Noise-Cancelling Headphones

Danylo Hirnyj
Final Project
Phys 406, Spring 2012

Motivation

In fall 2011, I studied abroad in Stockholm, Sweden. Besides living in Sweden, I was eager to travel around Europe. Unlike in the U.S. it is possible to find heavily discounted airline tickets, an awesome luxury for a college student. Over the course of the semester, I took at least 30 flights and many bus rides. It was often difficult to rest in these environments, largely due to the heavy noise pollution. Since then I have wanted to buy myself a pair of noise-canceling, though the price for these can run quite high in general.

Once I realized that this course allowed students to work on essentially any project related to acoustics, I had some difficulty trying to find one that would be interesting, feasible, and affordable. Given so many possibilities, I would often stumble across a possible idea that was interesting and affordable, but not feasible in a single semester, or interesting and feasible, but not affordable. After a few days of browsing the web and watching many different projects exhibited on YouTube, I came across a circuit outlining a basic design for noise-canceling headphones. It was not overly-complicated, so I was confident that even for someone with limited circuit-building experience, it would be feasible. Further, the project would be completed not only for this course, but I would be able to use what I built during my travels.

How It Works

In theory, the idea of noise-canceling headphones is quite simple once one understands the basics of sound waves. The simplest way to understand the mechanism behind noise-canceling headphones is to imagine a transverse sinusoidal sound wave at a single frequency. See Figure 1.
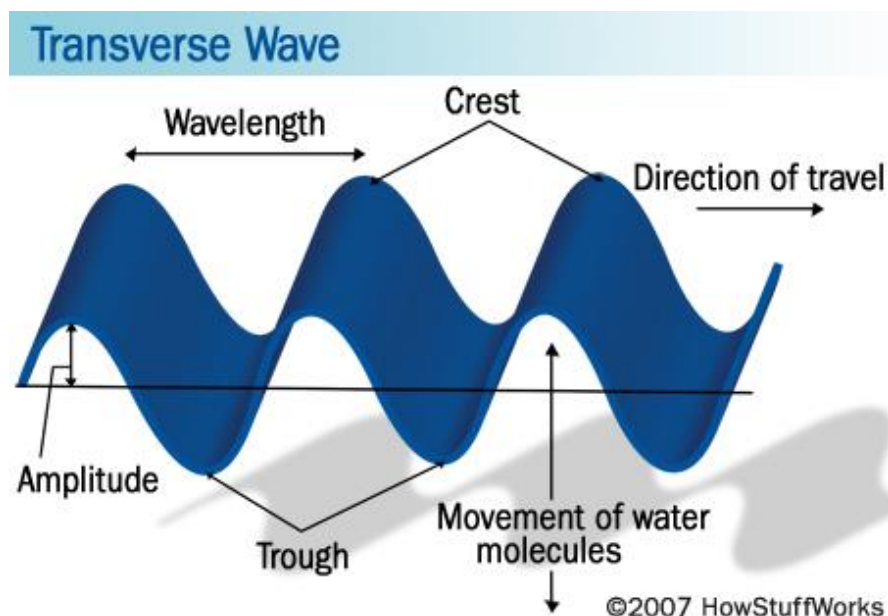

Figure 1

The principle of superposition tells us that, given an arbitrary number of waves, we can sum them up at the same position and time to obtain a new wave. See Figure 2.
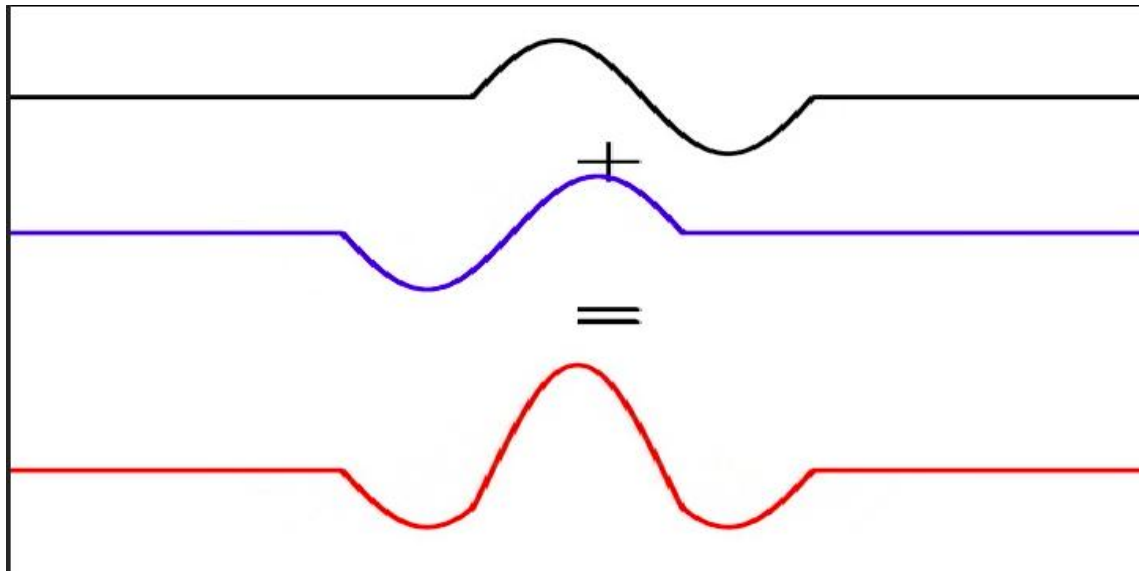

Figure 2

This simple idea turns out to be the only tool we need to develop the theory behind noise-cancellation. Since we want to cancel incoming sound, our goal is to create our own "noise" that can produce destructive interference with the noise coming from the environment. This is the same idea as above, except this time we have a specific end-result in mind: a waveform that is as close to zero amplitude as possible, minimizing the sound intensity we hear. This idea is illustrated in Figure 3.
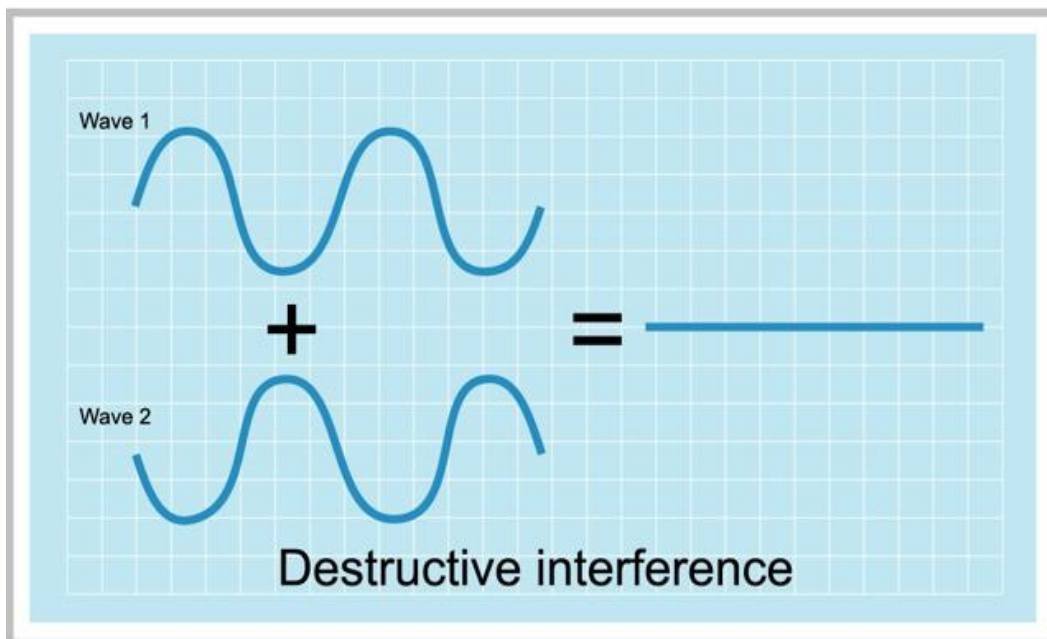

Figure 3

Notice that destructive interference occurs when a wave interferes with its exact inverted self. Hence, the design becomes clear. All we need is a way to capture the noise from the outside, invert the signal and output it in time such that we can match the crests of the original wave to the troughs of the inverted wave. For this, we require a microphone that can convert the noise outside into the form of an audio signal and circuitry that will invert the signal and output it through the headphones, illustrated in Figure 4.
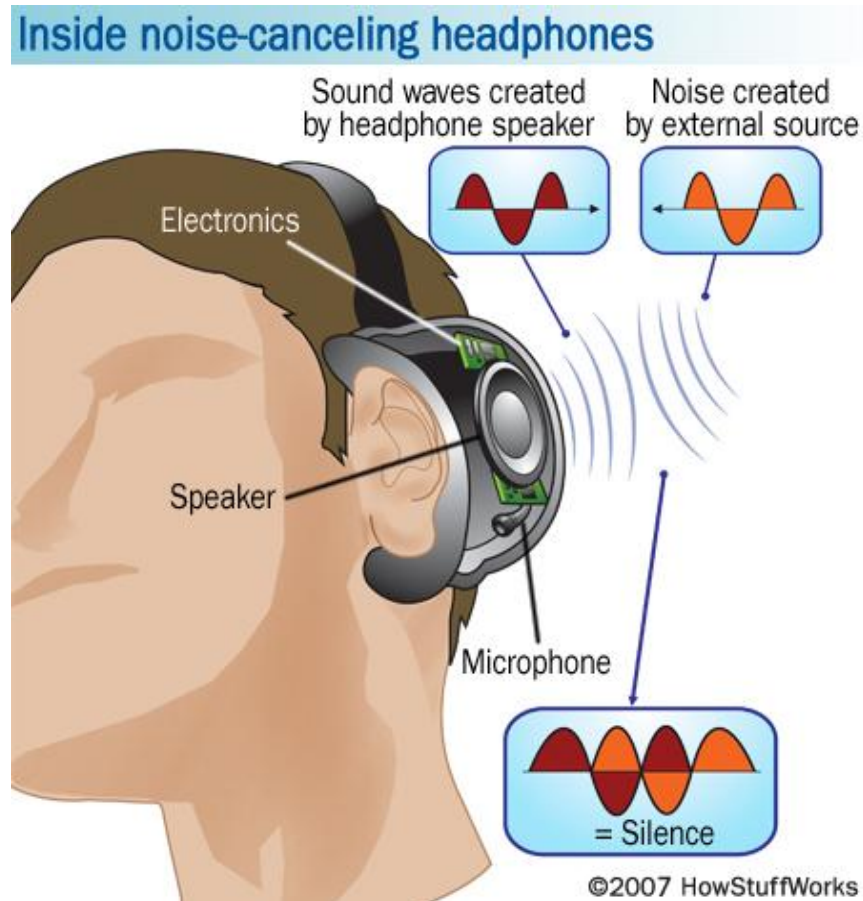

Inside noise-canceling headphones

Sound waves created by headphone speaker   Noise created by external source

Electronics

Speaker

Microphone

= Silence

©2007 HowStuffWorks

Figure 4

The Design

There are of course many different possible circuits that can be used to invert a signal. Further, it is important to understand the signal strength provided by the microphone and necessary to drive the headphones. Hence, the circuit is simple, but certainly not trivial. I thus decided that for my initial attempt at a design, it would be best to find an existing circuit-design and attempt to re-create it myself. A bit of browsing on the web led me to a project done by Jules Ryckebuschm during his years at Northwestern University. He provides a clear circuit-schematic (see Figure 5) and explains various components and design decisions.
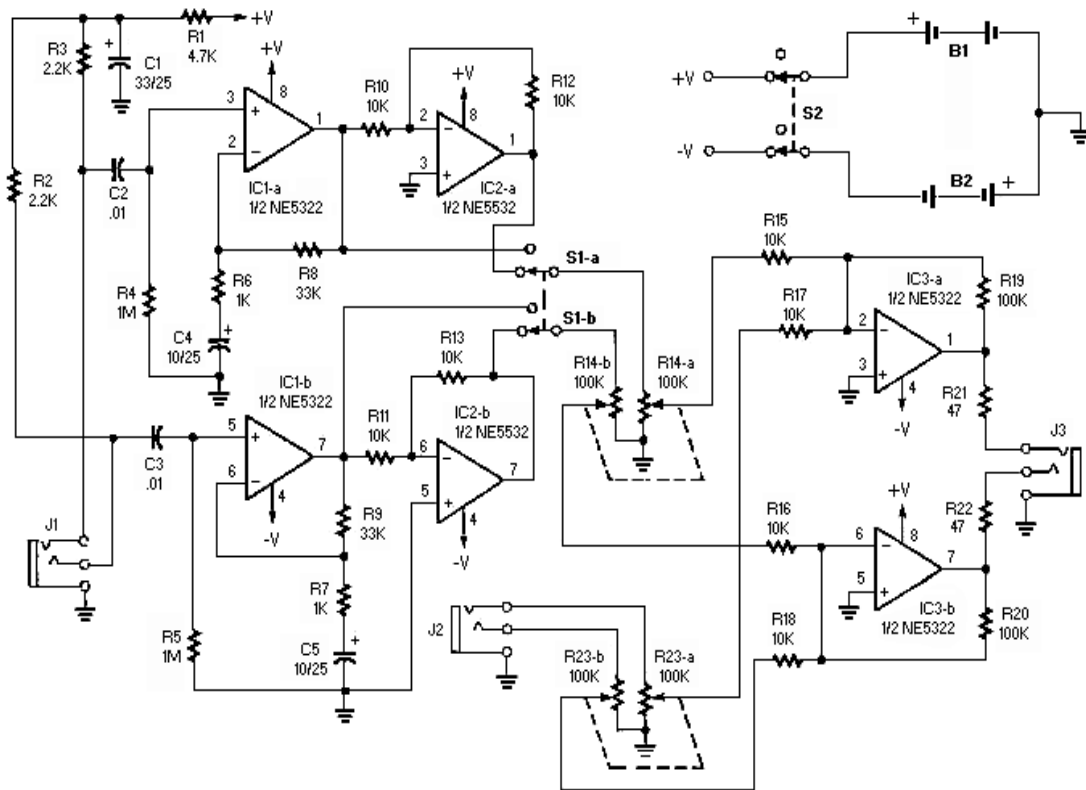
Figure 5

However, Jules did not provide a layout of this circuit on a perfboard. Given my limited experience working with circuits prior to taking this course, I decided that it would be best to proceed by coming up with a component layout schematic, so that I was sure there was enough space for everything and that everything was as clean as possible. To do this, I consulted another student with the same concern for his own project who recommended a program called Do It Yourself Layout Creator version 3.2.0 beta [http://code.google.com/p/diy-layout-creator/downloads/detail?name=diylc-3.2.0-beta.zip&can=4&q=].

Using the program I began to create a design for the circuit, including as many components as I could. After getting through about half of the circuit, I began to grow a bit frustrated since the program was quite buggy. It was also difficult to map many adjacent components since they would overlap and it was difficult to see what was going on. I decided that the amount of the circuit that I managed to lay out successfully was sufficient to start building. Figure 6 shows the layout I completed prior to abandoning this intermediate step of the project.
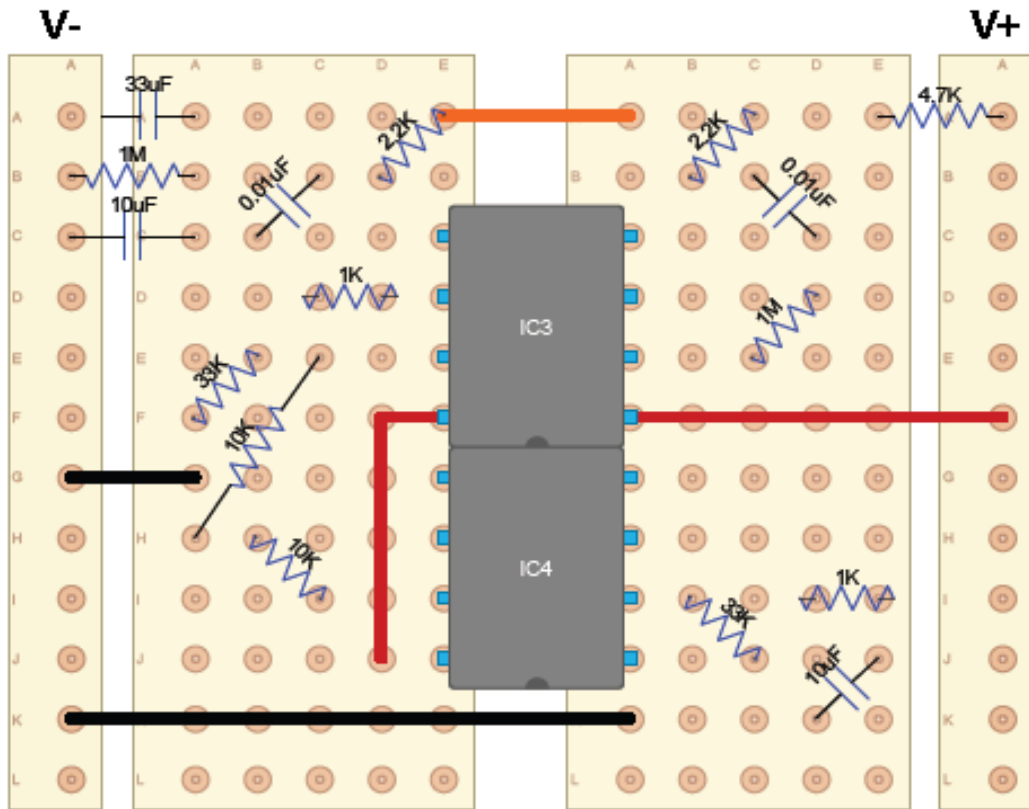
Figure 6

Given this figure, I could estimate how much space I needed in certain areas of the board which proved quite helpful for the actual build.

Circuit-Breakdown

The first thing to notice in Figure 5 is that it has two identical circuits parallel to each other. This is because we are dealing with a stereo system, and we must cancel each ear independently since the noise at each location is different.

The easiest way to understand the circuit is to break it down into sub-circuits and see what each part is doing. We will focus on the configuration of each op-amp in the top half of the circuit, representing one channel. The same applies to the other channel shown in the bottom half of the circuit.
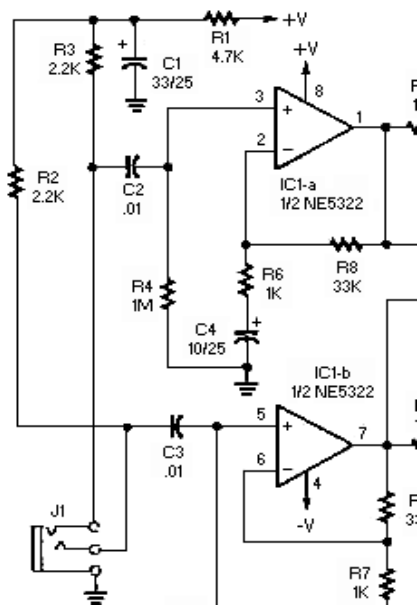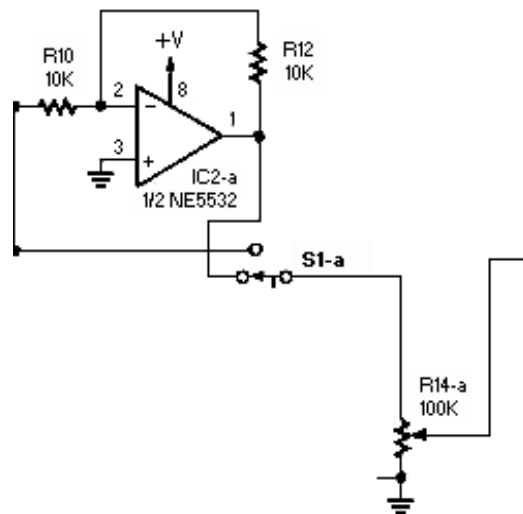
Figure 7

Figure 8

Looking at the connections for 1C1-a in Figure 7, we see an op-amp configured as a pre-amp of gain 34. The input of this op-amp comes from J1, which is the input jack for the microphone. It is important to amplify the signal of the microphone since the incoming signal is extremely weak compared to what we need to drive the headphones. The gain here is about 31dB. R4 is simply used to provide a ground reference. R2 acts as a pull-up resistor for the microphone. C2 and C4 act as a high pass filter. This is important for amplifying the signal by such a large amount because the noise will be amplified as well.

Now, taking a look at Figure 8, which focuses on the op-amp configured as a phase inverter. R10 in this image corresponds to the top-right corner of Figure 7. Hence, the output of the pre-amp is routed through the phase inverter. The gain here is 1. Notice that a separate connection of the output from the pre-amp goes to the same switch as the output of the inverting op-amp. The switch, S1-a allows us to choose between either the original amplified signal or the inverted signal. This is useful for testing/debugging and also provides the additional feature of using the noise-canceling as bionic headphones to enhance hearing.

R14-a is a potentiometer with a max resistance of 100K Ohms. This essentially acts as a volume control on the incoming signals from the switch.

The two sub-components described are essentially all that is needed for the noise-canceling part of the circuit. All that is left is a circuit that can mix this signal in with the audio signal coming out of an mp3 player such that we can listen to music while simultaneously canceling surrounding noise. The final sub-circuit does precisely this.
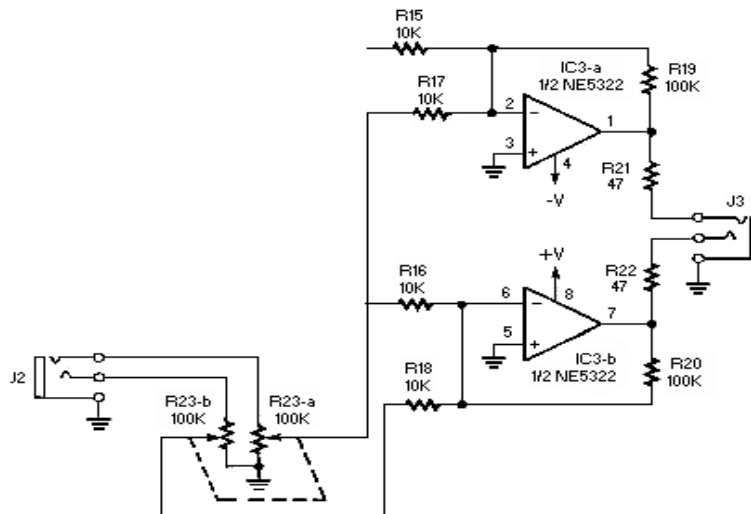
Figure 9

Once again, focusing our attention to the "top" part of the circuit, we notice that J2 feeds into IC3-a, which is configured as an inverting op-amp. The gain is given by R19/R15 for the microphone signal and R19/R17 for the audio signal, coming from J2. The potentiometer between J2 and the op-amp can be used to set the signal strength (or the volume) coming from the line in. The potentiometer is configured in such a way that it produces a nonlinear response. It acts as a voltage divider between the 100K variable resistor (R23-a) and the 10K resistor (R17). This has the effect of us hearing linear signal attenuation when in reality it is changing much more rapidly. This is important because human hearing works logarithmically. This is done to save on costs since linear taper potentiometers are less costly than audio tapered ones.

Understanding the circuit, the next challenge was to construct it from individual components on a board.

Build Results

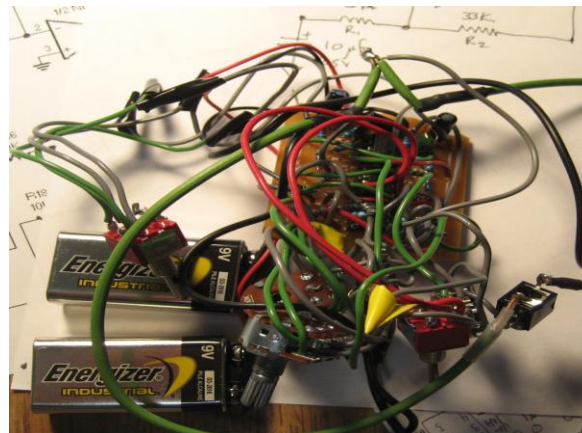A very nice schematic can easily turn into a quite messy circuit. See Figure 10.

Figure 10

This could have been made a bit cleaner, but the best improvement would have been to use a larger board and adopt some simple rules of thumb such as grouping specific wires and running others under the board. Still though, there was not enough time to create another prototype, and since it seemed to be working, it did not really matter how sloppy it looked on the inside so long as no debugging was needed.

Putting it into a nicely labeled box actually helps conceal this circuit as shown in Figure 11.

Figure 11

The headphones were created by taking an ordinary pair of headphones and using a hot glue gun to attach the microphones on the exterior of each ear. See Figure 12 and 13.

Figure 12

Figure 13

In order to save time, the microphone uses a cable from some old headphones from which I clipped off the ear buds. This also makes it much cleaner to work with the other cable than if one were to try to create one from individual wires and a plug.

Testing/Analysis

The first thing to do was to simply put on the headphones, flip the switch, and qualitatively assess what is happening to sounds at different frequencies. Doing this, I noticed that the "bionic" function of the headphones was working well. This is where the signal from the microphone is not inverted. However, the inverted signal seemed to have different behavior at different frequencies. It seemed to cancel lower frequencies, while it was difficult to assess what was going on at higher frequencies. Before any testing, I made sure that the circuit was working correctly by probing both the inverted and non-inverted signals with an oscilloscope to see if they were equal and opposite. This initial test showed promising results. See Figure 14.
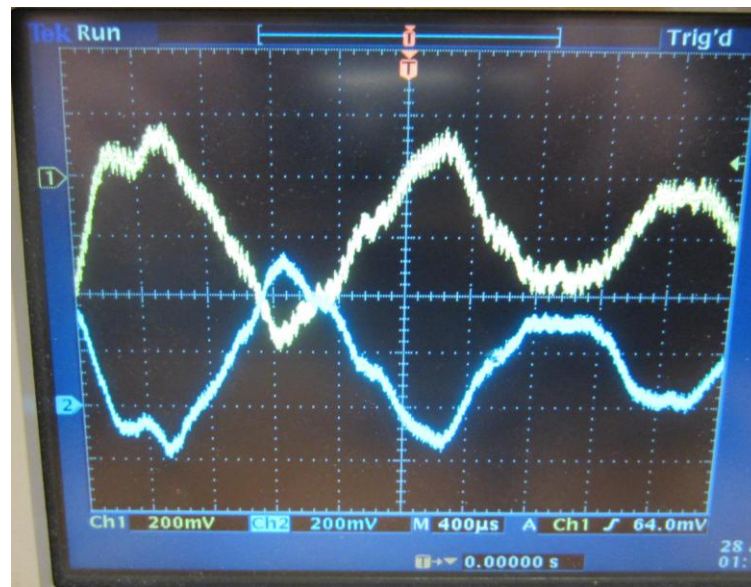

Figure 14

Next, we wanted to see how the volume of the region inside the headphones where cancellation should be taking place looks at various frequencies. For this, we generated constant white noise, and did a frequency scan using a spectrum analyzer. The following results are shown in Figure 15.
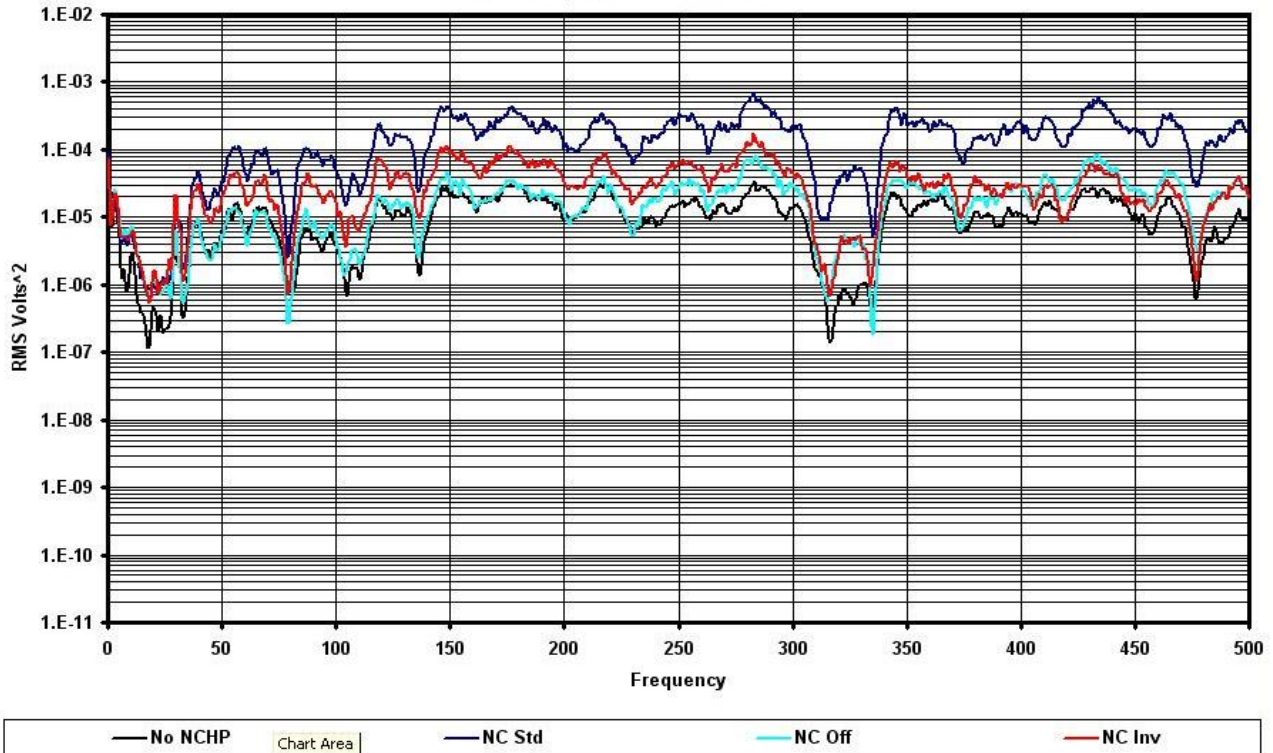
Figure 15

The graph shows loudness versus frequency. The spectrum analyzer scanned loudness in the frequency range of 0 to 500 Hz. In black, we see data for the microphone before the noise-canceling headphones covered it up. The next three sets of line plots show data with the noise-canceling headphones covering and with the microphone turned on with the non-inverted signal (NC Std), no signal at all (NC Off), and the inverted signal (NC Inv) coming through the speakers.

This experiment showed strange results since it suggests that not having the headphones over the microphone is actually quieter than having them over it, whether they are on or off. We would expect that even if the circuit is not working correctly, covering up the microphone with the headphones would reduce some of the sound. Ignoring this strange data and looking at the other plots, we see that the microphone detects least noise for the headphones being off completely.

Furthermore, turning the headphones on through the non-inverted part of the circuit produces the loudest atmosphere. On the other hand, the inverting mode is not as loud, but is louder than not having anything on. This suggests that something is wrong with the headphones. In terms of loudness, the cancellation feature is not working correctly since it seems to add noise, rather than cancel it at most frequencies that we scanned. One likely culprit to this problem is that although the circuit is working correctly, due to the finite space between the microphones and the headphone speakers, there is a time delay that is not accounted for, assuming the circuit processes the signal and drives the headphones instantaneously. We decided to investigate the time-delay to see if it would provide any clues to explain why the loudness results showed such poor performance.

The setup was to put a pressure microphone facing a loudspeaker that emitted a piezo horn pulse at 10 Hz. First, we compared the signals of the pressure microphone and one of the microphones on the headphones at the same location to how the signals should look when aligned. This is shown in Figure 16.
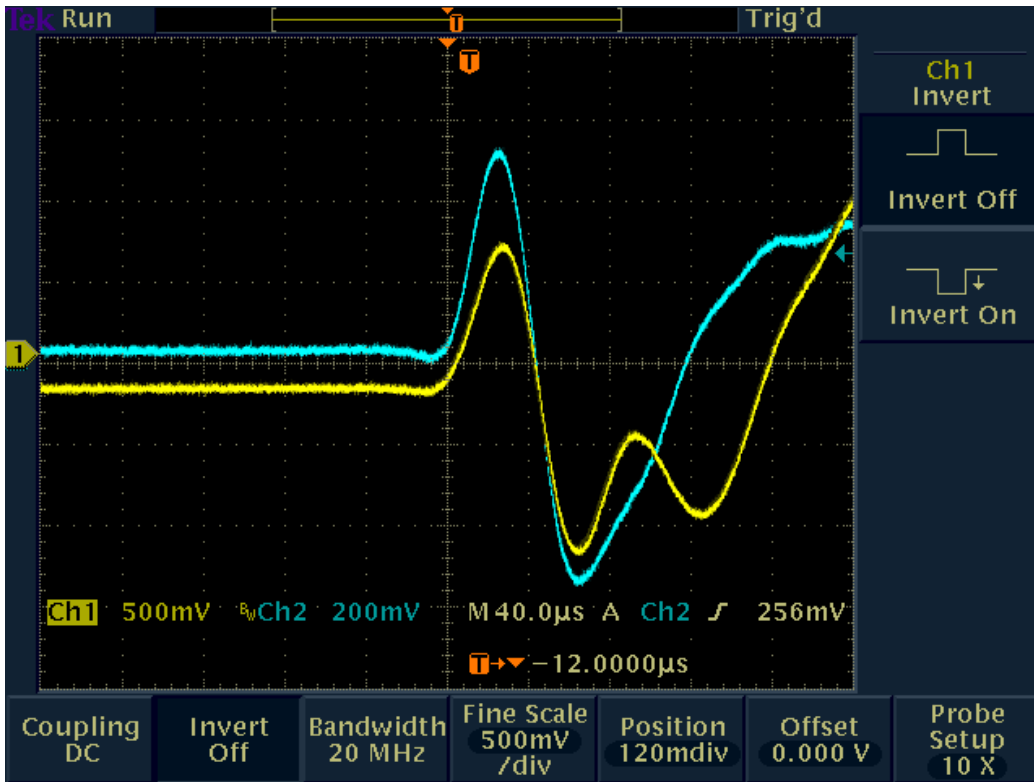

Figure 16

The signals are fairly close to each other. The discrepancy that we observe is likely due to the pulse dying down. We are only interested in the region where the two waveforms are roughly the same shape.

Next, we obtained oscilloscope data comparing the response of both microphones with the noise cancelling headphones clamped over the pressure microphone. We observed the results for the inverted signal, non-inverted signal, and no signal coming out of the headphone speakers. The results are shown in Figures 17-19.

Channel 1 shows the signal coming from the headphone-mounted microphone. Channel 2, from the pressure microphone inside the headphones.
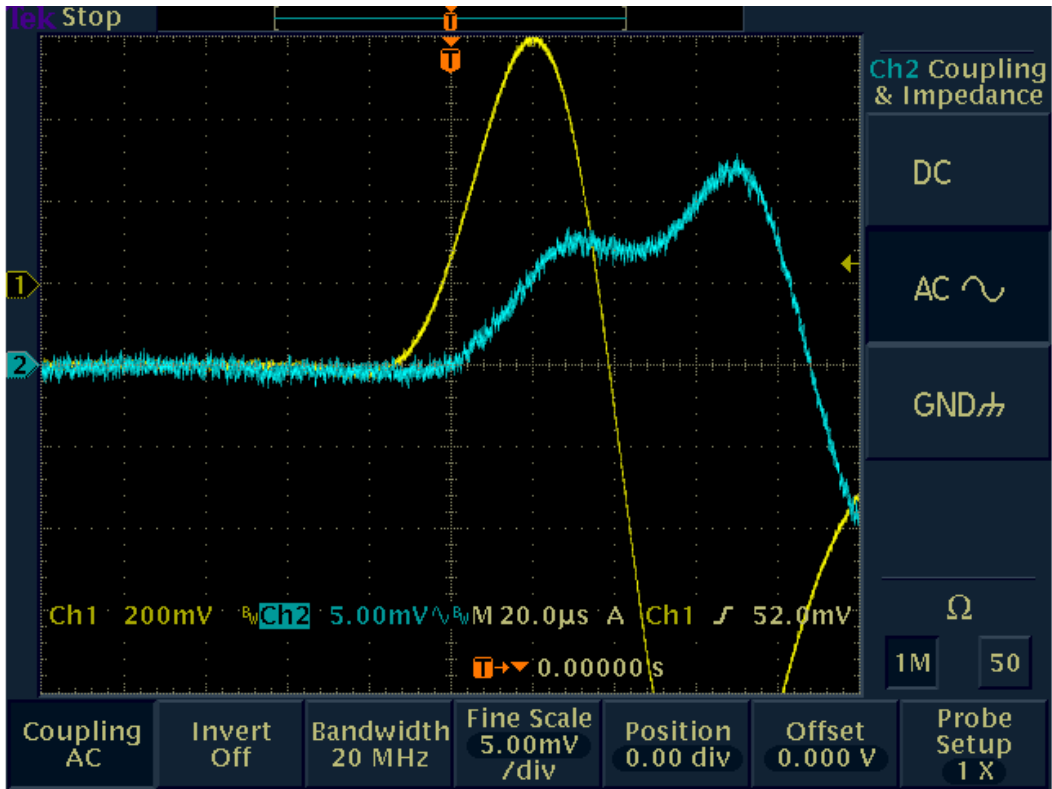
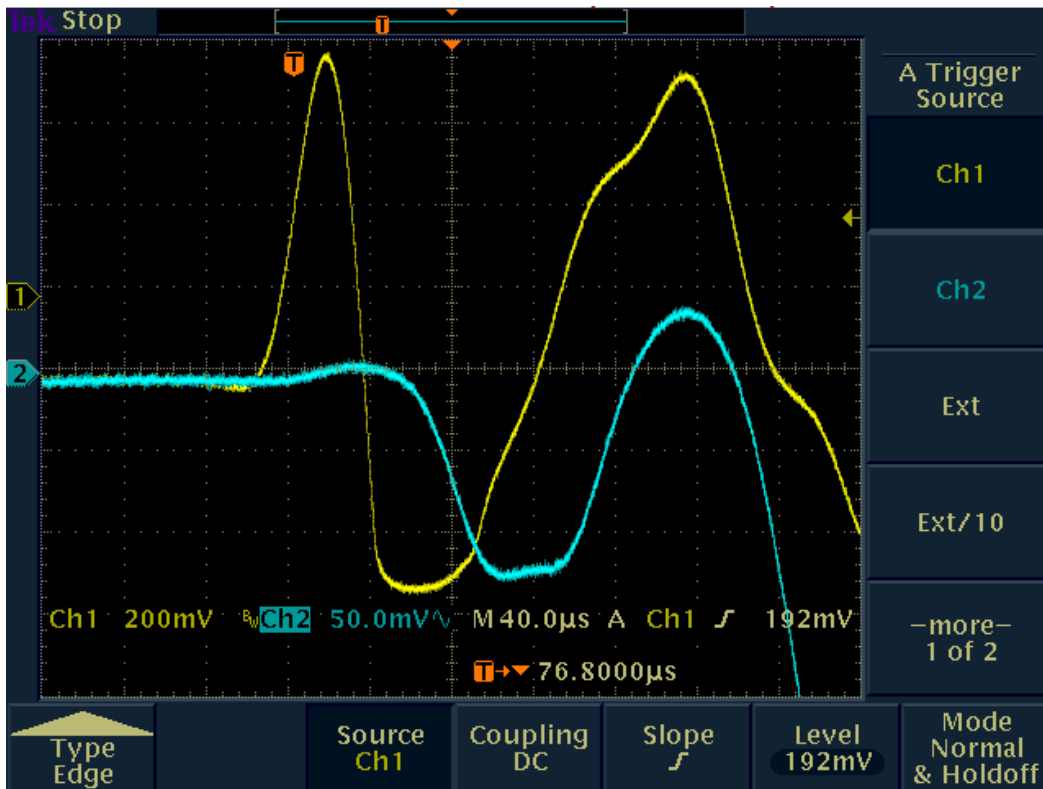**Figure 17 – Time-Delay Headphone Speaker Off (Δt 16 µs)**

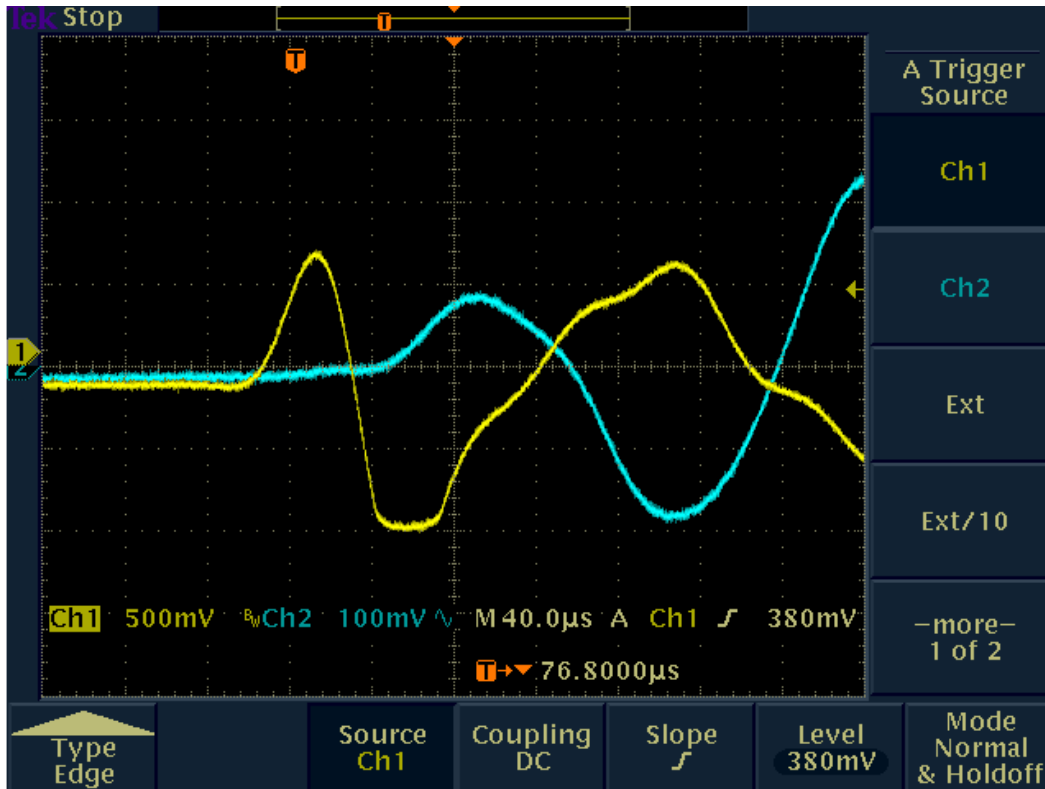**Figure 18 – Time-Delay for Inverted Signal (Δt 76 µs)**

Figure 19 - Time-Delay, Non-inverted Signal (Δt 72 μs)

The Δt values were obtained by printing the above figures out and counting ticks to approximate the time-difference between the rising edges of the signals. The non-inverted and inverted signal time-delays should be roughly the same. We expect that with more trials, the average of the two would be roughly equal.

The ultimate Δt that we are interested is the difference between the time-delay with the microphones on and off. Taking the inverted signal-delay and subtracting the delay with the headphone speakers off, we get 76 μs – 16 μs = 60 μs. Multiplying this by the speed of sound measured in our lab, approximately 345.8 m/s, we get that the microphone would have to move 2.075 cm toward the speaker to avoid the time-delay. Another way to alleviate this issue is by introducing circuitry that provides a time-delay of 60 μs before the signal from the microphone is outputted.

By inspecting the headphones, it becomes apparent that the microphone may be possible to move 2.075 cm if there are holes drilled into the headphones, something that I plan to look at in the future.

One circuit component that may help with this is a CCD Analog Bucket Brigade chip; however, the clock speed needed to achieve a delay of this magnitude is much too fast compared to that of standard available chips. The best solution is using a chip that allows one to provide a programmable audio delay. A combination moving the microphone and introducing a time-delay is probably the best solution.

Conclusion

Despite not being able to come out of the course with a pair of functional noise-cancelling headphones, I came out with much more than if I had simply bought a pair. I went from having virtually no soldering experience to being proficient at it. I learned how to use op-amps and how to create simple filters using resistors and capacitors. In general, I learned how to build a circuit on an empty board given a pile of individual resistors, capacitors, and integrated-circuit chips. The practical hands-on experience is a necessary skill for any kind of project since physical testing of the design should be the end-goal. However, the most important benefit to completing this project is the confidence I now have to tackle a seemingly complicated electronics project. Ultimately, my experiences in this course have inspired me to seek out and work through a more complicated project in the future.