EE-217 Final Project
The Hunt for Noise (and All Things Audible)
5-7-14

Introduction

Noise is in everything. All modern communication systems must deal with noise in one way or another. Different types of systems deal with noise in different ways, from digital systems/networks, analog communications, wireless and signal conditioning. Noise is natural and in some cases, unnatural, and noise suppression techniques are every part of a system as the controller that operates it.

This experiment attempts to extract an audio message from a noisy environment. Different noise sources were used to test the validity of using modern noise suppression techniques in attempting to filter out the noise in order to discern the audio message buried within. The noise sources used were a child's scream, a hair dryer, and a Dremel tool. Each source operates in a specific band of frequencies and presents challenges in filtering out an audio message signal.

Preliminaries

First, a test signal was used to test the FFT algorithm that would be used for the rest of the project. The signal was generated using a program called Audacity. This same program was also used to take the noise recordings. The test signal used was a 1KHz sine wave. Figure 1 below shows the FFT of this test signal. Refer to Appendix A.6 for the Matlab code.

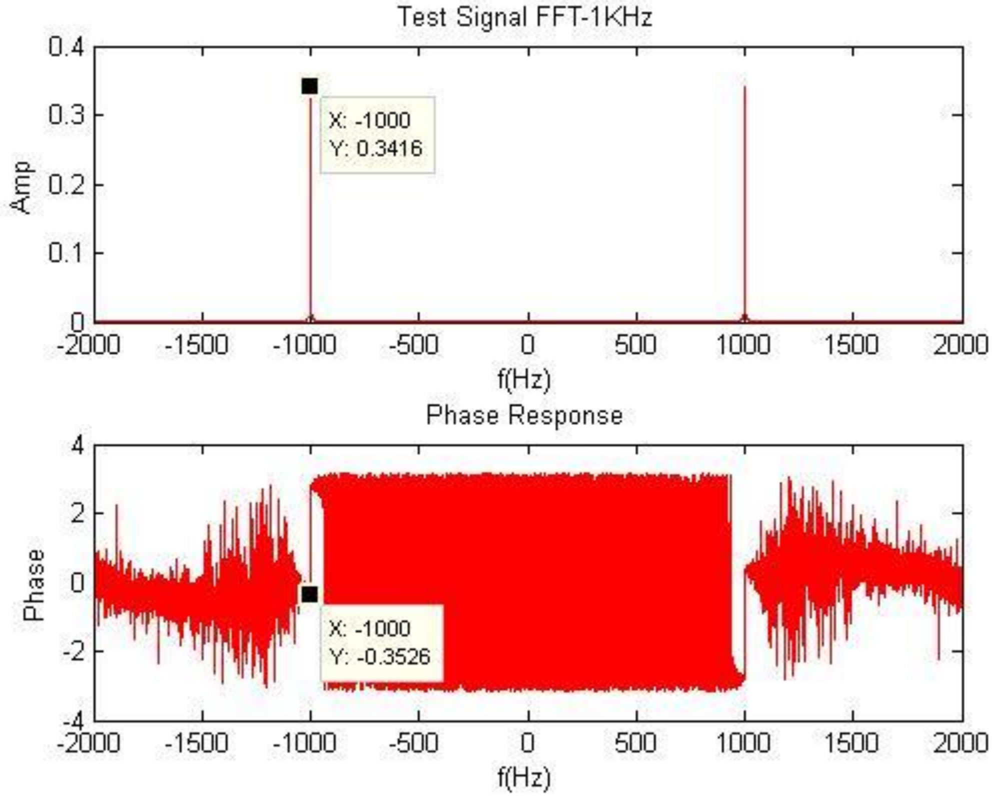

Figure 1: 1KHz Test Signal for System Calibration

Multiple voice and noise recordings were obtained to test the consistency of the frequency spectrum. The noise recordings were more consistent than the voice, which is what was expected. Based on these results, it was decided to use the same voice recording for the hair dryer and the Dremel tool. The scream recording uses the same scream but different voice recordings. For the scream analysis, the scream was consistent; using a recording at 2 feet from the recording device, and the same message was repeated at different distances, though separate voice recordings were used. Table 1 below shows a breakdown of the recordings used and the distances of the recordings.

Table 1: Recording Distances for each audio source

	Distance (feet)			
Audio				
Scream	2	2	2	2
Voice	2	6	12	18
Hair Dryer	2	6	12	18
Voice	2	2	2	2
Dremel	2	6	12	18
Voice	2	2	2	2

I. Noise Source: Scream

(Note: Refer to Appendix A.1 for the code)

The same audio message was used for all three noise sources. The audio is the author's voice and the message spoken is, "Come on feel the noise." Quiet Riot fans will get the meaning of this message immediately, but in this context, the meaning is something entirely different. This message was chosen for its simplicity. Figure 2 shows the spectrum of the audio message. The main components are at 119, 447, and 768Hz. There are many minor components in between these major parts, but most are less than 400Hz. For this part of the experiment, the scream was kept constant and was recorded at 2 feet from the recording device. The same scream recording was used for all samples. The voice recording however varied at different distances (see Table 1). The biggest challenge was trying to maintain the same volume and tone with each voice recording. The phase relationship between the recordings as well as the amplitude varied. This of course is not ideal when trying to extract a voice recording using another recording as the template, but fortunately for this experiment, most of the scream spectra fell well outside the voice spectra. The scream spectra is shown in Figure 3. Note the major components of this spectra are at or above 1586Hz.

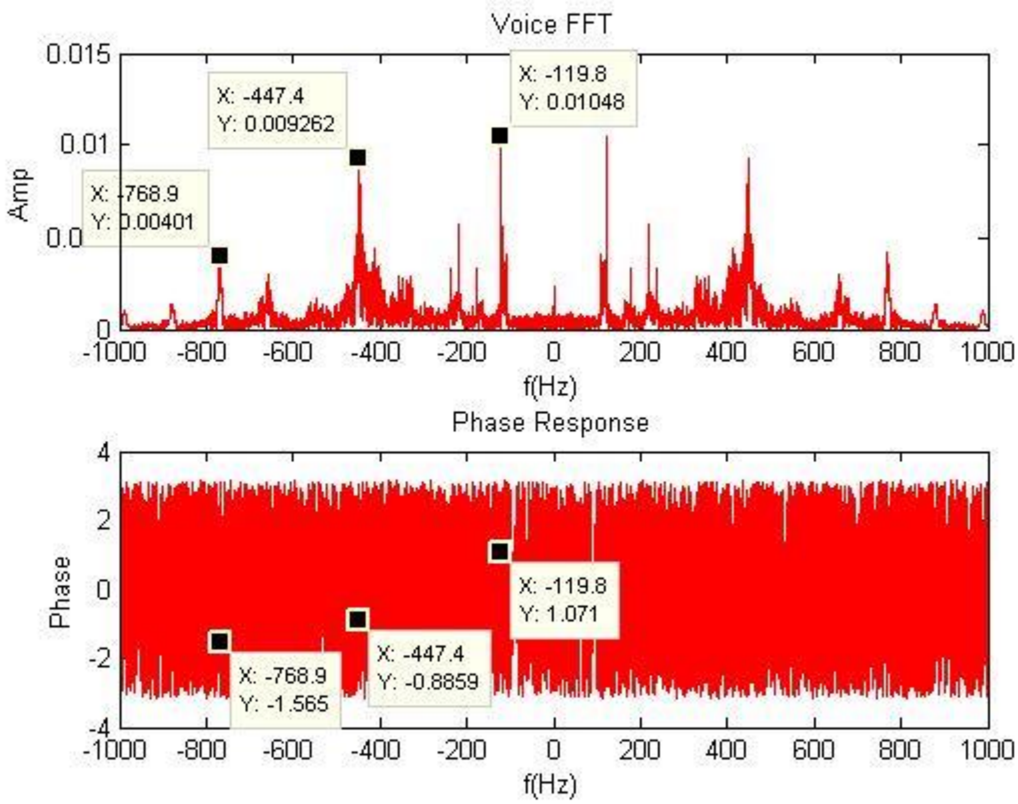

Figure 2: Frequency spectrum of Audio Message.

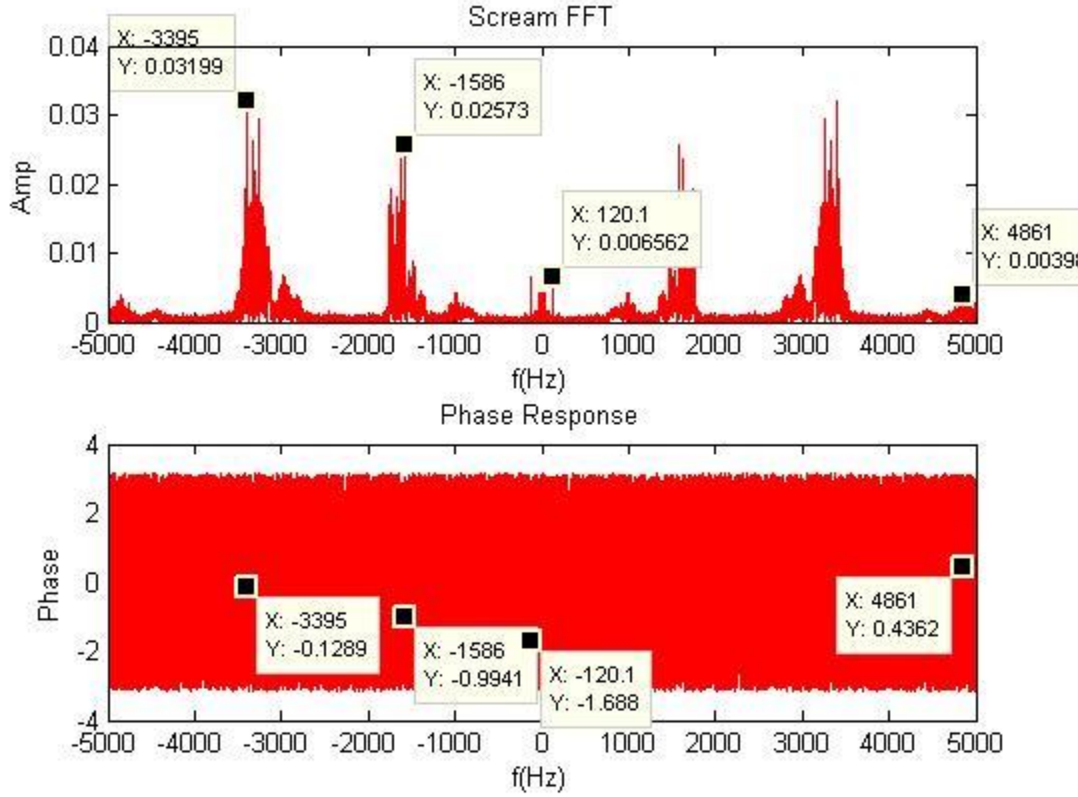

Figure 3: Frequency Spectrum of Scream.

Figure 4 shows the noisy signal spectra. Comparing this graph with Figures 2 and 3, the main components of the voice and scream are clearly identifiable. The voice spectra falls well below the frequency of the scream spectra. The scream was chosen for the high pitch, making it easier to filter out as a result, with very little interference of the scream with the voice message. The voice can be heard in the recording, but the message is not discernable.

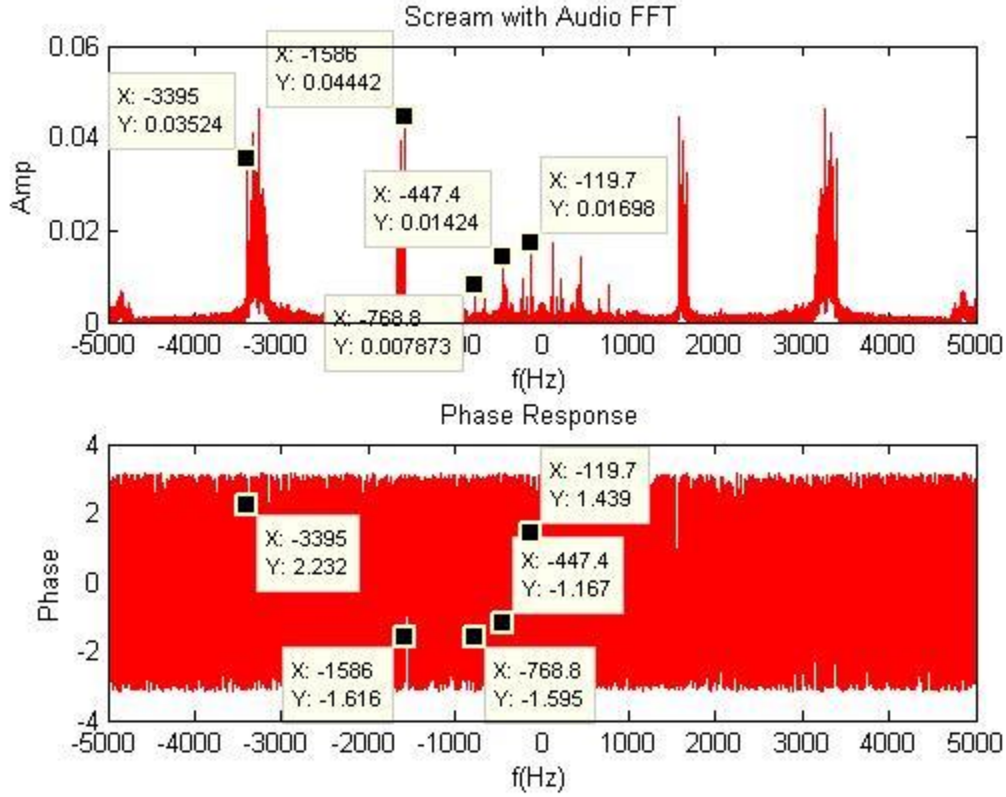

Figure 4: Frequency Spectrum of Noisy Signal (Voice and Noise at 2 feet).

For this part of the experiment, a low-pass filter (LPF) was used to filter out the noise of the scream. The spectra shown in Figure 5 is the result of the noisy signal being passed through the filter. It was found that the LPF model used was not an ideal filter, meaning there was some roll-off. As a result, there was quite a bit of the noise left over after being passed through the filter. It required multiple passes to obtain the spectra shown in Figure 5. Comparing this with Figure 2, the same frequency components can be seen and the noise portion is, for the most part, non-existent. There is a small signal of noise present at 1586Hz. The audio message post-filtering is still muffled, but more discernable and the noise is clearly suppressed.

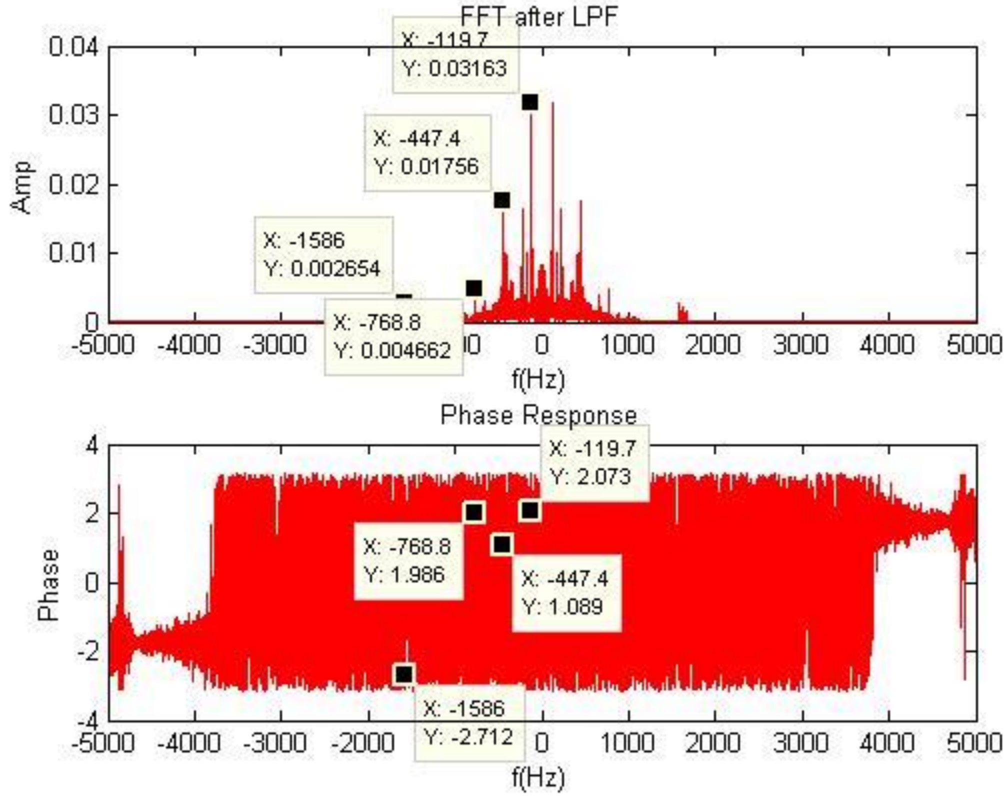

Figure 5: Spectra of “noisy” signal after LPF (Voice and Noise at 2 feet).

The spectras shown thus far are for the recordings at 2 feet for both the noise and the voice. As stated earlier, the voice distance varied (moved further away); thereby, decreasing in amplitude. Figures 6 and 7 show the spectra resulting from passing the signal through the LPF with the voice at a distance of 18 feet. The voice amplitude is clearly decreased in value compared with the noise; however, Figure 7 shows the same pattern after filtering as Figure 5. The LPF was set for a cutoff frequency of 200Hz.

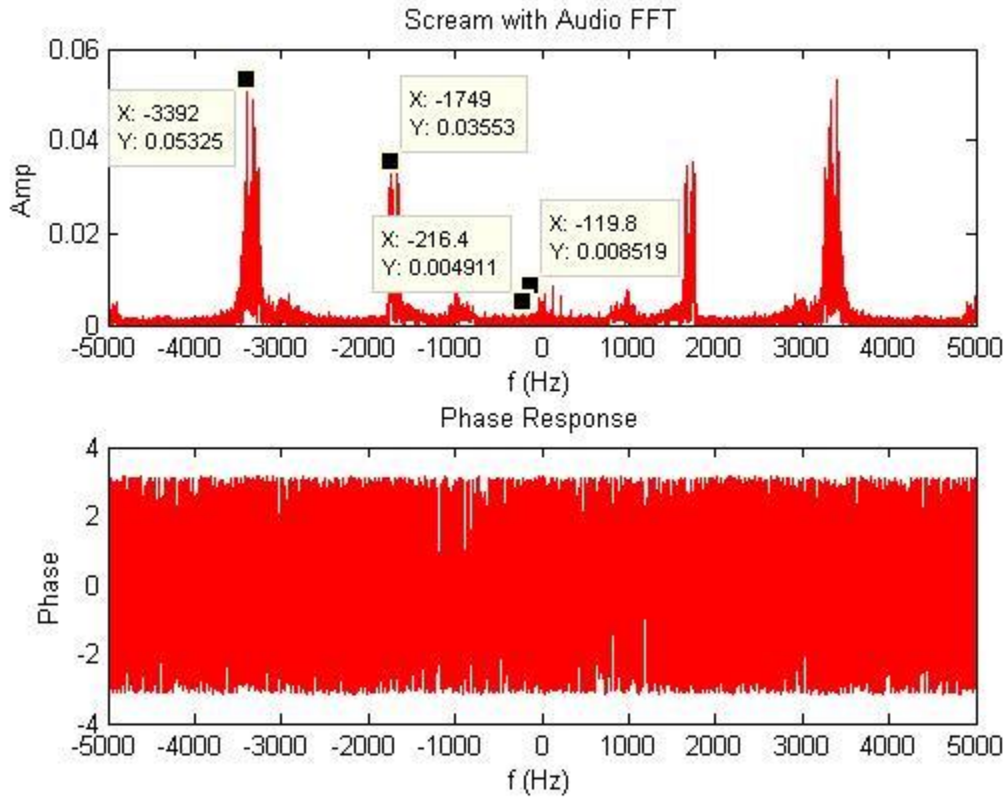

Figure 6: Noisy Signal with Scream at 2 feet and Voice at 18 feet.

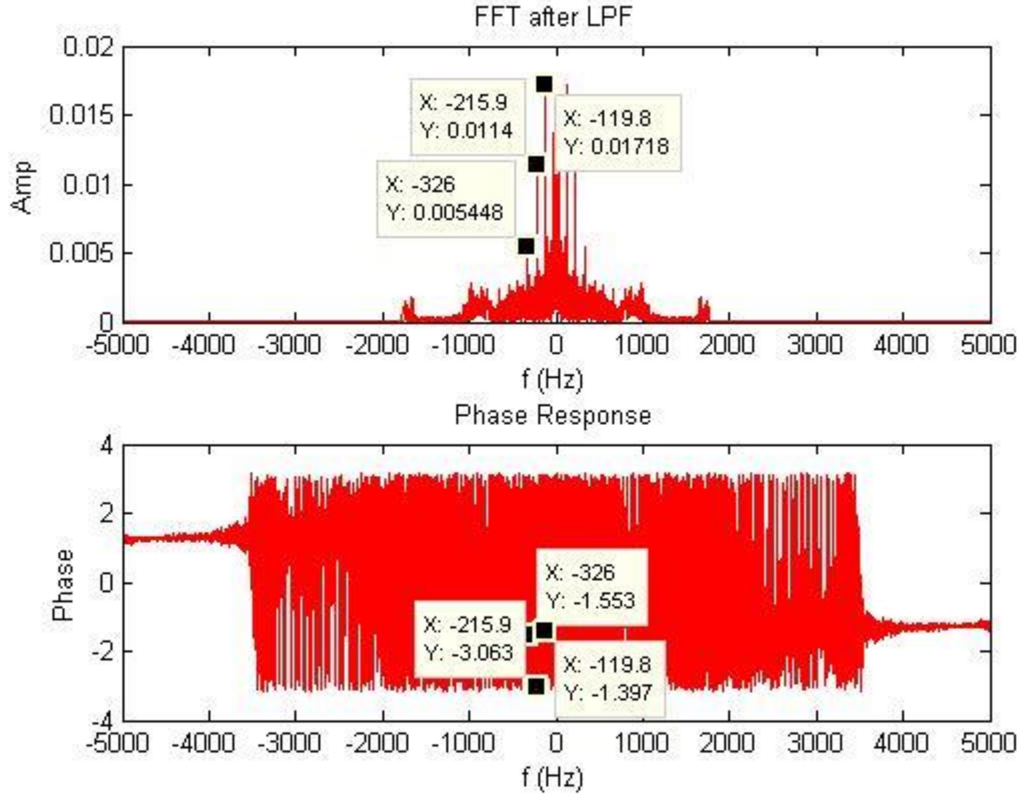

Figure 7: Noisy Signal with Scream at 2 feet and Voice at 18 feet, after the LPF.

II. Noise Source: Hair Dryer

(Note: Refer to Appendix A.2 for the code.)

For this experiment, the voice was kept constant and the noise distance was varied. This was done to provide a better SNR to aid in filtering out high-band spread that the hair dryer provided. Unlike the scream used earlier, the hair dryer violated the frequency band of the voice communication making it much more difficult to filter out as a result. The LPF was still used in this and the proceeding experiment to test the results of using a traditional means of filtering. Most of the noise could be filtered, but there is still left the challenge of filtering the lower frequency components. The next four figures (Figure 8-11) show the FFT of the hair dryer at the different test distances. The amplitude of the noise decreased with an increase in distance. This aids in filtering because of the increased SNR.

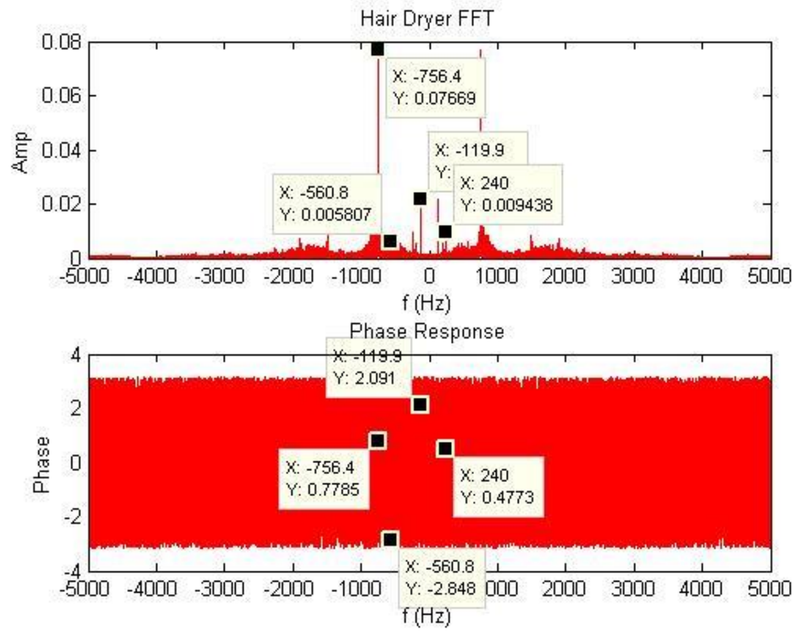

Figure 8: Hair Dryer FFT, at 2ft

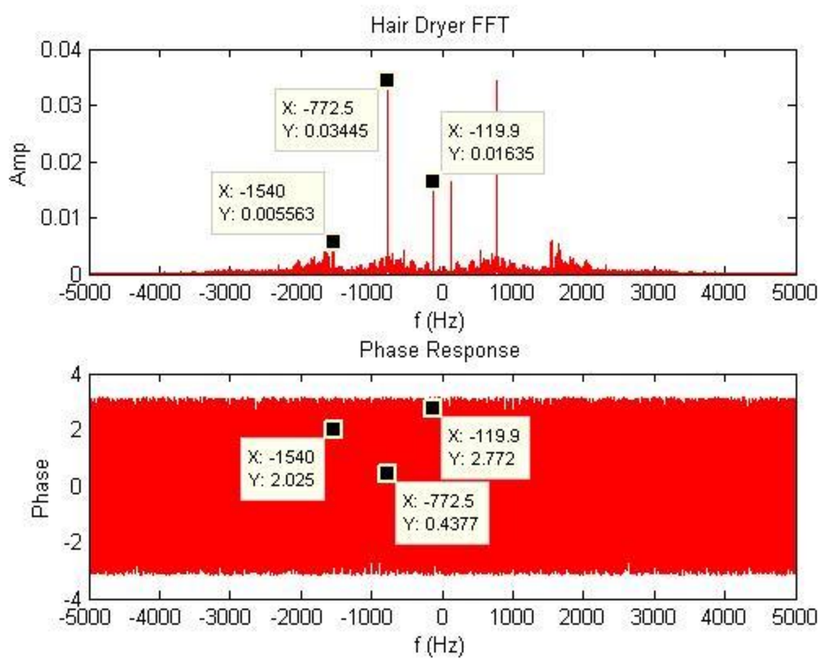

Figure 9: Hair Dryer FFT, at 6ft

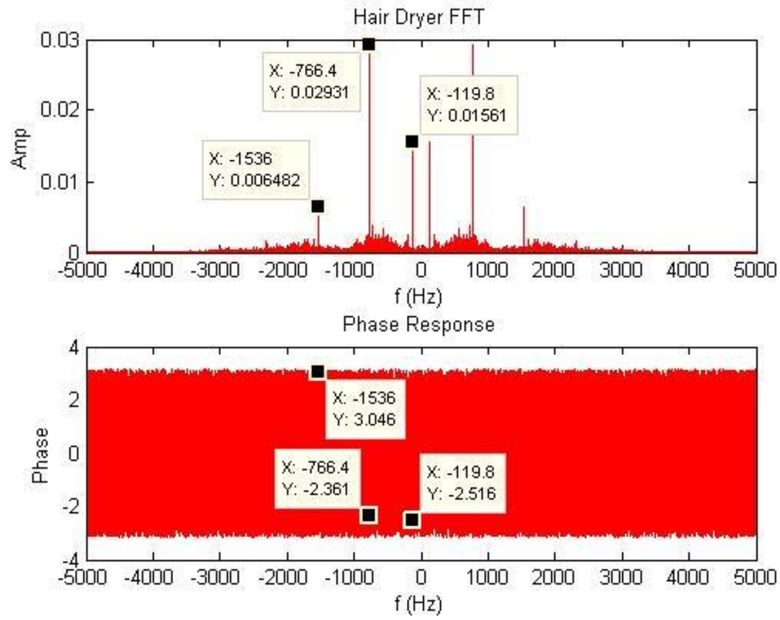

Figure 10: Hair Dryer FFT, at 12ft

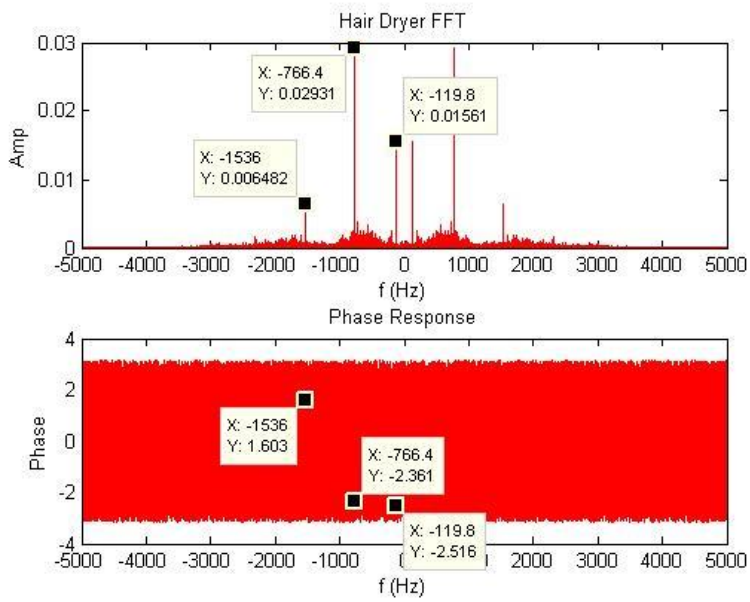

Figure 11: Hair Dryer FFT, at 18ft

Figure 12 below shows the spectra of the noisy signal at 2 feet. The hair dryer was over-powering at this distance. After passing through the LPF (Figure 13), the noise amplitude as attenuated, but there still remains high amplitude parts of the noise that remain within the band of the voice signal. Refer to the peaks at 756.1Hz and 119.7Hz. The amplitude has decreased by half, and the noise appears suppressed by listening to the audio, but it still has an obvious presence in the signal. Note in Figure 13 that most of the high frequency noise has been suppressed.

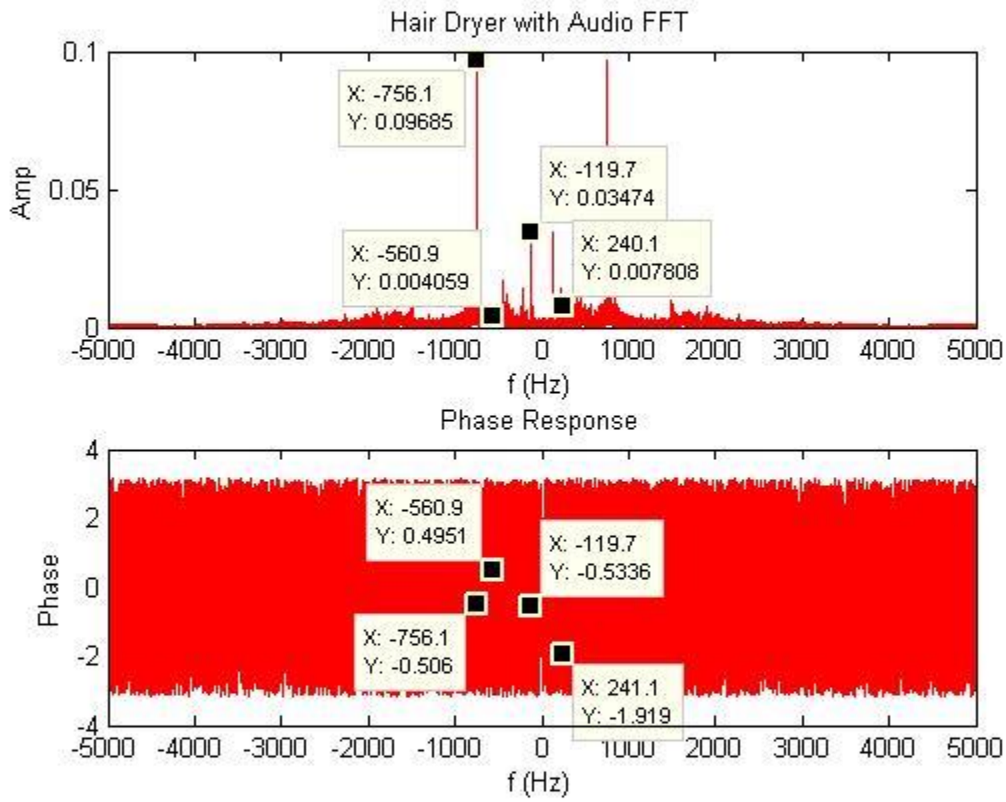

Figure 12: Noisy Signal: Voice and Hair Dryer at 2 feet.

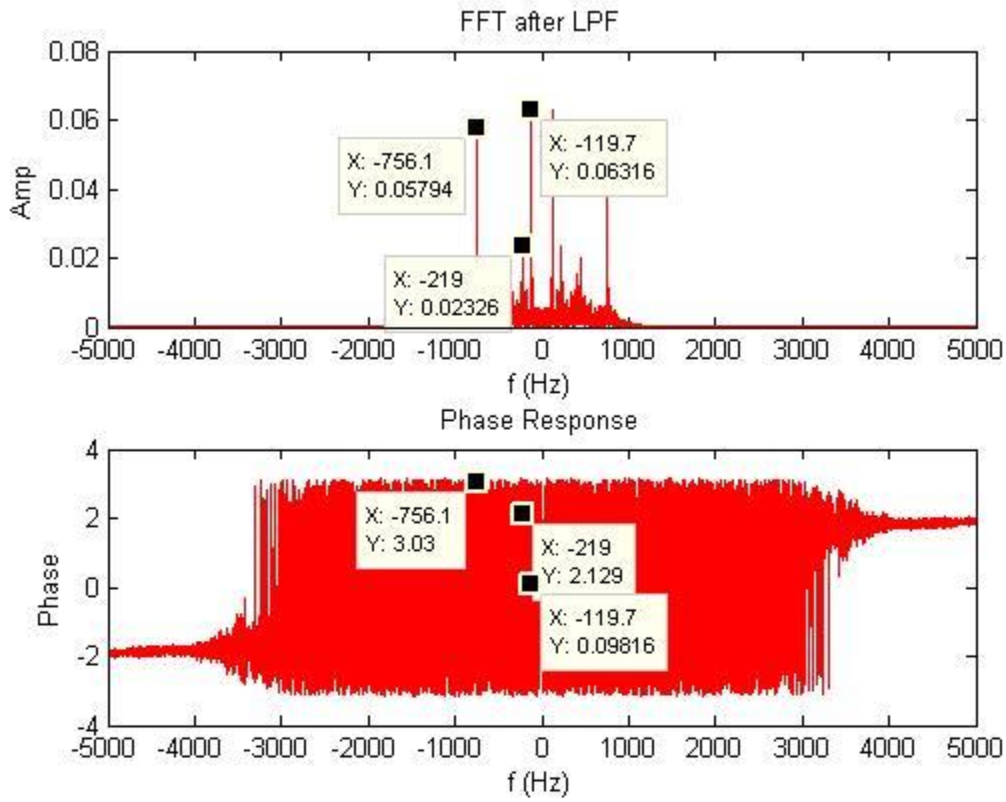

Figure 13: Noisy Signal after LPF: Voice and Hair Dryer at 2 feet.

Figures 14 and 15 are the spectras for the hair dryer at 6 feet. The amplitude of the noise has obviously dropped; although with slightly different frequency components. Now comparing Figure 15 to that of Figure 13, it is apparent that the noise has been suppressed even further compared to the adjacent voice components. The audio reveals a stronger voice presence than it did at 2 feet. The next four figures (16-19) are the spectras with the hair dryer at 12 feet and 18 feet. The same results were seen here.

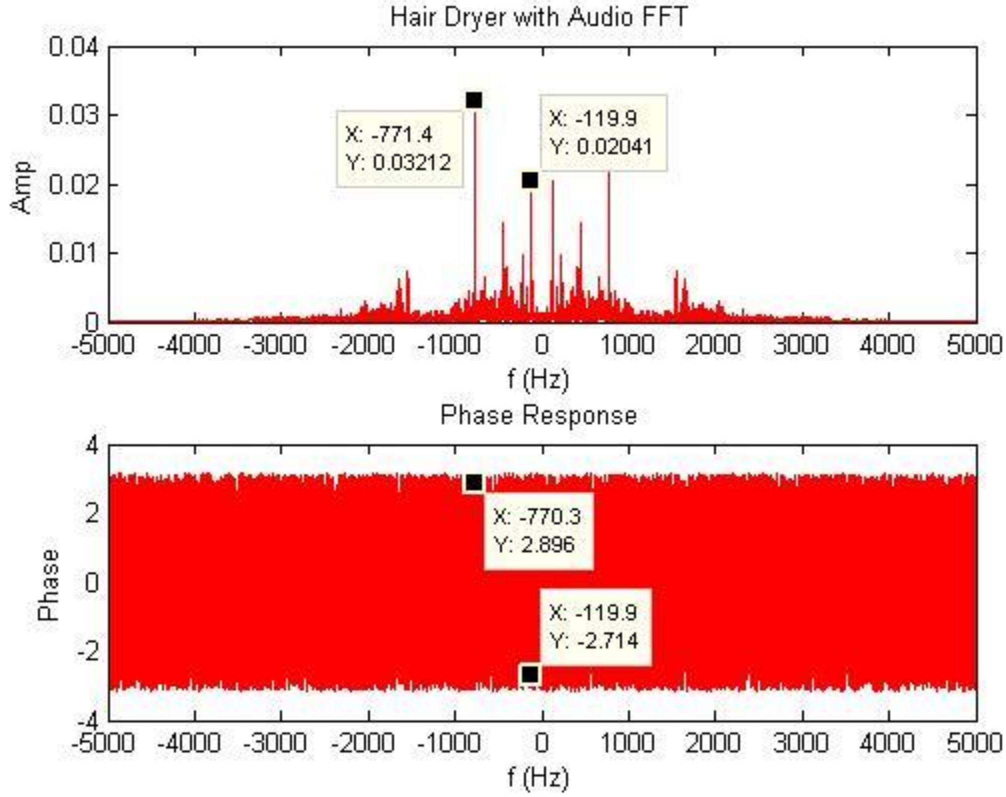

Figure 14: Noisy Signal: voice at 2 feet, Hair Dryer at 6 feet.

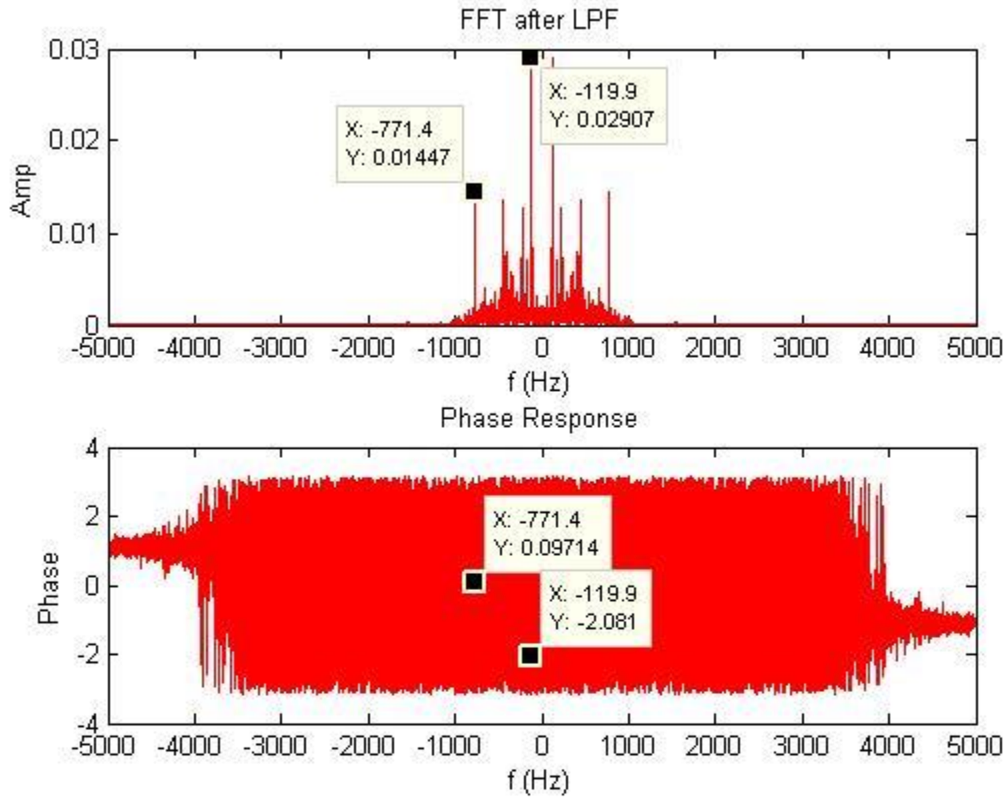

Figure 15: Noisy Signal after LPF: voice at 2 feet, Hair Dryer at 6 feet.

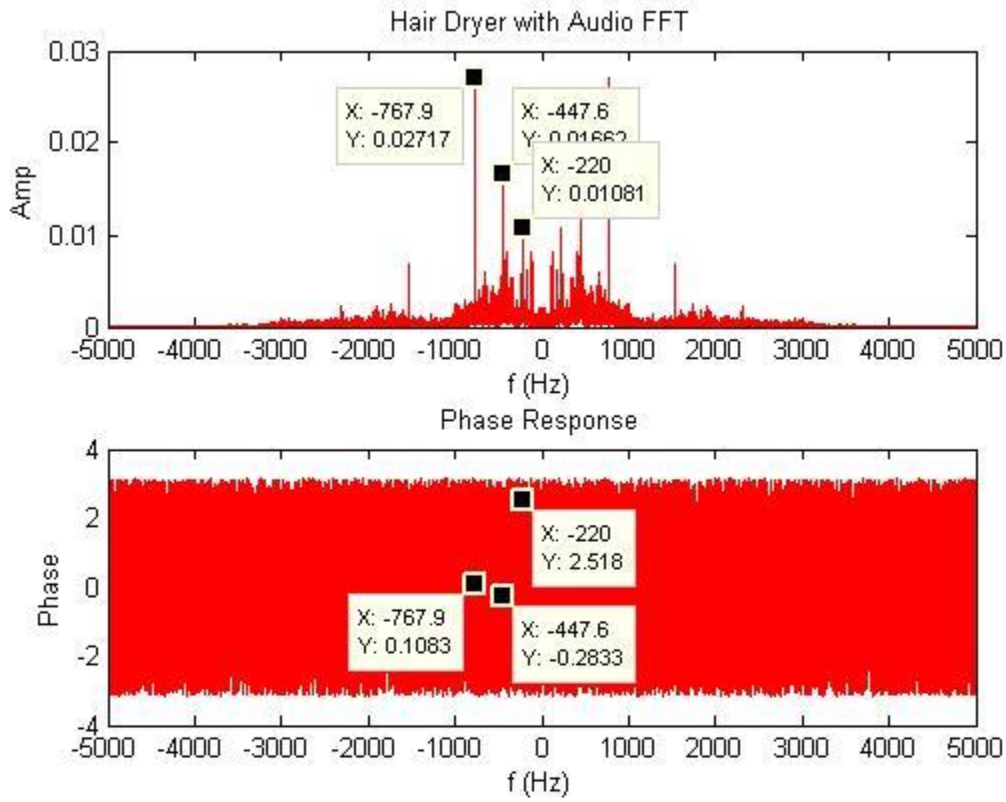

Figure 16: Noisy Signal: Voice at 2 feet, Hair Dryer at 12 feet

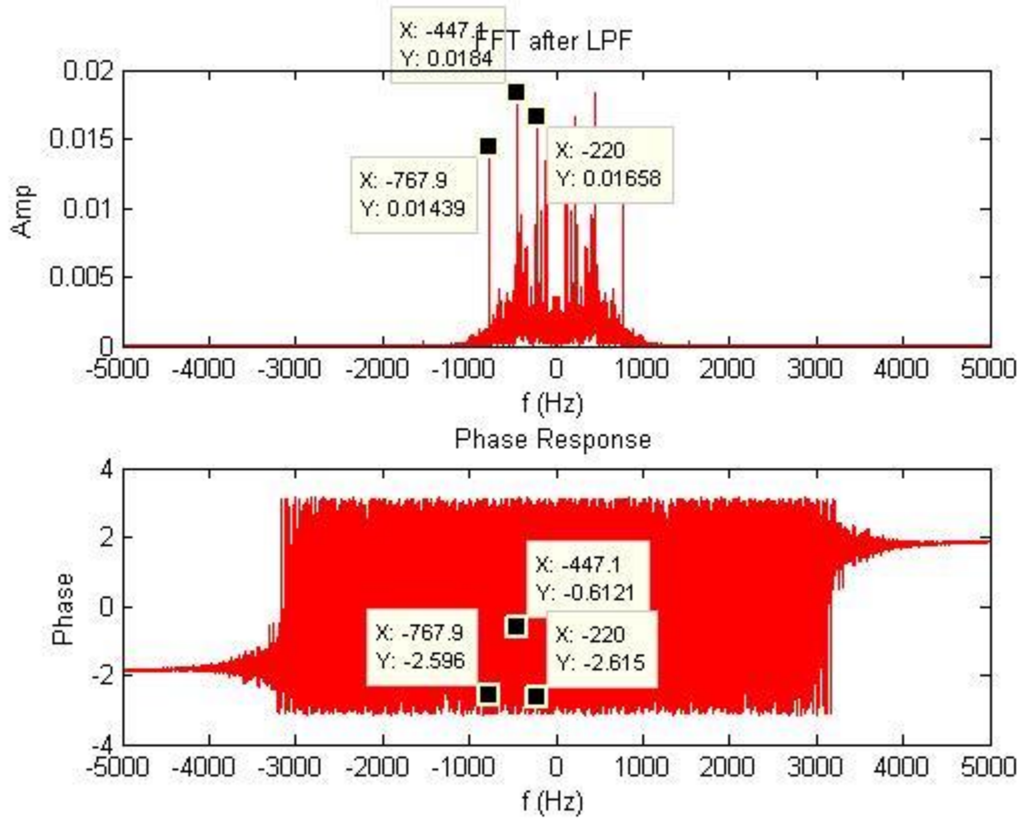

Figure 17: Noisy Signal after LPF: Voice at 2 feet, Hair Dryer at 12 feet

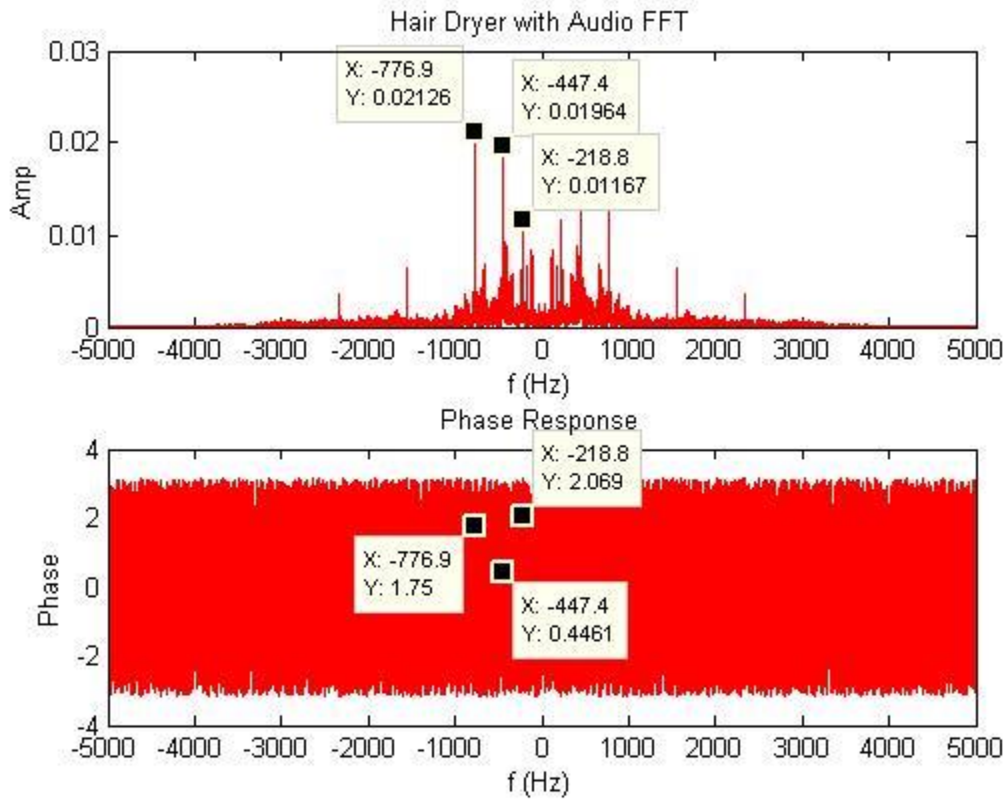

Figure 18: Noisy Signal: Voice at 2 feet, Hair Dryer at 18 feet.

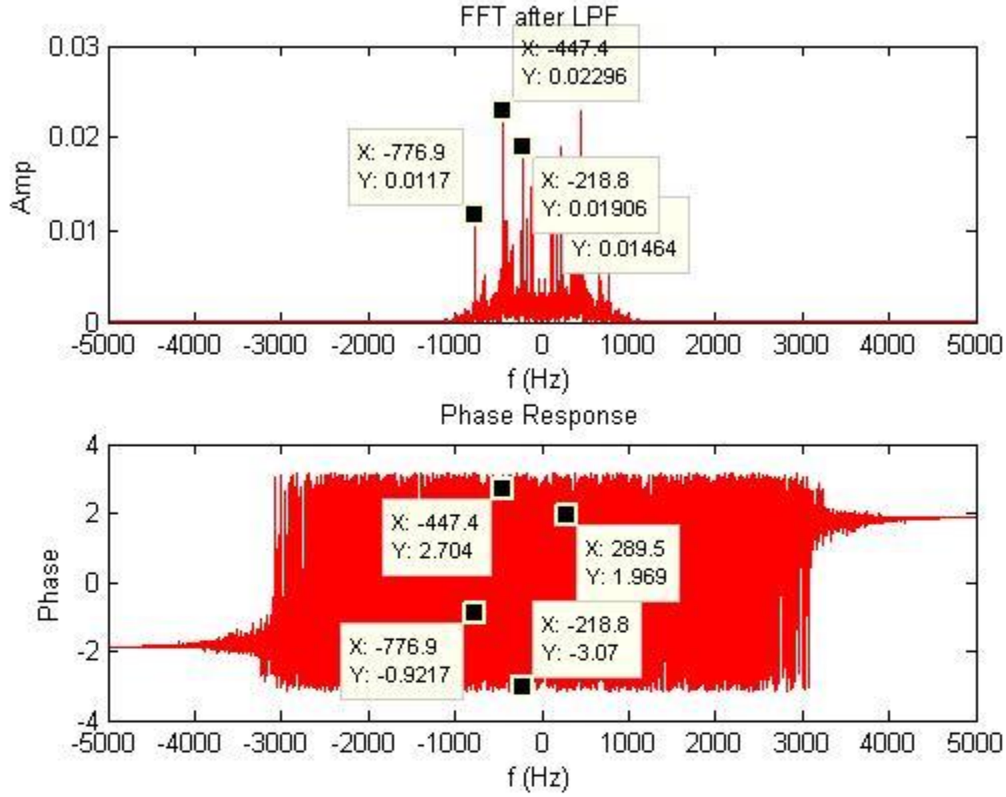

Figure 19: Noisy Signal after LPF: Voice at 2 feet, Hair Dryer at 18 feet.

III. Noise Source: Dremel Tool

(Note: Refer to Appendix A.3 for the code.)

A dremel tool was used for the third noise source. The motor of the dremel operates at a much higher RMP than a hair dryer, so it was expected that the spectra would reveal higher frequency components. The next four figures (Figures 20-23) show the dremel frequency spectra at the different distances. One interesting thing about Figures 20 and 21 are the frequency components present at 478.2Hz for Figure 20 and 487.9Hz for Figure 21. Note the amplitude in Figure 20. The frequency at 478.2Hz is the dominate magnitude, yet at 6 feet the magnitude is not the dominate component. This indicates the magnitude not only changes with distance, but certain frequencies carry further over distance. Also note that the frequency around 478Hz decreases in magnitude over distance but the lower frequency around 120Hz does not. It’s fairly consistent in magnitude.

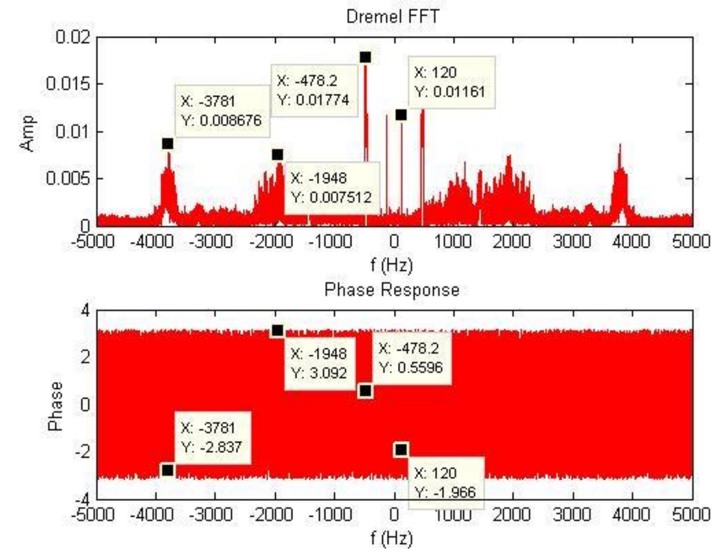

Figure 20: Dremel FFT, at 2 feet

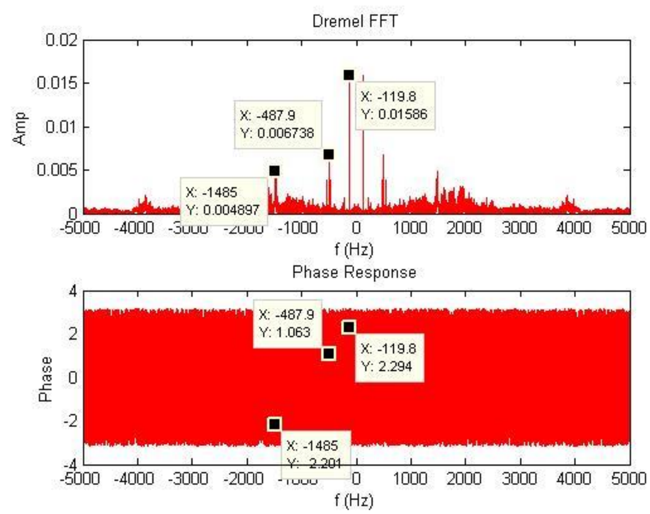

Figure 21: Dremel FFT, at 6 feet

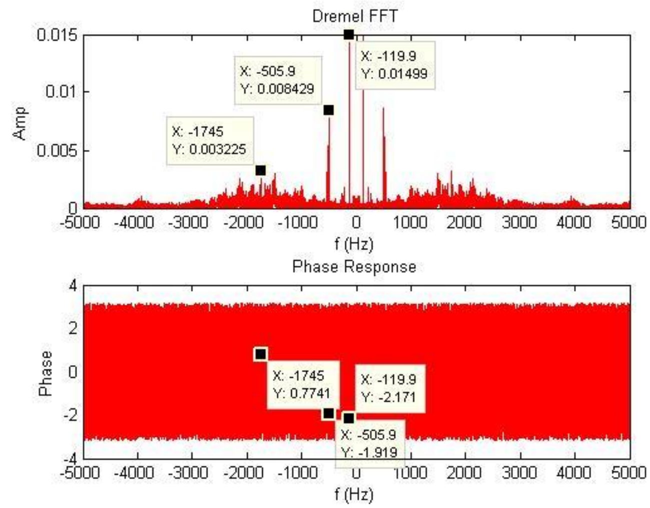

Figure 22: Dremel FFT, at 12 feet

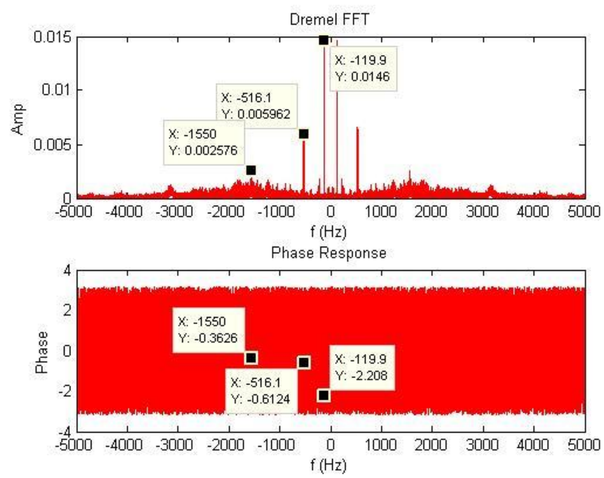

Figure 23: Dremel FFT, at 18 feet

Figure 24 below shows the FFT for the noisy signal. Note the wide frequency spread of the Dremel is much wider than that of the hair dryer. Most of the high frequency components can be filtered out, but there is still plenty of noise in the vocal band. Figure 25 shows the results after the signal passes through a LPF. Most of the frequencies above 1KHz are attenuated. Many of the noisy components remain in the lower frequency band.

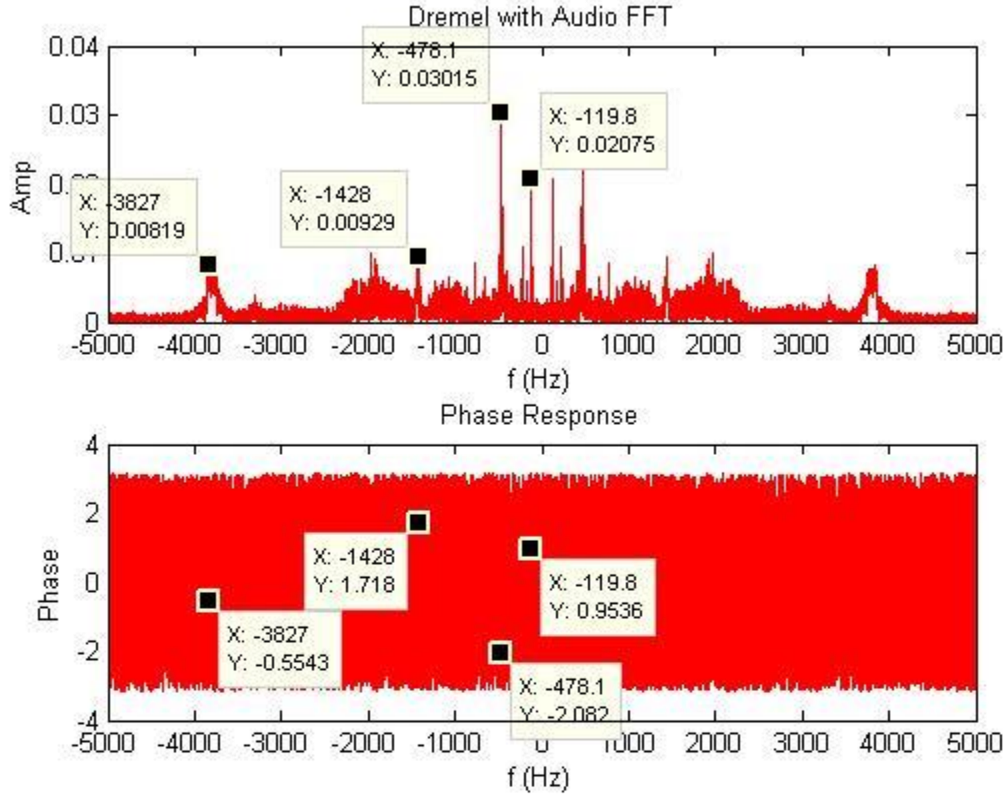

Figure 24: Noisy Signal: Voice and Dremel at 2 feet.

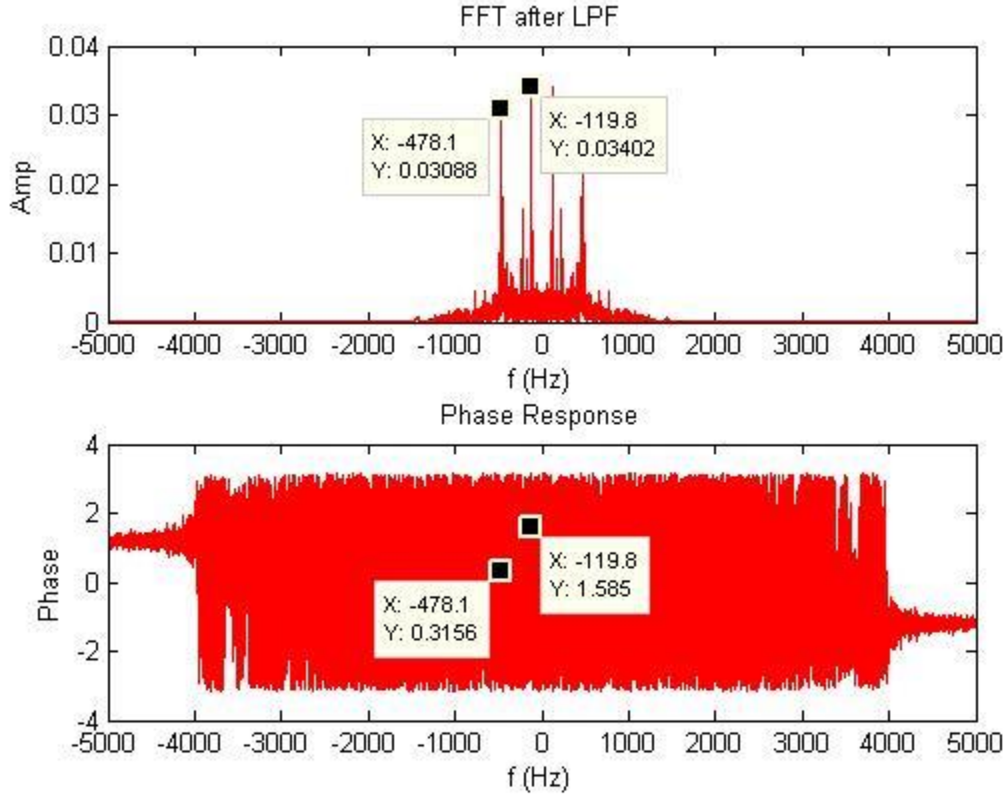

Figure 25: Noisy Signal after LPF: Voice and Dremel at 2 feet.

In Figure 26, the Dremel had been moved back to 6 feet. The amplitude of the noise is not as great resulting in a higher SNR. Figure 27 shows the results after the LPF. The noise is not as prevalent, but there is still some back ground noise. Figures 28-31 show the spectras at 12 feet and 18 feet. The results were similar to that of the hair dryer.

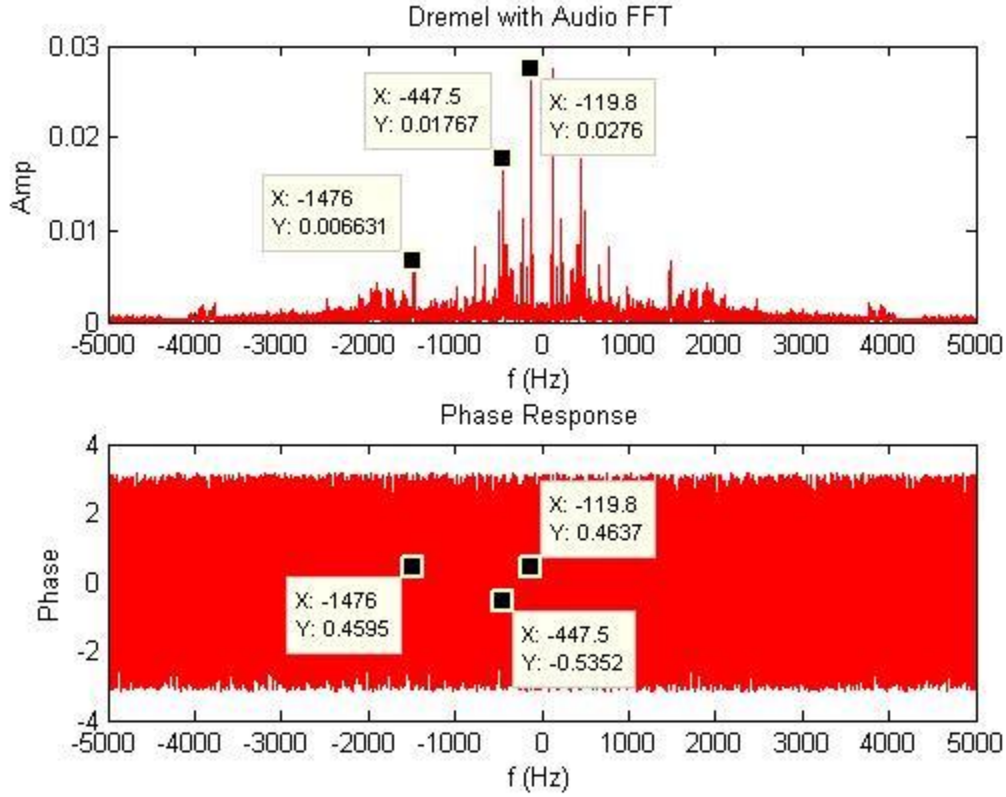

Figure 26: Noisy Signal: Voice at 2 feet, Dremel at 6 feet.

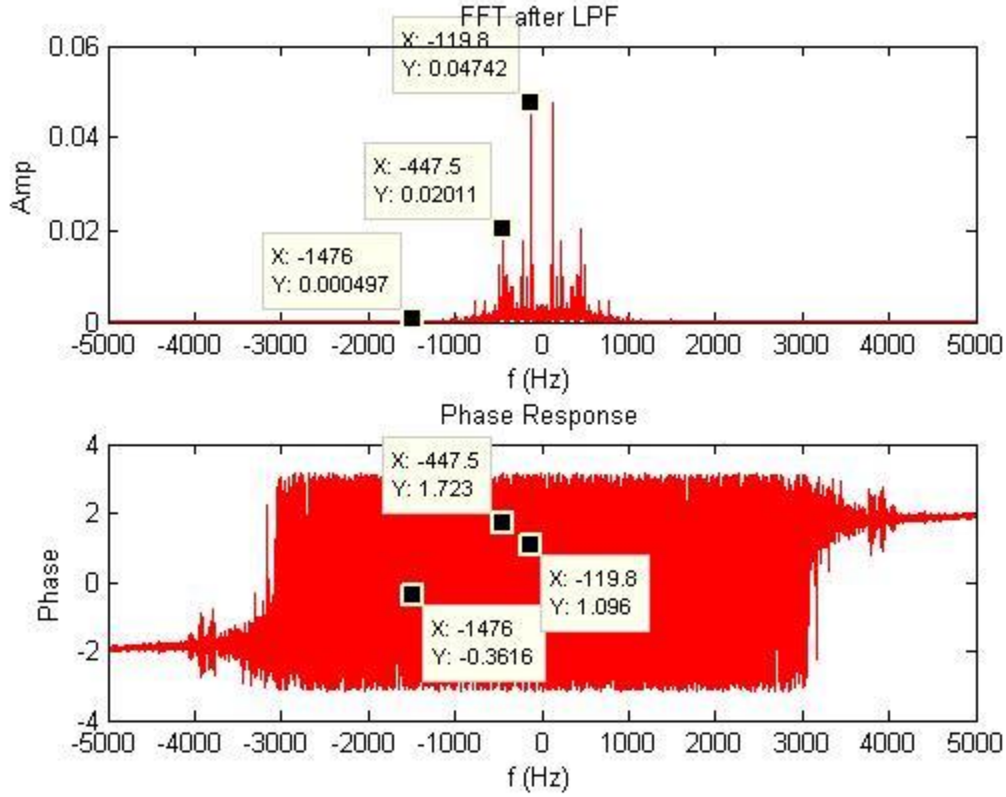

Figure 27: Noisy Signal after LPF: Voice at 2 feet, Dremel at 6 feet.

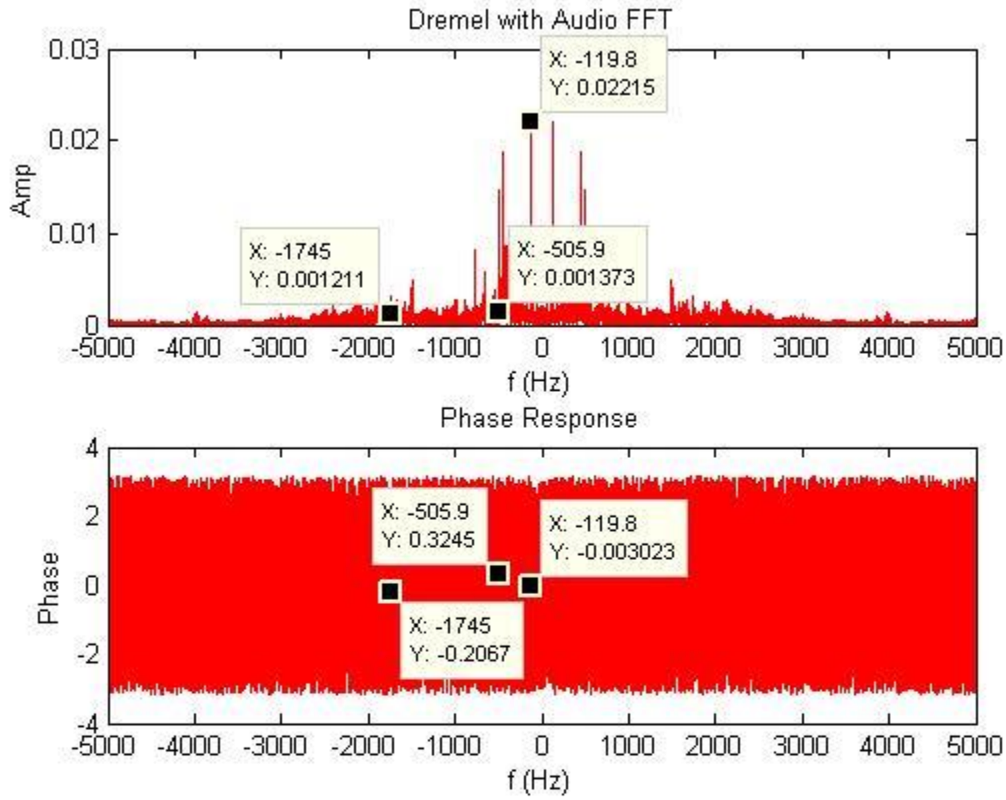

Figure 28: Noisy Signal: Voice at 2 feet, Dremel at 12 feet.

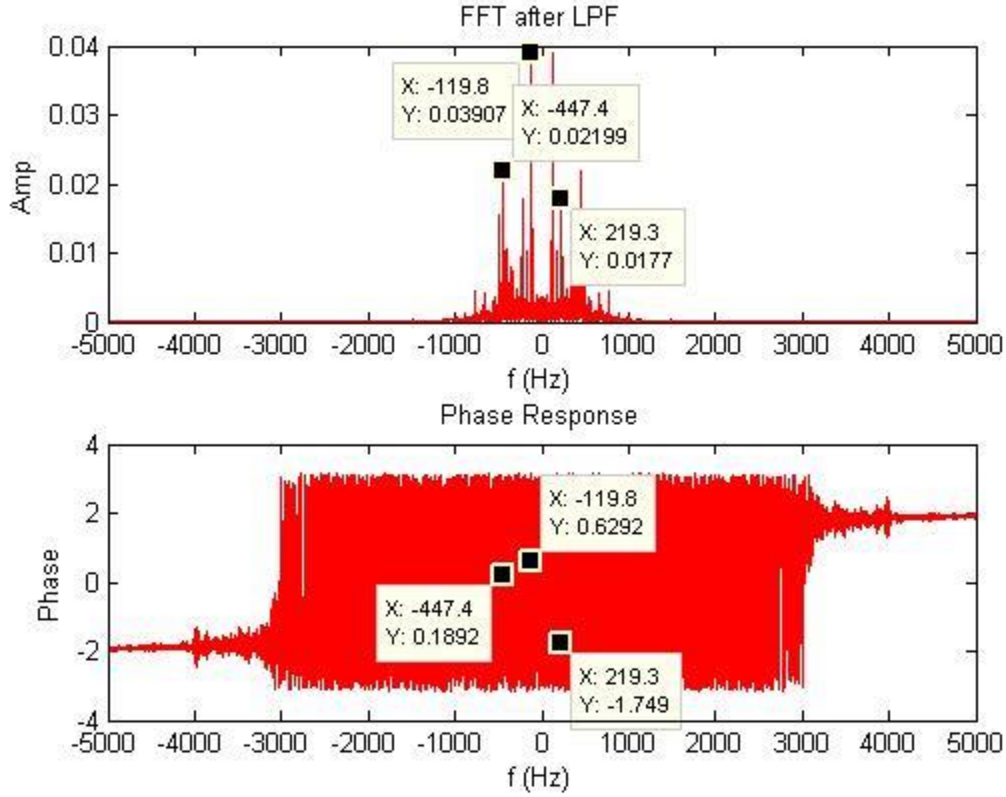

Figure 29: Noisy Signa after LPFl: Voice at 2 feet, Dremel at 12 feet.

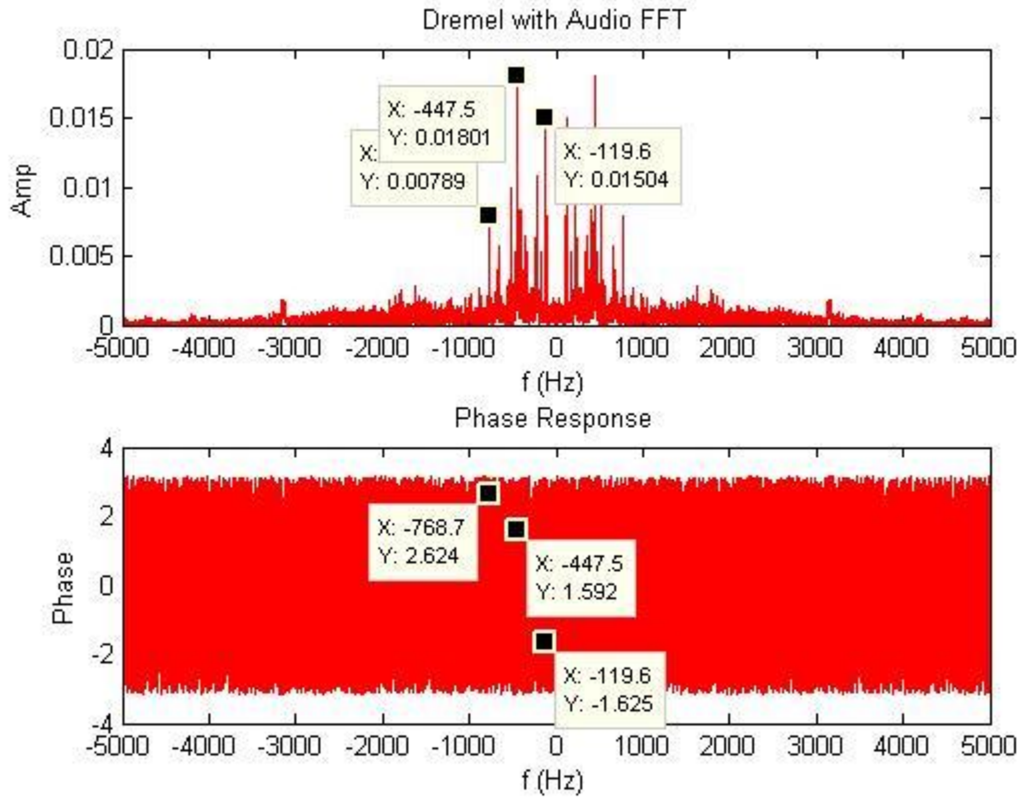

Figure 30: Noisy Signal: Voice at 2 feet, Dremel at 18 feet.

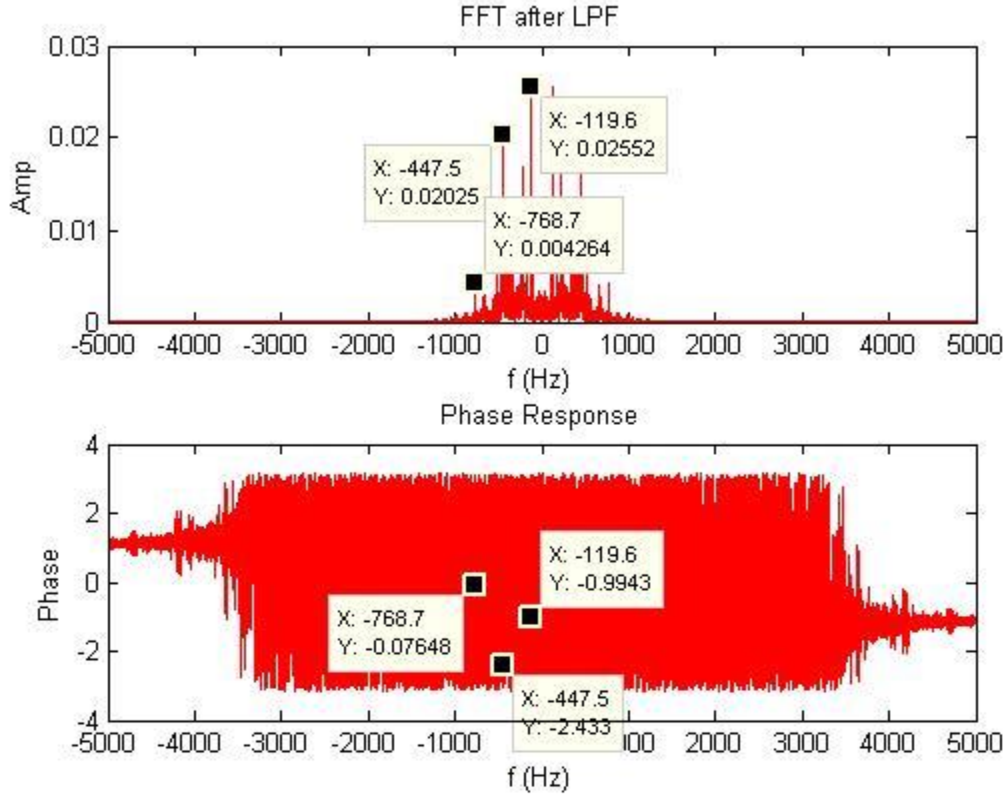

Figure 31: Noisy Signal after LPF: Voice at 2 feet, Dremel at 18 feet.

IV: Other Options

Template Matching

(Refer to Appendix A.4 for the code)

The voice signal in Figure 32 was mixed with the sound of a vaccum. The same message was broadcast as in the preceding experiments. A voice message at 6 feet was used to totally drown out the voice signal. The voice recording was used as a template to attempt to filter out the voice message from the noisy signal. The template is identical to the voice signal embedded in the noise. This way phase did not play a part in trying to extract the message out. Figure 33 shows the FFT of the voice signal. Figure 34 is the vacuum noise.

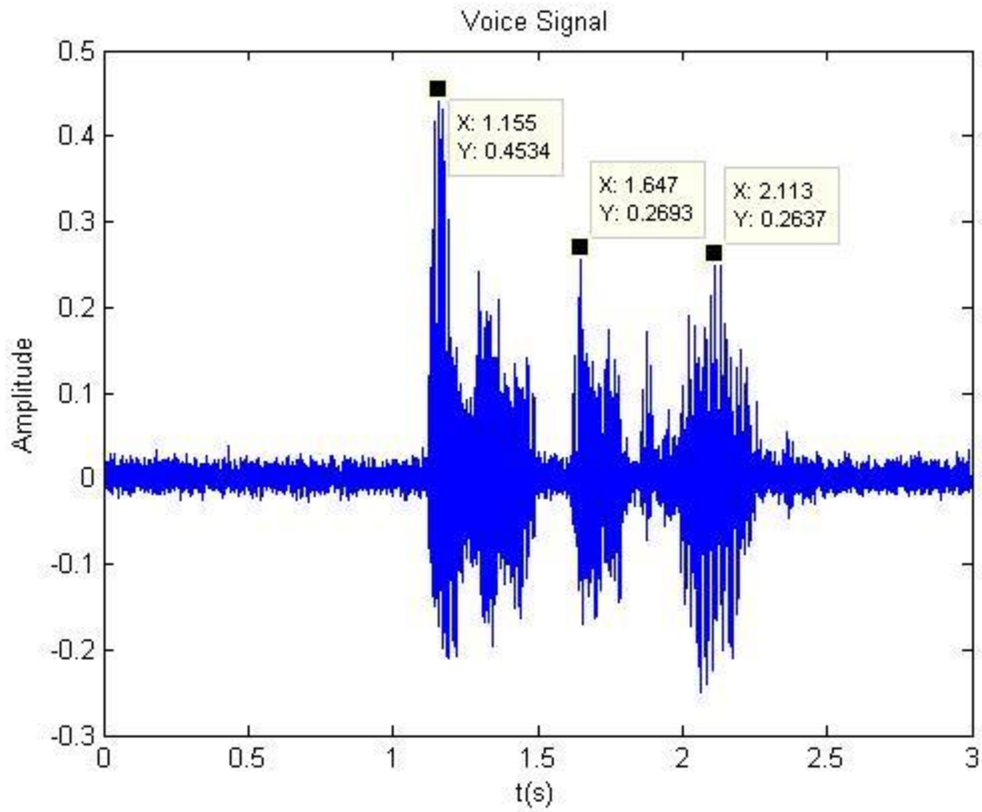

Figure 32: Voice Signal at 6 feet.

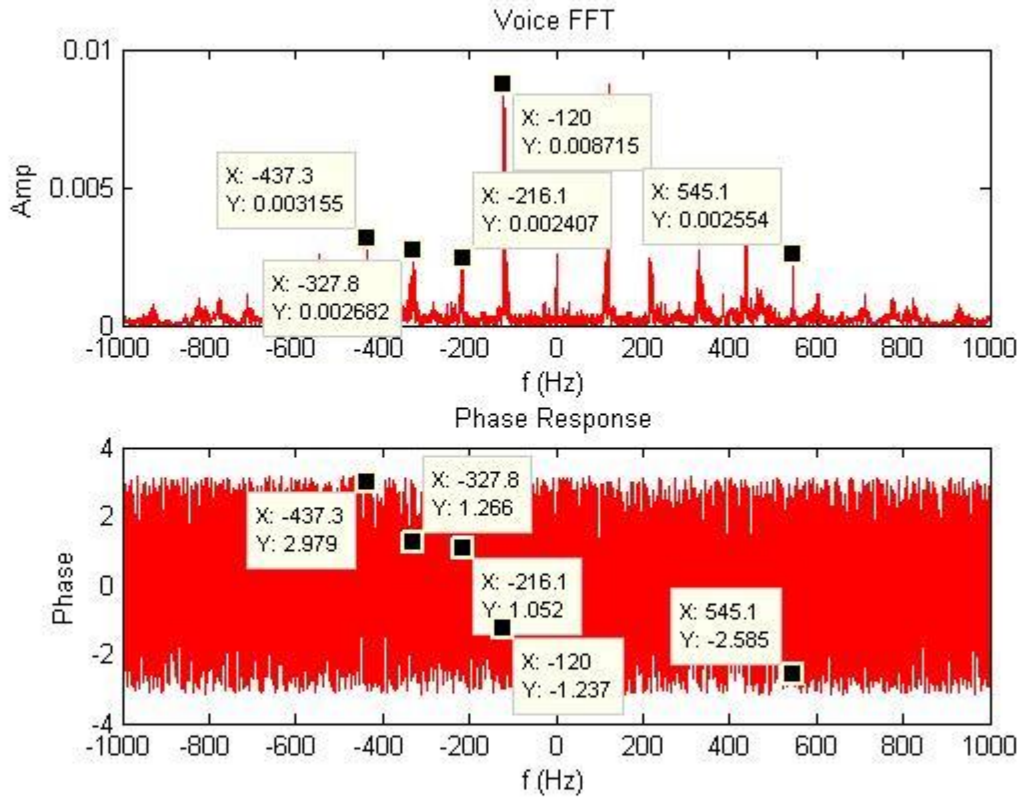

Figure 33: FFT of voice signal at 6 feet.

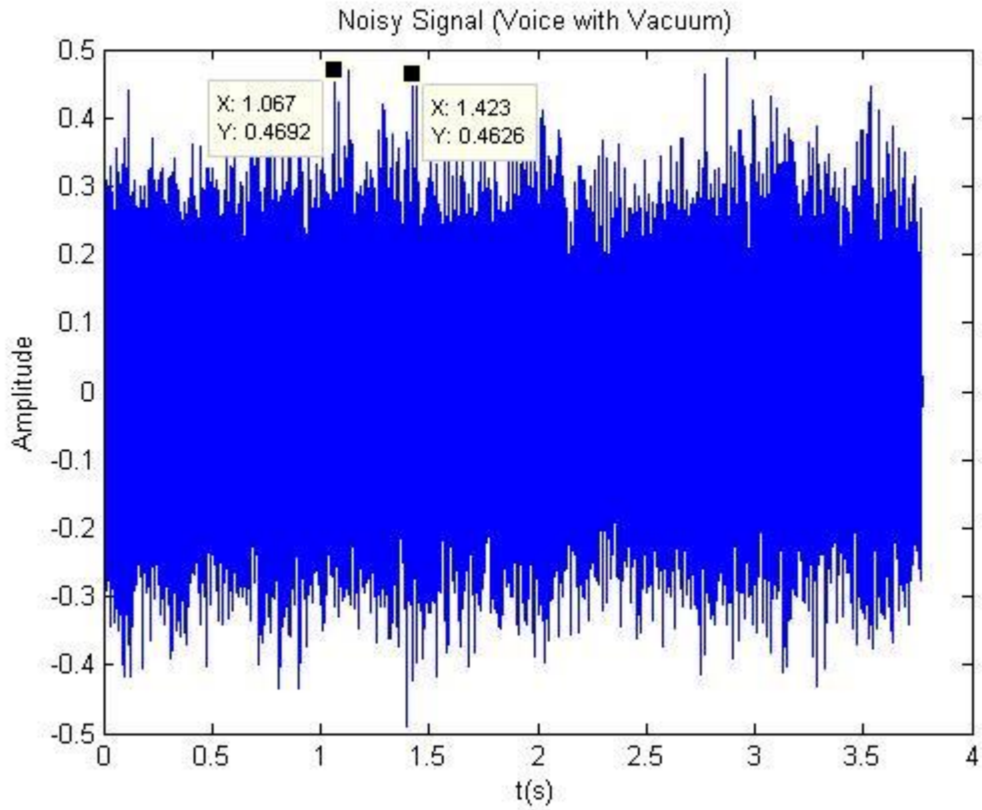

Figure 34: Noisy Signal in Time Domain

Figure 35 below is the FFT of the noisy signal. Most of the dominant frequencies are in the lower voice band. The noisy signal was passed through a For loop in an attempt to extract the message from the noisy signal using the voice template. Figure 36 is the result of this attempt. As can be seen by comparison to Figure 32, this attempt failed. The time domain signals look nothing alike.

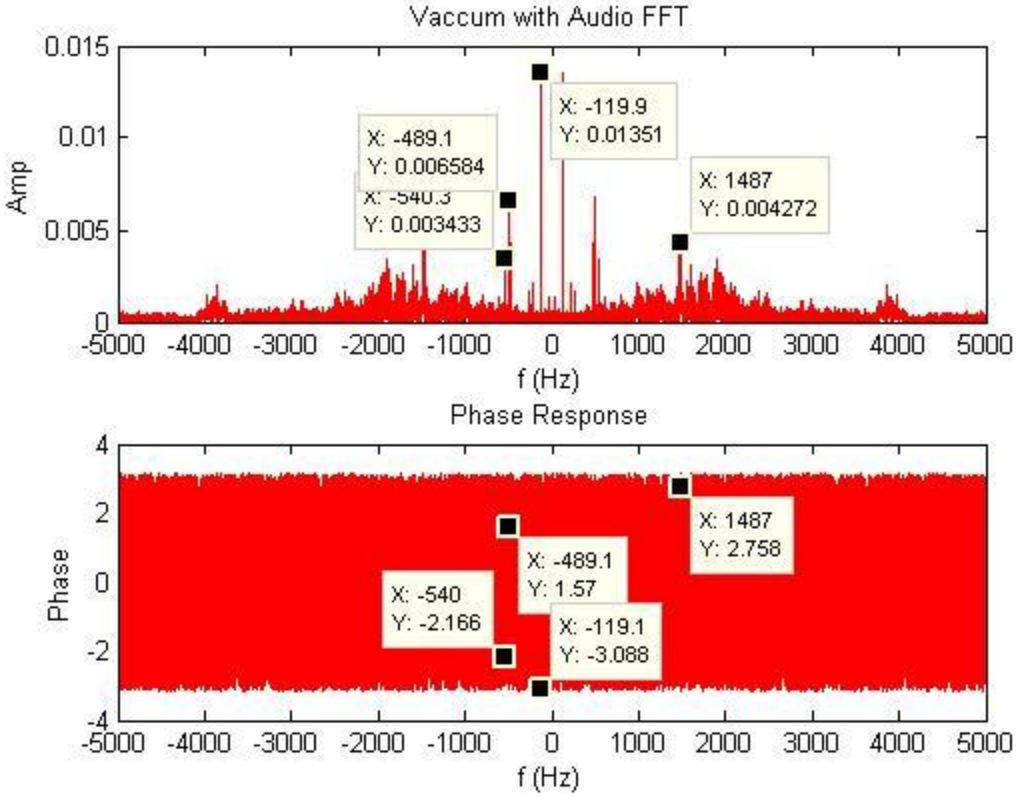

Figure 35: FFT of Noisy Signal

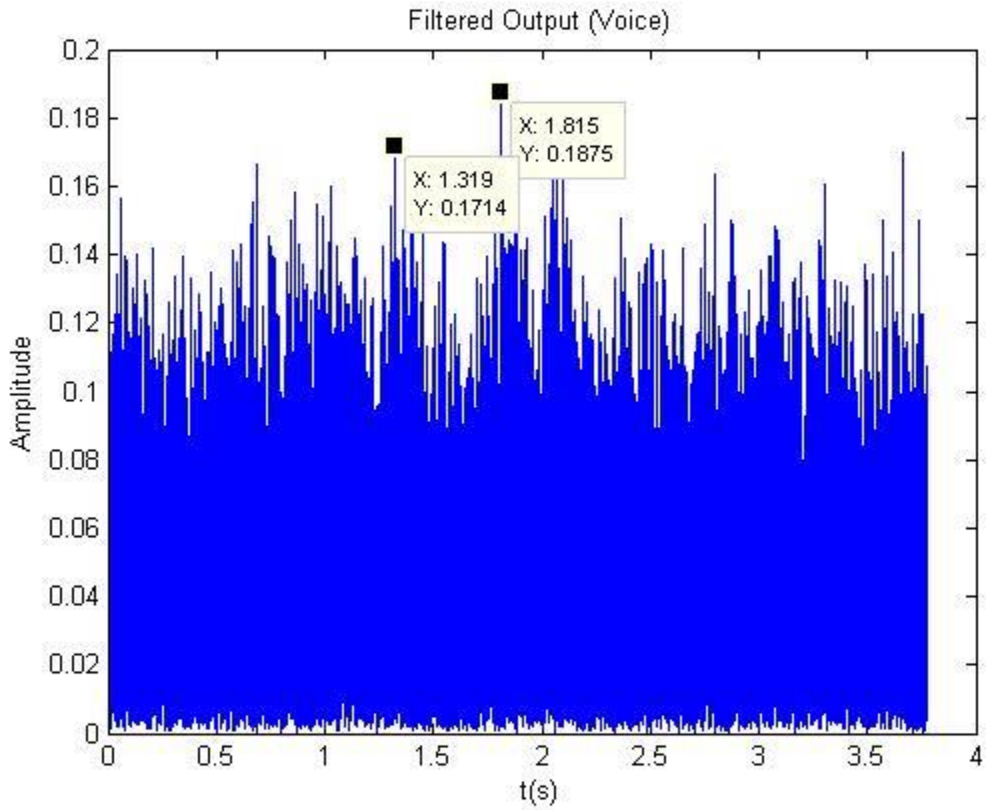

Figure 36: Output after Filtering

FIR1 Filter[1]

(Refer to Appendix A.5 for the code)

The FIR1 filter was used to take advantage of its flexibility. It supports all four basic filter designs, LPF, HPF, BPF, and BSF. As was expected the LPF had the same results as the LPF model used above. A band-pass filter was then used with cutoff frequencies of 15Hz and 500Hz. The sound of the noise was decreased substantially but the voice message was still a little muffled.

Conclusion

Given the right circumstances, noise can be easily filtered from a signal leaving only the intelligence intact. It has been proven through Fourier analysis that noise can be filtered out of a signal using conventional filtering techniques and conventional filters. The LPF was ideal for the first part of the project considering the low frequency band for voice and the high frequency components found in a child's scream. This project also shows the difficulty and challenges that await anyone who attempts separation. There are advanced filtering techniques that could assist in this matter, and with enough time and resources, better filtering can be achieved.   However, the scope of this project was limited and it's believed that the objectives and goals have been met satisfactorily; even though there was limited success at filtering the noise of the hair dryer and dremel tool. Finding out what doesn't work is just as important as finding out what does work. In this respect, success was achieved.