Theremin

6.2040 Final Project Report

Ryn Moore Sarah Coston

May 2023

Contents

1 Introduction

Abstract

A theremin is a musical instrument that changes pitch and volume based on the position of the player's hands relative to its antennae. The human body naturally has a capacitance, and a theremin takes advantage of that. By changing the distance between one's hands and the antenna, the amount of capacitance changes, which changes the resonant frequency of the circuit. This change in frequency, when amplified, is perceived by human ears as a change in pitch. A theremin is comprised of two major components, which can be broken down into smaller components: the volume circuit and the pitch circuit. The volume circuit can be broken down into an antenna, an oscillator, and an amplifier. The pitch circuit can be broken down into a reference oscillator, a pitch oscillator, and a circuit that combines those two signals. Alternatively, it is sometimes possible to simply generate a low enough frequency with a single oscillator that it is already in the audible range. There is also circuitry that combines the pitch and volume information into a single waveform. Time permitting, we will also include a tuner circuit that will provide the musician with visual feedback about how in tune they are.

1.1 Background

The theremin was invented in 1919 by Leon Theremin. It came about as a byproduct of Soviet research into proximity sensors. In 1928, Leon Theremin moved to the United States, bringing his innovative musical instrument with him. He eventually patented the invention and granted RCA the rights to produce it commercially. Poor timing in it's commercial release meant that it never became as popular as RCA hoped, though Clara Rockmore did much to popularize the instrument through her skilled manipulation of it's sounds. Ultimately, the development of new, easier to play electronic instruments, such as the synthesizer, saw the theremin fall into disuse. An interesting note is that the theremin very directly contributed to the invention and proliferation of the synthesizer. Robert Moog, the inventor of the groundbreaking Moog synthesizer, began his foray into electronic music by building theremins and theremin kits for people to purchase and assemble themselves. He directly credits this experience with aiding him in eventually creating the Moog synthesizer.

The 1990's saw a resurgence of interest in theremins from professional musicians. The theremin is now featured in orchestral music, pop music, film music, and even on Broadway. It is perhaps best known for being used to produce eerie, creepy sound effects for movies such as *The Day the Earth Stood Still*, *Spellbound*, and *Monster House*. However, even with this increase in popularity, it remains a niche instrument. It is arguably more popular among electronics hobbyists for it's interesting application of oscillators and amplifiers than among professional musicians.

1.2 Motivation

We would like to do our part to continue the usage of this unique and haunting instrument. In addition to producing sounds that other instruments struggle to reproduce, the playing of a theremin is a beautiful performance art. The way the musician interacts with the theremin is almost like a dance, and the shape and rotational position of the hand seems to have just as significant impacts as the distance from the antenna. Of course the question remains: what is the point of building on from the ground up if they are sold commercially? We believe that building one from scratch will be an extremely educational experience that will increase our understanding of analog electronics in general and reinforce the concepts we've been learning in class. Furthermore, since our report will be published on the class website, it will also be accessible to students who attempt similar projects in the future. Preliminary research into how we wanted to go about this project yielded surprisingly few results with detailed information about other people's implementations. Finally, the ground-up approach will provide us with thorough, intrinsic understandings of what factors impact the tone produced by the theremin. This will be valuable experience for future work in which we would like to implement some sort of visual feedback to the musician about their performance. We would like to implement the beginnings of such a system into this project, but we recognize the stringent time constraints we are under, and this may necessitate leaving that idea for a future iteration.

2 Overview

2.1 Block Diagram

The following block diagram depicts our original design, which is also the most common topology of a theremin that we found throughout our research.

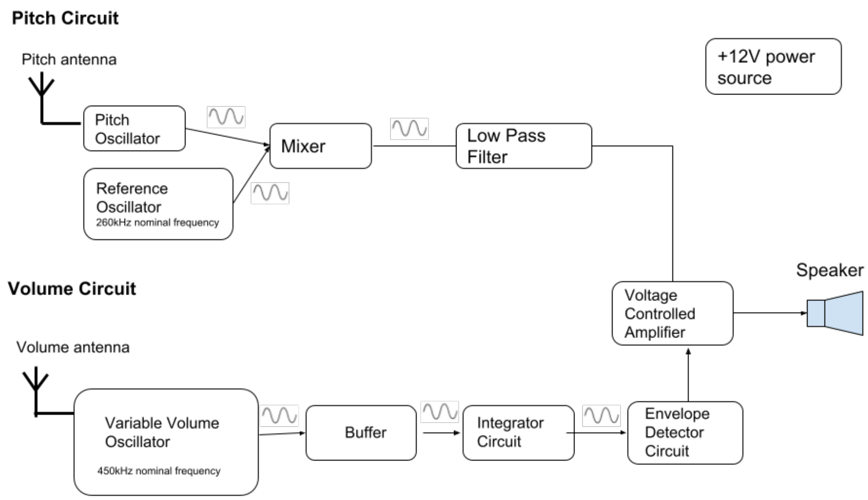

Fig.1 Proposed Theremin Block Diagram

The theremin circuit can be broken down into three main components: the pitch circuit, the volume circuit, and the amplification stage. As shown below, the pitch circuit consists of an antenna and two oscillators: a reference (the frequency of this oscillator varies widely with the particular implementation, and our research has yielded results as low as 125kHz and as high as 1MHz) and a variable pitch oscillator (typically with a bandwidth of 5kHz). The variable pitch oscillator changes frequencies as the user moves their hand closer and farther away from the antenna. These signals are combined in a mixing stage which leverages the properties of oscillating functions to produce an output frequency within the audible range (typically cited as being between 20Hz and 20kHz). The volume circuit, on the other hand, is a single oscillator, typically around 450kHz. The frequency variation of this circuit must be transformed into an amplitude variation, which can be accomplished through the use of an integrator and an envelope detector. Both of these outputs are sent to a complex amplifier that combines the frequency of the pitch circuit with the amplitude of the volume circuit. This amplifier produces the signal that is ultimately fed into the speaker.

For this project, given the time constraints and difficulties faced in getting the Colpitts oscillator topology to function, we decided to simplify everything down and use 555 timers to generate our oscillations instead.

This block diagram is a more accurate depiction of what our final product looked like.

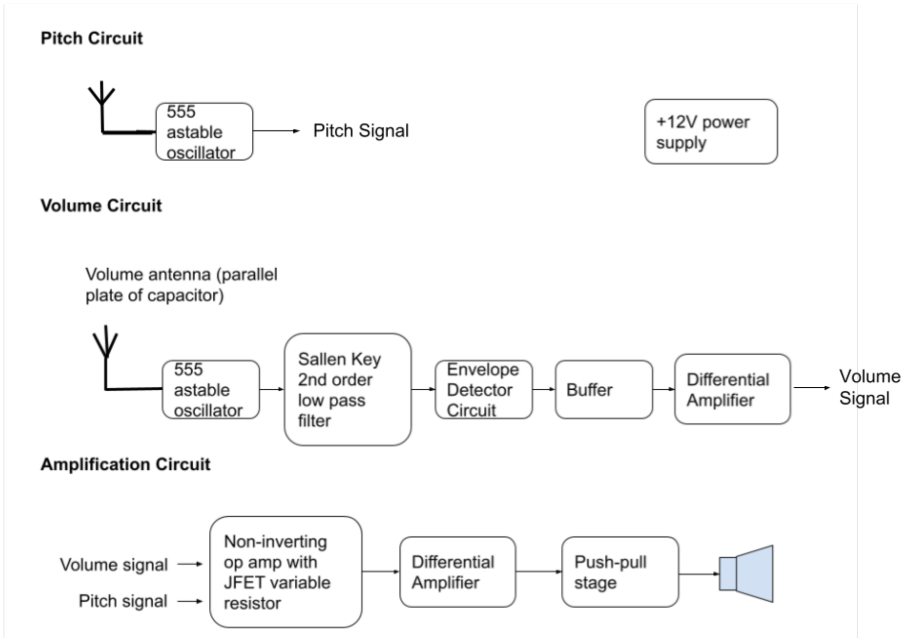

Fig.2 Final Theremin Block Diagram

We were able to get the 555 timers operating at a low enough frequency for the pitch circuit that it completely removed the necessity of a reference oscillator. The timer itself was operating within the audible range, so there was no need to subtract off a reference frequency. This drastically simplified the circuit by eliminating one of the oscillators, the need for a mixer, and the low pass filter. The volume circuit remained relatively unchanged. One big difference between the two iterations were that the output of the 555 timers were square waves instead of sine waves. This changes the timbre of the output, but in the scope of this project, this is a more than worthwhile tradeoff. The timbre of a noise is a description of what qualities it has besides the frequency at which it is oscillating, and is what allows humans to e.g. differentiate between different instruments playing the same note. However, the timbre difference between an electronically generated sine wave and an electronically generated square wave is so small, that most non-musicians would not even be able to tell the difference.

Further justification of the use of the 555 timer and explanation of its drawbacks can be found in section 3.3.

2.2 Detailed Schematic

This is what our final schematic looks like with each stage connected together.

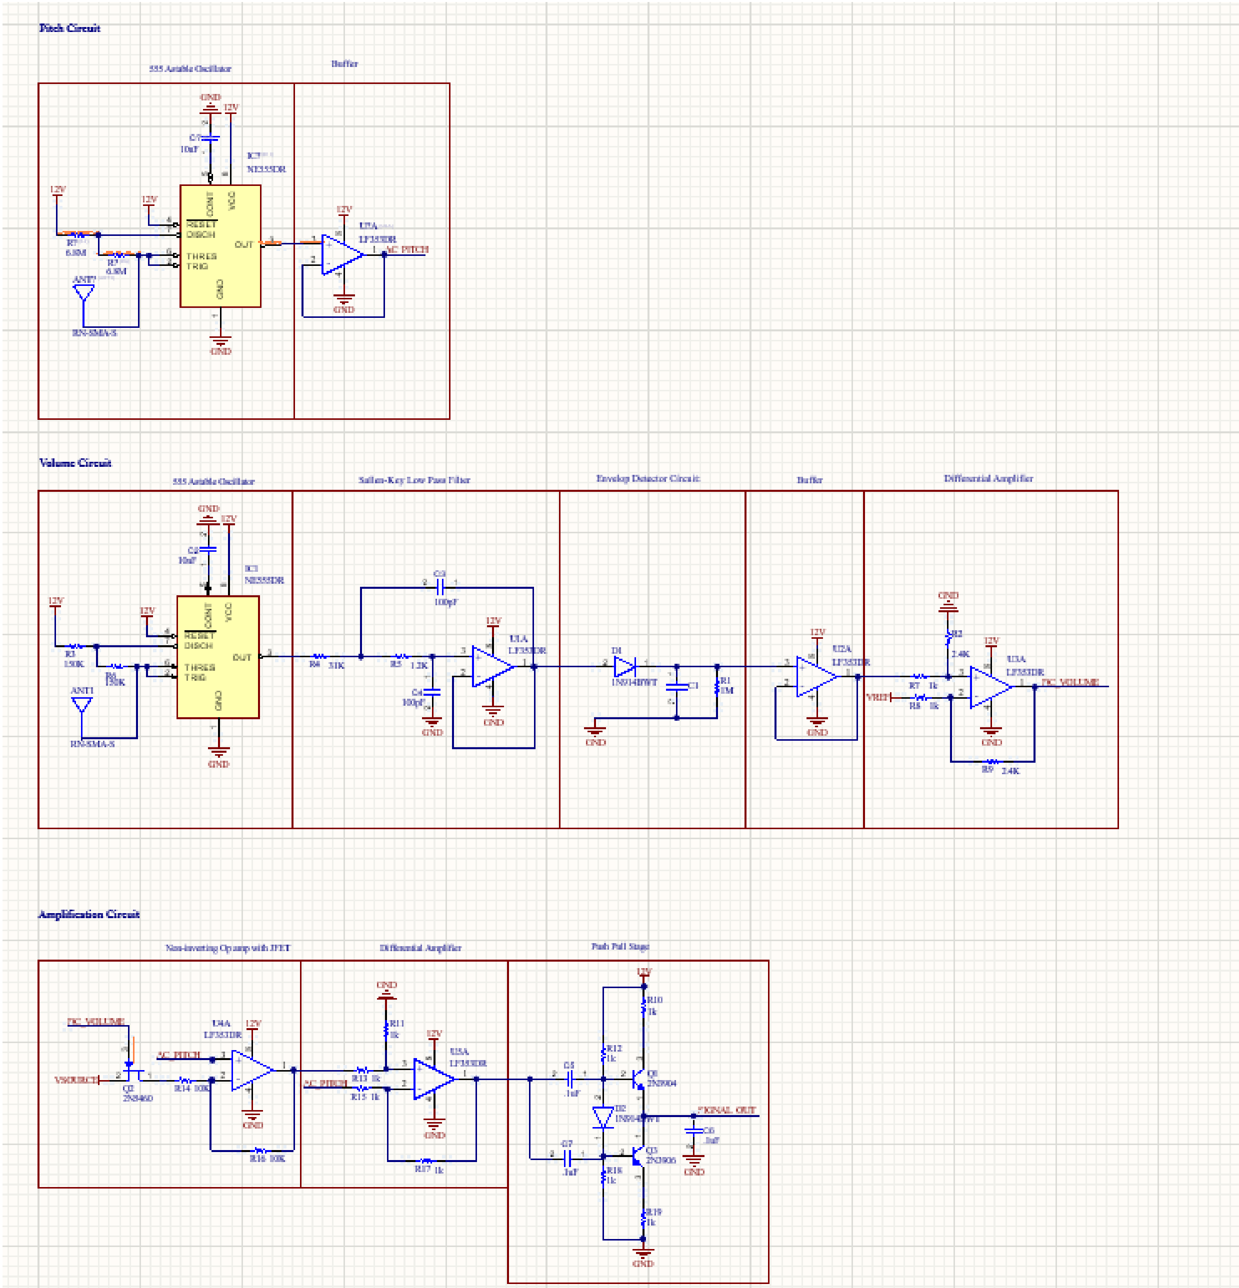

Fig.3 Final Theremin Circuit

3 Pitch Circuit (Ryn)

3.1 Original Design

Originally, I had planned to use Colpitts oscillators for both the reference oscillator and the variable oscillator of the pitch circuit (and Sarah planned to use the same type of oscillator for the volume circuit, which is covered in section 4). It seems that there are many theremins that use this topology with great success.

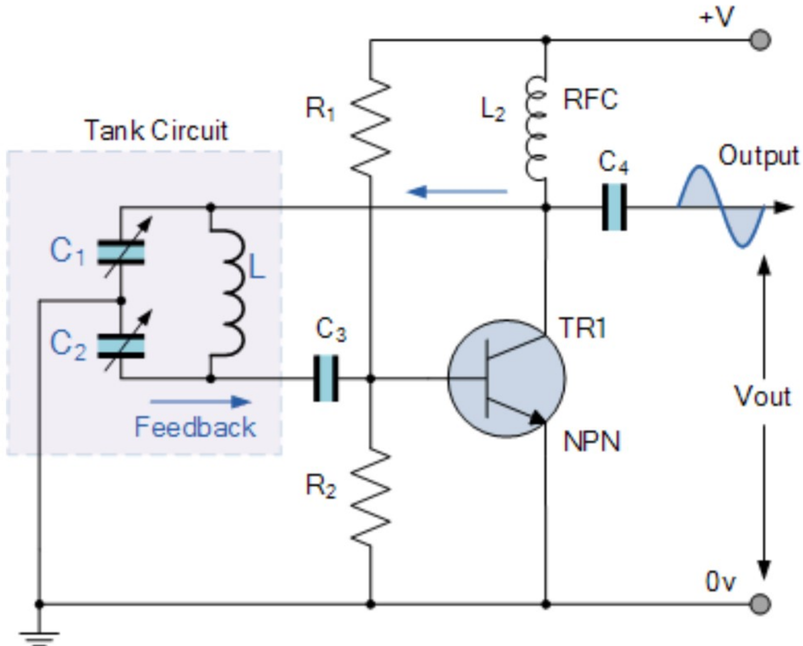

Fig. 4 Colpitts Oscillator

These oscillators take a DC signal and convert it into AC at a specific frequency. This is accomplished by having an LC tank circuit with two capacitors in series with the circuit grounded in between the two capacitors creating a voltage divider. Furthermore, the output of the tank circuit is fed into the base of a BJT, which has it's collector connected to the input. This creates feedback in the system, which produces a stable oscillation rather than one that decays. In the case of a theremin, the oscillators connected to the antennae have the antenna in place of one of the capacitors in the LC tank.

The complication to the pitch circuit is that a Colpitts oscillator is incapable of generating very low frequencies without significant amplitude loss. This means it is necessary to build two oscillators: one operating at a known frequency and one using the antenna as a variable capacitor that produces a frequency output with a minimum at the reference frequency of the other oscillator and a maximum at most 20kHz above that frequency. A mixer that multiplies two input signals can then be used to generate the audible range output frequency because of the follow trig identity:

$cos(\omega_1 t) * cos(\omega_2 t) = frac12 * (cos((\omega_1 + \omega_2) * t) + cos((\omega_1 - \omega_2) * t))$

It is then simple to filter out the signal comprised of the sum of the two frequencies leaving only the signal that is the difference between the two frequencies.

Unfortunately, I was unable to get any Colpitts oscillator to oscillate at all. Despite my best efforts, all I got was a constant output voltage that was the same as the input voltage. Sarah ran into the same trouble with the volume circuit. With time pressure growing, we both decided to abandon this stable oscillator in favor of a simpler design.

3.2 Modified Design

Due to our struggles with the Colpitts oscillator, I turned instead to the trusty 555 timer. 555 timer's are named such because of their characteristic three 5kΩ resistors. They can function in either a monostable or an astable configuration based on the resistors and capacitors connected to them. I used the astable configuration to produce a square wave output, the frequency of which varies with the capacitance produced by the antenna. This is the general layout of an astable 555 oscillator:

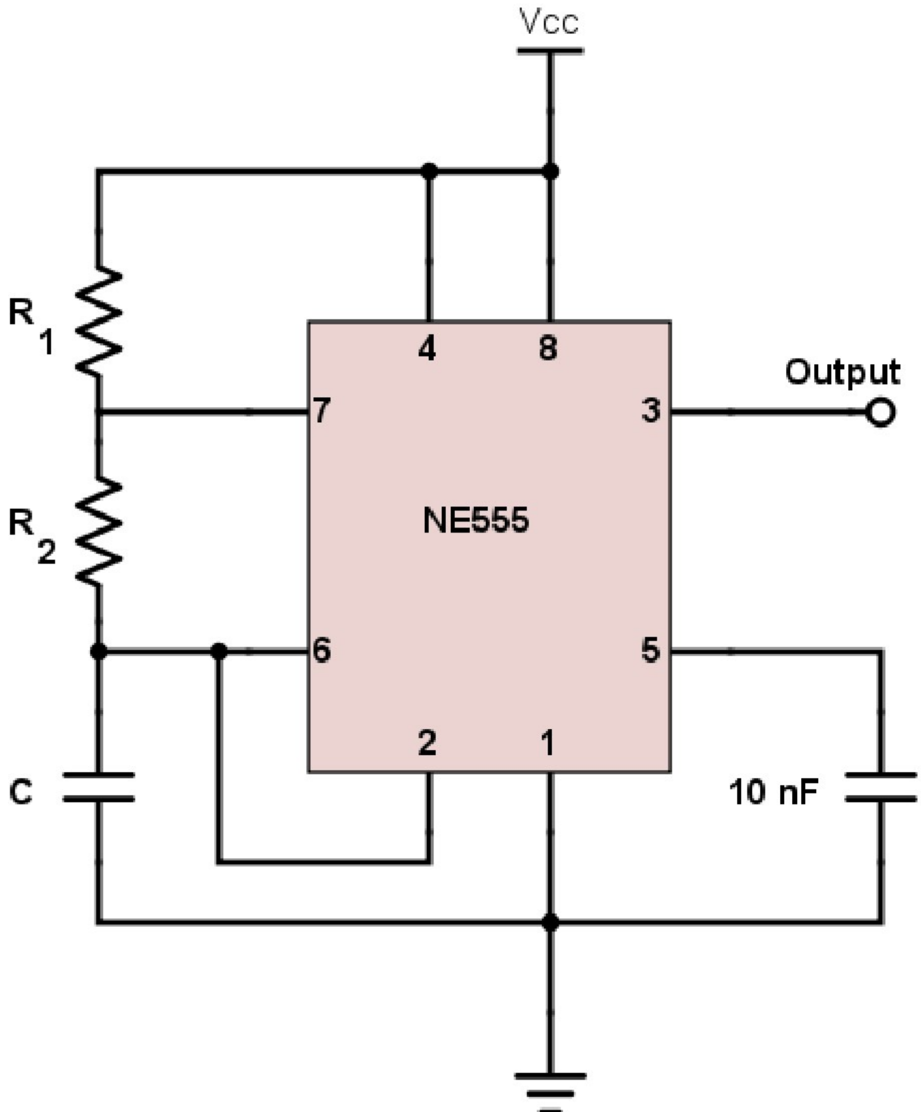

Fig. 5 Astable Oscillator Configuration

The following equations govern their operation:

$$f = \frac{1.44}{(R1 + 2 * R2) * C}$$

$$T1 = 0.694 * (R1 + R2) * C$$

$$T = \frac{1}{f}$$

$$DutyCycle = \frac{T1}{T}$$

With the correct ratio of resistors, even with the tiny capacitance from the antenna, I was able to generate an extremely low frequency. This led to the realization that the reference oscillator had become redundant. If the 555 timer could generate an audible frequency between 20Hz and 20kHz, having it operate at a higher frequency and then subtracting some of that frequency off added needless complexity.

3.3 Design Choices

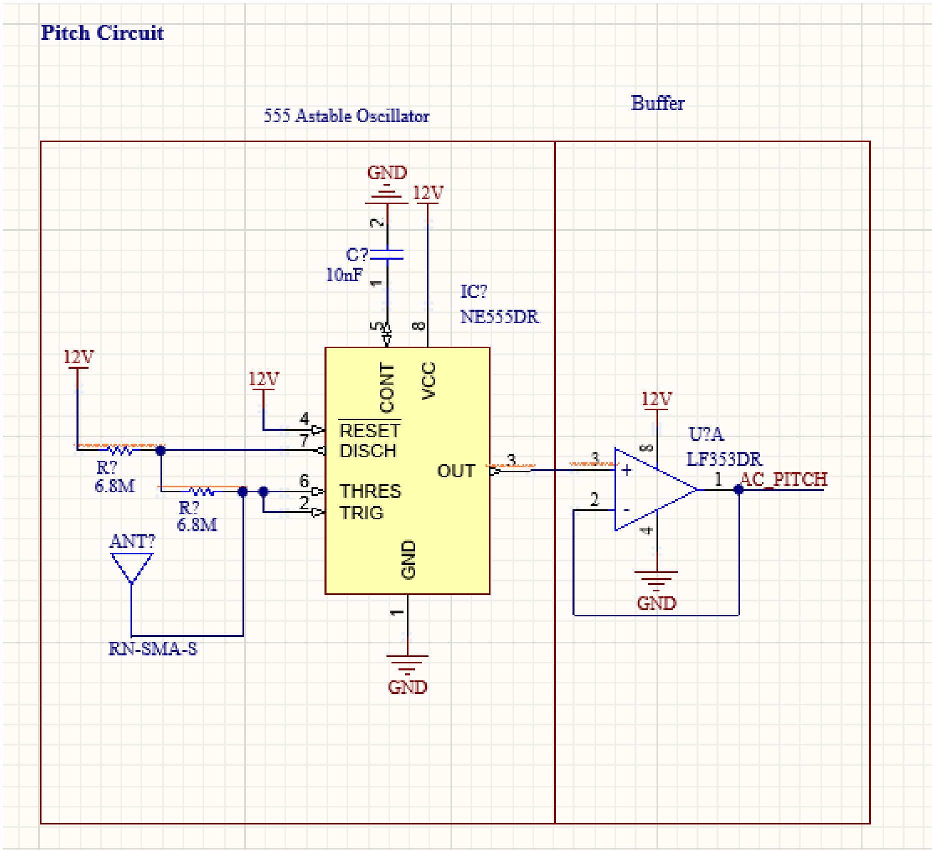

Fig. 6 Pitch Schematic

Whenever one is designing electrical system, it becomes a game of tradeoffs. Every decision that is made to optimize one aspect of the system necessarily acts to the detriment of other aspects of the system. In this section of the overall circuit, the biggest tradeoffs were between the center frequency and the frequency range. In order to keep the center frequency in the audible range (20Hz-20kHz), the resistors in the astable oscillator have to be massive. The ones ultimately used were each 6.8MΩ. I was further constrained by wanting to produce a sound that isn't earsplitting, which means it's best for the circuit to actually max out at more like 5kHz. With such a small capacitance (both research and empirical observation indicated it to be about 1-10pF), resistors that large aren't going to allow for much of a frequency swing. I wanted to make sure that the change in pitch was clearly audible. Thus, I had to find a middle ground between the resistors being big enough to pull the center frequency down to tolerable levels, but small enough to allow an audible change in frequency.

3.4 Buffer Stage

Over the course of constructing our circuits, Sarah concluded that, in order for the amplification circuit to work properly, the pitch signal needed to have a voltage offset of somewhere between 1.5 and 2.5 volts and the peak-to-peak voltage of about 6V. I attempted to add an offset by modulating down the 12V source, but this interfered with the signal. Even using an op-amp adder circuit, trying to add a 2.5V offset seemed to simply fix the voltage output of the circuit at 2.5V instead of offsetting the output oscillation. Thus, it seemed prudent to buffer the oscillating voltage.

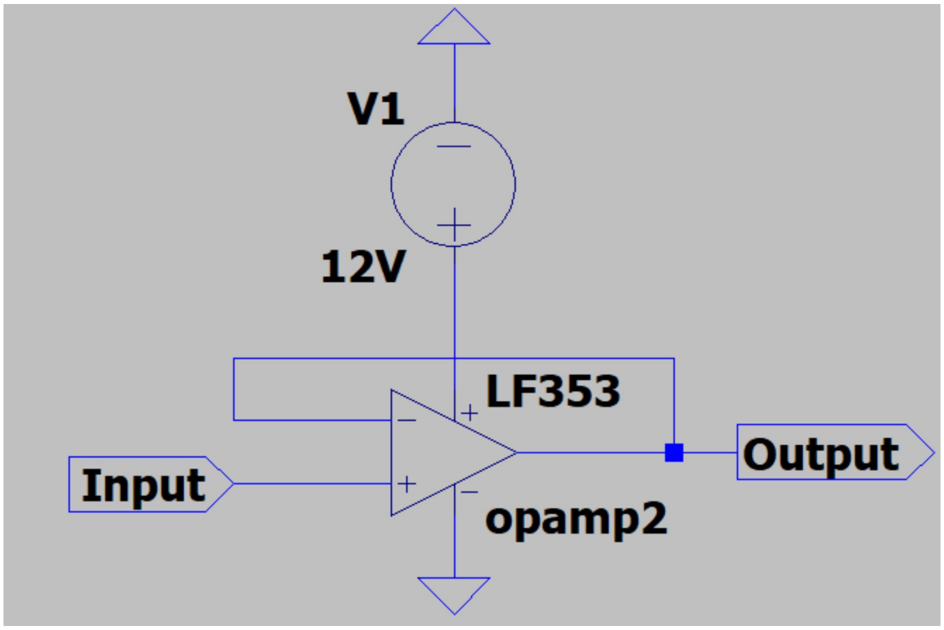

Fig. 7 Buffer Schematic (Not Used in Final Design)

I used a standard unity gain non-inverting amplifier configuration using an LF353. Surprisingly, this seemed to produce a constant output at 12V. The input signal had a duty cycle of about 65%, but the output had a duty cycle of 99%, essentially not oscillating at all. I then realized that the peak of the input was at about 13.2V despite the source voltage being only 12V. I never really figured out why that was happening, but I concluded that this was causing the op amp buffer to hit the rail for much of the cycle and thus interfering with the duty cycle. Putting a resistor between the oscillator and the buffer conveniently solved both the duty cycle problem and modulated the amplitude down to the desire 6V peak-to-peak value.

Ultimately, when I connected the pitch oscillator to the amplifier through the buffer, the amplifier had a much higher impedance meaning that there wasn't a high enough voltage going into the amplifier to produce a reasonable signal. That meant I had to remove the resistor between the oscillator and the buffer, reintroducing the railing issue with the buffer. When I plugged the oscillator directly into the amplifier circuit, it did seem to have the requisite offset without any adjustments, so the buffer circuit was deemed unnecessary and removed.

3.5 Results

The final result of this circuit was successful frequency modulation based on proximity to the antenna. However, the correlation between these two parameters was highly non-linear. It's not until the hand is within about three inches of the antenna that there is any response from the frequency at all, and about one inch to one sixteenth of an inch is where we get the largest, most perceptible change. I was able to get the frequency to vary between approximately 1.8kHz and 2.4kHz by the time the final video was filmed. For those unfamiliar with music, this ranges between about an A6 and a $D^{\#}7$.

Fig. 8 Pitch circuit signal when the musician's hand is far away
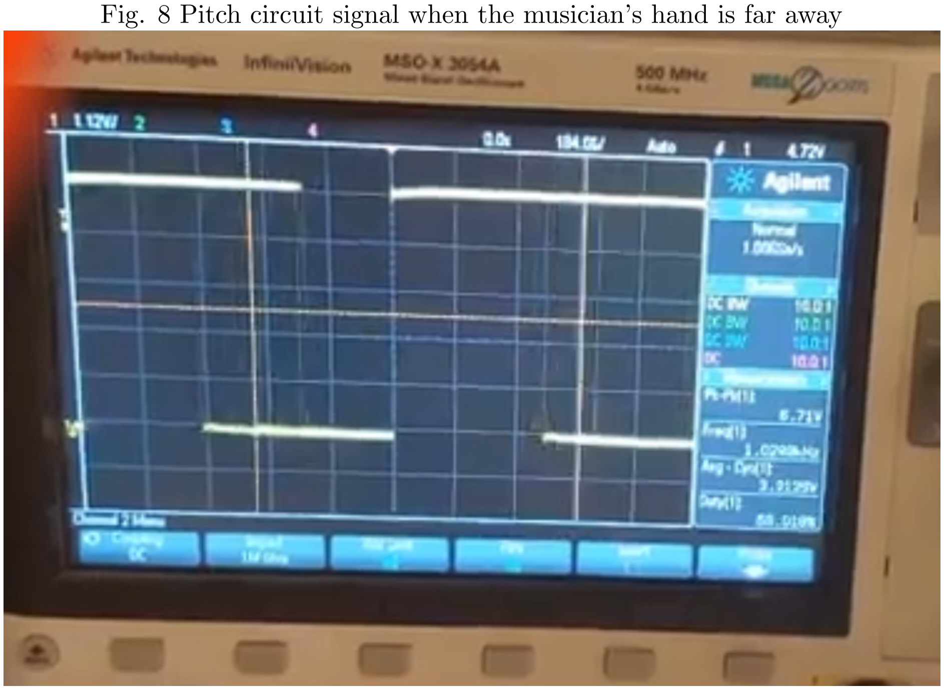

Fig. 9 Pitch circuit signal when the musician's hand is very close to the antenna

4 Volume Circuit (Sarah)

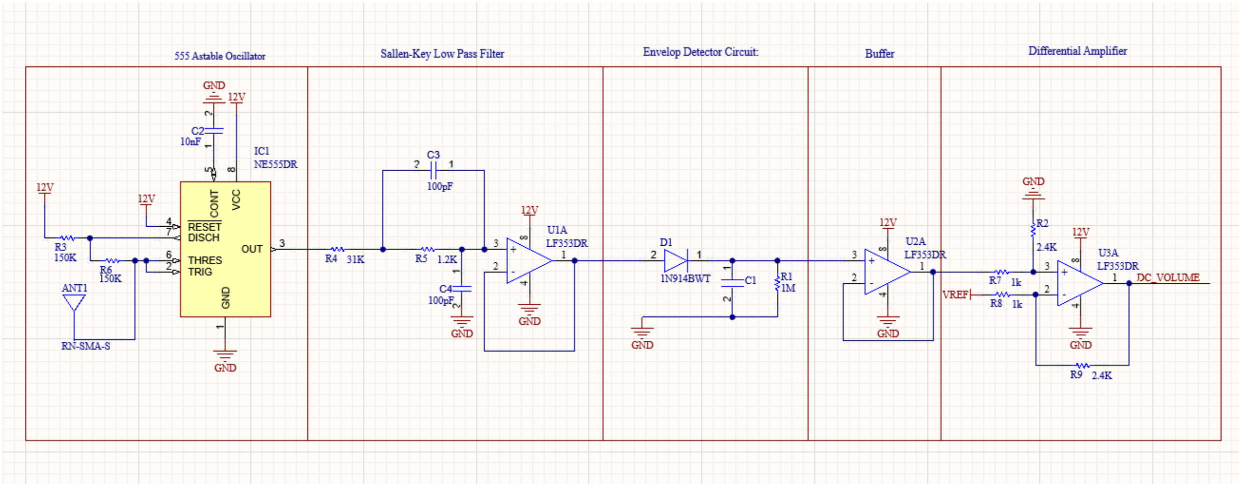

Fig.10 Volume Schematic

The volume circuit consisted of a 555 astable oscillator, a Sallen-Key second order low pass filter, an envelope detector, a buffer, and a differential amplifier.

555 Astable Oscillator

Current Theremin designs typically use a variable LC oscillator, as mentioned in the previous section, however I was unable to get the Colpitts LC oscillator to work. However I realized that the volume circuit did not specifically need a sinusoidal waveform in order to work, given that it was ultimately outputting a DC signal. As such, I determined that a 555 timer generating a square wave would be appropriate for my needs.

The 555 timer, as shown in the first block of Fig. 10, produced a variable frequency square wave that was between 70kHz and 200kHz with a duty cycle of around 70 percent and an amplitude of 12V. In my initial design, I thought to use a higher frequency (around 450kHz), given that a small change in capacitance would result in a larger change in frequency for an LC oscillator, but with the 555 I was able to produce a wider frequency range at a relatively low frequency, which I obtained empirically and using the following formulas to calculate the frequency and duty cycle.

$$f = \frac{1.44}{(R1 + 2 * R2) * C}$$

$$T1 = 0.694 * (R1 + R2) * C$$

$$T = \frac{1}{f}$$

$$DutyCycle = \frac{T1}{T}$$

We estimated that the capacitance of the antenna and the user's hand would be around 10pF, so I used a variable capacitor of 0.8pF to 20pF to test the circuit, before integrating the antenna. I used two small copper tubes to act as one half of the parallel plate. I tried several tube shapes in addition to a flat copper sheet, but found these tubes to generate the largest difference in capacitance. The capacitance was at a maximum when the user's hand is near the antenna, so the frequency was at a minimum, and vice versa.

Sallen and Key Low Pass Filter

This signal then went to a second order low pass filter which converted the variable frequency signal into a variable amplitude signal as shown in the second block of Fig. 10. The idea was that when the user's hand was close to the antenna, a lower frequency would be generated, resulting in a lower volume, and the opposite for when the user's hand was far away. The Sallen key filter was used due to its -40dB/decade slope which meant a small frequency change resulted in a larger amplitude change than that which could be achieved with a first order low pass filter, which typically has a -20dB/decade slope. I initially used a first order filter, but after simulating in LTSpice, I realized there was not a large enough change in amplitude, so I used a second order filter instead. The filter was designed to have a cutoff frequency low enough such that the varying frequency from the 555 timer would be in the most linear region of the bode plot. The cutoff frequency was calculated with the following formula:

$$f = \frac{1}{2\pi LC}$$

Envelope Detector, Buffer, and Differential Amplifier

This signal was then sent to an envelope detector which converted the AC signal into a DC signal, shown in the third block of Fig. 10. The diode in the envelope detector acted as a half wave rectifier, only passing positive voltage, while the capacitor on the output smoothed out the ripples to create a constant DC signal. This voltage was between around 5.85V and 8.9V. This signal was then fed to the gate of a JFET, which acted as a variable resistor used to control the amplitude of the pitch signal. However, I realized that the JFET gate needed between 1V and 6V. Therefore, a differential amplifier was used to step down the DC signal from the volume circuit down to the required range. The inverting terminal of the op amp was connected to a constant voltage which was tuned such that at the minimum capacitance, or maximum frequency, the output of the amplifier was 0. The output voltage was calculated using the following formula, assuming R1=R2 (resistors connected between input signals and inverting/non inverting terminals), R3=R4 (feedback and resistor between non inverting terminal and ground):

$$Vout = \frac{R3}{R1} * (V2 - V1)$$

During testing, I realized that the differential amplifier loaded the envelope detector so a buffer was added at the output of the envelope detector. The buffer was an op amp with negative feedback, shown in the fourth block of Fig. 10.

5 Amplification (Sarah)

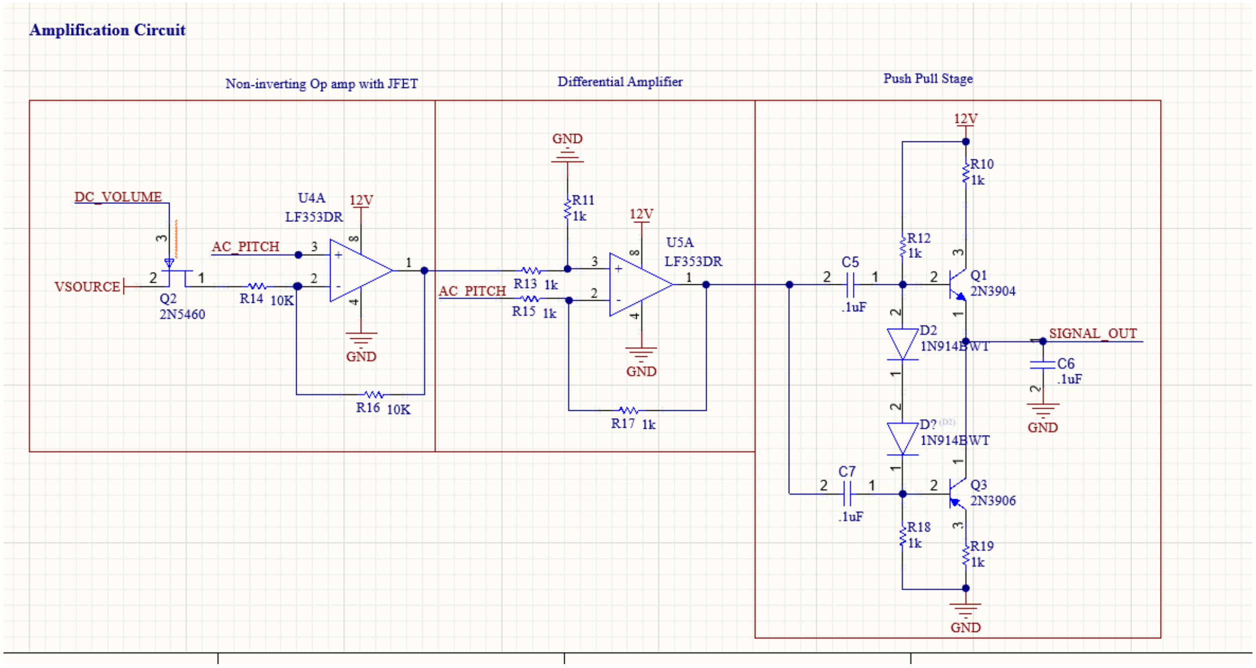

Fig.11 Amplification Circuit Schematic

The amplification circuit consisted of an inverting op amp with JFET, differential amplifier, and push pull stage.

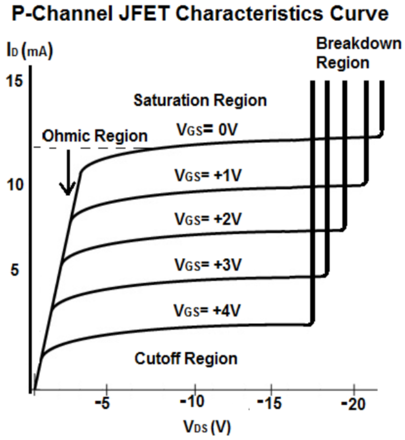

Fig.12 JFET characteristic curve

The p channel JFET operated in the ohmic linear region shown in Fig. 13, acting as a variable resistor, which changed the gain of the circuit, thereby modulating the amplitude of the pitch signal.

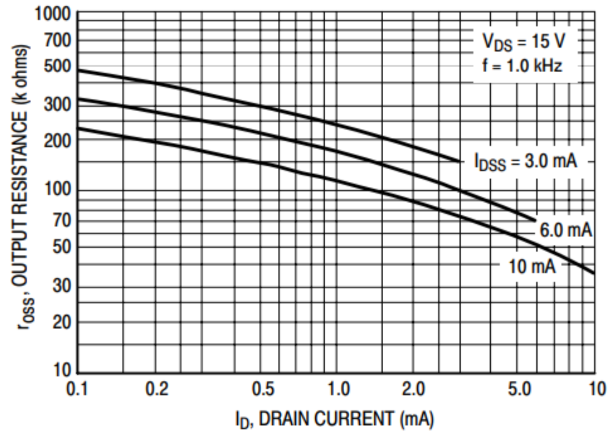

Fig.13 JFET 2N5460 drain current vs. gate source voltage

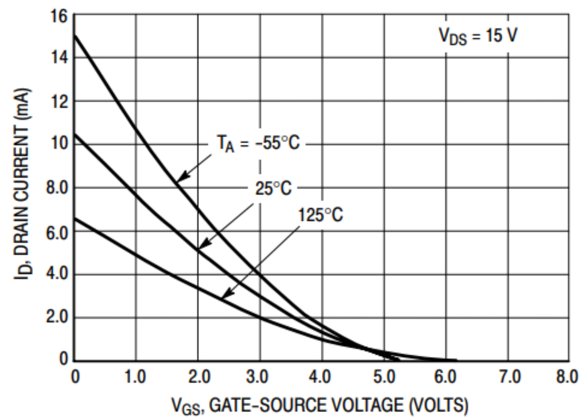

Fig.14 JFET 2N5460 output resistance vs. drain current

A p-channel JFET was chosen so that the gate voltage would be positive rather than negative. As can be seen from Fig 12 and Fig 13, when 0V was applied to the gate of the JFET, the current was at a maximum of 10mA, therefore the resistance was 40kohms. When 6V was applied to the gate, the current was 0mA, and the resistance was 250kohms. After testing in LTSpice and on a breadboard, a 10kohm resistor was added in series with the JFET in order to improve stability of the output waveform. After testing on a breadboard with the function generator, I determined that a square wave with an amplitude of around 3V and an offset of around 1.5V worked best to generate a clean output, so I requested that the output from the volume circuit meet those requirements. The circuit still worked outside of those measurements, but I found that a greater range of volume modulation could be achieved with those values.

In order to ensure that the volume was 0 whenever the user's hand wasn't near the antenna, and that the volume was at a maximum when the user's hand is close to the antenna, an additional differential op amp with unity gain was used. The two inputs were the original signal from the pitch circuit and the amplified version. Thus the output was the difference between the two. So when the 6V was applied to the JFET gate, the input and output signals of the differential op amp were equal, so the output to the speaker was 0, so no sound could be heard. However, when 0V was applied to the JFET, then the amplified version was at a maximum, so the output sent to the speaker was also at a maximum. Lastly, an output push pull stage was used to output a higher current for a speaker/buzzer.

6 Results (Sarah)

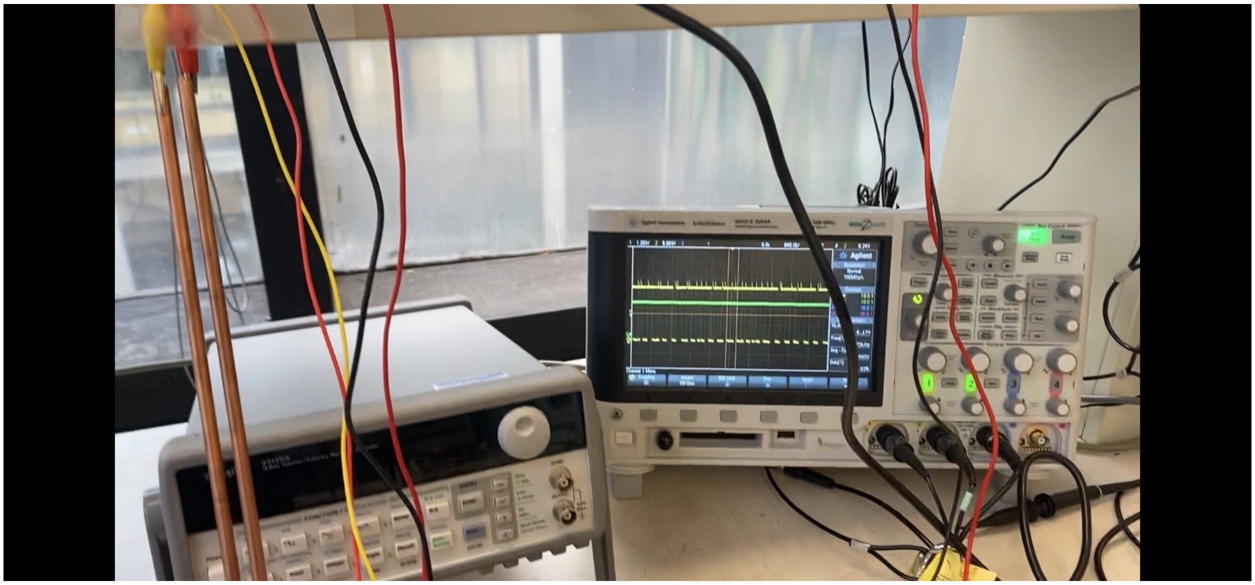

Fig.15 No volume when hand far away

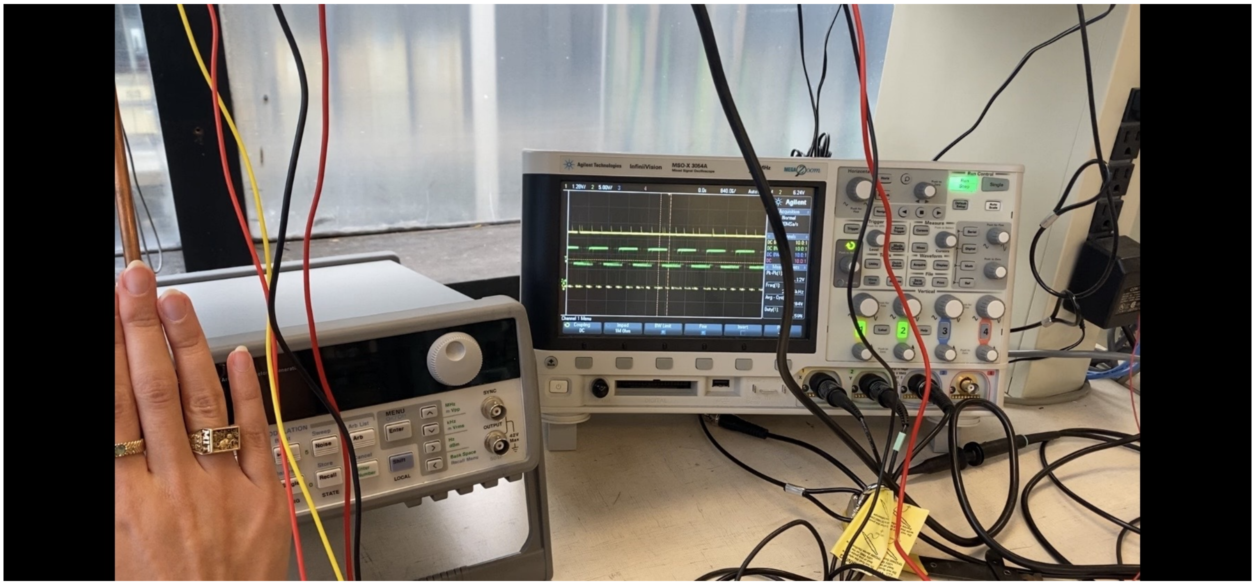

Fig.16 Maximum volume when hand close

I was successfully able to modulate the sound of a square wave as the user's hand moved closer and farther from the antenna. As shown by the green channel in the Fig. 14 and 15 above, when the user's hand was close to the antenna, the amplitude of the pitch signal was around 2.5V, and when the user's hand was far away from the antenna, there was no sound (0V).

7 Overall Results

7.1 Final Output

This link leads to a recording of the final output of our theremin.

This is far from the sound produced by a professional theremin, but this is not particularly surprising. The most obvious contributor to this is the speaker we were using, which was just as simple buzzer. However, the sound also wasn't very clear, indicating the signal had a substantial amount of noise in it. We could see on the oscilloscope output that the phase seemed to be oscillating about 0 rather than staying a stable place. We believe this is at least partially due to the instability of the capacitance from the pitch antenna because it seemed to improve when we reduced how large of an effect the antenna capacitor was having on the overall circuit by putting another capacitor in series with it. Of course, this also greatly reduced the range across which the frequency could vary, so it was deemed not worth it to leave in, but it was a good proof of concept for potential improvements in the future.

As was mentioned before, the final frequency output ranged from 1.8kHz (an A6) to 2.4kHz (a $D^{\#}7$). The volume circuit produced an increase in voltage up to 2V, resulting in a gain of 20 (if we assume Vinitial is 0.1V). This equates to a 26 dB audio gain.

7.2 Future Modifications

As with almost any project, we are not totally satisfied with our results, and can think of many modifications that, given more time, we would have chosen to implement. Perhaps the most compelling of these is the oscillator topology. We tried repeatedly to get the Colpitts oscillators working, and they worked great in simulations, but when we hooked them up in real life, we never got any sort of oscillation at all. It was just a constant output voltage. However, it is clear to us that the 555 timers are not the ideal oscillator to be using for the pitch circuit. Any oscillator that uses an LC tank, including the Colpitts, will have a more stable oscillation. This stability could make a big difference in the quality of signal we output, possibly helping considerably to alleviate the phase jittering. We also expect that we would have finer control over the frequency, and a more predictable output. Despite 555 timers theoretically being governed by well-defined equations as were listed earlier, working with them sometimes had unexpected outcomes that didn't follow those equations and for which we have not been able to account.

Furthermore, the use of an oscillator with an LC tank would mean that our output would be a sine wave, rather than a square wave. We had mentioned that for this project, the difference between a square wave and a sine wave was negligible, which is true, but as the circuit is refined and the output gets cleaner in general, the different timbre would become more obvious.

In all likelihood, returning to the Colpitts topology, or any LC tank topology, would make it difficult to directly generate a frequency low enough to be audible. As such, we would need to return to the idea of a reference oscillator and work through getting a mixer circuit to work. The nice thing about how low the desired frequency is compared to the likely reference frequency is that it would allow us to use an unbalanced mixer. An unbalanced mixer is the simplest kind of mixer because it doesn't do any of the filtering. This means that the output is actually a superposition of four different signals: the first input signal, the second input signal, the frequency sum of the two input signals, and the frequency difference of the two input signals. If we assumed that the reference frequency were approximately 125kHz (this was the lowest pitch reference volume we saw cited in our research), the expected output signals would be: 125kHz, 125-145kHz, 250-270kHz, and 0-20kHz. Therefore, we should be able to use a first order low pass filter to remove all the unwanted frequencies without the need for much precision in the filtering.

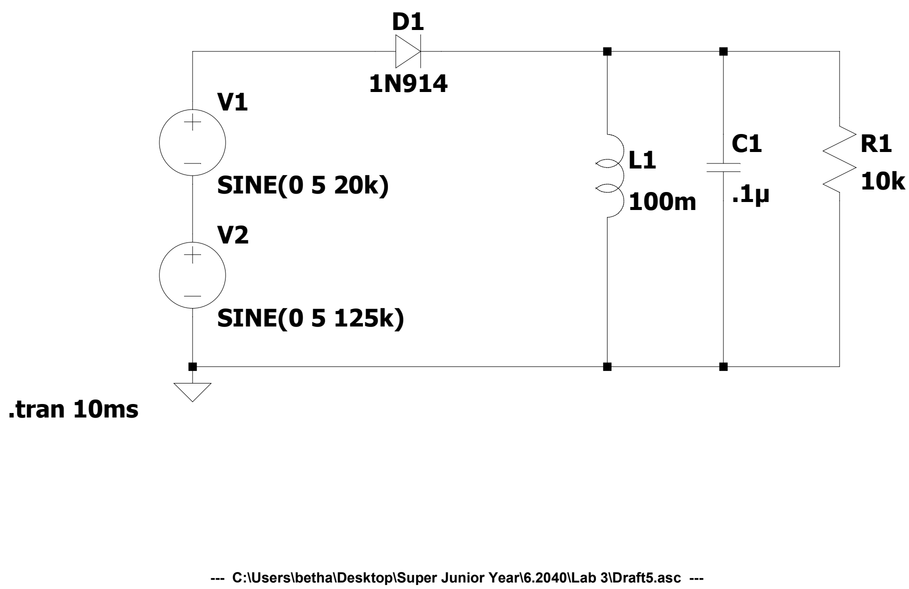

Fig. 17 The expected topology of an unbalanced mixer for a theremin using a reference oscillator

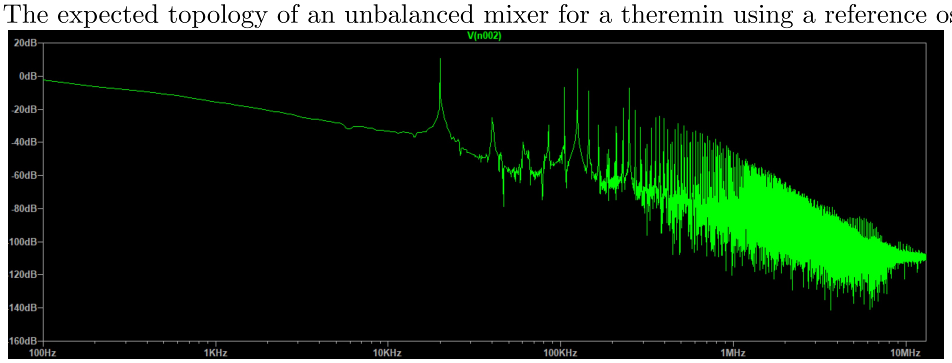

Fig. 18 The frequency domain output of an unbalanced mixer using the above topology

It also seems quite likely that, in order to achieve ideal results, we could spend many more hours fiddling with all the resistor and capacitor values. We found ones that worked and stuck with them because there was so much to do in so little time, but small improvements would likely make it possible to increase the frequency range. This is desirable because, as it stands, our theremin spans all of about 4 notes. This is satisfactory given the severe time constraints, but it won't be playing any concertos any time soon. Some of the resistors would likely be replaced with potentiometers because theremins are extremely sensitive to their environments. For example, Sarah initially constructed the volume circuit on the west side of the lab, and when she brought it over to the east side for the demonstration, she had to adjust things like the supply voltages in order for everything to continue working. When a theremin is taken into a new space for a performance, it must be carefully tuned to have the desired reaction to the musician's hands in that specific space. Using potentiometers instead of resistors would make this possible.

7.3 Future Additions

Throughout this project, we have generated many ideas for how to take this further. We've already mentioned in the previous sections the specific improvements we would like to make to the existing circuit, but there are also new components that we would like to add on.

7.3.1 Timbre Adjustment

One addition that we would like to make would be the addition of a dial that could change the timbre of the noise. As was mentioned earlier, timbre describes a set of qualities of a noise beyond the simple pitch (oscillation frequency). There are two main contributions to timbre: brightness and waveform. We would like to add dials to change both of these parameters.

Adjusting the waveform is a relatively easy addition. We saw in lab 5 the way that square waves can be turned into triangle waves, and we could use similar techniques to turn our sine wave into a square wave, triangle wave, or sawform wave. The tricky part of this addition would be how exactly we would make the output switch between these waveforms. It's relatively easy to construct circuits that generate each of these waveforms, but considerably more difficult to find a way to connect them all such that turning a dial (which are usually potentiometers) would change which one is outputting to the amplifier.

Adjusting the brightness might be much more difficult. Brightness refers to the strength of various harmonics within a sound. Harmonics are essentially the audio version of aliasing. The Fourier transform of a "brighter" sound will have high frequency components with larger amplitudes than a "darker" sound. Since we are intentionally filtering our signal to only include the pure frequency we're looking for, changing the brightness would be difficult with our current setup. We would either have to artificially recreate the harmonics or we would have to use a more complicated mixer that only filtered out specific frequencies.

7.3.2 Intonation Feedback

The other major addition that we were really hoping to be able to implement at least the first iteration of, but unfortunately ran out of time, would be feedback to the musician about how in tune their playing is. Theremins are notoriously hard to play because there is a lack of tactile feedback regarding the intonation like you would get from the frets of a guitar or even the fingerboard of a violin.

The simplest iteration of this is what we tried to implement, but ran out of time to. This would be an LED that indicates exactly one note. We would use a bandpass filter to send voltage to the LED only when the frequency was within a certain range that corresponds to the target note.

From there, we would branch off into two areas of exploration. The first would be giving feedback on how close to the note the musician is and the second is indicating multiple notes.

One way the first could be implemented would be by making the bandpass filter sufficiently narrow. This would mean that when the musician is close to the center frequency of the note, the LED would be getting a high voltage. As they move further away from it, the filter would begin to roll off and the voltage going to the LED would lower. Since we know that LEDs tend to dim as the batteries supplying them die, we can infer that the lower voltage applied would have the LED be dimmer. Thus, the strength of the light from the LED would be an indicator of how in tune the note was.

The second could be implemented using a series of bandpass filters. Each filter would correspond to a specific note. By sending the signal to each of the filters, only the one corresponding to the note closest to the frequency being played would light up. This would give the musician a sense of which note they are playing and how their hand movements are changing the sound.