The Development of a Novel Capacitive Water Conductivity Sensor

A THESIS SUBMITTED TO THE FACULTY OF THE UNIVERSITY OF MINNESOTA BY

KEIRAN CANTILINA

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF MASTER OF SCIENCE

ADVISOR: PETER MARCHETTO

MAY 2018

Acknowledgements

Many thanks to Pete who generously invited me to serve as one of his first graduate students. I am privileged to have benefited from his guidance, mentorship, and bad puns during my time at the University of Minnesota. Without Pete, this thesis would have never happened.

I am also indebted to Joe Magner and Andy Wickert who agreed to be my thesis committee members. Their willingness to take on yet another responsibility added to already full plates is much appreciated. Both of them also lent lab space and resources to this project, which proved to be extremely helpful.

Here's a shoutout to Dan Furuta who was an invaluable partner-in-crime on this project. Many thanks to him for having the wisdom to try a solution I had dismissed as unlikely to succeed (which then became a critical part of the design), and also for giving it to me straight when I'm wrong. I also appreciate the time and personally owned equipment he generously donated during the feasibility testing and characterization stages of this project. The success of this project is his bragging right at least as much as it is mine.

Thanks to Hanna Lin and Dane Kouttron for moral support, ideas, heckling, and that giant bottle of rum. The Ballmer Peak is a real thing guys. Special thanks to Hanna for patiently listening to me vent (the dishes are far cleaner than I deserve), and special thanks to Dane for the extra-special pat down I got from the TSA at MSP (what the hell did you have on your hands?).

If Bobby sees this, thanks to him for lending me space in Andy's lab to do my testing work. He endured my many questions and my attempts at small talk while he was doing SMT soldering work. sorry

To many unnamed people who were present at the Hackaday Superconference and at the IEEE Instrumentation conference in Torino, thank you for having good ideas and sharing them. Sometimes hearing others present good ideas inspires one to develop good ideas in turn.

To my parents, many thanks for supporting years of schooling and for impressing on me the value of pursuing truth. My career path is indebted to them for the hard work they put into raising me.

Last but not least, I must convey my sincere appreciation for the work done by Joe Warhol to keep things running smoothly in the background. Many thanks for his wisdom and advice in tough times, and for his insight on life's little mysteries (aka people). I'm getting better at remembering to go have some fun.

Dedication

This thesis is dedicated to my wonderful wife Ren, who has my sincere and everlasting appreciation. Her encouragement and support has made this thesis possible without the loss of my sanity. I am thankful for her immense patience through everything from midnight freak-outs to wet floors caused by me running out of the shower to scribble madly on the whiteboard. Also she told me about Mendeley, without which I would probably still be struggling to format my citations.

Abstract

A novel sensor was developed which measures water conductivity by tracking the capacitance of two insulated probes submerged in the water under test. Unlike commercially available water conductivity sensors based on 4-terminal potentiometric methods this sensor does not require metal to be in direct contact with water, and this sensor also has much higher sensitivity and lower power consumption compared to induction-based conductivity sensors. These differences make this sensor uniquely suited for use in distributed sensor networks, where resistance to corrosion and fouling, low power consumption, and reliability are necessary traits. In addition, this sensor is a good fit for use by citizen scientists or hobbyists due to the ease of assembly and the low price of components needed to construct the device.

Table of Contents

List of Figures

Introduction

PART I: Context of work

Salts in the environment can be useful for humans, but detrimental to natural water systems. Various chloride and acetate salts are used for winter road salting, and nitrate and phosphate salts are extremely important as fertilizers for agriculture.

According to 2008 paper published by the St. Anthony Falls Laboratory at the University of Minnesota, the Twin Cities Metropolitan Area sees more than 317,000 tons of road salt applied to roads every year (Novotny, Murphy, & Stefan, 2008). While road salt is absolutely necessary to maintain road safety during Minnesota's harsh winters, problems can arise when salt left over from the winter moves into surface water or groundwater during the spring melt. For watersheds associated with urban areas, this additional salt may seriously disturb freshwater ecosystems; in urban lakes, salt concentrations may be high enough to force stratification of the water column, thus preventing the turnover which would usually bring oxygen to the bottom of the lake (Novotny et al., 2008).

Similarly, runoff of fertilizers (which often take the form of nitrogen or phosphorus-containing salts) from agricultural fields or stockyards can have negative effects on freshwater systems. In fact, the most common impairment of surface waters in the United States is eutrophication caused by nitrogen and phosphorus pollution from nonpoint (usually agricultural) sources (Carpenter et al., 1998). Eutrophication can completely devastate a lake or river ecosystem;

the algal bloom that frequently accompanies eutrophication consumes nutrients, decreases dissolved oxygen in the water and may cause fish kills. Toxins (such as microcystin) produced by cyanobacteria in a bloom may produce water treatment problems and cause water to be unsuitable for drinking and recreation.

Fortunately, much research is being done to find solutions to the problems caused by human-applied salts in the environment. As of 2018, a quick search for “solutions to nonpoint source pollution” on Google Scholar reveals hundreds of highly-cited papers tackling the problem, with suggested solutions ranging from new governmental policies and regulations to novel devices and methods for capturing or eliminating the pollution. However, there are significant challenges to implementing these solutions: policy and regulation requires some way to enforce compliance and discover noncompliance, and mitigation methods usually require a way to identify specific targets. Both of these needs are satisfied if it is possible to spatially locate the source of the pollution.

Unfortunately, nonpoint pollution is exactly that – “nonpoint,” that is, a specific source of the pollution is difficult to spatially locate. Sometimes the cause of this difficulty is because the source of the salt pollution lies over a large area, or perhaps the source is intermittent. In other cases, the source might move, or the salt pollution might only be detectible after it has moved far from its source.

Typical methods for measuring concentrations of salts in water systems usually include taking water conductivity measurements. These measurements have been traditionally taken by hand at one location at a time, using a water conductivity sensing system (US EPA, 2012). Because of the limitations intrinsic

to taking measurements by hand at discrete points in both space and time, it is very hard to locate or measure/monitor pollution sources that vary over space and/or time. Thus, the obvious solution is to increase the distribution of the sensing system in order to keep up with the distribution of the salt pollution; if a nonpoint pollution source needs to be measured, perhaps a nonpoint sensing system should be used.

This “nonpoint” sensing system can take the shape of a distributed sensing network. In a network of this sort, many sensors are distributed over a wide area. These sensors all constantly take measurements over time, and the measurements taken together can be used to obtain an idea of the spatial distribution of the measurement parameter over time. In the case of water conductivity, this network could take the form of many independent conductivity sensors distributed over a water body or watershed, all recording measurements in parallel and transmitting the data back to a single base station. For example, if the task at hand is to monitor a nonpoint source polluting a lake, the sensor network could consist of many sensors installed along the shoreline of the lake. If the goal is to find a nonpoint source in a watershed consisting of rivers and streams, the network might consist of sensors placed at each major branch in the river or stream network. In all of these cases, the output would be data that can be used to create a spatial map of water conductivity as it varies over time.

The challenges in establishing a sensor network as opposed to the traditional by-hand method of taking measurements is twofold: it must be possible to deploy an appropriate number of sensors without prohibitively high

cost, and without prohibitively high maintenance requirements. Thus, for conductivity sensors, the sensors must be cheap, resistant to fouling or corrosion, and not require maintenance. Of course, to be useful the sensors must also be accurate, easy to calibrate and install, and sensitive over the range of expected water conductivity.

Currently, commercially available conductivity sensors are poor candidates for application in a distributed sensing network (as explained below).

PART II: Current state of the field

Extant commercially available conductivity sensors come in two varieties: 1) direct-contact electrode resistance sensors and 2) magnetic induction conductivity sensors (Ramos, Pereira, Ramos, & Ribeiro, 2008). Both of these types of conductivity sensors have weaknesses that render them inappropriate for use in a distributed sensing network.

Direct-contact electrode type sensors are mechanically and electrically very simple. The sensor consists of 2 to 4 metallic electrodes that are kept submerged in the water to be measured. In 2 electrode sensors, a voltage is generated between the electrodes and the current flowing between the electrodes is measured. In 4 electrode sensors, two of the electrodes are used to generate a constant current, while the other two electrodes measured the voltage difference between them. Using Ohm's Law:

$$V = IR$$

which dictates the relationship of voltage (V) with current (I) and resistance (R), the resistance of the water can be found. Because conductivity is resistance over distance, conductivity of the water can be calculated by dividing the resistance by the distance between the electrodes (Ramos et al., 2008).

While direct-contact type sensors have advantages mostly stemming from their mechanical and electrical simplicity, they have serious deficiencies in monetary and maintenance costs. Metal supporting an electrical charge in direct contact with water experiences various electrochemical effects, which can lead to corrosion, fouling, or other types of deterioration. To avoid these issues, high end direct-contact sensors (such as those used by most research institutions) use platinum wire for their electrodes. Although platinum has excellent anti-corrosion properties, it has fairly poor pecuniary properties. Low end direct-contact sensors forgo precious metals for stainless or high-nickel steels, but these metals are not completely impervious to electrochemical effects and must be recalibrated (and eventually replaced) regularly. Both low and high end direct-contact sensors are vulnerable to fouling from biological and mechanical sources such as algae growth or sediment deposition.

Magnetic induction sensors are far more resistant to fouling but are much more electrically complex. These sensors consist of two mechanically parallel coils of wire (usually coiled around a core) separated by a set distance. A sinusoidally varying voltage is applied across one coil, and the current flowing through this coil results in a magnetic field perpendicular to the plane of the coil. This magnetic field induces a measurable voltage in the second coil through the

principle of magnetic induction. The efficiency of this magnetic coupling between the two coils is related to the conductivity of the medium between the coils because the water *also* appears as a coil to the magnetic field and will steal some of the energy from the field (Wuliang Yin, Peyton, Zysko, & Denno, 2008). As the water conductivity increases, more of the magnetic field energy goes into inducing a voltage in the water, and proportionally less voltage appears on the second coil. This relationship is quantified by the following formula (Striggow & Dankert, 1985):

$$U_4 = \frac{n_4}{n_1} \frac{1}{1 + n_4^2 R_W \left(\frac{1}{R_A} + \frac{1}{j\omega L_{44}} \right)} U_1.$$

where U_4 is the voltage on the second coil, R_W is the resistance of the water, U_1 is the voltage on the first coil, n_4 is the number of windings on the second coil, n_1 is the number of windings on the first coil, R_A is the resistance of a resistor used to terminate the second coil, L_{44} is the combined inductance of the two coils, j is the square root of -1, and ω is 2π multiplied by the frequency of the sinusoidal voltage applied to the first coil (Striggow & Dankert, 1985). Again, once the resistance of the water has been determined, conductivity can be calculated if the distance between the coils is known.

As can be seen, interpreting measurements from magnetic induction conductivity sensors is a much more complex affair compared to direct-contact sensors. However, because magnetic induction occurs without direct metallic contact, inductive sensors can be completely enclosed which greatly improves

their resistance to corrosion and fouling compared to direct-contact sensors. Unfortunately, this improved ruggedness comes at the expense of sensitivity and power consumption. Inductive sensors have poor performance at low levels of water conductivity, and attempts to increasing sensitivity of these sensors involve either increasing the number of coil windings (which affects n_1, n_4, and L_{44}), or increasing the supply voltage (U_1). Both of these changes have the summary effect of vastly increasing power consumption to a degree such that magnetic induction sensors are only used for oceanographic purposes in seawater, which has a much higher nominal water conductivity than inland freshwater.

Thus, it can be seen that neither direct-contact type nor magnetic induction conductivity sensors are appropriate for the purposes of identifying nonpoint sources of salt-based pollution.

In this work, an attempt has been made to develop a third kind of sensor, based on the principle of capacitance. This sensor does not rely on direct contact of metals with water, but unlike the magnetic induction sensor this sensor retains usable sensitivity at the low end of the range of water conductivities normally encountered in freshwater studies and also does not have impractical power requirements. While there are a few cases of capacitive water sensors reported in the literature, these sensors sense only presence or absence of water and in some cases are used for water level (ie. Gauge) sensing (Reverter, Li, & Meijer, 2007). The development of the sensor reported in this work overcomes these and other limitations of currently available water sensors in order to produce a

non-contact capacitance-based water conductivity sensor appropriate for use in a distributed sensor network application.

Summary of Design Requirements

The list below is a summary of attributes that the design of the sensor must have in order to meet the requirements for use in a distributed sensing network.

- The sensor must feature no direct contact of metal with water
- The sensor must be useful across the entire conductivity range from DI water to mine runoff
- The sensor must be reasonably inexpensive (below $30/unit) so that many can be used together
- The sensor must not require an inordinately large enclosure
- The sensor must have sufficiently low power consumption (<< 1W) such that a sensing system might operate for months at a time on a single set of batteries.
- The sensor must be able to be fabricated using techniques easily available to hobbyists
- The sensor must be relatively easy to calibrate
- The sensor must not require complex or expensive circuitry to drive and read
- The sensor must be easy to integrate with networking or mesh networking hardware

Description of Project Goals

There were four goals for this project.

1) Design a sensing strategy and mechanism that fulfills the previously described design requirements.
2) Implement the strategy in the simplest, least expensive way possible.
3) Characterize properties of the sensor, including linearity of response, susceptibility to interference, stability of measurements over time, sensitive, and range.
4) Develop several possible approaches to making the sensor practically useful within a sensing system.

The fulfillment of these goals is described in the following sections.

Theoretical Design

The basic design of the water-conductivity sensor resembles a parallel-plate capacitor. The sensor consists of two legs, each made out of a conductive material coated in a thin, uniform layer of insulating material. This insulating material isolates the conductive material of the sensor from the water. The legs of the sensor are immersed in the water to be tested, and the capacitance of the sensor will vary depending on the conductivity of the water.

In their simplest form, capacitors consist of a sandwich of two electrical conductors with a dielectric substance between them. When a voltage difference is exerted between the two conductors, potential energy is stored in the capacitor in the form of a static electric field between the conductors as opposite charges accumulate on the conductors. The dielectric material serves to prevent charges from passing between the conductors.

In an ideal parallel-plate capacitor, the conductors take the form of two metallic plates with an insulator between them. The energy storing capability (capacitance) is described by the formula:

$$C = \frac{\varepsilon A}{d}$$

Where capacitance (C) is determined by the area of the conductors (A), the distance between them (d), and certain electrical properties of the substance between them (permittivity, ε).

In this conductivity sensor, A and d remain constant, while ε varies with the conductivity of the water. In reality, the situation is slightly more complex because of the presence of insulating material on the legs of the sensor between

the conductive elements and the water. Thus, there are two ε's; one describing the permittivity of the insulating coating, and another describing the permittivity of the water. However, since the insulating coating ideally does not vary in permittivity over time, the capacitance of the system should vary only in relation to the conductivity of the water.

Capacitance of this sensor can be measured by imbedding the sensor in an electrical circuit colloquially called a "tank circuit." This circuit consists of one or more capacitors and inductors connected together such that the assembly behaves as an electrical resonator. Such a circuit will have a characteristic resonant frequency (or multiple frequencies) determined by the values of the components in the circuit. In this case, the capacitance of the sensor will vary relative to the conductivity of the water between its legs, which will in turn vary the resonant frequency of the tank circuit.

Tank circuits have a characteristic frequency response; that is, a sinusoidal signal applied as an input to the circuit will produce an output dependent on the frequency of the input signal. Thus, if the frequency of the input signal is varied while the input magnitude stays constant, both the magnitude and frequency of the output signal will vary. Conversely, if the frequency and magnitude of the input signal both are held constant but the resonant frequency of the tank circuit changes, the magnitude of the output from the circuit will also change.

In summary, the design of this sensor is such that, if a sinusoidal signal (or signals) with static frequency and magnitude is fed to the sensor, the sensor will

produce an output signal with a magnitude proportional to the conductivity of the water between the legs of the sensor.

Implementation

From the previous section, it can be seen that there are several components critical to the electrical design of this sensor:

- Conductive material to form the legs
- Insulating material of very uniform thickness and stability to cover the legs
- Electrical components to form the tank circuit
- Circuitry to produce the sinusoidal signal fed to the sensor
- Circuitry to read the magnitude of the signal output from the sensor

The first two points were tackled by the use of magnet wire. Magnet wire is copper wire commonly used to form the windings on electromagnets or transformers. It is produced in a variety of thicknesses and with different insulating coatings to suit various industrial applications, and these properties are standardized and well-controlled by manufacturers. In addition, magnet wire is very cheap, easy to manipulate by hand, and relatively easy to solder to. Luckily for this project, there was magnet wire available at hand that came manufactured with a coating resistant to degradation by water.

For the third point, it was found by chance that the configuration of the magnet wire itself in the prototype sensor (in combination with a single resistor and the coaxial cables used to connect it to test equipment) had enough parasitic inductance to form a tank circuit with a reasonable resonant frequency when considered with the capacitance of the sensor legs.

The last two points were solved by the use of a Tsunami Board manufactured by Arachnid Labs. This board consists of an Arduino-compatible ATmega32u4 microcontroller ganged with a Direct Digital Synthesis chip (AD9838) and a high-speed comparator. Also on board are a reference clock built around a precision crystal, a USB interface, and full-size BNC connectors.

Using the DDS chip, this board can produce a sinusoidal signal of arbitrary frequency between 0 Hz and 2 MHz and it supposedly can vary the magnitude of said signal between 0V and 6V, which made it suitable for producing the input signal for the sensor. The comparator on the Tsunami was also deemed to be useful for reading the output signal from the sensor. The precision temperature-invariant reference clock means that the DDS output frequency is very stable, and the USB interface and microcontroller made for an easily programmable and configurable platform that was thought to be potentially very useful for integrating the sensor into a field-capable sensing system.

Additionally, the Tsunami board is relatively cheap ($60) and completely open source. This means that all aspects of the design necessary to replicate the device (schematics, source code, Gerber files required for manufacturing, etc) are freely available to the public via Arachnid Lab's Github pages. Since one of the project requirements was to make the sensor easy to replicate by anyone with hobbyist level skills, equipment, and budget, the open source nature of the Tsunami board was very attractive. Also, I happened to have one of these boards in my desk from another project and it was convenient to use.

Of course, there are also practical considerations which necessitated a few more components:

- Some means by which the various electrical parts are connected together
- A mechanical means by which the parts of the sensor are held rigidly and stably together
- Anyone which has ever dropped their phone in a toilet knows that there needs to be some means by which the electrical parts of the circuit that aren't meant to touch water are kept away from the water

The first point was solved internal to the sensor by using more magnet wire, and external to the sensor by using BNC-connector terminated RG6 coaxial cable. Coaxial cable, usually used for radio applications, was deemed necessary because the resonant frequency range of the sensor fell within the domain of high frequency radio.

The second point was solved in the prototypes by using 1/8 inch thick laser-cut acrylic plastic, cut and hot-glued together into a skeleton for the sensor to be built on. Acrylic plastic has the benefits of being very cheap, widely available, and easy to laser-cut. It has very good insulating properties (and thus does not interfere with the function of the sensor) and its mechanical and electrical properties are not much effected by water. In addition, acrylic plastic is transparent by default, which is useful for observing potential problems including cracks, leaks, corrosion of the magnet wire (should the insulation fail), and biofouling.

The legs of the sensor and some cabling were additionally secured by the use of plastic zip ties.

The last point was solved in a coarse and primitive (but very effective and fast) manner by potting the entire device (minus the legs) in two-part epoxy resin, using a chocolate candy form from Amazon as a mold.

The total cost of the most current prototype is around $70, including the Tsunami board. Excluding the Tsunami board, and considering the resources and skills of a typical electronics hobbyist (not including theft), the sensor could conceivably be constructed for as low as $0.

Testing and Experimental Methods

Four stages of experimental testing were used to evaluate the sensor's performance. An overview of the goals for each stage is presented below, followed by a more detailed description of the testing method.

Stage one evaluated basic operation of the sensor in order to confirm that the hypothetical sensing method was legitimate. Questions answered by stage one testing included:

- Can the sensor differentiate between water of different conductivities (i.e. does the sensor work at all in a way useful to users)?
- Do measurements remain stable or drift over time?
- Do different salts in the water effect measurements?
- What is an optimal operating frequency?

Stage two evaluated the sensor's susceptibility to interference of various kinds, along with basic waterproofness and robustness testing. Some questions answered by stage two included:

- Is the sensor subject to interference from:
    - Metal objects near the sensor
    - People
    - Vibration
    - Orientation of its input feed cables
    - Orientation of the sensor
    - Liquid-air boundaries

- Does the sensor leak if left for an extended period of time?
- How does the sensor hold up to mechanical impact and stress? (i.e. What happens if the sensor is dropped and then run over by a 55 gallon drum on a dolly?)

Stage three evaluated the response of the sensor to different levels of water conductivity and attempted to gather data to draw calibration curves. Stage three testing also featured comparison with a conventional commercial water conductivity probe. In addition, some prototype code on the Tsunami board was tested in an attempt to do data collection, processing (ie. conversion of capacitance measurements to conductivity values), and data logging all on a single platform in real time. Some questions answered by stage three included:

- What is the characteristic response of the sensor?
- What does a calibration function for the sensor look like?
- How does the sensor compare to a commercial resistance-measuring conductivity sensor?
- Is it feasible to control all aspects of a sensing system for this sensor on a single microcontroller?

Lastly, stage four directly measured electrical properties of the sensor using an impedance analyzer. The purpose of this stage of testing was to measure the impedance of the sensor in various configurations, and use this information to calculate equivalent lumped-element capacitance and inductance of the sensor. These values are important for fine-tuning the theoretical

computational model of the sensor. Some questions answered by stage four included:

- What is the series and parallel capacitance between the input and output of the sensor?
- What is the series and parallel inductance between the input and output of the sensor?
- What is the ohmic resistance between the input and output of the sensor?
- What are the values of the above parameters within the input and the output alone?

Below follows a further description of each testing stage.

STAGE 1: Jar Testing

Stage 1 testing was relatively unstructured. For all stage 1 tests, the sensor was driven by function generator which generated a 100ms sinusoidal sweep, typically from 100Khz to 5Mhz. The magnitude of the input sweep ranged from 1Vpp to about 10Vpp. The output from the sensor was read on a digital oscilloscope.

For the first set of tests, three pint-sized canning jars manufactured by Ball were filled with distilled water. A random quantity of table salt was dissolved in one of the jars. One jar of distilled water was designated as a wash jar, and the other distilled water jar was designed the negative control. The sensor was first rinsed in the wash jar, then dipped into the control jar. The shape and magnitude

of the output waveform was noted, then the sensor was dipped into the jar with the salt water. Observations were noted again. This process was then repeated several times, occasionally with different sweep parameters. This testing process confirmed that the output of the sensor does vary, at least to some degree, with water conductivity.

For the second set of tests, the sensor was suspended in a single jar filled with a table salt solution of random concentration. The setup remained static for 3 days, and the output of the sensor was noted at regular intervals. The liquid level in the jar was carefully monitored, and deionized water was occasionally added to the jar to make up for water lost by evaporation. This experiment was then repeated with distilled water, and then with a saturated salt solution. The testing process confirmed that the output of the sensor does not vary significantly with time, though the little variation that was found is consistent with the effect of temperature changes on water conductivity.

For the third set of tests, two jars of salt solutions were prepared. One jar contained a sodium chloride solution, and a second jar contained a sodium nitrate solution. As expected, the sensor showed a difference in output for each, compared with distilled water. It was possible to add sodium nitrate to the second jar until the output of the sensor was nearly identical to the output obtained with measuring the sodium chloride solution. This process confirmed the expectation that the sensor does not discriminate between dissolved ion species.

The final set of stage 1 tests was used to find an optimal input frequency for the sensor. The sensor was immersed in DI water and then fed with the

normal 100Khz to 5Mhz sinusoidal sweep, while a digital oscilloscope was used to monitor the output. A frequency was found within the sweep for which the sensor had the greatest output magnitude; this frequency was recorded as a frequency of interest for future tests.

STAGE 2: Barrel Testing

Stage 2 testing utilized the same signal chain as the stage 1 tests. However, instead of glass canning jars, all tests took place using a plastic 55 gallon drum filled with tap water. Unlike previous tests where the sensor was merely dipped into the test liquids, in stage 2 tests the sensor was fully submerged in the test liquid.

All stage 2 tests followed a similar protocol: The sensor was submerged in the water in the barrel, about 4 inches below the surface. The output of the sensor was then observed. The sensor was then treated to a test stimulus, and the output of the sensor in the presence of the stimulus was then observed. Below is a list of tested variables, along with a brief description of each test.

- Metal objects near the sensor
    - A steel butter knife was submerged in the barrel next to the sensor and moved around
- People
    - A hand was submerged in the barrel next to the sensor and moved around

- Vibration
    - The output of the sensor was observed during a prolonged period of vibrations created by a forklift adjacent to the laboratory
- Orientation of the device’s input feed cables
    - The position of the input feed cables were altered while observing any changes in the sensor’s output
- Orientation of the sensor
    - The orientation and position of the sensor was altered while observing the output of the sensor
- Liquid-air boundaries
    - The sensor was moved near the surface of the water and the sides of the barrel while the output of the sensor was observed
- Time
    - The sensor was left in the barrel for one week during a university holiday. Output of the sensor after the holiday was compared to output of the sensor before the break.
- Mechanical impact and stress
    - The sensor was dropped on a concrete floor from an elevation of about 5 feet. This was followed by treatment with an impact from a loaded drum dolly at low speed. The sensor assembly was then visually inspected for damage or deformation, after which an observation of the output of the sensor was then compared to the output of the sensor before the impact.

STAGE 3: Beaker Testing

Stage 3 testing utilized a more advanced signal chain compared to the stage 1 and 2 tests. The sensor was driven by the Tsunami board with a sinusoidal signal at 6Vpp (approx.) for 1 second. In some tests, this signal was followed by a second signal at a different frequency with a duration of 1 second. After a brief pause (10 ms), this would then repeat. The output of the sensor was read by the Tsunami board input circuitry. For each 1 second sense pulse, the Tsunami was programmed to obtain 100 magnitude readings spaced 10 ms apart. These readings were then averaged and output to a laptop in a human-readable format (LSBs) via the serial terminal.

All tests took place in a 3.5L Pyrex beaker at approx. 24 degrees Celsius. A Hach HQ40d Multi which provided standard conductivity and temperature measurements was used as a control and measurement reference point.

The object for the first set of tests in stage 3 was to obtain a rough idea of the response of the sensor to varying water conductivity. Thus, the sensor was set to read with a single probe frequency of 1.52 MHz in the first test, 1.0 MHz in the second test, and 0.8Mhz in the third test. The protocol for the tests was as follows:

The 3.5L beaker was filled to the 3.5L line with DI water, and an arbitrary quantity of Magnesium Chloride Hexahydrate road salt (kindly donated by the Tate Hall maintenance crew) was added to the water. The solution was stirred

until the salt completely dissolved, and then temperature and conductivity of the solution was measured using the Hach HQ40d Multi. The sensor under test was then submerged in the liquid, and the reading from the sensor recorded. The Hach HQ40d Multi and the sensor were never in the beaker at exactly the same time, as the 1.52 MHz probing frequency of the sensor interfered with the Hach HQ40d Multi. Each data point consisted of a conductivity measurement, a temperature measurement, a measurement from the sensor under test, and the frequency (or frequencies) used for probing the sensor. After measurements were recorded, an additional arbitrary amount of salt was added and the process repeated. This occurred until the conductivity of the water became very high, at which point the entire process was repeated from the beginning after washing all equipment with DI water. After a sufficient number of data points was gathered, the next test would begin.

The object for the second set of tests in stage 3 was to obtain a set of calibration curves for the sensor using two-frequency probing. The frequencies chosen for probing were 1.52 MHz and 0.8 MHz. The protocol for the second set of tests in stage 3 was identical to the protocol for the first set of tests (above), except that each data point contained two measurements from the sensor under test (one for each probing frequency).

The object for the last set of tests in stage 3 was to test software developed using the calibration curve drawn from the two-frequency probing data. This software was meant to let the sensor output a real conductivity value instead of a measurement in arbitrary units. This very brief period of testing

involved a similar protocol as previously described, except this time the Hach HQ40d Multi was used as a control against the sensor under test rather than as a calibration standard.

STAGE 4: Impedance measurements

Stage 4 testing involved the use of a HP 4192A impedance analyzer exhumed from the basement home of an occult anarchist. Although this device was of questionable provenance and required occasional profanity-laced percussive maintenance, reasonable measurements were probably obtained. It is likely that the last time the instrument was calibrated, the Berlin Wall was still standing. Unfortunately, budgetary constraints made the use of an alternative instrument impossible.

The protocol for stage 4 testing was unstructured because of the unpredictable nature of the instrument. Various configurations of the sensor were repeatedly measured at various frequencies using the machine and the measurements were written down by hand. Data points were frequency-dependent vector impedance measurements with magnitude in ohms, phase angle in degrees, and frequency in Hertz.

Results: Data and Characterization

The purpose of this chapter is to present both the results of design and fabrication work (in pictorial form) and the resultant data obtained from testing (in the form of graphs and tables). As such, this chapter is divided into three sections: Section A describes design and fabrication, Section B describes experimental testing data, and Section C describes output from impedance analyzer testing.

Tables of raw data may be found in the Appendix.

SECTION A: Design and Fabrication

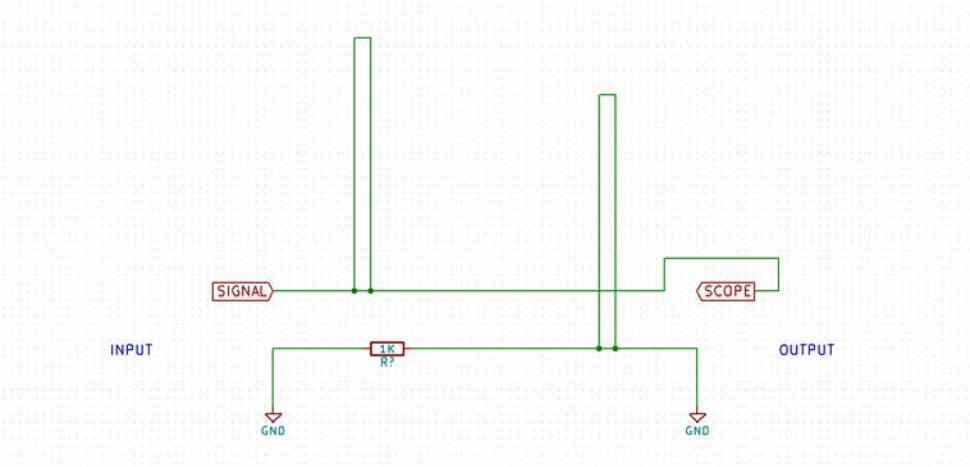

Figure 1*: Physical layout of the sensor*

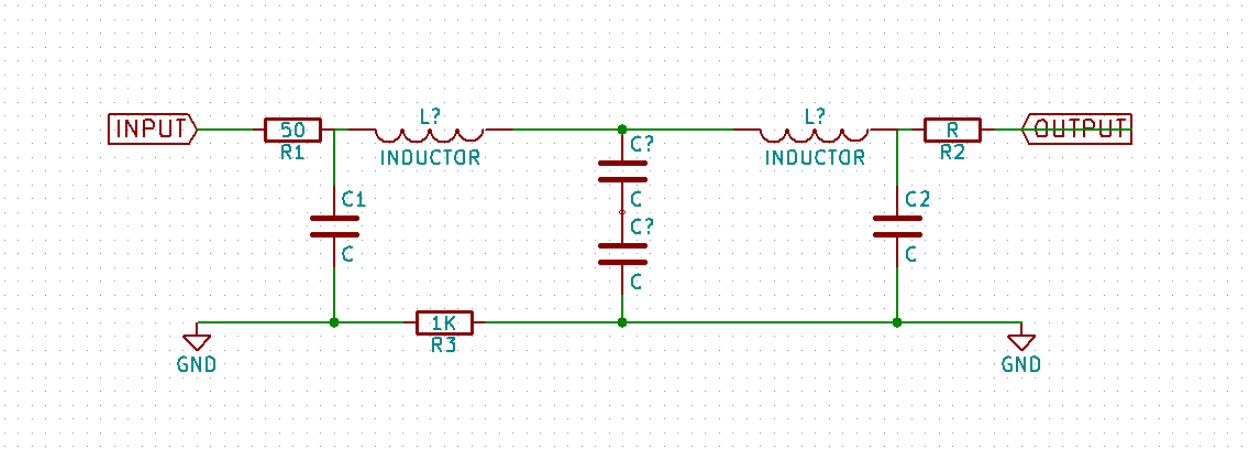

Figure 2: *Theoretical lumped-element schematic*

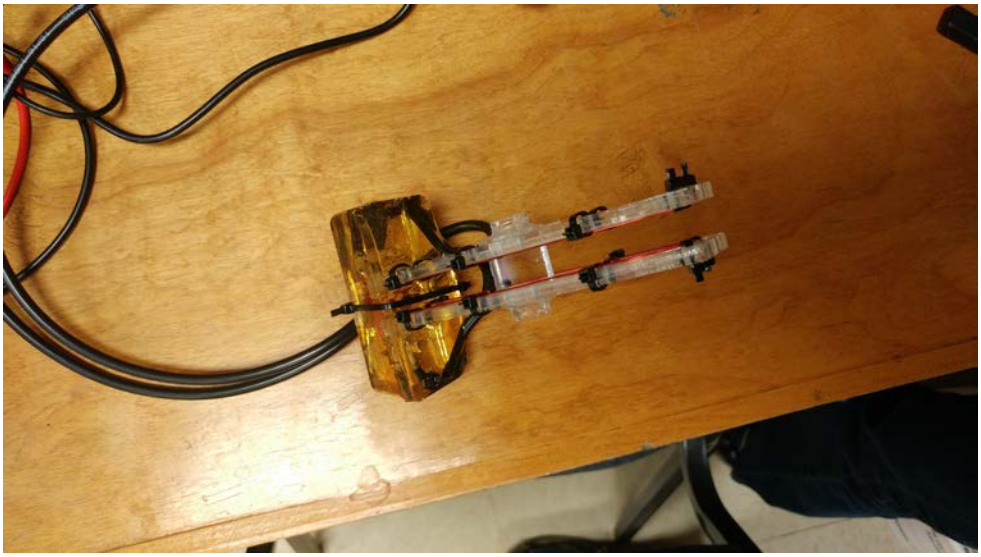

Figure 3: *Photograph of actual prototype sensor. Main components of the sensor can be easily identified: the yellow transparent block contains connections between the legs of the sensor (orange wire on*

transparent plastic) and the input and output coaxial cables

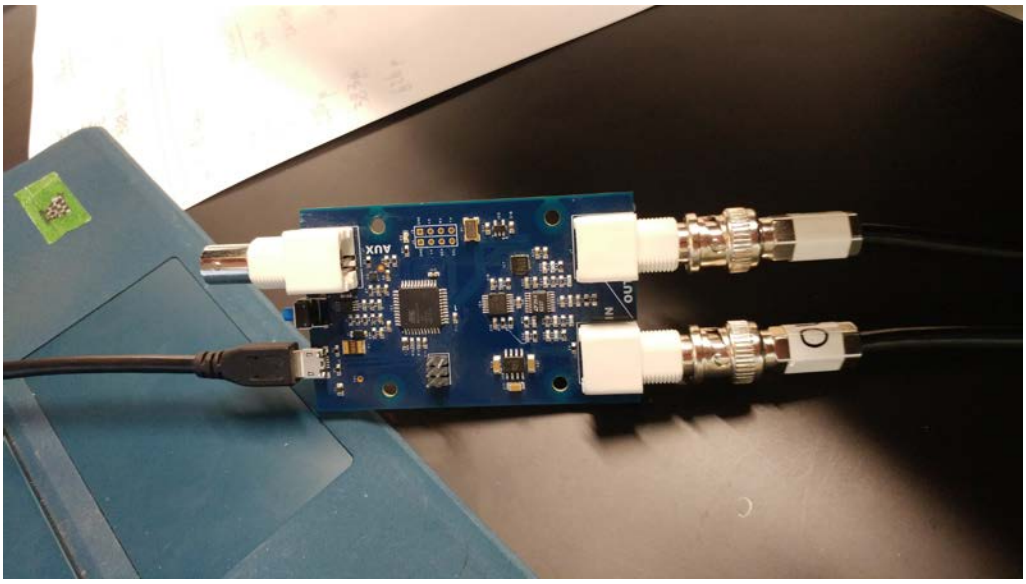

Figure 4*: Tsunami board. The input and output connections to the sensor can be seen on the right side of the board. The connector on the upper left is an unused trigger output. Power is supplied to the board via a 5V usb connection on the lower left.*

SECTION B: Experimental Testing

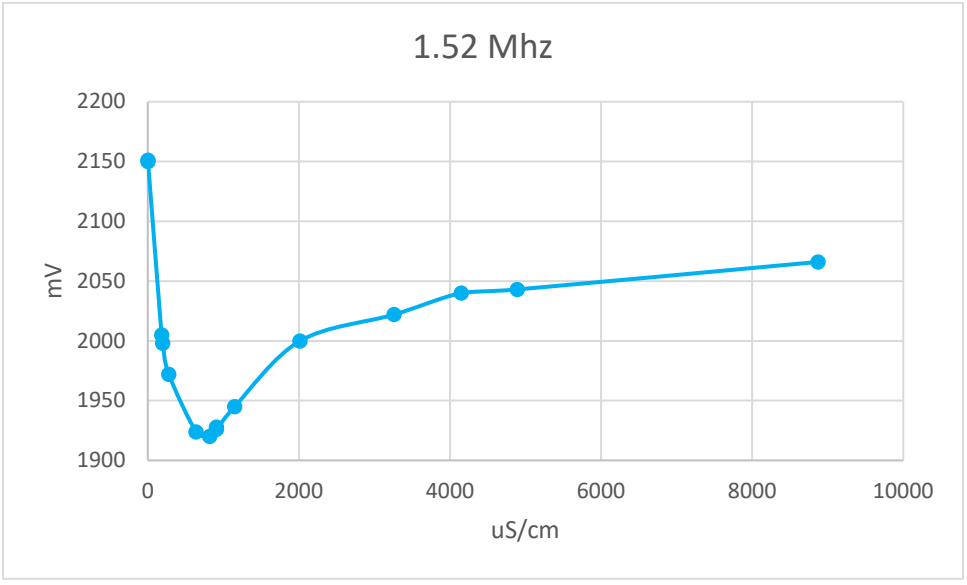

Figure 5*: Sensor output in millivolts dependent on water specific conductivity in microSiemens/centimeter. This data was obtained using code version 1 and a probing frequency of 1.52 Mhz.*

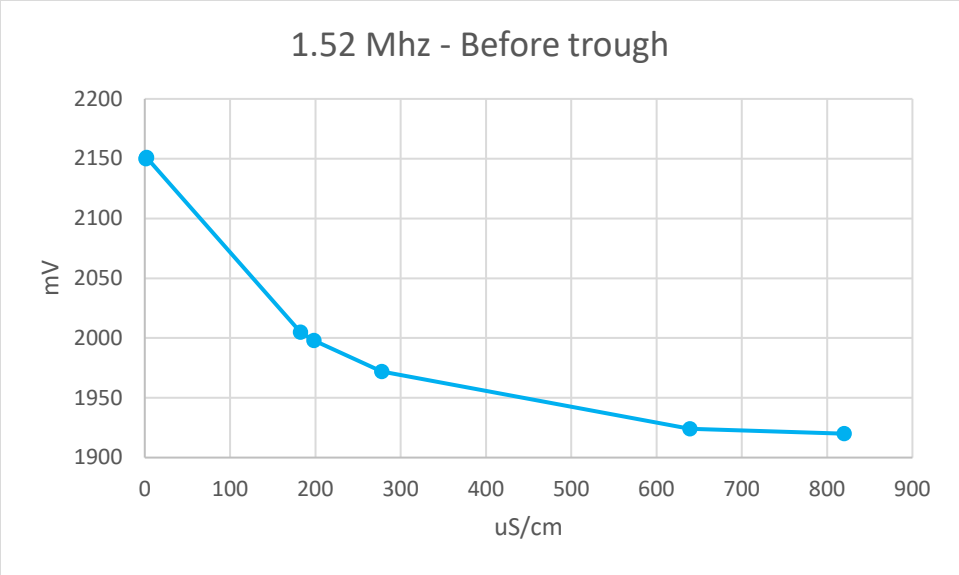

Figure 6: *Sensor output before trough, in millivolts dependent on water specific conductivity in microSiemens/centimeter. This data was obtained using code version 1 and a probing frequency of 1.52 Mhz.*

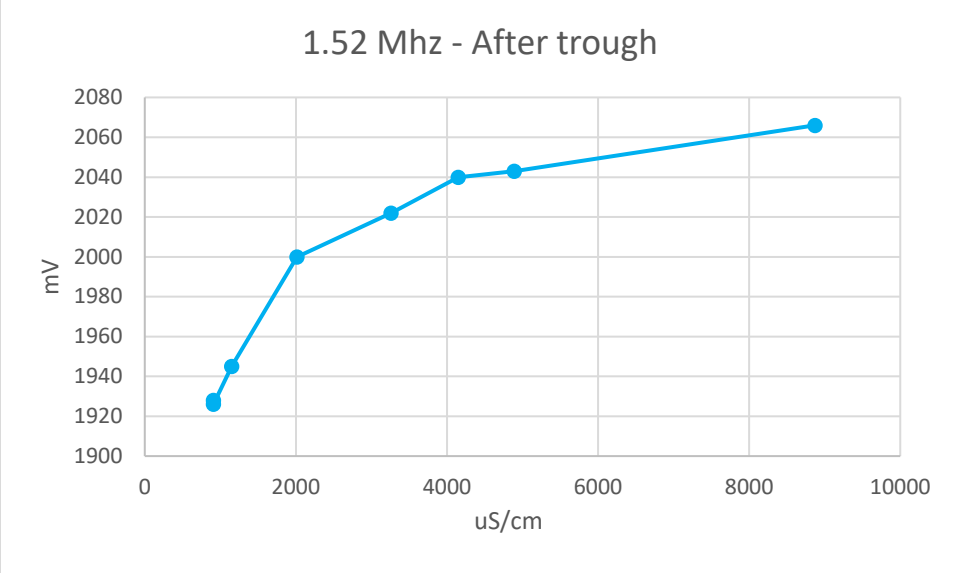

Figure 7: *Sensor output after trough, in millivolts dependent on water specific conductivity in microSiemens/centimeter. This data was obtained using code version 1 and a probing frequency of 1.52 Mhz.*

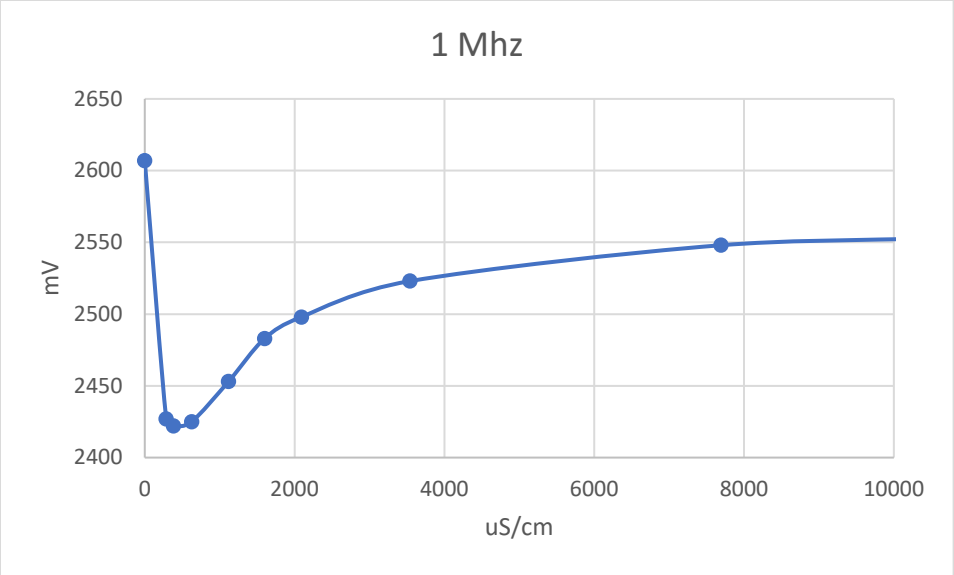

Figure 8: *Sensor output in millivolts dependent on water specific conductivity in microSiemens/centimeter. This data was obtained using code version 1 and a probing frequency of 1.0 Mhz.*

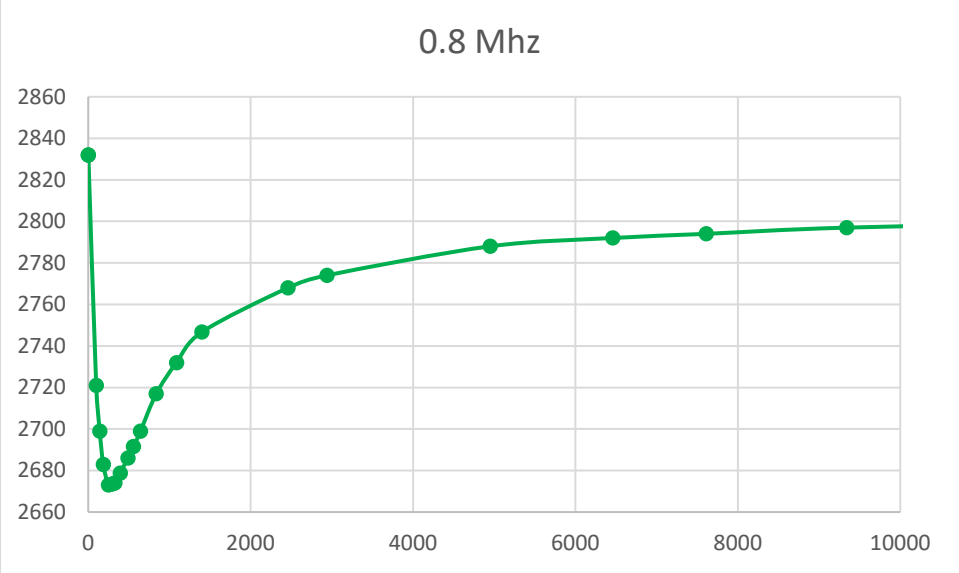

Figure 9: *Sensor output in millivolts dependent on water specific conductivity in microSiemens/centimeter. This data was obtained using code version 1 and a probing frequency of 0.8 Mhz.*

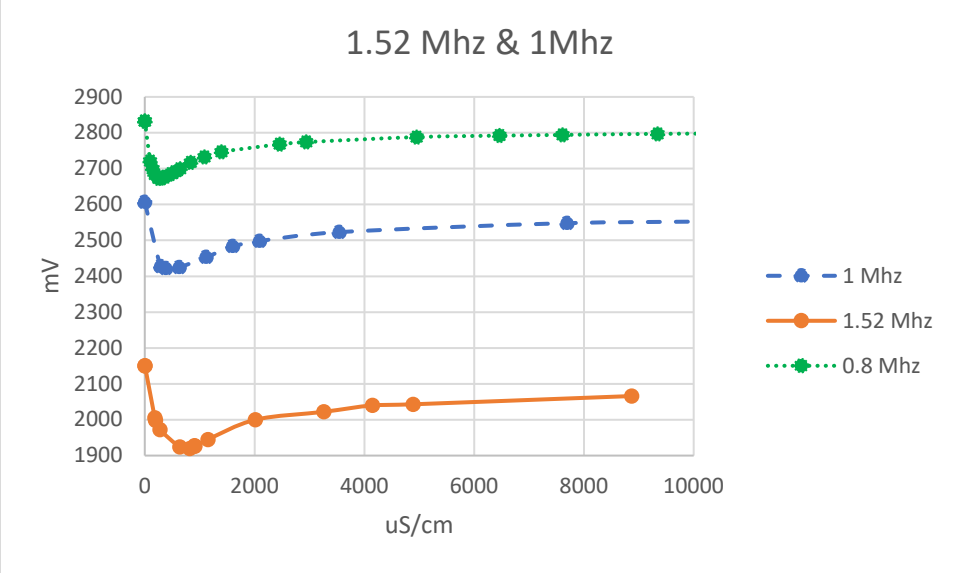

Figure 10: *Comparison of sensor output in millivolts dependent on water specific conductivity in microSiemens/centimeter using 1.52 Mhz, 1 Mhz, and 0.8 Mhz probing frequencies. This data was obtained using code version 1. The three datasets were obtained sequentially.*

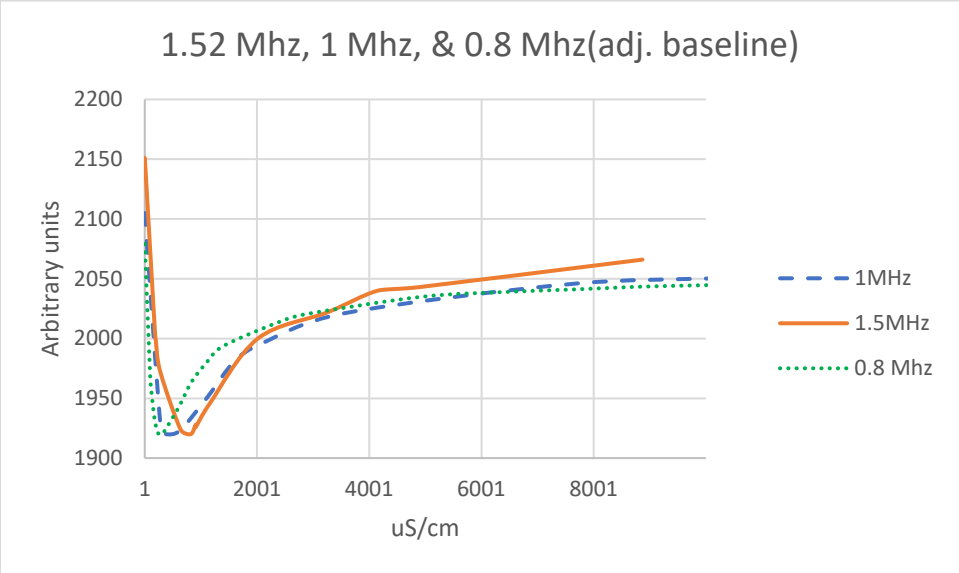

Figure 11: *Comparison of sensor output in arbitrary units dependent on water specific conductivity in microSiemens/centimeter using 1.52 Mhz , 1 Mhz, and 0.8 Mhz probing frequencies. This data was obtained using code version 1. The two datasets were obtained sequentially.*

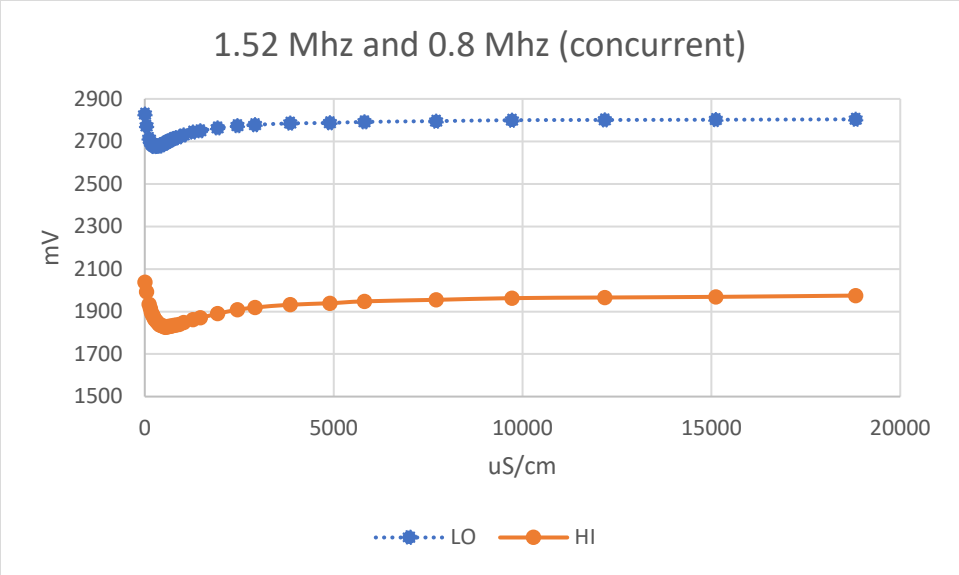

Figure 12: *Comparison of sensor output in millivolts dependent on water specific conductivity in microSiemens/centimeter using 1.52 Mhz and 0.8 Mhz probing frequencies. This data was obtained using code version 2. The two datasets were obtained concurrently. The data has not been adjusted in any way.*

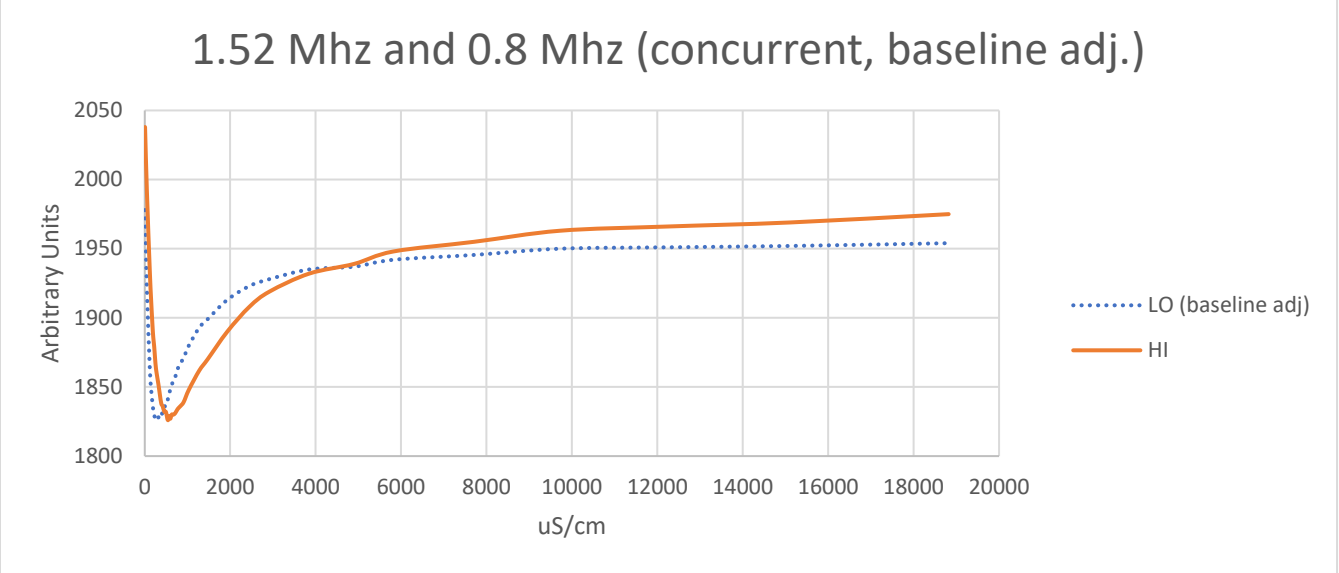

Figure 13: *Comparison of sensor output in arbitrary units dependent on water specific conductivity in microSiemens/centimeter using 1.52 Mhz and 0.8 Mhz probing frequencies. This data was obtained using code version 2. The two datasets were obtained concurrently. The data for 0.8 Mhz measurements has been baseline adjusted using the 1.52 Mhz data. Data point icons have been removed for clarity.*

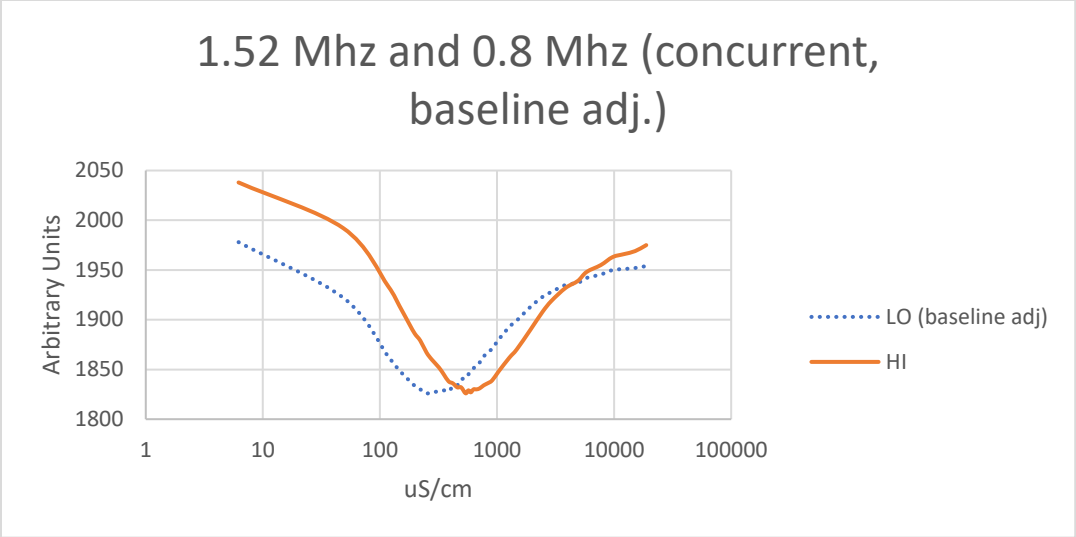

Figure 14: *Same as Figure 13, but with logarithmic x-axis*

SECTION C: Instrumented Characterization

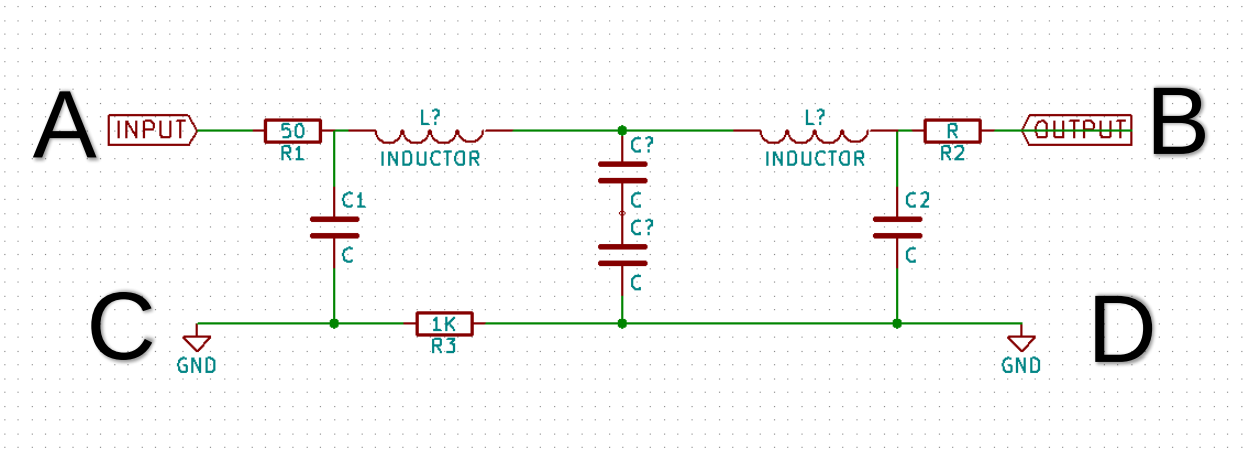

Figure 15: *Duplicate of Figure 2 with added labels to notate test points*

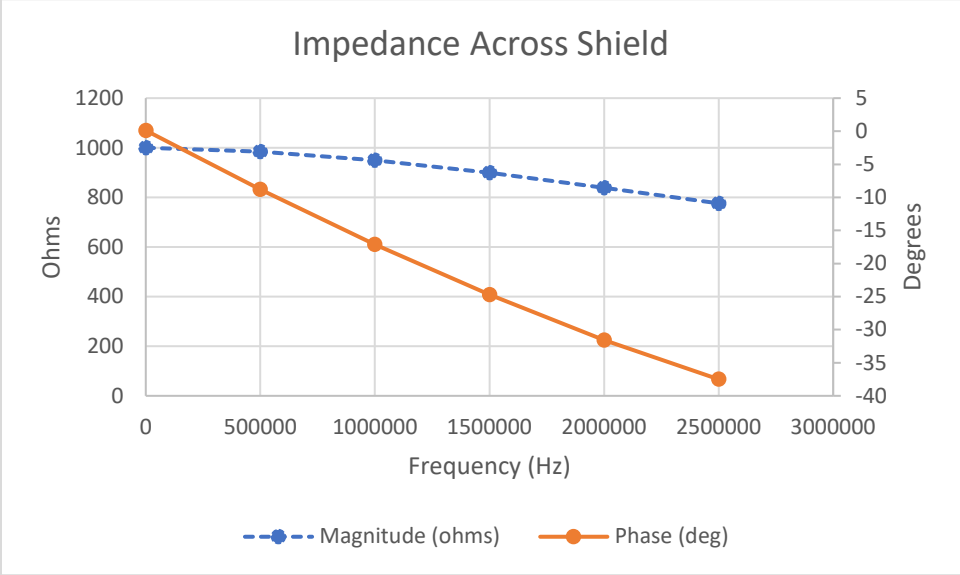

Figure 16: *Impedance measured across the ground shield of the sensor, between point C and point D.*

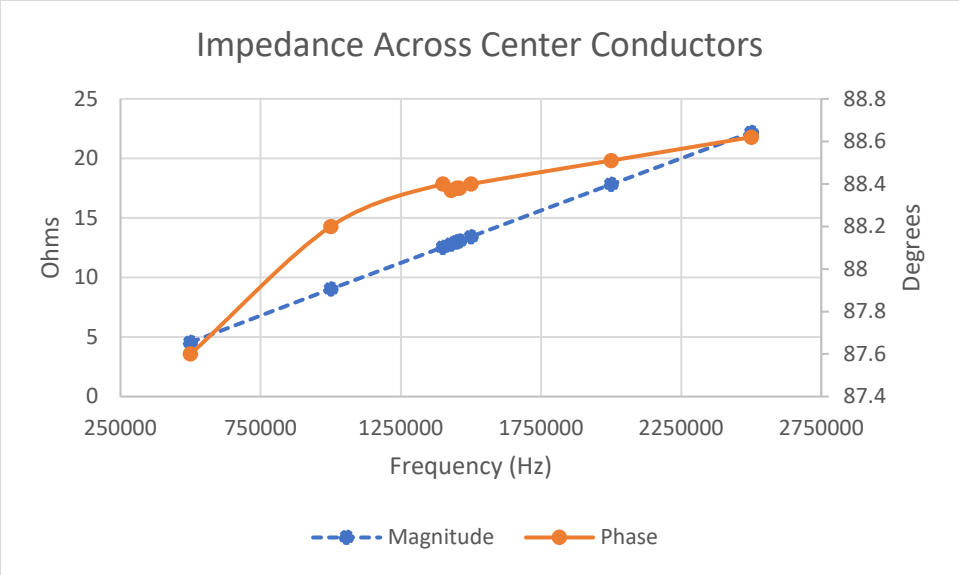

Figure 17: *Impedance measured across the center conductors of the input and outputs, between point A and point B.*

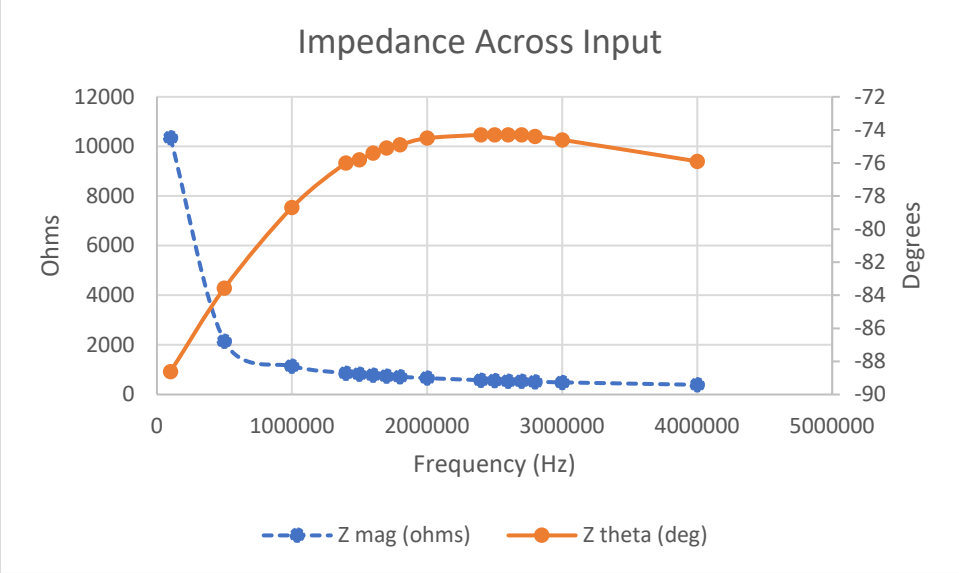

Figure 18: *Impedance measured across the center conductor and shield of the sensor input, between point A and point C.*

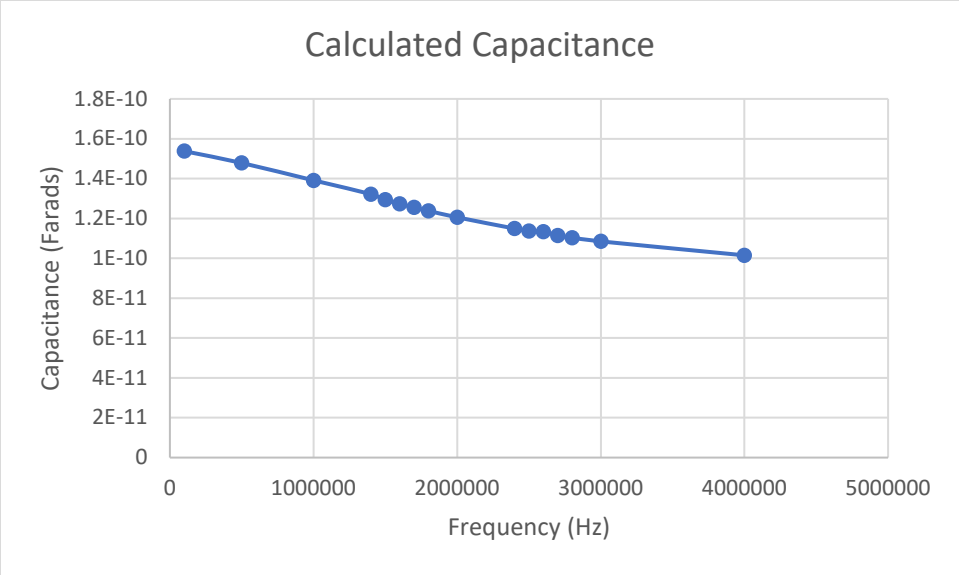

Figure 19: *Calculated capacitance across input (A and C), based on Figure 18 data*

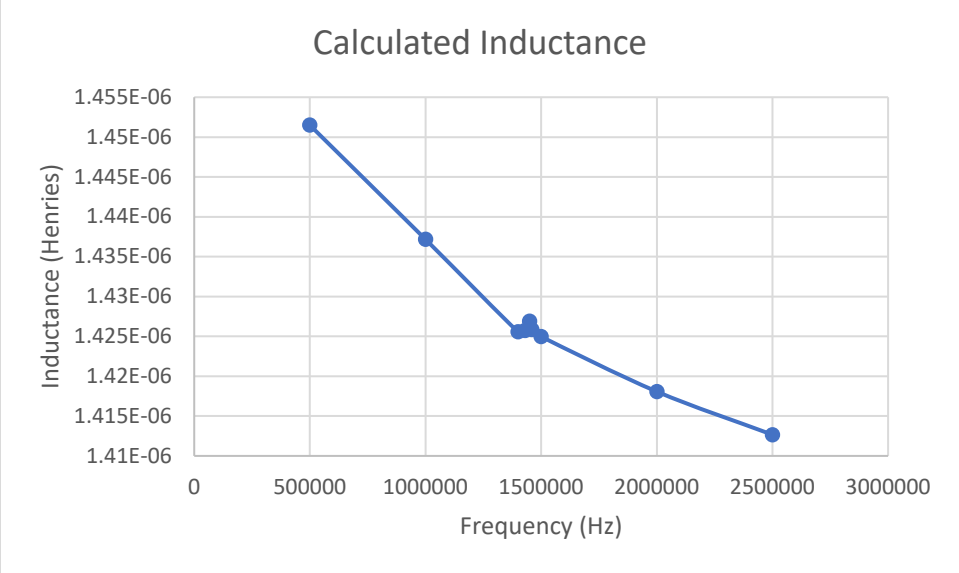

Figure 20: *Calculated inductance across center conductors (A and B), based on Figure 17 data*

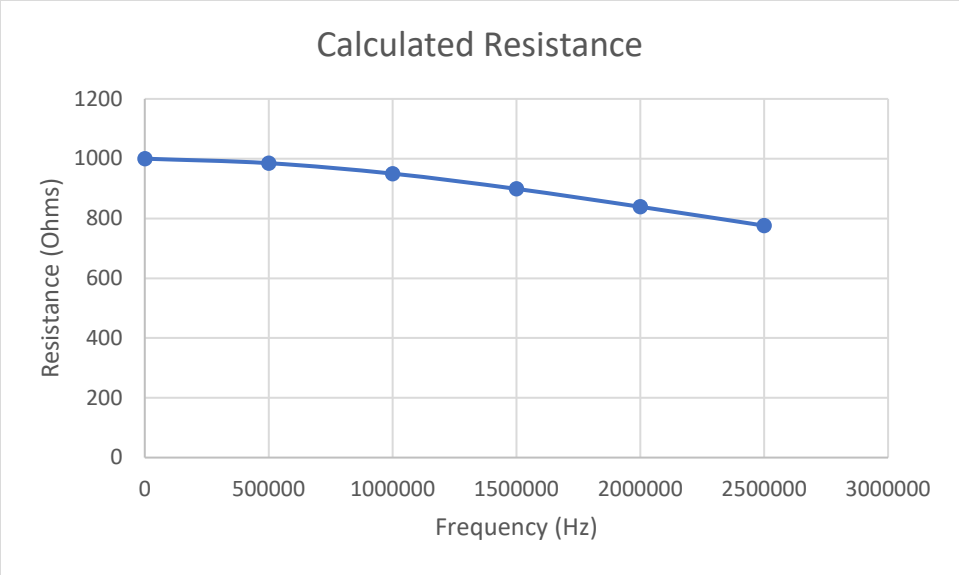

Figure 21: *Calculated resistance across shield (D and C), based on Figure 16 data*

Discussion of design properties and application strategies

As can be seen from Figure 5, the sensor response curve appears as one would expect given the design of the sensor as a tank circuit behaving as a band-stop filter. As a reminder, a sinusoidal signal of fixed amplitude is applied to the sensor as an input, and the output from the sensor is the same signal shape as the input signal, but with a change in amplitude relative to water conductivity.

To confirm that the sensor is indeed a tank circuit, the impedance of the sensor (out of water, in open air) was measured in various configurations across a range of frequencies (Figures 16, 17, and 18). From this data, the intrinsic capacitance (C), inductance (L), and resistance (R) of the sensor can be calculated (Figures 19, 20, and 21).

In air, at an operating frequency of 1.0 MHz:

- The sensor intrinsic resistance (R3 in Figure 15) is 950 Ohms.
- The sensor intrinsic inductance (the two L?s added together in Figure 15) is 1.44E-06 Henries.
- The sensor intrinsic capacitance (C1 and C2 added together in Figure 15) is 1.39121E-10 Farads.

If we model the intrinsic capacitance, inductance, and resistance of the sensor as connected in series, we obtain a value of around 12 MHz as the resonant frequency of the sensor in air.

Although obvious, it is important to note that, in air, the value of capacitance contributed by water conductivity (the two C? in Figure 15) is assumed to be 0 Farads.

As can be seen from the Figure 19-21 series of charts, the values for L, C, and R vary with frequency: while ideal inductors, capacitors, and resistors do not have frequency independent values, the sensor is a real-world example of these devices and so these supposed constants exhibit some frequency dependence. However, this variance is fairly insignificant because the operating frequencies of the sensor are fixed so these values will be static.

To summarize, as input frequency remains constant, the total capacitance of the sensor changes relative to water conductivity and the output magnitude changes. In the development of this sensor, the output magnitude was measured by the Tsunami board.

For every operating frequency, there is one water conductivity value that corresponds to a minimum output response. From Figures 10 and 11 it is also easy to see that this minimum (or “trough”) is frequency dependent: The minimum appears at a lower water conductivity value as the frequency decreases. Thus, the whole response curve shifts directly dependent on frequency.

Unfortunately, the response curve of the sensor is a many-to-one function, which by definition does not have an inverse function. An inverse function, (such as one that might be useful as a calibration curve to calculate conductivity from

the sensor's output), would fail the vertical line test because a given input would have multiple outputs. Thus, for some of the possible output values from the sensor there are two corresponding possible conductivity values.

Two solutions to this problem have been considered. The first and most obvious solution is to select an operating frequency such that the trough occurs at the minimum expected water conductivity (which can be assumed is 0 µS/cm in most cases). Thus, if we ignore parts of the response curve that correlate with negative water conductivities, the response curve can be considered a one-to-one function instead of a many-to-one function. An implementation of this solution would involve greatly lowering the operating frequency of the sensor.

The problem with this approach can be deduced from Figure 11. As operating frequency decreases, the response of the sensor output becomes *less sensitive* (especially at higher conductivity levels); that is, the response function becomes flatter as frequency decreases. Besides decreasing dynamic range, this loss of sensitivity would also increase relative error. If the magnitude of uncertainty in the measurements remains constant , a shrinking dynamic range would increase the proportional uncertainty of the measurements. For example, let's say that the sensor has a frequency independent uncertainty of ±3 units. At frequency A, the sensor has a dynamic range of 100 units (that is, the lowest possible measurement is 0 units, and the highest possible measurement is 100 units) and so the uncertainty in measurement is ±3%. At frequency B, if the sensor has a dynamic range of 50, the uncertainty in measurement now becomes ±6%.

This problem is further compounded by the nonlinear sensitivity of the sensor. If 60% of the dynamic range of the sensor represents only the first 25% of the expected range of measurable water conductivity values, uncertainty becomes greatly magnified as water conductivity increases.

The second solution considered avoids the previously mentioned pitfalls by maintaining the operating frequency around its most optimal value. Instead, the response curve of the sensor is split into two independent functions around the trough (example: Figures 6 and 7). Because they are one-to-one, each of these half-functions can be inverted to form two respective calibration functions. However, this does not change the fact that many (if not all) values for water conductivity will correspond with two values – one for each half-function. There must be a way to decide which half-function is to be used for a given measurement, and by logical necessity this way must involve a third variable.

This ideal third variable is phase. In an ideal band-stop filter, there is a phase inversion at the resonant frequency of the filter. For this conductivity sensor, we would expect this phase inversion to occur at the trough. In real world applied terms, this means that we would expect a negative phase shift between the sinusoidal signal input and the output from the sensor to the left of the trough, and a positive phase shift to the right of the trough. Thus, it becomes easy to decide which half-function to use is phase is also measured: if phase is negative, the half-function to the left of the trough is used. If phase is positive, the half-function to the right of the trough is used.

Sadly, due to budget and time constraints, it was not possible to implement phase measurements that effectively coordinated with the magnitude measurements made by the Tsunami board. The makers of the Tsunami board claims that it is able to measure phase, but in practice phase is only reported as an absolute magnitude without sign. This fact renders the phase-measuring capabilities of the Tsunami board useless for this application.

A less ideal third variable is frequency. Luckily, the Tsunami board is able to control and measure frequency with a useful degree of accuracy. Because the response curve of the sensor changes with input frequency, two independent responses can be obtained from the sensor by sequentially probing the sensor with two signals of different frequencies. These two measurements can each be correlated to two conductivity values (again, one for each half-function for a total of 4 values all together) using their own calibration curves, but we expect that two of these values will be identical (barring error). The identical values should be the real conductivity.

This two-frequency measurement scheme was the purpose for gathering the data shown in Figures 12, 13, and 14. Two frequencies were chosen such that the response curves would be sufficiently different, but that the lower frequency was not so low as to impede sensitivity of the sensor. Of course, the 2 MHz upper limit of the Tsunami board constrained the range of the higher frequency. As can be seen, the response curves for 1.52 MHz (HI) and 0.8 MHz (LO) frequencies are quite different. An attempt was made to fit calibration functions to the data corresponding to the inverse half-functions of the HI and LO

responses using the free online curve fitting software MyCurveFit. An example of these fitting attempts using symmetric sigmoidal functions can be found in the Appendix. Code written for the Tsunami board that applied these calibration functions to measurements from the sensor (including an algorithm for generating and detecting the consensus values) can also be found in the Appendix.

Unfortunately, these attempts at using the two-frequency measurements were not successful for a variety of reasons. The possible reasons include:

- The Tsunami board output amplitude is unstable and is effected by output frequency, microcontroller instruction load, and power supply source.
- The two frequencies chosen have response curves that are too similar relative to error and uncertainty in the measurements
- The free curve fit “software” is insufficiently powerful to produce good calibration curves

These factors probably all contributed to failure. Instability of the Tsunami output amplitude creates error that may not be random or independent, which can compound problems from poor fit of the calibration curve.

It is likely that success may be found by trading some loss of sensitivity in the LO curve for a much bigger difference between the two response curves. In addition, the use of higher quality nonlinear regression software and better hardware than the Tsunami board would be of much help. It might also be beneficial to use quality test equipment to characterize the effect of temperature,

output frequency, load, and power source on the Tsunami’s output. Unfortunately, these resources were not available during this research.

Future Work

Potential future work on this topic include improvements in hardware design and construction, enclosure design, testing, and calibration methods.

- Hardware design
    - Purpose-designed frequency generation and measurement circuits. This may include:
    - Signal generator with control loop to stabilize output amplitude and frequency
    - Precision instrumentation rectifier to make accurate amplitude/magnitude measurements
    - Battery power system
    - Data logging or transmission
- Hardware construction
    - Custom ordered PCBs and waterproof connectors
- Enclosure design
    - Waterproof enclosure for circuitry
    - Mechanically more robust cable enclosure
    - Waterproof connectors
- Testing
    - Use of better electrical test equipment, including a digital impedance analyzer, vector network analyzer, etc.
    - Better mechanical testing
    - Field tests

    - Usability and interface tests
- Calibration
    - Construction of automated calibration method.
    - Use of commercial curve fitting software

Citations

Carpenter, S. R., Caraco, N. F., Correll, D. L., Howarth, R. W., Sharpley, A. N., & Smith, V. H. (1998). NONPOINT POLLUTION OF SURFACE WATERS WITH PHOSPHORUS AND NITROGEN. *Ecological Applications*, *8*(3), 559–568. https://doi.org/10.1890/1051-0761(1998)008[0559:NPOSWW]2.0.CO;2

Ho, E. J., Surenkok, G., & Zayas, V. (2014). Explicit but Not Implicit Mood is Affected by Progressive Social Exclusion. *Intergroup Relations and Identity*, *7*. Retrieved from https://www.researchgate.net

Novotny, E. V., Murphy, D., & Stefan, H. G. (2008). Increase of urban lake salinity by road deicing salt. *Science of The Total Environment*, *406*(1–2), 131–144. https://doi.org/10.1016/J.SCITOTENV.2008.07.037

Ramos, P. M., Pereira, J. M. D., Ramos, H. M. G., & Ribeiro, A. L. (2008). A Four-Terminal Water-Quality-Monitoring Conductivity Sensor. *IEEE Transactions on Instrumentation and Measurement*, *57*(3), 577–583. https://doi.org/10.1109/TIM.2007.911703

Reverter, F., Li, X., & Meijer, G. C. M. (2007). Liquid-level measurement system based on a remote grounded capacitive sensor. *Sensors and Actuators A: Physical*, *138*(1), 1–8. https://doi.org/10.1016/J.SNA.2007.04.027

Stogryn, A. (1971). Equations for Calculating the Dielectric Constant of Saline Water (Correspondence). *IEEE Transactions on Microwave Theory and Techniques*, *19*(8), 733–736. https://doi.org/10.1109/TMTT.1971.1127617

Striggow, K., & Dankert, R. (1985). The exact theory of inductive conductivity sensors for oceanographic application. *IEEE Journal of Oceanic Engineering*, *10*(2), 175–179. https://doi.org/10.1109/JOE.1985.1145085

US EPA, O. of W. (2012). Water Monitoring and Assessment: Conductivity. Retrieved May 5, 2018, from https://archive.epa.gov/water/archive/web/html/vms59.html

Wuliang Yin, Peyton, A. J., Zysko, G., & Denno, R. (2008). Simultaneous Noncontact Measurement of Water Level and Conductivity. *IEEE Transactions on Instrumentation and Measurement*, *57*(11), 2665–2669. https://doi.org/10.1109/TIM.2008.926054

Appendix

In order to save the trees (and to prevent the file size of this document from becoming ridiculously large), the contents of this appendix exist online at the following links:

For more photos, see Flickr Album at: https://flic.kr/s/aHsmj7PTsQ

For additional documentation, raw data, and code, see Github repository at https://github.com/KeiranCantilina/Contactless-Conductivity-Sensor

Hardcopy or offline copies of this data can be requested from the author at canti021@umn.edu.