores the application of portable HNSW sensors to field robotics for NDE/SHM and presents a novel robot design combining the technologies.

The designs of this thesis are validated via a series of laboratory experiments carried out on a variety of materials. The HNSW equipped robot is shown to be able to traverse metallic surfaces and detect simulated defects of varying depths. A mathematical model was developed that

allows for the mapping of features from HNSW sensor waveforms to material thickness measurements. Results show that incorporation of the model into the design of the robot greatly improves the performance of HNSW sensing and allows for the robot to successfully detect outliers in the thickness of the plate.

# Table of Contents

|   |            |
|---|------------|
| <b>Preface.....</b>                                       | <b>xii</b> |
| <b>1.0 Introduction.....</b>                              | <b>1</b>   |
| <b>1.1 Background of HNSW's.....</b>                      | <b>2</b>   |
| <b>1.2 Previous Work on Pipe Robots .....</b>             | <b>4</b>   |
| <b>1.3 Research Contributions .....</b>                   | <b>5</b>   |
| <b>1.4 Organization of Thesis .....</b>                   | <b>6</b>   |
| <b>2.0 Hardware.....</b>                                  | <b>8</b>   |
| <b>2.1 Previous Work .....</b>                            | <b>8</b>   |
| <b>2.1.1 First Generation Wireless HNSW Device .....</b>  | <b>8</b>   |
| <b>2.1.2 Results .....</b>                                | <b>10</b>  |
| <b>2.1.3 Second Generation Wireless HNSW Device .....</b> | <b>13</b>  |
| <b>2.1.3.1 Signal Conditioning .....</b>                  | <b>14</b>  |
| <b>2.1.3.2 External ADC Operation .....</b>               | <b>15</b>  |
| <b>2.1.3.3 Power Management .....</b>                     | <b>16</b>  |
| <b>2.1.4 Experimental Setup and Results.....</b>          | <b>17</b>  |
| <b>2.2 Construction of Custom Sensor .....</b>            | <b>19</b>  |
| <b>2.3 Robot Hardware .....</b>                           | <b>22</b>  |
| <b>2.3.1 Printed Circuit Board.....</b>                   | <b>24</b>  |
| <b>2.3.2 Microcontroller .....</b>                        | <b>25</b>  |
| <b>2.3.3 Motion Hardware.....</b>                         | <b>26</b>  |
| <b>2.3.4 Sensors for Position Estimation .....</b>        | <b>27</b>  |

|                                      |    |
|--------------------------------------|----|
| 2.3.5 Wireless Communication .....   | 27 |
| 3.0 Software .....                   | 28 |
| 3.1 Control.....                     | 28 |
| 3.2 Estimation .....                 | 35 |
| 3.3 Waveform Collection.....         | 37 |
| 4.0 Experimental Methods .....       | 40 |
| 4.1 Experimental Setup .....         | 46 |
| 5.0 Results .....                    | 48 |
| 6.0 Future Work and Conclusions..... | 53 |
| 6.1 Future Work .....                | 53 |
| 6.2 Conclusions .....                | 54 |
| Bibliography .....                   | 55 |

## List of Tables

|   |           |
|---|-----------|
| <b>Table 1 - Table of results from week-long test .....</b>   | <b>18</b> |
| <b>Table 2 - The intended depths alongside the emasured range of depths. A,B,C,D from Figure<br/>14 correspond to rows 1,2,3,4 in thsi table.....</b> | <b>41</b> |
| <b>Table 3 - N-Way ANOVA table from dataset. ....</b>   | <b>42</b> |
| <b>Table 4 - Average ToF/RefA for different depths of corrosion .....</b>   | <b>44</b> |

## List of Figures

|  |           |
|--|-----------|
| <b>Figure 1 - A basic schematic of an HNSW sensor .....</b>  | <b>2</b>  |
| <b>Figure 2 - A sample HNSW from a steel plate. The first highlighted peak is the primary wave (PW), and the second highlighted peak is the secondary wave (SW).....</b>   | <b>3</b>  |
| <b>Figure 3 - On the left, a pipe robot meant for 4” pipes designed by Versatrax [8]. On the right, a pipe robot meant for 2” pipes designed by Honeybee [9]. .....</b>  | <b>4</b>  |
| <b>Figure 4 - The developed PCB, with a filter, a MCU, a driver, and a voltage regulator (VR) and a bluetooth module [6]. .....</b>  | <b>9</b>  |
| <b>Figure 5 - Schematic of the control circuit used to drive the direct current through the solenoid. <math>V_{dd}</math> represents the 14.8V power source while the coil directly below <math>V_{dd}</math> represents the electromagnet used to lift the topmost ball [6]......</b> | <b>9</b>  |
| <b>Figure 6 - The a) average waveforms and b) ToF distribution collected during the experiment from steel.[6].....</b>   | <b>11</b> |
| <b>Figure 7 - The a) average waveform and b) distribution of ToF's collected from foam[6].</b>   | <b>11</b> |
| <b>Figure 8 - The a) average waveform and b) distribution of ToFs collected from granite[6]. .....</b>   | <b>12</b> |
| <b>Figure 9 – The improved PCB with a transducer .....</b>   | <b>13</b> |
| <b>Figure 10 - The second generation HNSW device, with separate subsections labeled. ....</b>  | <b>14</b> |
| <b>Figure 11 - A picture of the signal conditioning circuit used to allow the external ADC to measure a HNSW.....</b>  | <b>14</b> |
| <b>Figure 12 - Box plot showing results of samples collected from both sets of transducers. ..</b>   | <b>17</b> |

|  |           |
|--|-----------|
| <b>Figure 13 – Picture from Fusion 360 showing complete design for both halves of the sensor casing.....</b>   | <b>20</b> |
| <b>Figure 14 - Addition of tabs to casing. The three tabs are highlighted in blue. ....</b>  | <b>20</b> |
| <b>Figure 15 - A picture of the two 3d printed halves of the sensor casing, with one of the halves containing the materials used in the sensor.....</b>  | <b>21</b> |
| <b>Figure 16 - a) A side-by-side comparison of the robot and a more traditional HNSW sensor. b) A close up of the robot. c) The robot driving itself across a plate.....</b>   | <b>23</b> |
| <b>Figure 17 - A system diagram of the electronics on the robot. ....</b>  | <b>24</b> |
| <b>Figure 18 - The completed PCB, designed in Altium and annotated. ....</b>   | <b>25</b> |
| <b>Figure 19 - Velocity vs time for the three trials.....</b>  | <b>30</b> |
| <b>Figure 20 - Both the averaged velocity waveforms and the exponential fit to the averaged waveform .....</b>   | <b>32</b> |
| <b>Figure 21 - Controller based on the ideal velocity profile.....</b>   | <b>34</b> |
| <b>Figure 22 - Example waveform collected on steel .....</b>   | <b>38</b> |
| <b>Figure 23 - A waveform collected from the same sensor and material, but with a different microcontroller.....</b>   | <b>39</b> |
| <b>Figure 24 - Machined metal plate.....</b>   | <b>40</b> |
| <b>Figure 25 - Multiple waveforms taken from an uncorroded section of the plate layered on top of each other.....</b>  | <b>43</b> |
| <b>Figure 26 - Picture of the side of the plate that the robot drove on .....</b>  | <b>46</b> |
| <b>Figure 27 - A picture of the underside of the plate. The two corroded sections are outline in blue, and the robot drove on the opposite of the pictured side. Also included is a picture of the robot driving on the uncorroded side.....</b> | <b>47</b> |

**Figure 28 - Plot of ideal thickness vs distance and measured thickness vs distance from the machined plate ..... 48**

**Figure 29 - A second graph of thickness vs distance obtained using the alternate formula. 50**

**Figure 30 - Graph of thickness vs distance obtained from the larger plate. .... 52**

## Preface

This thesis exists because of me, but I exist because of my parents, so I think I should thank them first. There's a chance I might have made it to this point without their support, but it would have been much harder. In particular, a daily reminder to eat well, while sometimes annoying, was very appreciated in a time when human interaction was scarce because of a pandemic.

The next most significant reason this thesis exists is Dr. Rizzo. Some time ago, he wrote the grant application that resulted in me being able to get a Master's in the first place, and his vision has been invaluable both in producing good results and presenting them properly. I think this is also the best spot to mention Hoda, Rizzo's graduate student and a PhD candidate, who taught me much about the theory regarding HNSW's and was an excellent team member in general.

The next most significant reason after that is Dr. Dickerson. Dr. Rizzo approached Dr. Dickerson to collaborate with him, which resulted in Dr. Dickerson approaching me. Dr. Dickerson has been consistently supportive as an advisor, making sure I had what I needed to get my work done and being tolerant of my many flaws. My thesis could have been on any topic I chose related to HNSW's, but he encouraged me to pursue the robot and is part of the reason I managed to see this through to the end, despite random parts of my robot breaking/my results not "behaving" the way I expected them to. He has also been a role model in many aspects.

I want to thank all of Pitt RAS from the time I was an undergrad, from the officers and team leads to just the ordinary members. I wouldn't be half as technically skilled as I am now if it were not for the experience and knowledge I gained from working alongside that group of people. I also want to thank the graduate students I shared a lab space with in Schenely Place for general camaraderie and helping with some of the classes we took together.

I also want to thank Bill and Jim, the maintainers of SERC and the IT leads for all of Pitt ECE. Strictly for this thesis, they let me use a 3d printer so that I could prototype my custom sensor, but outside of that, they have taught me a lot about the practical aspects of electrical engineering.

Next I want to thank Dr. Mao and Dr. Hu for taking time out of their schedules to be a part of my committee. In addition, Dr. Mao's enthusiasm in the classroom is part of the reason I wanted to make a robot in the first place.

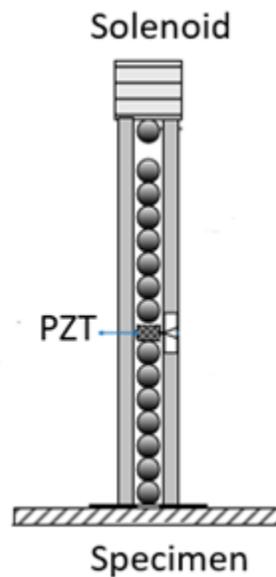
Finally, I want to thank Jason Xuu, a sales rep from ALLPCB. This probably seems out of place here seeing as every other entity I've acknowledged has a personal connection with me, but my graduation depended on the manufacturing of a few PCB's and being able to Skype Jason made it a lot easier. No, I am not being paid to type this here.

## 1.0 Introduction

Structural health monitoring (SHM) and nondestructive evaluation (NDE) are important topics of research because of their importance in the management of almost any engineering structure subject to corrosion or wear of any kind. The field of NDE is rich in both the number of potential applications and the diversity of techniques investigated in research and used commercially. For water and wastewater pipes alone, acoustic emission, eddy currents, magnetic flux leakage, and sonar have all proven to be useful[1]. Applications of NDE range from the transmission testing of aerospace laminates to pulse-echo flaw detection in plastic packaging materials[2]. In recent years, a novel SHM/NDE approach based on the propagation of highly nonlinear solitary waves (HNSWs) has emerged as a valuable tool for the monitoring or inspection of a variety of structures and materials. This mechanism involves inducing an incident solitary wave (ISW) at one end of a chain of spherical particles. As a result of the ISW being induced at one end of the chain, a compact non-dispersive compressive wave propagates along the chain and reaches the opposite end. This single pulse, when reflected off of the material of interest, results in one or two reflected waves. The amplitude and ToF of these reflected waves can reveal some characteristics of the material/structure of interest. These sensors have been proven to be useful to characterize concrete, steel, and even tennis balls[1, 3, 4]. As of now, these sensors have mostly been kept stationary, and been presented as a way to monitor the health of a structure. In the one instance that could count as an exception, the sensor was simply mounted on a 4 wheel cart with no motors[5]. The potential of this sensor to be used for an autonomous, one time inspection by transporting it across a structure has never been explored, to the best of the author's knowledge. The purpose of this thesis is to prove the viability of attaching this sensor to a mobile platform and

carrying out an inspection with said platform. The construction of a smaller HNSW sensor and a robot platform for the sensor is discussed. An algorithm to classify materials in real time is introduced. Finally, the results of an experiment in which the robot inspects a metal plate with artificially induced corrosion are shown to prove the viability of the sensor-robot pairing.

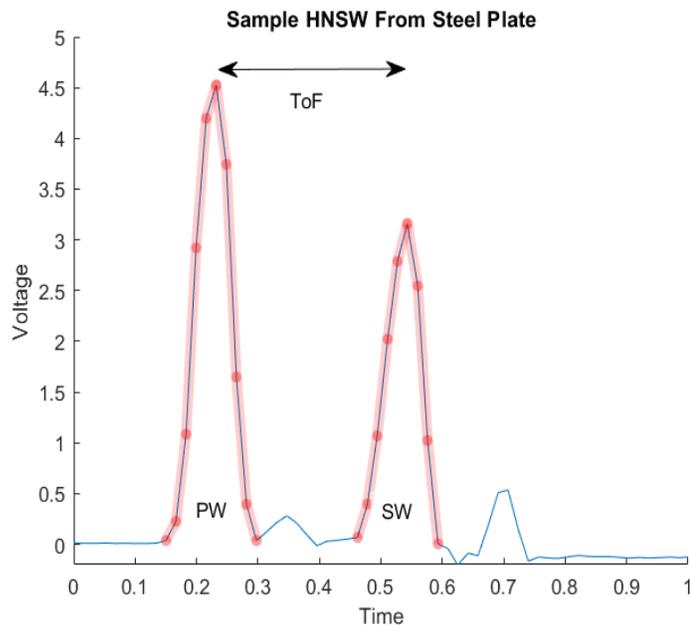
### 1.1 Background of HNSW's



**Figure 1 - A basic schematic of an HNSW sensor**

The basic scheme of an HNSW-based NDE/SHM is shown in Figure 1. One end of an array of identical spheres is in contact with the structure to be inspected/monitored. However, the topmost ball is made of a ferromagnetic material, unlike the rest of the balls in the chain. This allows the solenoid at the top of the chain to lift the topmost ball, given that there is a significant amount of current flowing through it. The dropping of the topmost ball induces an ISW, which results in a reflected wave referred to as a secondary wave (SW). The PZT element in the middle

serves to turn the ISW and the reflected waves into a waveform that can be measured by suitable electronic equipment. A graph of a sample HNSW that was taken from a steel plate is shown in Figure 2.



**Figure 2 - A sample HNSW from a steel plate. The first highlighted peak is the primary wave (PW), and the second highlighted peak is the secondary wave (SW).**

The difference in time between the arrival of the PW and the SW is referred to as the time-of-flight (ToF). The ToF of an HNSW is often a key statistic used to characterize the material of interest, such as the aforementioned concrete and steel [3, 5]. In most applications explored by civil engineers, this waveform is measured by an oscilloscope or some other traditional data acquisition system, such as custom desktops built by Labview [5]. The actuating of the solenoid is usually performed by a power supply controlled by the data acquisition system [6]. The aforementioned hardware was useful for characterizing these sensors and investigating applications in which they

might be useful, but in order for this technology to be practical, the apparatus for actuating the solenoid and measuring the induced signal had to be minimized to the bare necessities.

## 1.2 Previous Work on Pipe Robots

Every pipe robot is fundamentally a platform to carry a detector through a pipe, where the detector used depends on the application. Developing a robot that can inspect the inside of a pipe offers many advantages, such as increased speed of inspection, the safety of those who would otherwise be responsible for inspections, more consistent examinations, and the ability to inspect spaces that humans would otherwise be unable to access. Robots can perform in hours what would take humans weeks [7].

Pipe robots have existed for decades, and numerous products exist on the market for the inspection of pipes ranging from a diameter of 2 inches to 42 inches. Products from companies such as Inuktun and Honeybee sell industrial-quality inspection robots that are ready for delivery, as shown in Figure 3.



**Figure 3 - On the left, a pipe robot meant for 4” pipes designed by Versatrax [8]. On the right, a pipe robot meant for 2” pipes designed by Honeybee [9].**

Robots have been developed for corrosion and radiation inspection, to name just a few applications. RadPiper, a robot developed at Carnegie Mellon, uses a sodium iodide scintillation

detector to find trace quantities of U-235 in pipes[10]. A robot designed at the University of Strathclyde used a non-contact ultrasonic inspection system to find defects in pipes[11]. This work in particular was very similar in spirit to this thesis, as they investigated a new kind of ultrasonic coupling as well as the hardware necessary to make this novel coupling method mobile on a robotic platform. However, the aluminum plate they inspected to prove the functionality of their sensor was only 1 mm thick, whereas the plate used in the experiment discussed in this thesis is 0.12 in (~3mm thick) which is a much more realistic thickness to prove the viability of a sensor for pipe inspection on. Other robots designed for corrosion testing often use ultrasonic sensors, an well-established NDT method for finding irregularities in materials. However, some ultrasonic sensors suitable for being attached to a robot can cost up to \$450[11]. Proving the viability of an alternative NDT sensor could serve to drive down the cost of future pipe inspection robots while maintaining the same functionality. In our case, the sensor we have developed is a maximum of 50 USD worth of parts.

### **1.3 Research Contributions**

Previous publications have used dedicated signal acquisition equipment to collect HNSW's from transducers, such as the National Instruments PXI [3]. This thesis outlines the minimal electronics necessary to perform the same functionality as one of these much more complex machines. Given that the possibility of moving this kind of sensor around to perform a one-time inspection has never been explored, it then builds upon this work to design a robot and a custom sensor that the robot is capable of moving around. It is hypothesized that the combination of the

robot and the custom sensor will be capable of distinguishing between differing levels of corrosion. The total research contributions of this thesis are as follows:

- ***Hardware design for the electronics necessary to capture and measure a HNSW:***  
The electronics described in this thesis are designed to replace relatively expensive lab equipment normally used by civil engineers to validate manufactured HNSW transducers. This functionality is tested by using both the designed electronics and more complicated equipment to collect samples from three different materials.
- ***The design of a robot that can transport the sensor, actuate the solenoid, and capture the waveform:*** The aforementioned electronics are used as a starting point for a robot designed to traverse a surface and inspect it automatically.
- ***A model that allows for the extraction of features from a waveform and maps these features to a thickness measurement, allowing for the detection of corrosion:*** A plate with manually induced defects is used to collect training data, which is in turn used to find a correlation between features from collected waveforms and the corresponding thickness of the sample. Results from a traversal of the robot are obtained using a model derived from this data, which in turn proves the viability of this sensor-robot pairing.

## 1.4 Organization of Thesis

Section 2.0 discusses the hardware of previous modules built for the capture and transmission of HNSW's, as well as the hardware design of the robot. Section 3.0 discusses various aspects of the software that were necessary to make the robot capable of moving around

autonomously and accurately, so as to allow accurate corrosion measurement/position pairings. Section 4.0 describes the procedure of the experiment used to verify the robot, and Section 5.0 discusses the results from this experiment.

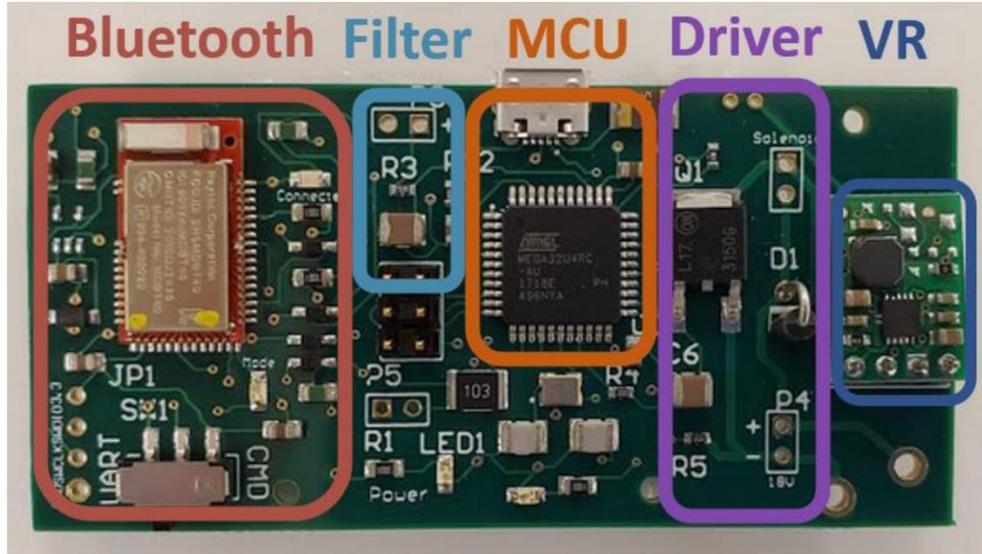
## **2.0 Hardware**

There has been a significant amount of previous work on the hardware necessary to induce an HNSW, capture it, and extract a ToF. The way in which this work laid the groundwork for the robot developed for this thesis will be elaborated upon. The construction of the custom sensor used for the robot will also be explained.

### **2.1 Previous Work**

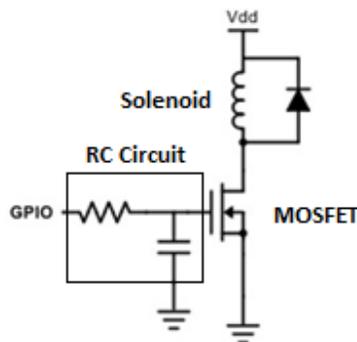
#### **2.1.1 First Generation Wireless HNSW Device**

In the past year, a PCB was created by the author of this thesis for the purpose of proving the viability of inducing and measuring an HNSW using the minimal electronics necessary to do so[6]. This custom solution used an ATmega32u4 to control the induction and sampling of the signal. The ATmega microcontroller (MCU) provided an on-board analog-to-digital converter (ADC) that could sample the signal and the capability to control the driver, which in turn controlled the flow of current through a solenoid. A step-up voltage regulator converted 3.7V from a single cell LiPO battery into 5V for the MCU. A picture of the custom solution is shown in Figure 4.



**Figure 4 - The developed PCB, with a filter, a MCU, a driver, and a voltage regulator (VR) and a bluetooth module [6].**

The driver was a circuit designed to control the flow of current through a solenoid based on the state of a digital output from the MCU. A schematic of the driver and the solenoid is shown below.



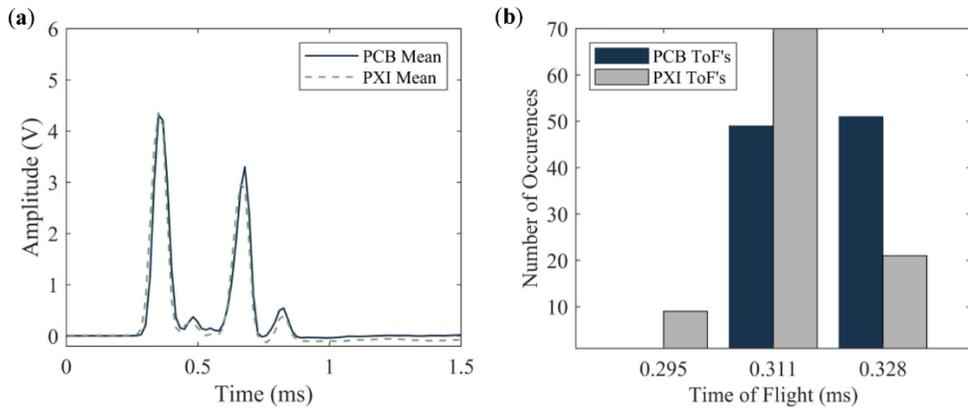
**Figure 5 - Schematic of the control circuit used to drive the direct current through the solenoid.  $V_{dd}$  represents the 14.8V power source while the coil directly below  $V_{dd}$  represents the electromagnet used to lift the topmost ball [6].**

The RC circuit was added to prevent the worst case scenario of the MCU's ADC sampling the PZT sensor after the PSW had passed. When the MCU turns its GPIO pin on, the

capacitor charges up to 5V (the logic level of the MCU), which turns the N-channel MOSFET on, allowing current to flow through the solenoid and resulting in the ball rising until it was blocked by the solenoid. Turning the MOSFET off would result in the capacitor discharging itself through the resistor into the GPIO pin, and the voltage at the gate would steadily fall until the MOSFET turned off. Increasing the amount of time it takes for the ball to fall after the pin is set to low helps guarantee that the MCU can sample the ADC in time to catch the PSW. The diode in parallel with the MOSFET is a 1N4003 diode. In this configuration, it is referred to as a flyback diode, and it is added to prevent the voltage spike that results from turning off the solenoid from damaging the metal–oxide–semiconductor field-effect transistor (MOSFET), which would in turn shorten the life of the overall system [12].

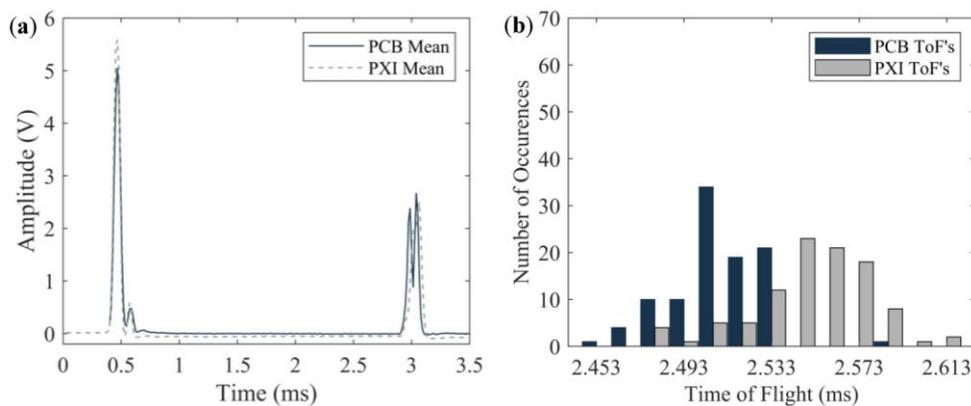
### **2.1.2 Results**

The PCB pictured in Figure 4 was validated against a PXI module by taking samples using the same transducer on 3 different materials; steel, granite, and foam. 100 waveforms were collected from both devices, and the waveforms from each device/material pair were aligned using the incident solitary waves and plotted against each other to compare the average waveforms collected from both devices. The average waveforms and ToF distribution collected from both devices from a sample of steel is shown in Figure 6 .



**Figure 6 - The a) average waveforms and b) ToF distribution collected during the experiment from steel.[6]**

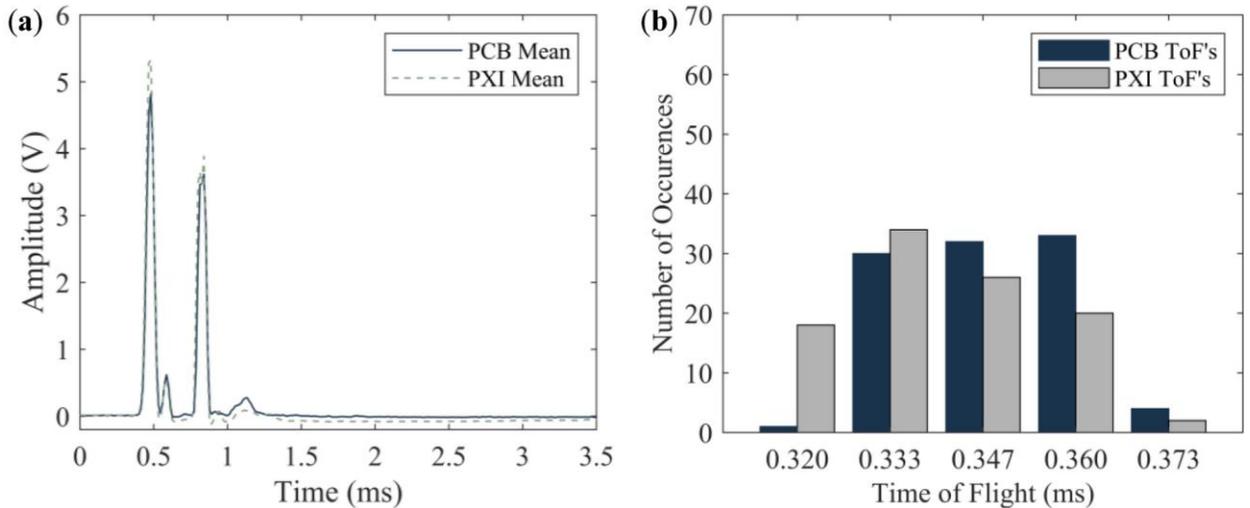
As seen in Figure 6, the average waveform and distribution of ToF's were very similar. In order to quantify the differences between the measured ToFs from each system, Figure 6b shows the number of occurrences for a given ToF for both setups. Three values, namely 295, 311, and 328  $\mu$ s were measured. It is noted here that the width of each bar is equal to the sampling period, which, for the steel plate test, was equal to 16.4  $\mu$ s.



**Figure 7 - The a) average waveform and b) distribution of ToF's collected from foam[6].**

Figure 7 shows the average waveforms collected from both systems and the distribution of ToFs from both systems from foam. The ToFs collected in Figure 7b have more variance than the ToFs collected the steel test. This is attributable to the soft nature of the foam specimen. Additionally, the occurrences of the ToF measured with the wireless platform are skewed to the

left, i.e., the PSW traveling along the transducer connected to the PCB seems faster than in the transducer connected to the PXI. While the cause of this skew is not clear, it is believed that what was observed is more related to the dynamic interaction of the waves with the foam rather than a variability associated with the transducers [6].



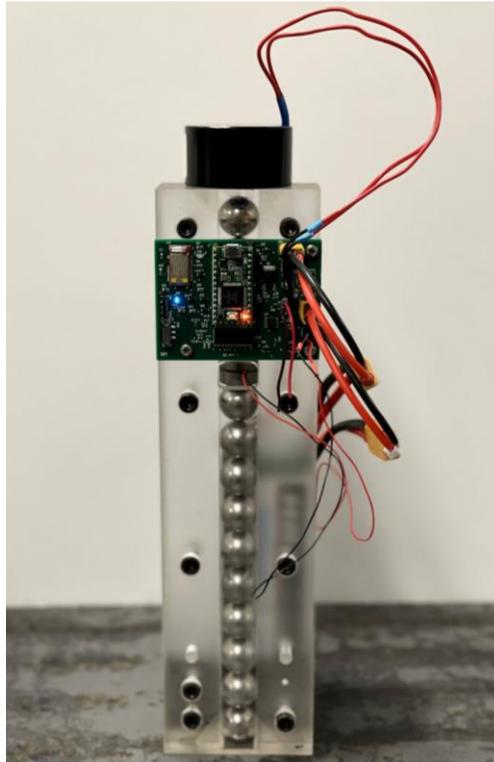
**Figure 8 - The a) average waveform and b) distribution of ToFs collected from granite[6].**

Figure 8 shows the average waveforms collected from both systems and the distribution of ToFs from both systems from granite. Like the results from steel, the ToFs were very similar and not spread apart, suggesting that there is a correlation between variance in measured ToFs and the softness of the material sampled.

Across all three of the materials tested, the largest percent error between the mean ToF collected by both devices was 3.90%, which was impressive considering the existence of mechanical differences between the two transducers used as well as the sheer difference in electrical complexity between the PXI computer and the PCB. These results served as proof that this thesis was a viable undertaking.

### 2.1.3 Second Generation Wireless HNSW Device

Another PCB was made that had an improved sampling frequency, better power efficiency, the capacity for local storage. The PCB was also designed alongside a transducer that it was designed to be attached to. This PCB along with a transducer is shown in Figure 9.



**Figure 9 – The improved PCB with a transducer**

This PCB used the same solenoid driver as the last version of the PCB. Like the last PCB, it also has a microcontroller and a wireless module that can be used to transmit samples. However, it also had an external ADC that allowed it to sample much faster than an ADC onboard a microcontroller would. The PCB with its separate components is shown in Figure 10 - The second generation HNSW device, with separate subsections labeled..

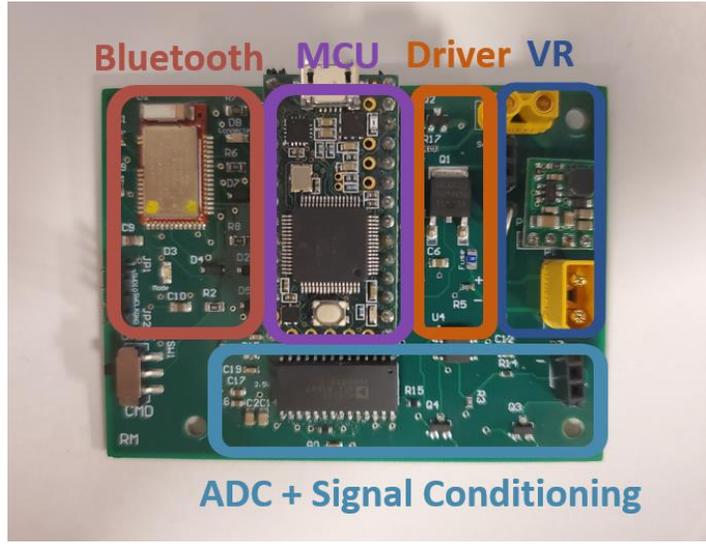


Figure 10 - The second generation HNSW device, with separate subsections labeled.

2.1.3.1 Signal Conditioning

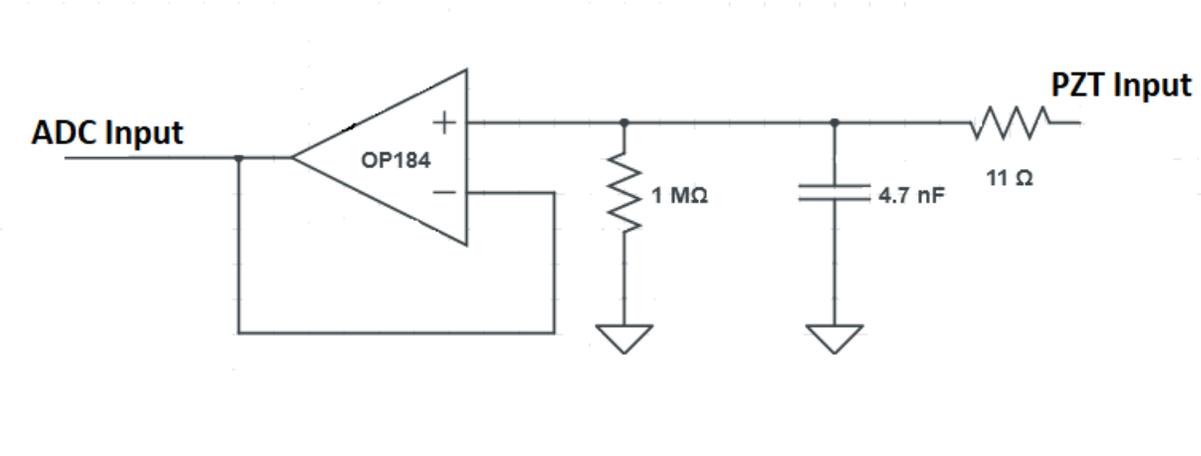


Figure 11 - A picture of the signal conditioning circuit used to allow the external ADC to measure a HNSW.

A passive low pass filter, a resistor that acts as an impedance matching network, and a rail-to-rail op amp were combined to provide the external ADC with minimal output impedance and anti-aliasing. A diagram of the filtering circuitry is shown in Figure 11 - A picture of the signal conditioning circuit used to allow the external ADC to measure a HNSW.. The maximum sampling

frequency of the system was measured to be 8 MHz, thus the Nyquist frequency is 4 MHz. The low pass filter shown in Figure 11 was designed to have cutoff of 3 MHz. A relatively high cutoff frequency was used because it is undesirable to eliminate too much information below the Nyquist frequency of 4 MHz. At the same time, the filter provided anti-aliasing above the Nyquist frequency.

### **2.1.3.2 External ADC Operation**

After the signal is filtered, it is sampled by an external ADC. The external ADC selected was an AD9220, which has a maximum sampling frequency of 10 MHz and uses a parallel data interface. The MCU that interfaces to the ADC can read the digital state of 8 pins in the span of a clock cycle. The ADC also has a clock input, which acts as a rising-edge trigger for the ADC to take a sample. It was chosen a loop was implemented on the MCU that first induces a rising edge on the ADC clock input by turning the corresponding digital pin off and on. Then, it reads an 8 bit number from parallel data interface on the ADC and places this sample into an array that is stored in RAM. In testing the data acquisition software, it was determined that the highest sampling frequency that this MCU/external ADC approach can achieve is 8 MHz. An additional benefit to controlling the clock and sample collection entirely through firmware (instead of using an external oscillator to control the clock) is that the sampling frequency can easily be configured by inserting delays in the code that controls the clock output to the ADC. The ability to add clock cycle delays results in sampling periods can be adjusted with a granularity of  $1/72$  us, a unique advantage of this approach to triggering the ADC in this way. The MCU system clock frequency is 72 MHz and the data acquisition sampling frequency is 8 MHz. Thus, the software loop described takes 9 processor clock cycles per iteration. Individual clock cycle-level delays (NOOPs, or no operation

instructions) added to this loop would decrease the sampling frequency to  $72/(9 + n)$  MHz, where  $n$  is the number of added no-op CPU instructions.

### **2.1.3.3 Power Management**

One of the goals for this system is for it to monitor the structural health of a material for an extended period. To accomplish this goal, it is vital that the power consumption be optimized in periods of time when the system does not have to take measurements. The chosen MCU had an extremely efficient sleep mode, capable of running on only up to 40  $\mu$ A, until it is woken (either through an external trigger or through an internal timer). However, the circuitry used to make analog measurements is not quite as efficient. The ADC consumes 30 mA whether it is being used to collect samples or not, and the operational amplifier constantly consumes 10 mA. If a micro SD card is utilized for data-logging, it can also constantly take up to 30 mA. These three components combined require much more power than a sleeping MCU, so additional circuitry to turn these three components off while the MCU was sleeping was necessary. N-channel MOSFETs were added that can cut off all three components' connections to ground. A separate plane within the dedicated ground plane under the ADC and op amp was connected to the drain of our MOSFET, and the source of the MOSFET was connected to the rest of the ground plane. The source terminal of the MOSFET was connected to the rest of the ground plane. The ground pin on the micro SD card slot was also connected to the drain of the MOSFET. The gate of the MOSFET was controlled by the MCU, enabling the MCU to simply control the total current consumption of the board through a single digital output. The ability to turn these components greatly reduced the minimum current consumption from 70 mA to 5 mA, dramatically extending the amount of time the device can operate without human intervention.

### 2.1.4 Experimental Setup and Results

In order to investigate the power efficiency of the device and its capacity to make accurate measurements over time, 3 PCB's were assembled with transducers as portrayed in Figure 9 – The improved PCB with a transducer. These three transducers were set to collect 10 measurements every 20 minutes on a large metal plate. Three other transducers were connected to a National Instruments PXI and set to collect measurements at the same interval from a slightly thicker, but mostly similarly sized plate. Figure 12 shows a box plot that represents the span of ToF measurements taken from all 6 transducers. P1 through P3 represent distributions of the samples taken from the PXI, while W1 through W3 represent samples taken from the wireless module. Each box represents at least 5,040 measurements with the exception of W1, which failed halfway through the week-long test. It was later found that this failure was simply due to a cold solder joint on one of the power connections.

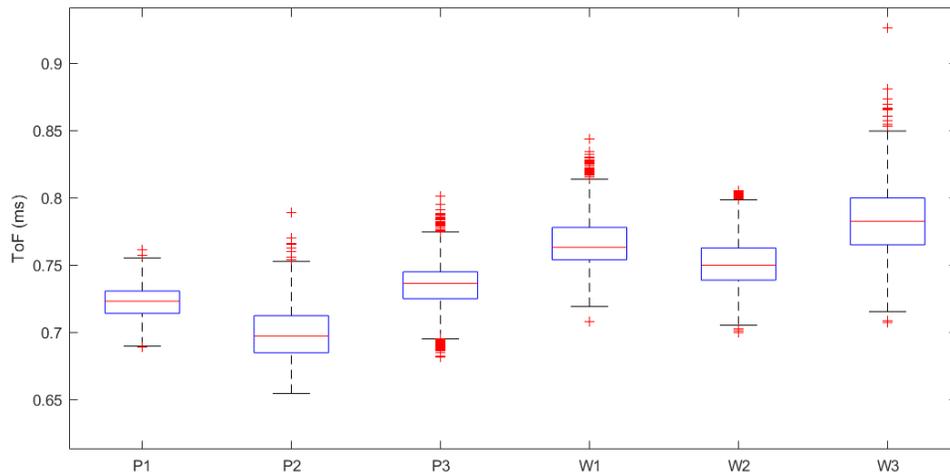


Figure 12 - Box plot showing results of samples collected from both sets of transducers.

Ignoring the outliers, which are represented by the red markers, on average the ToFs collected from the PXI seem lower than the ToFs collected from the wireless modules. However, this is to be expected, as the two metal plates used were of different thicknesses. The wireless modules were placed on the thinner plate, so observing a larger ToF on average agrees with previous studies that show that there is a negative correlation between thickness and ToF [13]. The goal of this experiment was to show that the wireless module was capable of making reliable measurements as the PXI was across the span of a week. In order to prove this, the ranges of collected ToFs as well as the variances of the total ToFs will be compared. Table III shows these statistics. It should be noted that the maximum and minimum fields exclude outliers denoted by red markings on the box plot. The criteria for being an outlier was defined as being more than 2.7 standard deviations from the mean.

**Table 1 - Table of results from week-long test**

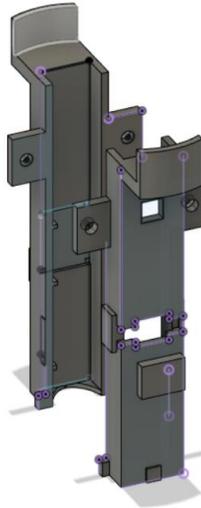
|                         | Device |      |      |      |      |      |
|-------------------------|--------|------|------|------|------|------|
|                         | P1     | P2   | P3   | W1   | W2   | W3   |
| <b>ToF Maximum (ms)</b> | 0.76   | 0.75 | 0.80 | 0.81 | 0.79 | 0.84 |
| <b>ToF Minimum (ms)</b> | 0.69   | 0.65 | 0.68 | 0.71 | 0.70 | 0.71 |
| <b>Variance (us)</b>    | 1.7    | 3.25 | 2.56 | 7.3  | 3.1  | 0.5  |
| <b># of outliers</b>    | 4      | 19   | 123  | 30   | 23   | 14   |

The ranges of ToFs from the PXI range from 0.07 ms to 0.12 ms, and the ranges of ToFs from the wireless module range from 0.09 ms to 0.13 ms. The maximum range of ToFs from the wireless modules is therefore 0.01 ms more than that of the PXI, but this is an insignificant difference. The variance for W1 is an outlier relative to the variances of the other samples, but W1 also yielded a fewer number of samples due to stopping halfway through the week. The calculation for the variance of a set of samples involves dividing by the number of samples, so it is likely that

the fewer number of samples here is the dominant factor in the variance being significantly larger than the others. The number of outliers for each set of measurements seems to be completely random, and so it is concluded that the presence of outliers is likely to be correlated with characteristics of mechanical flaws within the transducers instead of the devices making the measurements.

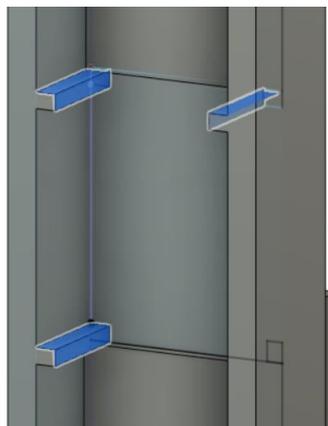
## **2.2 Construction of Custom Sensor**

In making an HNSW sensor (also referred to as a transducer) that is suitable for a robot, there are several things that must be addressed. The spherical particles and the sensor particle should ideally have a very small amount of friction against the walls of the casing, such that they do not become misaligned (i.e. become dislocated from the formation of a straight chain), but not so much friction that they cannot vibrate freely. The top spherical particle especially should not be restricted so that it can be freely lifted by the solenoid at the top of the chain. The sensor particle, ideally, should maintain a constant orientation throughout the operation of the robot so that the collected ToF's do not drift over time. As a starting point, a casing was designed in Fusion 360 that could hold spherical particles  $\frac{3}{8}$ ths of an inch in diameter, as well as a slot for a solenoid on top. The complete design for both halves of the sensor casing is shown in Figure 13.



**Figure 13 – Picture from Fusion 360 showing complete design for both halves of the sensor casing**

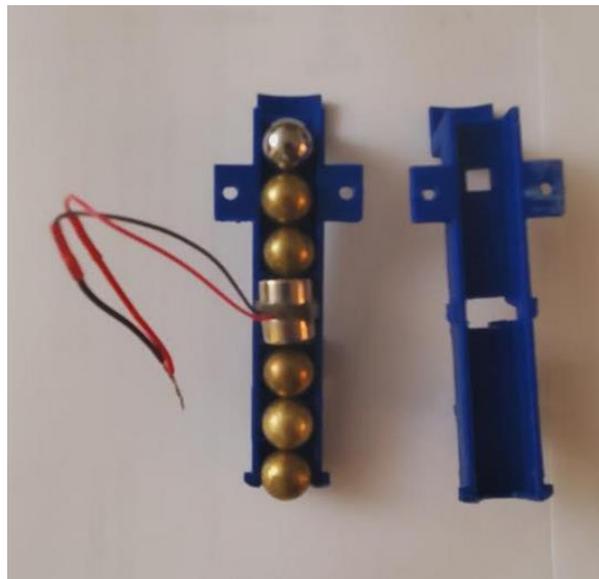
In order to keep the sensor particle level, one half of the sensor had 3 tabs added to it that would keep the sensor particle in place.



**Figure 14 - Addition of tabs to casing. The three tabs are highlighted in blue.**

Finally, a sensor particle was made with two magnetic disks  $\frac{3}{8}$ ths of an inch in diameter, a piezoelectric disk, a piece of foam, and Kapton tape. The magnets were first heated in an oven at 275 degrees Fahrenheit for 10 minutes to partially demagnetize them. This was done just to make the final assembly more feasible. The piece of foam was necessary because the piezoelectric

disks that were purchased had presoldered wires attached, making it impossible to assemble a level sensor disk without adding some sort of cushioning. The sensor behaves the most ideally if it is exactly the same weight and size as the rest of the particles, however, this was not readily achievable due to a lack of commercially available materials. However, as shall be proven later, this does not completely impede the sensor's ability to detect corrosion. A steel sphere of diameter 3/8" was used for the striker particle, and five brass spheres of diameter 3/8" were used for the rest of the particles in the chain, excluding the sensor particle. A final picture of the sensor assembly is shown in Figure 15.

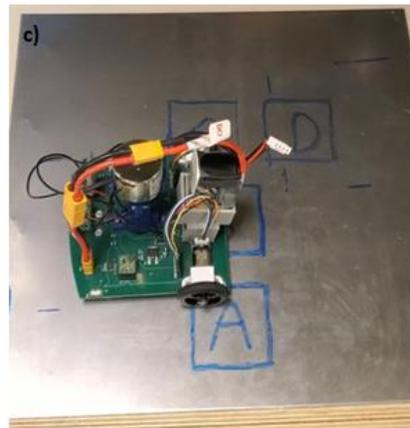
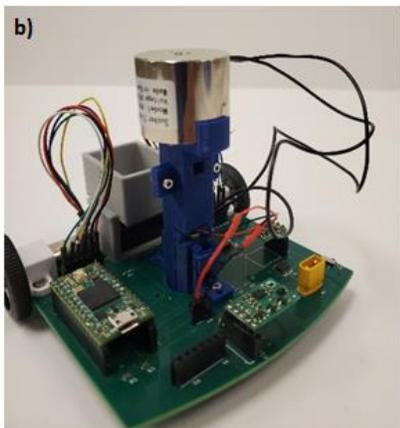
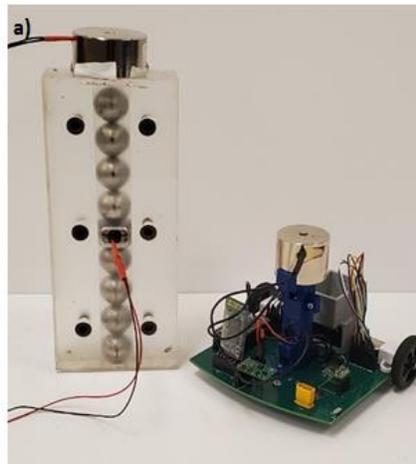


**Figure 15 - A picture of the two 3d printed halves of the sensor casing, with one of the halves containing the materials used in the sensor.**

After constructing the sensor and finding an appropriate solenoid, it was found that at least 10V needed to be applied across the solenoid to lift the striker particle. This motivated the choice of power supply for the robot, which will be discussed later.

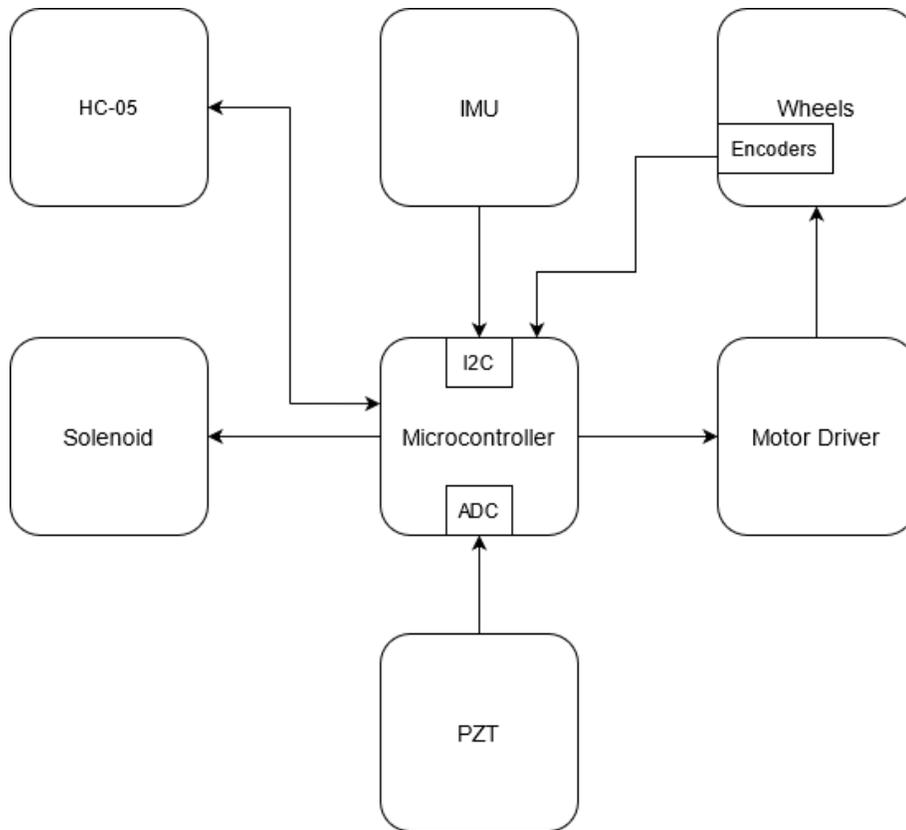
### **2.3 Robot Hardware**

In order to transport the sensor, a basic platform that was capable of moving the sensor and the power source for the solenoid had to be constructed. For this, a PCB was manufactured that allowed for the wiring of a microcontroller to the sensor that waveforms would be collected from and to a DC motor driver that allowed the microcontroller to control the movement of the robot. The mechanical design for the robot was simplified by allowing the PCB to double as a mechanical base for the motors and the sensors. Finally, a gyroscope, IMU, and compass were added to provide information during the operation of the robot that would help with localization. The final robot is shown in Figure 16a next to a more conventional transducer in order to display the extent of the miniaturization of the custom sensor mounted on the robot.



**Figure 16 - a) A side-by-side comparison of the robot and a more traditional HNSW sensor. b) A close up of the robot. c) The robot driving itself across a plate.**

Figure 17 shows a basic system diagram for all of the different components of the robot. Each individual component will be expanded upon in the following sections.



**Figure 17 - A system diagram of the electronics on the robot.**

### **2.3.1 Printed Circuit Board**

The printed circuit board created for this robot creates the connections between each subcomponent that allow the robot to function, and doubles as the mechanical base for the robot. The PCB was first designed in Altium, keeping in mind the size of the constructed sensor, the motors, and the wheels. It had 4 layers; 2 of them were dedicated to signals on the board, and the other 2 were dedicated power planes. Figure 18 displays an annotated picture of the final PCB design.

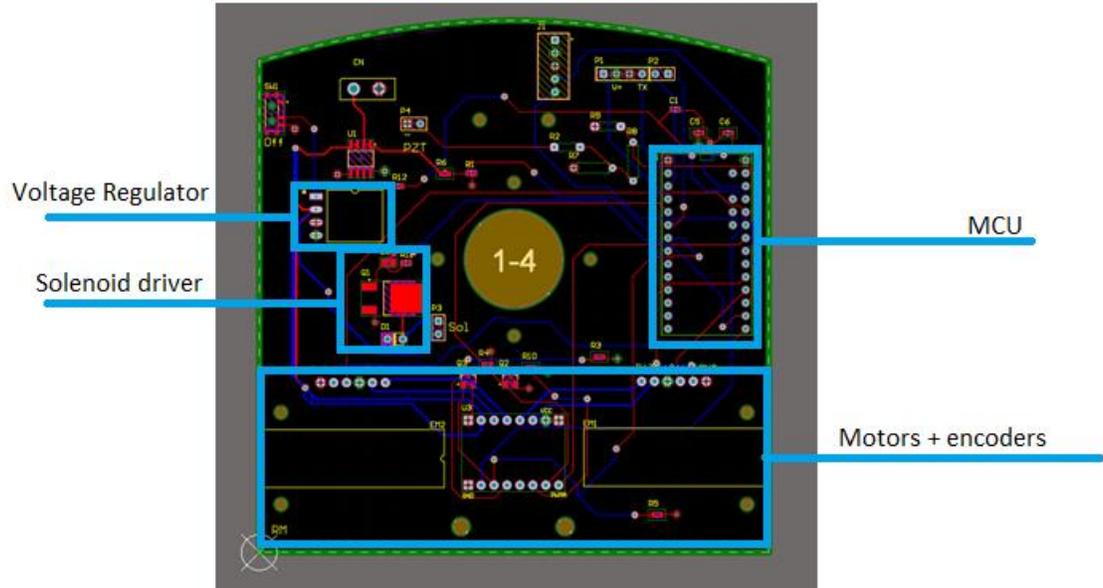


Figure 18 - The completed PCB, designed in Altium and annotated.

### 2.3.2 Microcontroller

The robot needed a central computer to carry out the main tasks of waveform collection and robot control. For this purpose, we chose to use the Teensy 4.0. The Teensy 4.0 has support for multiple communication protocols, such as I2C and SPI, allows the user to select one of multiple base clock frequencies (ranging from 24 MHz to 1 GHz), and has a floating point unit, which makes the kind of calculations necessary for position estimation execute much faster. The robot also needs to be able to measure the waveforms produced by the piezoelectric sensor. It was found through experimentation that the Teensy 4.0's ADC was capable of accurately measuring the signal from the custom transducer without any external circuitry. In short, it is a very flexible, powerful microcontroller that was more than capable of handling the requirements for this project.

### 2.3.3 Motion Hardware

In order to make an electronic system capable of moving itself, one usually needs a power source, some kind of motor, and an electrical system that provides an interface between the motors and a computer that allows the computer to control those motors. The power source was going to have to be able to provide at least 10V, given that the solenoid needed 10V and more than one battery would unnecessarily weigh the robot down. For this reason, a 3S LiPo battery with a nominal voltage of 11.2V was chosen. For motors, two micro metal gearmotors with a gear ratio of 30:1 from Pololu were chosen to drive a left wheel and a right wheel on the robot. Two separate quadrature encoder attachments designed specifically for those motors were also added to provide position measurements. The encoders provided 12 counts per revolution of the main shaft, but because the gear ratio was 30:1, the main shaft completed 30 revolutions for each revolution of the output shaft. Therefore, the encoder resolution was a total of 360 counts per revolution of the wheel. Wheels 32 mm in diameter were used, and by calculating the circumference and dividing by 360, the derived resolution of our encoder based position measurements is 0.28 mm per encoder count. Finally, the TB6612FNG dual H-Bridge motor driver was used to provide the microcontroller with an interface to the motors. The TB6612FNG has 4 outputs meant to be wired to the input and output of two DC motors, and can control bidirectional current through both DC motors given two digital inputs determining the direction of the current. The TB6612FNG also has pulse-width modulation (PWM) inputs that determine how often voltage is applied across the motors, which allows a microcontroller with a PWM output to directly control the speed of the motors by setting a duty cycle on said PWM output. In order to know what amount of power will be delivered for a given duty cycle, the robot also has to be able to measure the voltage of the battery powering the motors. A simple resistor divider was used to make this possible.

### **2.3.4 Sensors for Position Estimation**

The ability to detect corrosion in the context of a robot is less effective if it is not accompanied with the ability to discern the robot's current position. Besides the aforementioned encoders, additional sensors were needed to increase the accuracy of position measurements. The MiniIMU v9 was chosen to provide the robot with more information that would help with position estimations. The MiniIMU v9 provides an I2C interface that allows a microcontroller to obtain values from a gyroscope and a compass. While encoder values can provide accurate measurements during turns, additional information from a gyroscope or compass can be fused with estimates from the encoder to increase the accuracy of angular movements.

### **2.3.5 Wireless Communication**

A way to collect data remotely was necessary, as though a cable could be used to transmit data, it would affect the robot's dynamics and make it much harder to control. Turns would become especially difficult. Two HC-05 modules were used to allow the robot to send data without having to be tethered to the main computer through a USB cable. An HC-05 module is simply a device that allows for remote communication with other HC-05 modules that are configured to be paired with it. The HC-05 was connected to the Teensy through a serial interface.

### **3.0 Software**

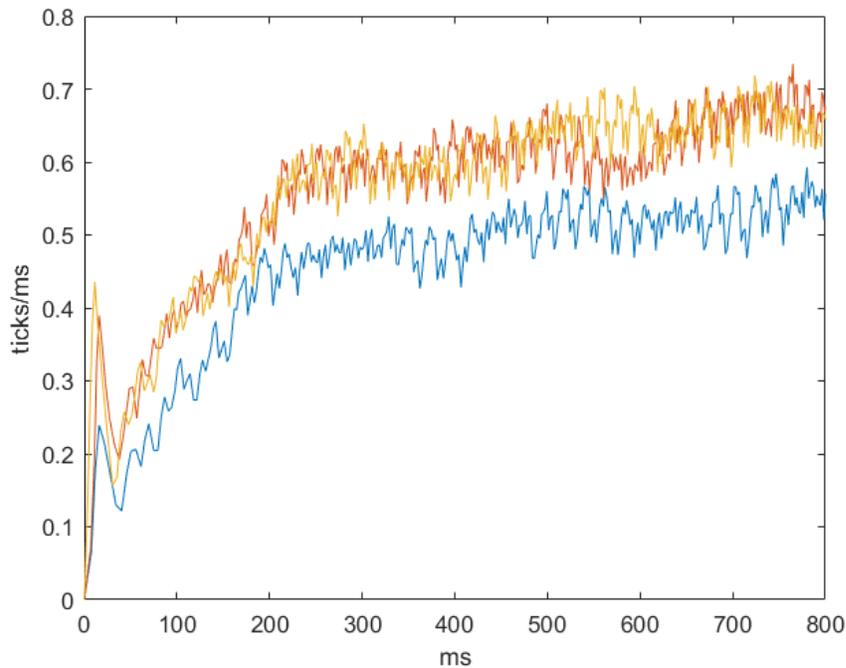
With the hardware completed, software needed to be developed for controlling the robot. In this section, the software for control, estimation, and waveform collection and transmission will be discussed.

#### **3.1 Control**

It was decided in advance that the experiment that would be used to prove the robot worked would only require straight movements and point turns, so as to simplify what would be a difficult motion planning problem since the focus of this thesis is on the robot-sensor pairing. In both of these movements, the goal would be explicitly to move as slowly as possible while not letting the motors stall so as to prevent slippage in the wheels. Any slippage would reduce the accuracy of position estimates derived from the encoders.

The first objective was to derive a model of how the motors behaved at a comfortable velocity. This future model required that there be a way to accurately measure the velocity first. Two interrupt service routines (ISR's) were set up that triggered whenever the digital state from one of the two encoders changed. These ISR's were programmed to store the amount of time since the encoder value last changed, which allowed for a reasonably accurate velocity estimate [14]. First, a script was written that slowly incremented the PWM value sent to both wheels until the robot went from stalling to moving forward. The battery voltage at this point was measured, and

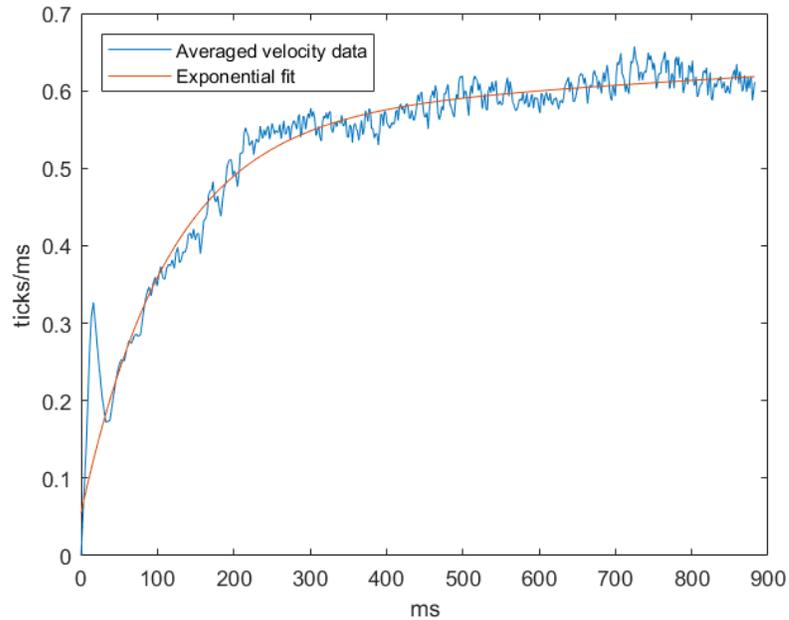
then multiplied by the duty cycle percentage to obtain a voltage measurement, assuming the law of averages [15]. This served as an experimental way of deriving the average voltage that had to be applied across the motors to overcome the force of static friction. This minimum voltage was incremented by 0.5V to simply obtain a voltage that was more guaranteed to make the robot move, in the case of any small changes to the weight of the robot or a change in the force of static friction given a different surface that the robot was driving on. This final ideal voltage was 1.8V. Finally, this voltage was applied across the motors for one second, and as the robot drove forwards it continually polled the encoders, recording the times at which the encoder value for both the left wheel and the right wheel changed. This trial was run three times, and a corresponding graph of velocity vs time for both the right and left wheels for each of those three trials was obtained by differentiating the encoder readings vs time. Figure 19 shows the velocity vs time for just the left wheel, as including a plot for the right wheel would be redundant as it was very similar to Figure 19.



**Figure 19 - Velocity vs time for the three trials.**

There are a few things to note here. There is an initial spike in velocity, most likely because the wheel spins freely for a moment before the internal gears engage. After that, the velocity seems to rise steadily, then plateaus. The final velocity is not consistent across trials, indicating that this is likely a time varying system. The battery voltage was measured at the start of each trial and was found to be the same to the tenth of a millivolt, so though it might not be time varying and might instead be varying across changes in the supply voltage, if these changes in the supply voltage that cause variations in our final velocity are too small to measure it doesn't matter. Therefore, a purely model-based approach might not be reliable, as a model that was derived one day might prove to be inaccurate the next day. For this reason, the decision was made to create an ideal velocity profile and make a controller that would make a best effort to make both wheels follow that velocity profile. In order to obtain this velocity profile, it was decided that the three aforementioned velocity

measurements would be averaged. However, the velocity measurements were made by differentiating the instances where the encoder value changed (as this would result in more accurate velocity measurements than polling the encoders at a set period). Therefore, the timestamps of the velocity measurements between trials were not aligned. In order to average these waveforms, it was assumed that the wheel velocity was smooth between the intervals in which it was measured, and a new graph of velocity vs time was obtained for two trials using the timestamps of the first trial by linearly interpolating between known time-velocity pairings to get two new graphs of velocity vs time that used the same timestamps as the first trial. This allowed the three velocity profiles to be averaged. The resulting graph of velocity vs time was fit to an exponential equation of the form  $v(t) = a * e^{bt} + c * e^{dt}$ . The values that were found for a,b,c and d resulted in a final equation to represent an ideal graph of velocity vs time:  $0.594e^{0.022t} - 0.0076e^{-2.09t}$ . Because 0.022 is almost 0, the equation was simplified to  $0.594 - 0.0076e^{-2.09t}$ . The graph of the resulting exponential fit along with the averaged velocity profile is shown in Figure 20.[16]



**Figure 20 - Both the averaged velocity waveforms and the exponential fit to the averaged waveform**

It was decided that this profile would be used for both of the wheels, as the left wheel was on average slower than the right wheel, assuming the same voltage was applied across both. Additionally, as noted before Now that there existed a model for how the velocity of one of the wheels would behave over time, the model had to be incorporated into the robot. A PID controller was implemented on the robot that measured the velocity on both wheels and attempted to control the velocity of both wheels to match the velocity profile derived earlier. However, this resulted in uneven movements, and the final encoder counts for the left wheel and the right wheel would almost never match. In order to remedy this, the velocity PID controller was cascaded with a PID controller that used the current encoder counts as a setpoint instead, attempting to make sure the encoder count of the left wheel always matched the encoder count of the right wheel [17]. The final controller architecture is pictured in Figure 21. The “zero” block feeding into the left of both the left wheel controller and right wheel controller is simply there to

emphasize that the controller's goal is to keep the difference between the left wheel's encoder count and the right wheel's encoder count at zero. The picture of the graph at the top and the bottom of the diagram is simply a symbol for the ideal velocity profile derived earlier in this section. The "xMotor-Encoder" blocks represent the input to the motors, and the resulting readings from the encoders.  $\frac{1-z^{-1}}{T}$  is simply the z transform equivalent of "s" in the Laplace domain, which represents differentiation. In simpler terms, that block represents the obtaining of a velocity estimate through the discrete differentiation of the encoder readings.

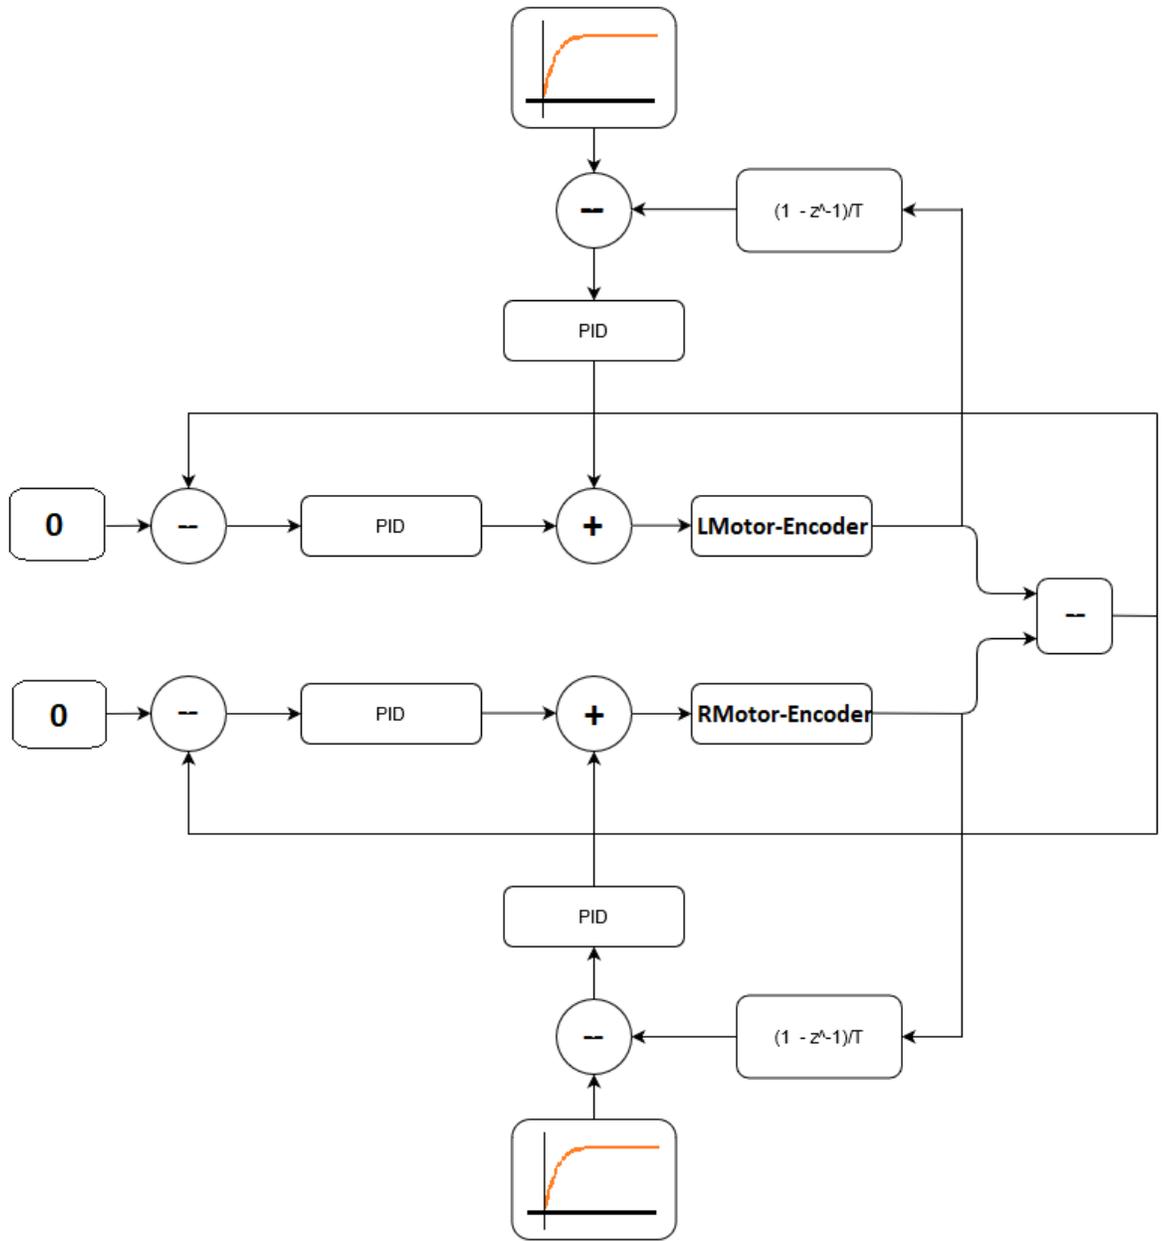


Figure 21 - Controller based on the ideal velocity profile

### 3.2 Estimation

Because of the metal samples that were within the possession of the author, it was known in advance that the robot would have to make accurate turns. For an idea of why the final experiment would require that the robot make turns, refer to Figure 26. The task of making accurate turns called for more sophisticated estimation of the current heading and/or the angular velocity than simply relying entirely on encoders for estimates. Accurately changing the heading of the robot is a much more critical function than moving straight, as incorrect headings cause an accumulation in error over time. It was decided that information from the gyroscope would be fused along with estimates from the encoders to obtain more accurate heading estimations as the robot was performing a turn. In order to implement this, a simple Kalman filter was the method of choice. A Kalman filter is a predictor-corrector type estimator that minimizes the estimated error covariance of states that are being observed. It involves a prediction stage, where the current state and a model of how the system evolves over time are used to predict the next state, and an update stage, where any measurement of the states of interest is used to refine the estimate. The implemented Kalman filter estimated two states; the current angular position and the angular velocity. The gyroscope directly provided an angular velocity estimate, and the encoders provided an angular position and an angular velocity estimate. Since turning was performed by driving one of the two wheels forwards and the other wheel backwards, a formula was derived that took encoder values obtained during this movement and calculated an angular position. Another formula was derived that took the velocity estimates from both encoders and calculated a total angular velocity estimate. These two formulas are:

$$\frac{(\Delta LEnc - \Delta REnc) * 0.28mm/enc\#}{diameter} \quad \text{Equation 3.1}$$

$$\frac{(LVel - RVel)}{diameter} \quad \text{Equation 3.2}$$

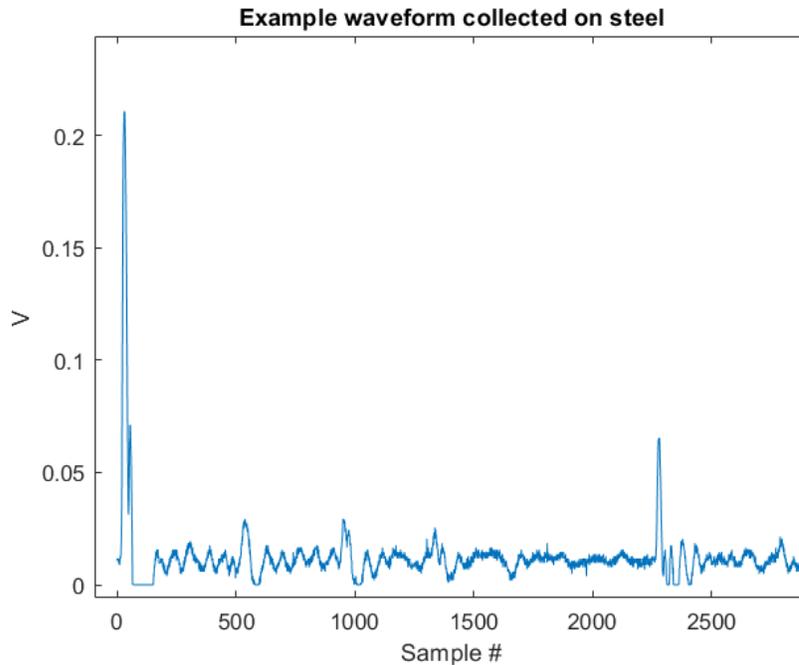
In Equation 3.1,  $\Delta LEnc/\Delta REnc$  are the change in the left/right wheel encoder count from any given starting point, and the equation yields the degree change in radians from the starting point. In both equations, the diameter is the distance between the center of the left wheel and the right wheel. The sensor variances used for the Kalman filter had to be chosen carefully in order to guarantee that the output velocity and position estimates were appropriate. To estimate the variance in the gyroscope output, the robot was kept still, and the gyroscope values were printed over serial to a computer. The recorded values had a slight offset from zero and varied +/- 0.02 radians/second around this offset. The units we used inside the Kalman filter were in degrees or degrees/second, so we set the variance of the gyroscope estimate to 2 degrees/second (rounded up from 1.14 degrees, which is equal to 0.02 radians). The variance of the angular position estimate was technically bounded to the quantization of the encoder or the error in the measurement of the moments when the encoder changed values. The code measuring the amount of time elapsed between encoder ticks did so down to the microsecond, so it shall be assumed that that error is negligible. The diameter of the robot from the left edge of the left wheel to the right edge of the right wheel was measured to be 111 mm. Since the resolution of the encoders is 0.28 mm, we shall assume that the maximum error in an encoder position estimate is this quantization interval, or

0.28 mm. If one wheel is forward 0.28 mm more than expected and the other wheel is backward 0.28 mm more than expected, then the robot's angle is  $(0.28/(111 * \pi)) * 1/360 = 2.2e - 06$  degrees off from the estimated angle. This would theoretically make the variance for the angular velocity estimate from the encoders much less than the previously found variance for the estimate from the gyroscope. However, the encoder estimate would become very inaccurate if there was any sort of slippage, whereas the gyroscope estimate, while far less accurate on average, would remain as accurate as it was before. If there was enough slippage, it could be possible that the gyroscope would become a more accurate measurement of the angular velocity than velocity estimates derived from the encoders. However, the robot would never be able to take advantage of this if the sensor variance for the encoder-based angular velocity estimate was far lower than the gyroscope-based angular velocity estimate, since the Kalman filter would always heavily weigh estimates from the encoders if the variances were set according to each sensor's "ideal" ceiling of error. For this reason, the sensor variance for the encoder-based velocity estimate was simply set to half of the variance for the gyroscope estimate.

### 3.3 Waveform Collection

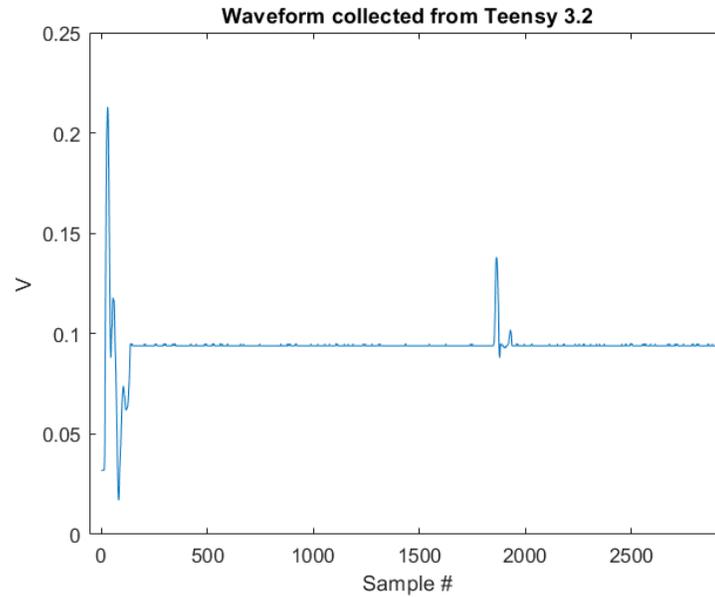
First, an array with space for 6000 integers is preallocated. All it took to induce the production of the initial solitary wave was turning on the solenoid, waiting 600 milliseconds (ms), and turning off the solenoid. This lifted the topmost ferromagnetic particle, held it in place for 600 ms minus the time it took to rise to make contact with the solenoid, and dropped it. It should be mentioned that it did not take 600 ms for the ball to rise to its topmost position, but turning the solenoid on for shorter periods (such as 150-300 ms) resulted in much more inconsistent ToF's.

After the solenoid had been turned off, the first 1000 samples of the preallocated array are treated as a circular buffer, and samples from the ADC that the piezoelectric element is connected to are continually read into this buffer. If a positive trend in the measured voltages is noticed, the program transitions to freely filling the rest of the array with samples of the ADC. After the array is filled, the collection of the waveform is completed and the waveform exists as discrete samples taken from the analog waveform in the memory of the microcontroller. An example of a waveform collected using this method from a Teensy 4.0 is shown in Figure 22.



**Figure 22 - Example waveform collected on steel**

One immediately notices the jitter present immediately after the incident solitary wave. In an attempt to find the source of this jitter, the Teensy 4.0 was switched out for a Teensy 3.2, a different microcontroller board that happens to be pin compatible with the Teensy 4.0, and recorded another waveform with the same program.

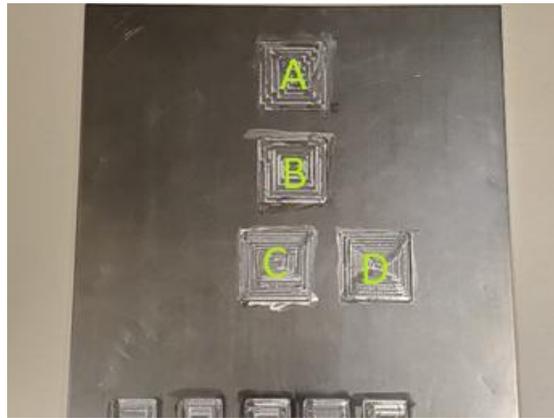


**Figure 23 - A waveform collected from the same sensor and material, but with a different microcontroller**

This makes it clear that the seemingly random appearance of small peaks in Figure 22 is a result of either some flaw within the Teensy 4.0's ADC, or due to some source of internal noise within the Teensy 4.0. The Teensy 4.0 was run at 450 MHz while the Teensy 3.2 was run at 72 MHz, so it is entirely possible that the extra noise seen on the Teensy 4.0 is a result of jitter from the faster clock [18]. In any case, it will be shown that this noise does not make it impossible to extract the information that is relevant to discerning thickness from these waveforms.

## 4.0 Experimental Methods

First, the sensor itself had to be verified. A metal plate 1ft x 1ft x 0.12in (length x width x height) had 4 pockets of varying depths machined into it, as well as a few smaller slots at the bottom that were used to verify a different instrument. Each individual pocket was meant to be a uniform depth. The final machined plate is shown in Figure 24. The 5 slots pictured at the bottom were only used to verify a depth gauge and are not relevant to any other aspect of the experiment.



**Figure 24 - Machined metal plate**

The goal with choosing depths was to test the extent to which the sensor was capable of discerning between different depths. The final depths of the individual squares were not entirely accurate, mostly due to slight warping during the machining process. The range of depths was measured by using a iGaging Digital Depth Gauge on many different points on each of the individual squares. A table with the intended depths and the actual depths is shown.

**Table 2 - The intended depths alongside the measured range of depths. A,B,C,D from Figure 24 correspond to rows 1,2,3,4 in this table.**

| <b>Intended depth</b> | <b>Actual range of depth</b> |
|-----------------------|------------------------------|
| 0.014 in              | 0.009 – 0.015 in             |
| 0.028 in              | 0.031 – 0.041 in             |
| 0.042 in              | 0.038 – 0.042 in             |
| 0.056 in              | 0.054 – 0.058 in             |

The most problematic pocket is the one in the second row, which unfortunately overlaps in depth with the pocket listed in the third row. A set of 4 distinctive depths would provide a greater amount of training and/or calibration data, but 3 distinctive depths would do. The ToF and amplitude ratio (AR) are the two most investigated variables when it comes to corrosion detection, so it was decided that the information provided by the waveform would be “compressed” into three variables; the amplitude of the ISW, the amplitude of the PSW, and the ToF. The combination of these three variables will hereafter be referred to as a “feature vector”. The robot was then placed on A, C, D, and one other uncorroded portion of the plate next to square B. 20 feature vectors were collected after the robot was placed each time. The location of the PSW was obtained simply by finding the maximum voltage measured by the ADC after the ISW had passed by. In cases where there can be likely more than one reflected wave, sometimes this approach will catch a secondary reflected wave instead of the first one, resulting in inaccurate ToF measurements. However, the custom sensor would very rarely generate a secondary reflected wave when sampling from metal (in other words, the waveforms were always similar to the waveforms from Figure 22, so this relatively naïve approach was sufficient. The ToF was calculated by subtracting the previously calculated index of the reflected wave from the index of the peak of the ISW. A good amount of previous work in HNSW’s shows that there is usually a correlation between ToF and/or amplitude ratio (the amplitude of the ISW divided by the amplitude of the PSW) and thickness. In order to

check that this correlation had been observed, an N-way ANOVA test was run on the collected data, treating the feature vector as the factors and the thickness as the response. The N-way ANOVA table is shown in table 2. In this table, “RefA” is short for “reflected wave amplitude”, and “ISWA” is short for “incident solitary wave amplitude”.

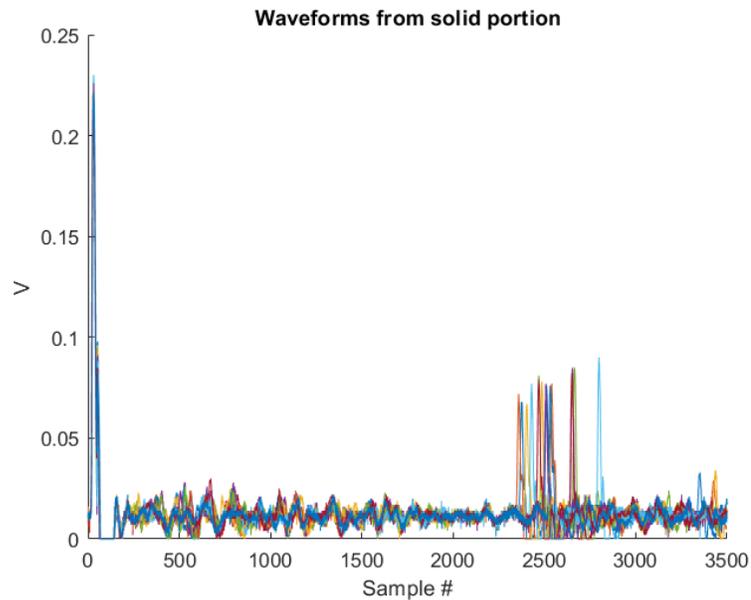
**Table 3 - N-Way ANOVA table from dataset.**

| <b>Source</b>    | <b>Sum Sq.</b> | <b>F</b> | <b>Prob&gt;F</b> |
|------------------|----------------|----------|------------------|
| <b>ToF</b>       | 0.000486       | 2.5763   | 0.11252          |
| <b>ISWA</b>      | 0.00064        | 3.3956   | 0.069168         |
| <b>RefA</b>      | 0.000387       | 2.0561   | 0.15559          |
| <b>ToF*ISWA</b>  | 0.000475       | 2.5224   | 0.11628          |
| <b>ToF*RefA</b>  | 0.001314       | 6.9709   | 0.010003         |
| <b>ISWA*RefA</b> | 0.000308       | 1.64     | 0.2048           |

It is noted that the most variable with the lowest p value (“Prob>F” is the same thing as the p-value) is the interaction variable between the ToF and the amplitude of the reflected wave. The general interpretation for this result is that the effect of the thickness of the sample on the ToF or the amplitude of the reflected wave depends on the value of the other variable [16]. In short, a method that used both variables to discern a thickness would likely perform much better than a method that used one of the two alone. This represents an interesting departure from what is usually seen in HNSW literature, as it is very common for either the ToF or the amplitude ratio alone to be correlated with some feature of interest [13, 19, 20]. This also proves that the noise observed in Figure 22 is not detrimental to the performance of this sensor. However, while it is now known that there is a correlation between the collected feature vectors and differing levels of corrosion,

this table tells us nothing about the nature of this correlation. Finding the nature of this correlation to discern the thickness of a point on the metal plate using waveforms collected from that point.

Figure 25 shows multiple waveforms collected from an uncorroded section of the plate laid on top of each other.



**Figure 25 - Multiple waveforms taken from an uncorroded section of the plate layered on top of each other.**

**Each peak observed in the latter half of this graph is from a different waveform.**

From this figure, it is apparent that a higher reflected wave amplitude is correlated with a larger ToF. It was hypothesized that this relationship might be approximately linear, and so the result of  $ToF / (\text{Reflected wave amplitude})$  for the samples that were used to conduct the N-Way ANOVA test were compared with each other.

**Table 4 - Average ToF/RefA for different depths of corrosion**

| Average ToF/RefA (Sample #/V) | Depth of corrosion (in) |
|-------------------------------|-------------------------|
| 21.07                         | 0 in                    |
| 20.52                         | 0.012 in                |
| 21.23                         | 0.04 in                 |
| 22.21                         | 0.056 in                |

There is clearly a correlation, but it was not strong enough to be convincing. However, it could be used as a starting point for finding a function that would result in a strong correlation between our feature vectors and the corresponding thicknesses. The task of creating a function of our feature vector that extracts the thickness of the sample at that point is a daunting one, as there is an infinite number of possible solutions. For this reason, it was decided that a candidate function would be optimized for a linear correlation between the output of the function and the thickness corresponding to the feature vector. The Pearson correlation coefficient (PCC) is a well-known measure of linear correlation. In statistics, the PCC is generally used to measure the linear correlation between two variables, and the formula for PCC (assuming a pair of variables, X and Y) is  $\frac{cov(X,Y)}{\sigma(X)\sigma(Y)}$ . The coefficient can be comprised between +1 and -1, where 1 is a total positive linear correlation, 0 is no linear correlation, and -1 is a total negative linear correlation [21]. The candidate function in this case was  $ToF^X ISW^Y PSW^Z$ , where X, Y, and Z would be the variables that would be optimized. This formula was chosen simply because of the moderate success of the simpler formula used to create Table 3. The variable that would be optimized for would be the calculated Pearson correlation coefficient between the output of this function for feature vectors

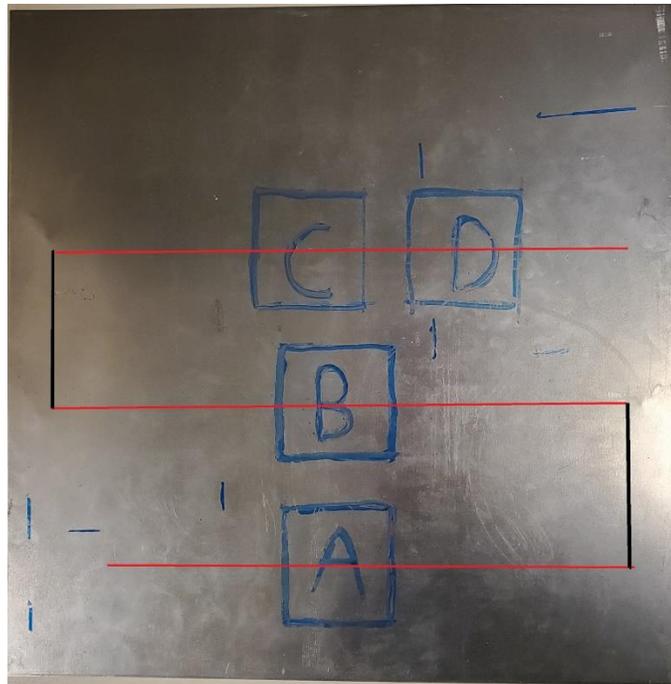
collected from corroded (namely squares A, C, and D)/uncorroded parts of the plate and the thicknesses of those corresponding spots. Succeeding in doing so would guarantee that linear regression could be used to find a linear relationship between this output of this function and the true thickness. Because the state space was not that large, a brute force approach was taken where every value between -1 and 1 (with an increment of 0.1) was considered for X, Y, and Z, using the same dataset used for the N-way ANOVA test discussed earlier. So, the problem could be formulated as:

$$\max \frac{\text{cov}(ToF^X ISW^Y PSW^Z, T)}{\sigma(ToF^X ISW^Y PSW^Z)\sigma(T)} \quad \text{subject to } -1 < X, Y, Z < 1 \quad \text{Equation 4.1}$$

In Equation 4.1, T refers to the thickness of the sample corresponding to the features in the waveform. There were two extrema found after running this brute force search; X = 0.5, Y = 0.7, and Z = -0.3 resulted in a correlation coefficient of -0.9251, and X = 0.6, Y = 1, and Z = -1 resulted in a correlation coefficient of 0.7486. In other words, the negative correlation found seemed like it would be a more accurate representation of the relationship between our collected feature vectors and the measured thicknesses. Linear regression was used to obtain coefficients for both “solutions” to our optimization problem, and these coefficients were saved so they could be used to analyze the results. In short, there were now two equations of the form: *thickness* = *a* \*  $ToF^X ISW^Y PSW^Z$  + *b*. Now that the sensor had been characterized, the next goal was to set up a suitable experiment.

## 4.1 Experimental Setup

Figure 26 is a picture of the plate before it was used for the experiment. Squares corresponding to the corroded sections were manually drawn on the unmachined side of the plate so that the author could place the robot at an appropriate starting point. The image has had red and black lines digitally added to it. A red line indicates a path the robot was programmed to take while collecting samples roughly every 10 mm. The black lines indicate “sample-less” traveling. The robot would start at the bottom left (before square A) and make its way to square D, stopping once it reached the edge of the plate. Each transition from a red section to a black section (or vice versa) necessitated a 90 degree turn in place.



**Figure 26 - Picture of the side of the plate that the robot drove on**

In this case, “collecting a sample” meant that the robot actuated the solenoid 10 times, collected 10 waveforms, and sent 10 corresponding feature vectors using the HC-05 module back to a laptop hooked up to another HC-05 module.

The robot was also programmed to travel across a different metal plate that was 0.18 inches thick. This metal plate had been selectively corroded by being placed in a tank filled with a saltine water solution. There were 6 squares of material that were exposed to this material, however 4 of these squares were corroded all of the way through, so they were unsuitable for inspection with the robot. However, there were two 1”x1” squares that were not corroded all of the way through, so the robot was simply programmed to drive across both of them. The distance between the centers of these two squares was 8”. A picture of both corroded sections is shown in Figure 27.



**Figure 27 - A picture of the underside of the plate. The two corroded sections are outline in blue, and the robot drove on the opposite of the pictured side. Also included is a picture of the robot driving on the uncorroded side.**

## 5.0 Results

For the first experiment on the machined plate, the robot sent 45 sets of 10 feature vectors. The earlier optimized candidate function that had been found to have a significantly negative correlation with the measured thicknesses and the corresponding coefficients found through linear regression were used to generate a set of 10 thickness estimates for each spot that the robot collected samples on. However, the first thickness estimate (which ideally should have been 0.12 in) was -1.66 in. To remedy this, an offset of 1.78 in was added to every thickness estimate. The starting point of the robot, the CAD file submitted to the machine shop to generate the metal plate in the first place, and the thicknesses measured using the digital depth gauge were used to obtain an ideal plot of minimum and maximum thickness vs distance. The two are plotted in Figure 28.

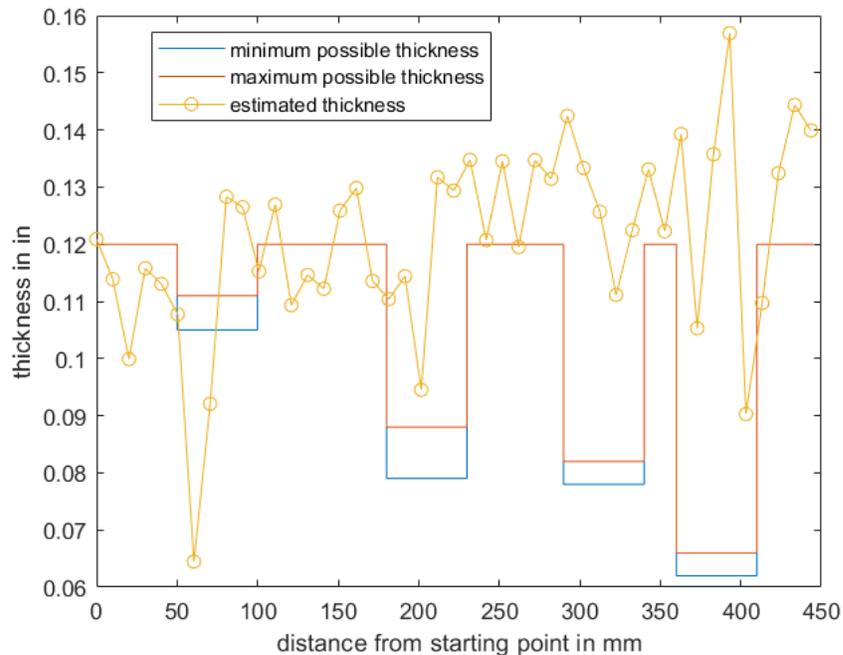
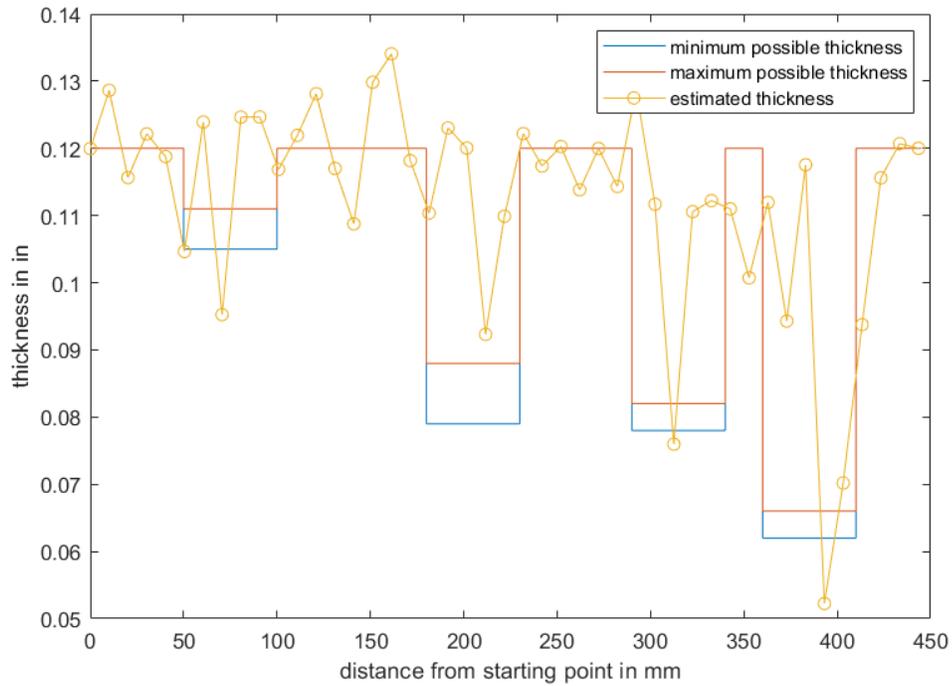


Figure 28 - Plot of ideal thickness vs distance and measured thickness vs distance from the machined plate

Clearly, the robot was capable of determining that the thickness of the plate had changed at certain points, but these observations did not seem to be proportional to the amount by which the thickness changed whatsoever. To further investigate, the same procedure outlined above was used to obtain another plot of thickness vs. distance, but instead the coefficients of X, Y, and Z that resulted in a positive correlation and the coefficients from the resulting linear regression were used to obtain a thickness estimate instead. As mentioned earlier, the negative correlation coefficient found was -0.9251, and the positive correlation coefficient found was 0.7486. In general, a correlation coefficient greater than 0.8 is considered to be strong evidence of a linear relationship, and the positive correlation coefficient clearly does not meet this criteria[22]. So, one would not expect the formula derived from the positive correlation coefficient to work as well. Similarly to the last set of thickness estimates, the first thickness estimate was not 0.12 in, so all of the thickness estimates had an offset of 0.39 inches added to them to get a set of thickness estimates that made sense. Figure 29 shows the resulting graph of thickness vs distance.

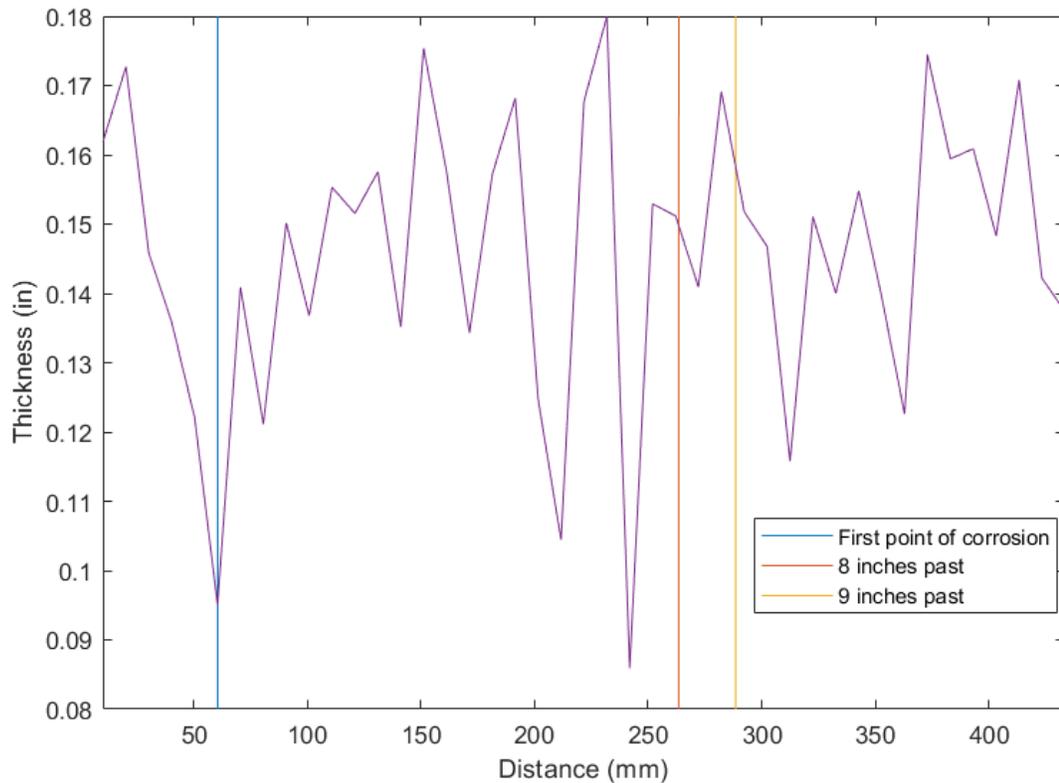


**Figure 29 - A second graph of thickness vs distance obtained using the alternate formula.**

Clearly, the alternate formula did a better job of distinguishing properly between the differing levels of corrosion. The maximum thickness measured is also much more reasonable, as in Figure 28 the measured thickness climbs to 0.16 inches, though the plate itself is no more than 0.12 inches thick.

For the larger plate that had been corroded by the saline solution in a tank, there was no guarantee that our model would be extensible to a thickness of 0.18 in, and the exact type of steel might have been different as well. More importantly, there was no source of a ground truth for the thickness of the corroded sections. For this reason, the only goal is to observe two outliers on the plot that are somewhere between 8 and 9 inches away from each other. The robot collected 43 sets of 10 feature vectors every 10 mm as it traveled. The heading had to be corrected as it drove towards the second corroded section, not because the robot itself veered off course, but because

the initial heading was not set precisely enough. The coefficients of X, Y, and Z that resulted in a positive correlation coefficient (and also were more successful in distinguishing between depths, as shown in Figure 29) were reused to turn each feature vector into a “thickness coefficient”. The two largest and smaller coefficients from each set of 10 were thrown out to prevent the effect of potential outliers. Unlike the experiment on the machined plate, the coefficients derived from the linear regression performed earlier were not used. Instead, the largest thickness coefficient measured from the set of samples was assumed to correspond to 0.18 in, the smallest thickness coefficient was assumed to correspond to 0.09 in, and the rest of the measured coefficients were simply linearly mapped to thicknesses between those two values. Again, we have no ground truth for the actual “ideal” depth that should have been measured at a point, since the corrosion was induced naturally, so these assumptions are mostly meaningless since the only goal is to find two outliers that are the proper distance (8 or 9 inches) from each other.



**Figure 30 - Graph of thickness vs distance obtained from the larger plate.**

The blue line in Figure 30 represents the estimated point where the sensor on the robot first makes contact with the first corroded section of the plate it was meant to drive over. The orange line is drawn 8 inches after the blue line, and the yellow line is drawn 9 inches after the blue line. Ideally, the two large downward spikes would appear in that region. However, it is possible that the robot was moved forward more than intended when the heading was corrected during the experiment. Additionally, the start point of the first region of corrosion could be as much as 10 mm earlier than estimated on this graph.

## 6.0 Future Work and Conclusions

### 6.1 Future Work

However, its performance in determining relative depths seemed to depend on the choice of model, when both models should have been suitable. The output of both models also needed an offset in order for the results to make sense. This brings up an important point; anytime the robot underwent a significant mechanical disturbance, it was likely that the sensor became a “different” sensor. The behavior of the particles and the sensor particle depends on a number of very delicate mechanical factors; the friction of each individual particle against the walls of the sensor, the angle of the sensor particle within the casing, and possibly the orientation of each individual particle (assuming that they are not perfect spheres). The robot moving itself likely induced a mechanical disturbance within the sensor casing, meaning the sensor could have been fundamentally different after every movement. The effect of this disturbance on the thickness calculations can be thought of as noise, and the real changes in thickness in the plate can be thought of as a signal. The reason moving the sensor around manually and taking measurements manually is not sufficient as a proof of concept is because the only way to find the limitations that the signal-to-noise ratio imposed upon the viability of this concept was to build this robot. The fact that one of two models seemed to result in accurate thickness measurements from the plate suggests that there should be a way to create a model that can account for mechanical disturbances. Perhaps multiple sensors could be used to create different sets of training data, which could then be used to figure out how to account for mechanical differences. The resulting model might be resistant to mechanical disturbances. Simple mechanical improvements to the sensor could also go a long way in terms of increasing

the reliability of the sensor. There might also be other hidden features in the waveform that have not been accounted for thus far; for example, maybe the width of the incident solitary wave and/or the reflected wave have some relationship with other properties of materials. It was also obvious that the thickness of the corroded sections was “clearest” when the sensor was placed near to the center of the corrosion. If it had been known in advance that there would be a result like that, the plate would have been machined to have corroded sections of different sizes. Perhaps the response of the sensor is more dependent on an amount of supporting material surrounding the point of impact, instead of just the thickness of the plate at the point of impact.

## **6.2 Conclusions**

In conclusion, three separate devices designed to capture HNSWs were discussed in this thesis; two modules designed for monitoring applications, and one robot intended to act as a proof of concept for the possibility of the sensor being used for a one-time inspection. These devices were validated with three separate experiments. The first generation wireless device was proven to be able of taking measurements from transducers and using them to differentiate between materials. The second generation wireless device was capable of running for much longer periods of time and displayed an increase in sampling frequency by a factor of ~80 compared to the first generation wireless device. It also experimentally displayed the capability to make reliable measurements across the span of a week. The robot traversed a plate that had pockets drilled into it to simulate corrosion, and was proven to be capable of detecting outliers within the plate.

## Bibliography

- [1] P. Rizzo, "Water and Wastewater Pipe Nondestructive Evaluation and Health Monitoring: A Review," *Advances in Civil Engineering*, vol. 2010, p. 818597, 2010/05/19 2010.
- [2] W. A. Grandia and C. M. Fortunko, "NDE applications of air-coupled ultrasonic transducers," in *1995 IEEE Ultrasonics Symposium. Proceedings. An International Symposium*, 1995, vol. 1, pp. 697-709 vol.1.
- [3] A. Nasrollahi, W. Deng, P. Rizzo, A. Vuotto, and J. M. Vandebossche, "Nondestructive testing of concrete using highly nonlinear solitary waves," *Nondestructive Testing and Evaluation*, vol. 32, no. 4, pp. 381-399, 2017/10/02 2017.
- [4] A. Nasrollahi, M. S. Orak, A. James, L. Weighardt, and P. Rizzo, "A Nondestructive Evaluation Approach to Characterize Tennis Balls," *Journal of Nondestructive Evaluation, Diagnostics and Prognostics of Engineering Systems*, vol. 2, no. 1, 2018.
- [5] B. Zheng, P. Rizzo, and A. Nasrollahi, "Outlier analysis of nonlinear solitary waves for health monitoring applications," *Structural Health Monitoring*, vol. 19, no. 4, pp. 1160-1174, 2020/07/01 2019.
- [6] R. Misra, H. Jalali, S. J. Dickerson, and P. Rizzo, "Wireless Module for Nondestructive Testing/Structural Health Monitoring Applications Based on Solitary Waves," *Sensors*, vol. 20, no. 11, 2020.
- [7] H. Jones *et al.*, "Automated Analysis, Reporting, and Archiving for Robotic Nondestructive Assay of Holdup Deposits," *ArXiv*, vol. abs/1901.10795, 2019.
- [8] EddyFi. Available: <https://eddyfi.com/en/product/linetrax-small-pipe-inspection-crawler>
- [9] T. Tours. Available: <https://turnstile-tours.com/from-the-brooklyn-navy-yard-to-mars-honeybee-robotics/>
- [10] N. Jog, "Highly Miniaturized Robots for Inspection of Small Nuclear Piping," 2019.
- [11] G. Dobie, R. Summan, S. G. Pierce, W. Galbraith, and G. Hayward, "A Noncontact Ultrasonic Platform for Structural Inspection," *IEEE Sensors Journal*, vol. 11, no. 10, pp. 2458-2468, 2011.
- [12] S. N. Easwaran and R. Weigel, "1.3A,-2V Tolerant Solenoid Drivers for Pedestrian Protection in Active Hood Lift Systems," in *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2018, pp. 1-4.
- [13] H. Jalali and P. Rizzo, "Highly nonlinear solitary waves for the detection of localized corrosion," *Smart Materials and Structures*, vol. 29, no. 8, p. 085051, 2020.

- [14] Q. Torgerson, L. Burner, and A. Miller. (2014). Available: <https://github.com/QuentinTorg/EncoderMod>
- [15] A. Ayachit and M. K. Kazimierczuk, "Averaged Small-Signal Model of PWM DC-DC Converters in CCM Including Switching Power Loss," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 66, no. 2, pp. 262-266, 2019.
- [16] (2021). *Main Effects and Interaction Effect*. Available: <https://stats.libretexts.org/@go/page/2904>
- [17] L. Zhang, P. Slaets, and H. Bruyninckx, "An open embedded industrial robot hardware and software architecture applied to position control and visual servoing application," *Int. J. of Mechatronics and Automation*, vol. 4, pp. 63-72, 01/01 2014.
- [18] T. Instruments. (2020). *Fundamentals of Precision ADC Noise Analysis*. Available: <https://www.ti.com/lit/eb/slyy192/slyy192.pdf>
- [19] P. Rizzo, A. Nasrollahi, W. Deng, and J. M. Vandenbossche, "Detecting the Presence of High Water-to-Cement Ratio in Concrete Surfaces Using Highly Nonlinear Solitary Waves," *Applied Sciences*, vol. 6, no. 4, p. 104, 2016.
- [20] T. Singhal, E. Kim, T.-Y. Kim, and J. Yang, "Weak bond detection in composites using highly nonlinear solitary waves," *Smart Materials and Structures*, vol. 26, no. 5, p. 055011, 2017.
- [21] P. Di Lena and L. Margara, "Optimal global alignment of signals by maximization of Pearson correlation," *Information Processing Letters*, vol. 110, no. 16, pp. 679-686, 2010/07/31/ 2010.
- [22] H. Akoglu, "User's guide to correlation coefficients," *Turkish Journal of Emergency Medicine*, vol. 18, no. 3, pp. 91-93, 2018/09/01/ 2018.