“EXTENDED REACH” ECHOLOCATION SENSOR

By

Matthew Lurie

Kyle Spesard

Final Report for ECE 445, Senior Design, Fall 2012

TA: Lydia Majure

8 December 2012

Project No. 18

Abstract

We designed and built a device for a person to sense their surroundings using sonar by analyzing the reflected signals of ultrasonic chirps we send out.

The “Extended Reach” Echolocation Sensor sends out a chirp at ultrasonic frequencies and receives ultrasonic signals reflected off objects in the vicinity of the device. The device then filters the signal and pulls out characteristics of the signal that will be able to indicate distance, shape and other properties of objects that reflected the signals. These characteristics then are exported so that they can be used by the “Extended Reach” Haptic Array project.

Unfortunately, the entire system did not function properly when put together, but a majority of the individual components function properly. This leaves room for a good amount of future work to be done.

The cost of this project was kept sufficiently low, showing that this design could be a viable product in the future.

Table of Contents

1. Introduction

1.1 Project Statement

We chose to develop a way for a person to sense their surroundings using sonar by analyzing the reflected signals of ultrasonic chirps we send out. This project is a modular design, which works with our sister project the "Extended Reach" Haptic Array. We chose this project because it has the possibility to lead to some groundbreaking studies, and can be a prototype for an invisible cane for blind people. Not only will it aid people in way a variety of applications, it provides unique signal processing challenges as well as circuit design experience.

1.2 Objectives

The "Extended Reach" Echolocation Sensor will be able to send out a chirp at ultrasonic frequencies and receive ultrasonic signals reflected off objects in the vicinity of the device. The device will then be able to filter the signal and pull out characteristics of the signal that will be able to indicate distance, shape and other properties of objects that reflected the signals. These characteristics will then be exported so that they can be used by the "Extended Reach" Haptic Array project.

Benefits to the Customer

- It may be used to sense the surrounding environment without the need of light or without the use of sight.
- It will be a virtually weightless alternative to a cane
- Can be used in applications where people need to reach something that is out of eyesight
- Reasonably priced
- User friendly, and fun to learn

Product Features

- Ability to sense distance to an object within a maximum distance
- Ability to process multiple characteristics of reflected audio signal
- Ultrasonic frequency transmission from 22kHz – 40kHz
- Fast processing speed allowing continuous updates of sensing feedback
- Safe and reliable design

2. Design

2.1 Block Diagram

The full block diagram can be seen in Appendix B, figure b.1.

2.1.1 Block Descriptions

Processing:

The processing block is the heart and soul of the echolocation sensor. For the sensor processing, we will be using an ST Discovery DSP (digital signal processing) board. The Discovery board will do three main things: send out the chirp waveform, process the incoming reflected signal, and output the signal characteristics to the output serial link. The chirp waveform will be linear with a frequency bandwidth from 22kHz to 40kHz. The incoming reflected signal will be analyzed using cross-correlation. The reflected signal will first be high-pass filtered in order to remove frequencies in the audible range. Then the signal will be cross-correlated with the outputted chirp waveform. Times (and therefore places) where the cross-correlation leads to larger values can be interpreted as important and will tell us information about the object reflecting the signal. We will utilize one of the programmable push-buttons on the Discovery board to yield two different output characteristic patterns. The first pattern will be in the time-domain, and the second pattern will be an FFT (Fast Fourier Transform) of the time-domain output pattern.

Transmitter Circuitry:

The transmitter circuitry will mainly handle amplification of the emitted chirp signal from the Discovery board. We want the signal to be amplified enough that the outgoing chirp will have a sound intensity of around 100-120dB (Sound Pressure Level). The reason for this is that this is the range that bats use in their echolocation and also the range that most ultrasonic rangefinders use.

Transmitter Transducer:

The transmitter transducer will send out the ultrasonic chirp. Since the chirp will have a bandwidth of about 22 kHz to 40 kHz, it ideally the transducer will have a flat frequency response up to about 60 kHz. We are trying not to use specially-built ultrasonic transducers because they can cost between $100 and $500. Cheaper ultrasonic transducers have too narrow of a bandwidth for our application. Instead we will be using audio transducers that have shown to have a flat response higher than their rated 20 kHz. We have decided on using Pyramid TW-125 tweeter.

Receiver Transducer:

The receiver transducer will receive the reflected ultrasonic signal. The specifications of the receiver transducer will be the similar to the speaker. Ideally the receiver microphone will have a flat frequency response up to 60 kHz and will be able to output a large enough signal for the ADC (analog to digital converter) to measure.

Receiver Circuitry:

Once the emitted chirp is reflected and received by the receiver microphone, it needs to be modified for proper use with the Discovery board. Specifically, it needs to be low pass filtered in order to prevent aliasing, and it needs to be biased and scaled in order to fit the specifications of the ADC on the Discovery board. The signal will go from being around 5 V peak-to-peak to a 0-3.3 V range.

Output Link:

Packets of information regarding the characteristics of the room will be sent through USART. The Discovery board has pins that can be specially assigned for USART capabilities.

Power Circuitry:

The power circuitry will supply power to the transmitter circuitry, the receiver circuitry, and the processing unit. Power will mainly come from a wall to mini-USB adapter. The mini-USB will power the ST Discovery board and the rest of the circuitry will draw power from the necessary pins on the board.

2.3 Hardware Schematics

The hardware schematic can be found in figure B.2. Note: the LMC6484 op-amp was chosen because it can operate down to a supply rail voltage of 3 Vdc and it can operate to within 100 mV of its supply rail voltage on both input and output.

2.1 Software Flowchart

For the software flowchart, see figure B.4 in Appendix B.

2.2 Calculations and Simulations

2.2.1 Hardware

Receiver Circuit

The receiver circuit can be broken into three functions: biasing, scaling, and filtering. Initially, we thought that this would necessitate three operational amplifiers (op-amps)--one per function. It was later conceived that all three functions could be handled using only one op-amp using symmetry and a common ground. The common ground is achieved by applying a DC bias to the positive terminal of the op-amp. The 1.65 V DC offset is produced from the 3.3 V Vdd by use of a voltage divider circuit. The DC offset is given by the equation

$$Voffset = \frac{Rref1}{Rref1 + Rref2} Vdd \qquad \{1\}$$

where Voffset is the DC offset in Volts, R1 and R2 are the same resistance given in Ohms, and Vdd is the positive reference voltage from the Discovery Board in volts. A differential op-amp circuit was selected as the basis of the design in order to allow for calculations of the DC offset to be independent of the circuit input gain. The input gain G is given by

$$G = -\frac{Rfb}{Rin} \qquad \{2\}$$

where G is in Volts/Volts and the resistors are measured in Ohms. The -3 dB cutoff frequency is given by

$$f_c = \frac{1}{2\pi RfbCfb} \qquad \{3\}$$

where fc is the cutoff frequency measured in Hertz, Rfb is the feedback resistance in Ohms, and Cfb is the feedback capacitance in Farads. Putting this all together leads to an equation for the magnitude of the voltage outputted to the ADC pin of the Discovery Board. The equation is as follows

$$V_{ADC} = \frac{-\frac{Rfb}{Rin}}{\sqrt{1+(\omega RfbCfb)^2}} V_{in} + Voffset \qquad \{4\}$$

where V_{ADC} is the voltage outputted to the ADC in Volts, Vin is the voltage from the receiver transducer measured in Volts, and is the frequency of the voltage from the receiver transducer measured in radians.

The pspice plot shown in figure B.6 demonstrates the biasing and scaling feature of the circuit functioning properly. The filtering aspect of the circuit was also demonstrated using pspice. The resulting simulation is shown in figure B.5 with the -3 dB cutoff frequency labeled.

Transmitter Circuit

The transmitter circuit's primary function is to provide gain. A non-inverting op-amp circuit was chosen in order to preserve the shape of the transmitted chirp and because the non-inverting op-amp circuit has a high input impedance which will keep the experimental gain close to the calculated theoretical gain. The equation for the gain in this circuit is given simply by the following

$$Vout = (1 + \frac{R4}{R3})V_{DAC} \qquad \{5\}$$

where Vout is the voltage going out to the transmitter transducer measured in Volts, the resistors are measured in Ohms, and VDAC is the voltage inputted to the circuit from the DAC pin on the Discovery Board. The plot in pspice shown in figure B.7 demonstrates the transmitter circuit operating properly.

2.2.2 Software

Chirp

The chirp is a linear frequency modulated waveform with an ultrasonic frequency range from 22 kHz to 40 kHz. A linear frequency modulated waveform is defined by:

$$x(t) = cos(\pi \frac{\beta}{\tau} t^2) \quad 0 \leq t \leq \tau \qquad \{6\}$$

Thus the instantaneous frequency of this waveform sweeps linearly across a total bandwidth of hertz during the -second pulse duration. A linear frequency modulated waveform was chosen over a simpler one-frequency sinusoidal pulse in order to increase the Rayleigh resolution of the cross-correlated signal. The length of the pulse was chosen to be the time it takes sound to travel one foot. This effectively becomes the smallest distance that the sensor can detect objects at. The chirp was created using the (aptly named) MATLAB function chirp and using a sampling frequency of 200 kHz. A plot of the chirp is shown in figure B.8.

Digital High-Pass Filter

After the chirp is sent, the ultrasonic sound is reflected on the various surfaces in the room and sent back to the receiver transducer. Once the reflected signal is captured by the Discovery Board, the reflected signal is high-pass filtered in order to remove frequencies from the audible sound range. The high-pass filter being implemented is a second-order IIR (infinite impulse response) digital butterworth filter using Direct Form 2 transpose structure. A diagram of the digital filter is shown in figure B.9. The filter coefficients were found through the MATLAB command butter and setting the -3 dB cutoff at 22 kHz. The frequency response of the designed filter is shown in figure B.10. The filter will be implemented on the Discovery Board by use of the Direct Form 2 IIR filter function in the CMSIS DSP Software Library.

Cross-Correlation/Waveform Matched Filtering

If the device is in the cross-correlation mode, the next step is to pick out the parts of the reflected signal where the chirp appears. This is done through waveform matched filtering. The reflected signal can be thought of as a collection of targets at various times plus noise. We would like a filter with an overall frequency response H(Ω) that will maximize the signal-to-noise ratio (SNR). To achieve this, consider that the spectrum of the desired output, y(t), will be Y(Ω)=H(Ω)X(Ω), where X(Ω) is the received signal. We want to maximize the SNR at a specific time T_M. The power of the output at that instant is

$$|y(T_M)|^2 = \left| \frac{1}{2\pi} \int_{-\infty}^{\infty} X(\Omega)H(\Omega)e^{j\Omega T_M} d\Omega \right|^2 \qquad \{7\}$$

Now consider the case where the noise interference is white noise with a power spectral density of $N_o/2$ watts per hertz. The total output noise power is thus

$$n_p = \frac{1}{2\pi}\frac{N_o}{2} \int_{-\infty}^{\infty} |H(\Omega)|^2 d\Omega \qquad \{8\}$$

and

$$SNR = \frac{|y(T_M)|^2}{n_p} = \frac{\left| \frac{1}{2\pi} \int_{-\infty}^{\infty} X(\Omega)H(\Omega)e^{j\Omega T_M} d\Omega \right|^2}{\frac{1}{2\pi}\frac{N_o}{2} \int_{-\infty}^{\infty} |H(\Omega)|^2 d\Omega} \qquad \{9\}$$

The choice of H(Ω) that will maximize SNR can be determined via the Schwartz inequality under the form

$$\left| \int A(\Omega)B(\Omega)d\Omega \right|^2 \leq \left\{ \int |A(\Omega)|^2 d\Omega \right\} \left\{ \int |B(\Omega)|^2 d\Omega \right\} \qquad \{10\}$$

with the equality if and only if B(Ω)=A*(Ω). Applying this equation to the numerator of the SNR equation gives

$$SNR \leq \frac{\left(\frac{1}{2\pi}\right)^2 \int_{-\infty}^{\infty} |X(\Omega)e^{j\Omega T_M}|^2 d\Omega \int_{-\infty}^{\infty} |H(\Omega)|^2 d\Omega}{\frac{N_o}{4\pi} \int_{-\infty}^{\infty} |H(\Omega)|^2 d\Omega} \qquad \{11\}$$

Thus SNR is maximized when H(Ω)=X*(Ω)e^(-jΩT_M) or h(t)=x*(T_M-t). This is called a matched filter because the waveform and the filter needed to maximize SNR are a "matched" pair. If one considers the waveform x(t) and the target-plus-noise signal x'(t), the output of the matched filter at time T_M would be the convolution

$$y(t) = \int_{-\infty}^{\infty} x'(s)h(t-s)ds \qquad h(t) = x^*(T_M - t)$$
$$y(t) = \int_{-\infty}^{\infty} x'(s)x^*(s + T_M - t)ds \qquad \{12\}$$

The second line is known as the cross-correlation of the target-plus-noise signal with the transmitted waveform. For our sensor, we will be cross-correlating the received signal with the transmitted chirp. Since our signals will be digital and therefore discrete, the discrete time-domain cross-correlation function becomes

$$y[n] = \sum_{m=-\infty}^{\infty} f[m]g[n+m] \tag{13}$$

where f is the chirp and g is the reflected signal. There is a frequency-domain method that utilizes the quick operation time of the Fast Fourier Transform (FFT). The equation for this is

$$Y[k] = F^{*}[k]G[k] \tag{14}$$

where G[k] is the FFT of the reflected signal, F[k] is the conjugate of the FFT of the chirp, and Y[k] is the FFT of the cross-correlated output. The time domain output can be found by simply taking the Inverse Fast Fourier Transform.

Power Spectrum

If the device is in power spectrum mode, the next step is to compute the power spectrum of the reflected signal. First, the reflected signal will be multiplied by a Hamming window in order to remove the effects of Gibb’s Phenomenon. A Hamming window looks like the following

$$w[n] = 0.54 - 0.45cos(\tfrac{2\pi n}{N-1}) \qquad 0 \leq n \leq N-1 \tag{15}$$

After the window is applied, the FFT of the signal is taken. The power spectrum is simply the square of the magnitude of the FFT, so after the FFT is taken, the magnitude of the FFT is computed and then squared.

MATLAB Simulation Results

We created a MATLAB simulation in order to show the software algorithm in action. First, a reflected signal needs to be generated. This was done by first placing scaled versions of the chirp in an array at various points. The result is shown in figure B.11. Next, noise was added to the reflected signal to properly model environmental conditions. The result is shown in figure B.12.

After the reflected signal was generated, it went through the high-pass filter and the initial chirp was removed. The result is shown in figure B.13. We then tested the cross-correlation using both time-domain and frequency-domain algorithms. The cross-correlation of the reflected signal with the chirp using the time-domain algorithm is shown in figure B.14.

At this point, it is easy to see where in the reflected signal the chirps were actually reflected. For the frequency-domain algorithm, it was noted that the Discovery Board has a max FFT size of 2048 for the built-in optimized FFT function. Thus, the reflected signal was windowed with a 2048-point Hamming window and was done using 50% overlap in order to remove scaling issues due to the window. The resulting cross-correlation is shown in figure B.15.

Comparing the two algorithms, one can see that both methods show where the chirp was reflected. Also, it both outputs have very similar SNRs (signal to noise ratio). After converting both algorithms into C code, we found that the frequency-domain algorithm completed twice as fast as the time-domain algorithm. For this reason, we will most likely use the frequency-domain method. We then simulated the power spectrum mode by applying a hamming window and taking the magnitude-squared of the FFT of the reflected signal. The result is shown in figure B.16.

2.3 PCB Design

We designed the PCB (printed circuit board) using Eagle, we had to create two new parts in Eagle, one for the LMC6484 Op-Amp chip, and another for the transducers. We used a through hole pin array to connect the board to the DSP discovery board. All the through-holes do not connect the bottom layer to the top layer; therefore, all traces to the DSP board are on the bottom. This is because it is not possible to solder the pins to the array at the top so the traces will not be connected otherwise. This design was sent to the ECE parts shop where it was constructed. The PCB design is shown in figure B.3.

2.4 Soldering

Once the PCB was finished, we soldered all parts onto the board. The vias did not connect the top layer to the bottom, so a wire was put through and soldered on both sides to connect it. The transducers were connected to wires which were then soldered to through-holes, so that they can be oriented perpendicular to the board.

3. Design Verification

Our goal was to ensure that our project modules were able to work independently and then would work together as a whole.

3.1 Hardware

3.1.1 Transmitter amplifier verification

To verify the functionality of the transmitter amplifier circuit we needed to verify that the signal from the DAC becomes amplified with a gain of 2.2.

To do this we supplied the circuit using the DSP board, and attached a function generator at the DAC (digital to analog converter) pin. We set it to output about 20 kHz, and attached an oscilloscope across the output of the transducer. The output is in green, and the input is in yellow. As you can see in figure B.20, the input waveform is amplified by 2.07 which is a close value to the simulated gain.

3.1.2 Transmitter (speaker) Verification

We wanted to ensure that the output of the speaker was loud enough to obtain accurate results when reflected off of an object at a maximum distance of 20 ft. A level of 100 dB is sufficient for distinguishing the reflected signal when reflected over our maximum distance.

To verify this we used a speaker that output 94 dB at 1 kHz as a reference to our microphone and measured the amplitude of the voltage at the output of the microphone. Then we output a signal at 1 kHz through our speaker from the function generator and did a linear sweep of the voltage of the signal. From this we were able to determine the voltage at which the speaker output 100 dB. The voltage of the function generator signal at which the input of the microphone was twice the amplitude of our reference signal is the point at which the speaker is outputting sound at 100 dB. This is because a signal at 94 dB has half the amplitude of a signal at 100 dB. We found that at the highest voltage we are able to output our amplified signal we do not quite reach 100 dB. This is shown in figure B.18.

3.1.3 Receiver pre-amp circuit

To verify the functionality of the receiver circuit we needed to make sure the input signal from the microphone was properly biased to the center of the ADC range, which is 1.5 and that the signal becomes inverted. We needed to show that the signal was amplified to map to the largest amount of

ADC values possible for greatest signal resolution. We also needed to show that the low pass filter has a 3 dB cutoff at 60 kHz.

To test the receiver circuit bias and gain we used a function generator to input a 20 kHz signal at output of the microphone and measured the input of the ADC using an oscilloscope. The signal was properly biased and amplified. We further tested this by outputting sound to our microphone at 20 kHz. We found that we needed a larger gain to get a large enough signal at the input of the ADC for resolution, and so we then maximized the gain by trying different resistive values for our differential amplifier and acquiring the resulting signal at the ADC. The maximum amplification we were able to achieve was 18 dB.

Then we tested the low-pass filter by using a function generator to simulate a range of frequencies from 20 kHz to 70 kHz. The output of the function generator was hooked up to the output of our microphone, and then we looked at the input of the ADC using an oscilloscope. We were able to show that the signal becomes diminished properly as the frequency of the inputted signal increases. This is shown in figures B.17 and B.19.

3.2 Software

3.2.1 Input Buffer

In order to verify that the input buffer worked independently, we applied a known signal to the ADC (analog to digital converter) pin of the STM32F4 Discovery Board. We then looked at the memory location of the stored data in the Keil uVision debugger and found that the stored values did indeed match up with the known inputted signal.

3.2.2 Output Buffer

The output buffer was tested by first loading the buffer with a known waveform. The waveform was sent out via the DAC (digital to analog converter). We connected an oscilloscope to the DAC pin on the STM32F4 Discovery Board and viewed the outputted waveform on the oscilloscope. We verified that the waveform was correct.

3.2.3 High-Pass Filter

The high-pass filter was tested by loading an array with known data, putting the array through the high-pass filter function and viewing the output array in the debugger. Unfortunately, the filter function did not work. The filter function relied on the CMSIS DSP function library which did not end up working with the STM32F4 Discovery Board.

3.2.4 Cross-Correlation

The cross correlation function was tested by sending in known arrays with a known expected result and comparing it to the actual array outputted by the function. The two were shown to agree with each other.

3.2.5 Power Spectrum

Unfortunately, the power spectrum function relied heavily on the CMSIS DSP function library. Without the use of this function library, this function could not work.

3.2.6 Output Link

The output link was tested by outputting known data and connecting the output pin to the oscilloscope and comparing the waveform on the oscilloscope to the known data. Unfortunately, accessing the memory to output the data caused the STM32F4 Discovery Board to fail. Thus, the output link did not work.

4. Costs and Schedule

4.1 Cost Analysis

Labor:

Employee	Labor Costs
Matthew Lurie	$35/hr * 15 hrs/week * 12 weeks * 2.5 = $15,750
Kyle Spesard	$35/hr * 15 hrs/week * 12 weeks * 2.5 = $15,750
Total	**$31,500**

Parts:

Item	Part Number	Quantity	Cost Per Unit	Total Cost Per Item
STM32F4 Discovery Microcontroller	STM32F407	1	$14.55	$14.55
Ultrasonic Transducer	WM-61A	2	$1.92	$3.84
PCB		1	$20.00	$20.00
LMC6484 Op-Amp	LMC6484	1	$3.31	$3.31
1kohm Resistor	40T1807	2	$.034	$.034
10kohm Resistor	18R3504	1	$.05	$.05
12kohm Resistor	89M6846	1	$.012	$.012
18kohm Resistor	59M6955	2	$.077	$.154
30kohm Resistor	59M6981	2	$.012	$.024
150pF Capacitor	75-562R5GAT15	2	$1.62	$3.24
3.3uF Capacitor	DT-35	1	$.17	$.17
Female Header Pin Array (2x8)	11P4167	1	$1.53	$1.53
USB 2.0 A to mini-USB B connector		1	$3.96	$3.96
Wall to mini-USB power adapter		1	$4.99	$4.99
Total		**19**		**$55.86**

Total cost: $31,500 (Labor) + $55.86 (parts) = **$31,555.86**

5. Conclusion

5.1 Accomplishments

We have made a circuit allows the control of the output though a speaker as well as the ability to read the reflected signal from the speaker. This includes an amplification circuit for the speaker, a pre-amp circuit for the microphone receiver, and a ST DSP discovery board. The speaker amplification circuit takes the output of the DAC from the DSP board, and amplifies it to produce a loud enough sound to be distinguishable when received after being reflected at our specified maximum distance of 20 feet. After the reflected signal is converted by the microphone to a voltage, the pre-amp takes the signal from the microphone, amplifies, biases, low pass filters, and then pass the signal to the ADC. The low pass filter prevents aliasing, and the gain and the bias map the signal to the middle of its range. This allows for the greatest signal resolution for the ADC. Our DSP discovery board is then able to store the signal digitally. The ST DSP discovery board is also able to perform the calculations for a cross-correlation in the time domain. This altogether covers a quite an array of functions that our design can accomplish.

5.2 Uncertainties

We are not able to characterize all of the different parts of the design that add distortion to the signal. We also have come across some uncertainties with the compatibility of the DSP discovery board with certain libraries functions. We are aware that our speaker adds frequency dependent distortion, and diminishes the sound pressure at higher frequencies. Our speaker also has a few resonant frequencies, which are at frequencies lower than our application. The microphone we are using to receive acts as low pass filter with a cutoff at approximately 31 kHz, but it is not clear if it acts as a filter on lower frequencies because we do not have a perfect speaker to test it. We have also been trying to figure issues when using some libraries with floating point numbers, but have not been able to determine the cause. Another question we have not been able to answer is why the DSP board has been overwriting some parts of memory, while there is nothing in code that we see causing the error. These uncertainties are areas we wish to address with more testing and analysis.

5.3 Ethical Considerations

We will make sure to adhere to the following IEEE Code of Ethics standards in the following ways:

1. to accept responsibility in making decisions consistent with the safety, health, and welfare of the public, and to disclose promptly factors that might endanger the public or the environment;

> Since our sensor will emit chirps at around 100 dB in the ultrasonic frequency range, we will make sure not to activate our device around any animals that can hear in this frequency range. The chirps would be very irritating and possibly damaging to their hearing.

3. to be honest and realistic in stating claims or estimates based on available data;

We will be cautious and make sure the characteristics of the room we send out are the actual characteristics of the rooms and are not distorted in any way. We will not falsify any data received or transmitted by the sensor.

6. *to maintain and improve our technical competence and to undertake technological tasks for others only if qualified by training or experience, or after full disclosure of pertinent limitations;*

This project will require skills and knowledge that we have developed after years of studying in and out of the classroom. It is possible that we will come across problems that test the boundaries of our knowledge. We will make great efforts to gain this knowledge and credit the sources whenever possible.

7. *to seek, accept, and offer honest criticism of technical work, to acknowledge and correct errors, and to credit properly the contributions of others;*

Our project group consists of two people, so we will make sure to be critical of each other's work in order to prevent errors. We are also receiving guidance from two members of Dr. Doug Jones's DSP lab, Erik Johnson and David Jun. We will make sure to credit them for their contributions.

5.4 Future Work

There is a good amount of work that can be done on this project in the future. The first thing step would be to get the CMSIS DSP library working with the STM32F4 Discovery board. With those functions working, the high-pass filter and the power spectrum could work properly. Another alternative might be to use a different DSP board. An arduino board of some kind might be a good alternative because a lot of user support exists online for those boards. Also, it would be a good idea to build some sort of housing for all of the components in the project. The exposed circuit board and wires makes the system less reliable for general use. Additionally, a new speaker should be found that can output ultrasonic frequencies without distortion. Lastly, Additional circuitry should be added so that the system can run off of batteries. This way the system could be run without having to be tethered to the computer via USB.

References

1. "Cross-correlation." Wikipedia. Wikipedia, September 26, 2012. Web. 30 Sep 2012. <http://en.wikipedia.org/wiki/Cross-correlation>.

2. "IEEE Code of Ethics." IEEE. N.p., 2012. Web. 30 Sep 2012. <http://www.ieee.org/about/corporate/governance/p7-8.html>.

3. "LMC6484 Datasheet." Texas Instruments. N.p., n.d. Web. 30 Sep 2012. <http://www.ti.com/lit/ds/symlink/lmc6484.pdf>.

4. "Panasonic WM-61A Datasheet." Panasonic. N.p., n.d. Web. 30 Sep 2012. <http://www.panasonic.com/industrial/components/pdf/em06_wm61_a_b_dne.pdf>.

5. Richards, Mark A. Fundamentals of Radar Signal Processing. McGraw-Hill, 2005. Print.

6. "RM0090 Reference Manual." ST. N.p., n.d. Web. 30 Sep 2012. <http://www.st.com/stonline/stappl/st/com/TECHNICAL_RESOURCES/TECHNICAL_LITERATURE/REFERENCE_MANUAL/DM00031020.pdf>.

7. "STM32F407xx Datasheet." ST. N.p., n.d. Web. 30 Sep 2012. <http://www.st.com/internet/com/TECHNICAL_RESOURCES/TECHNICAL_LITERATURE/DATASHEET/DM00037051.pdf>.

8. "STM32F4 User Manual." ST. N.p., n.d. Web. 30 Sep 2012. <http://www.st.com/internet/com/TECHNICAL_RESOURCES/TECHNICAL_LITERATURE/USER_MANUAL/DM00039084.pdf>.