Laser Harp (v2.0)

Briana Chavez, Henry Cheung, Mira Partha

May 18, 2017

6.101 Final Project Report

Table of Contents

I. Introduction

A laser harp is an electronic musical interface that replicates the functions of a harp, using laser beams as strings and analog circuitry to generate tones. Our project consists of three main subsystems: the harp interface with the laser beam "strings," a tone synthesizer, and a class-D audio amplifier. The harp contains six strings tuned to play a C major pentatonic scale. The harp is capable of switching between three different well-tuned key signatures, producing different waveshape audio signals, and adding a tremolo effect to each tone. Henry was responsible for the laser harp interface, envelope generator, and additional functionalities; Briana was responsible for the synthesizer; and Mira was responsible for the class-D amplifier.

II. Block Diagram

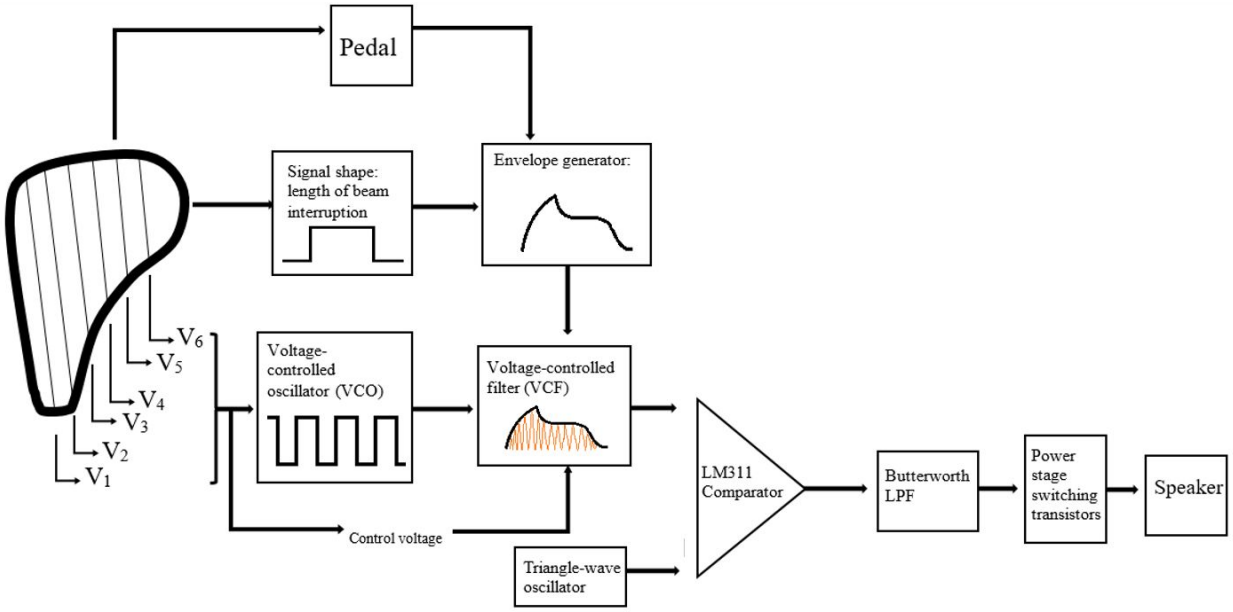

Figure 1. A simplified block diagram of the overall system.

Our entire system is powered by a ±15V power supply, specifically the one available at the 6.101 lab benches. Specific signal amplitudes/frequencies, power requirements for specific modules, and details of signal flow through each subsystem are explained in the sections below.

III. Harp interface - Henry

A. Overview

The harp interface and the overall design of connecting the action of hitting a laser beam and generating a tone to be played through a speaker is by its nature a more digitally inclined circuit. As such, the circuits behind it uses highly digital signals, versus a more analog approach, simply due to the fact that there was no need for an analog approach. The main challenge is

implementing the design multiple times over on numerous inputs, and in such a way as to minimize the number of discrete components used.

The harp interface consists of three two-by-four (referring to a common dimension applied to unprocessed lumber) pieces of wood, with the top horizontal section embedded with 18 holes drilled in rows, as shown in the figure below

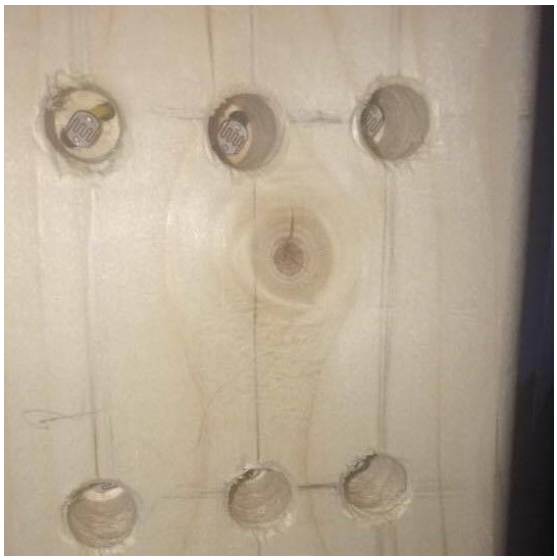

Figure 2A. The top row of photoresistors in the laser harp

Figure 2B. The bottom row of laser diodes embedded in the laser harp

The 5mW red laser diode were soldered to multiple pin headers, and connected to a breadboard, powered by a 5V rail. The top row of photoresistors were all soldered onto wires and extended down to a breadboard with multiple pin headers as well, and simply interfaced with just like regular resistors. Interfacing with them like resistors was one particular reason to choose

photoresistors, since a light detecting circuit simply requires the photoresistor in a voltage divider, and depending on where the photoresistor is placed in series, the voltage created will either rise or fall when light hits or is blocked from the detector.

B. Control Voltage Generator

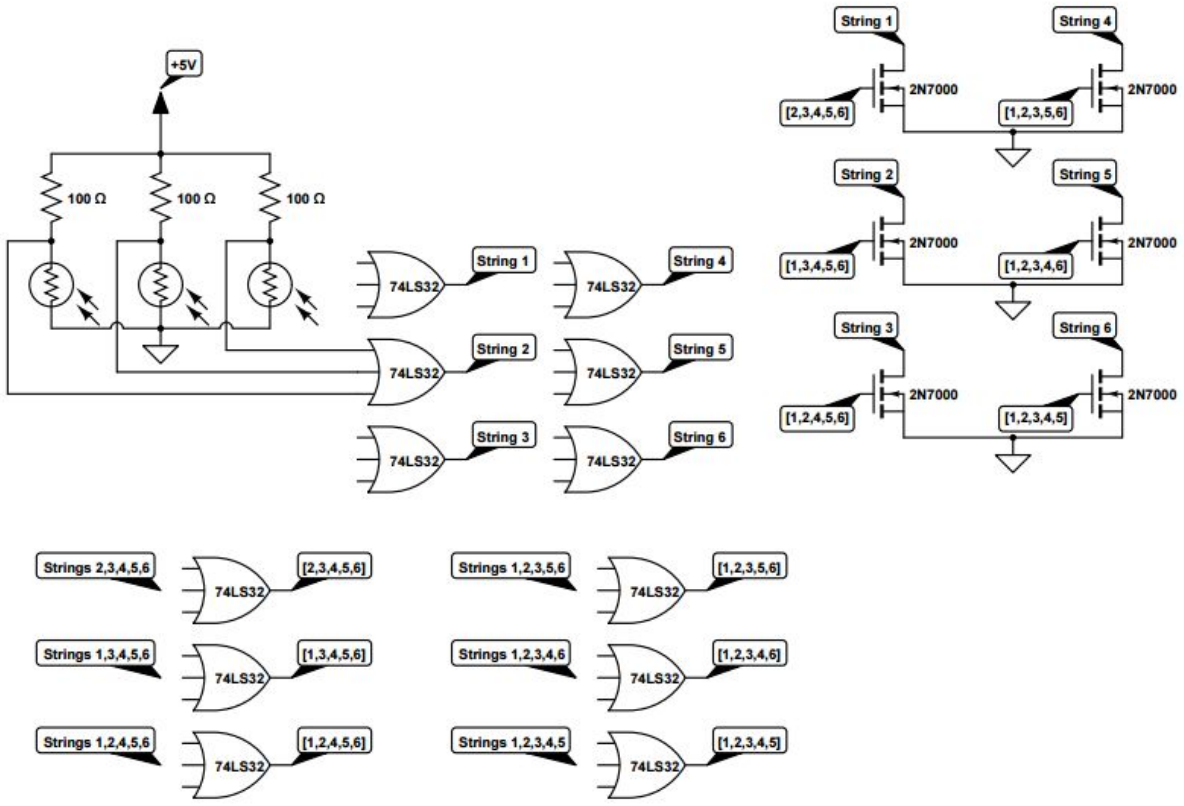

Figure 3A. Three photoresistors OR'ed together, representing one of six strings. The sets of 6 OR gates and N-FETs correspond to multiples of the same circuit for each string.

The three photoresistors in the circuit diagram above represents one string, of six possible strings. For the sake of brevity, those circuits were not drawn, but were the same with simply their respective signals. The photoresistors are placed in series with a sufficiently small resistor, in this circuit: 100Ω, and the output of the voltage dividers will rise to approximately 3.5~4.5V. For the laser harp, while three laser beams represent a string, only a single laser beam has to be hit in order to trigger a control voltage. As such, the three voltages were OR'ed together using 74LS232 logic OR gate chips. The output of the respective OR gate would represent that string, and if high, then the string had been "plucked".

The output of each strings' OR gate are then OR'ed together with all the other strings' OR gate outputs, except for one, done six times over. For example, the signals representing

strings 2, 3, 4, 5, and 6 are OR'ed together, where the output of that gate represents whether any of the other strings have been hit.

The output of each strings' OR gate, representing whether that particular string was hit, are tied to the drain of a 2N7000 N-FET, with the source tied to ground and the gate tied to the output of the OR'ed signals of all the other strings. When one of the strings is hit, then the output of all the OR gates with that string signal will go high, and pull all the other strings' signals down to ground. This was designed so that only one string can possibly trigger a control voltage, without having a summing effect of a control voltage being generated for multiple strings. The voltage-controlled oscillator and the laser harp itself is designed to only take and play one note at a time.

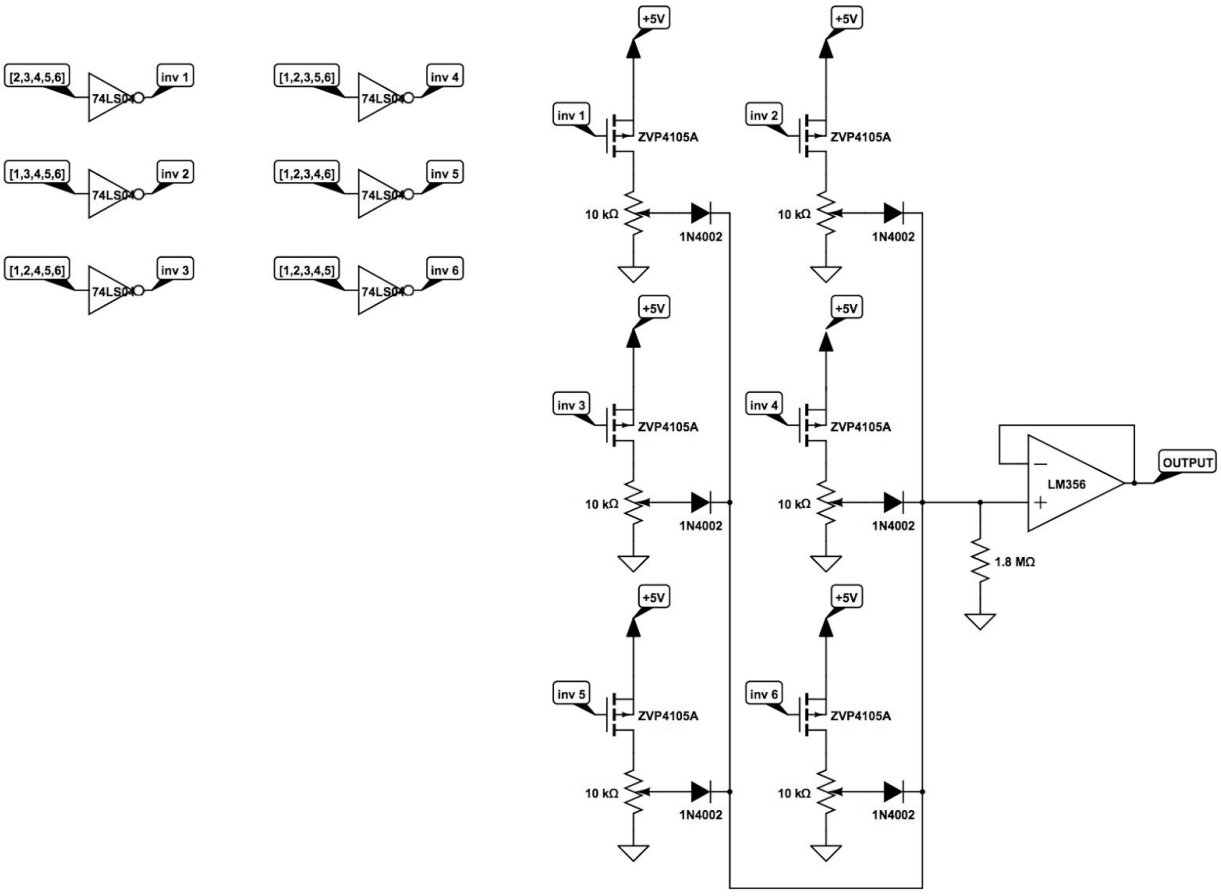

Figure 3B. Final stage of control voltage generation. The sets of 6 inverters and P-FETs represent the duplicity of the same circuit for each string.

The signals that represent whether all but one string are inverted, using the 74LS04 inverter logic chips. The inverted signals once again represent the string that was omitted in the signal going into the inverter. If a string was hit, then all other inverter inputs will be low except for the inverter with the signal omitting that specific string. The output of the inverter will rail

high and is tied to a P-FET's (ZVP4105A) gate, with the source tied to the +5V rail. When a string was hit, the P-FET's channel will open and drop +5V across the 10kΩ potentiometer to ground. The voltage divider output of the potentiometers can then be tuned to generate the exact control voltage required by the VCO.

Each voltage divider output sums into a single node through a 1N4002 diode, preventing backflow of the voltage into the other voltage divider outputs, which could pull a desired signal to ground or to some other voltage. Then, the node is tied to the non-inverting input of a LM356 op-amp voltage buffer and the output of the voltage buffer will be fed into the octave switching circuit, and then the voltage-controlled oscillator. As a caveat, a 1.8 MΩ resistor to ground is required at the node before in the non-inverting input as the diodes do not provide a path to discharge the voltage at that node, and without a path to ground, the control voltage output would look like an exponential decay of the voltage signal as it slowly leaks back to ground.

As long as the laser beam of string is continually blocked, the control voltage will continue to be generated, and when no longer hit, the control voltage will drop back down to ground.

C. Octave Switch

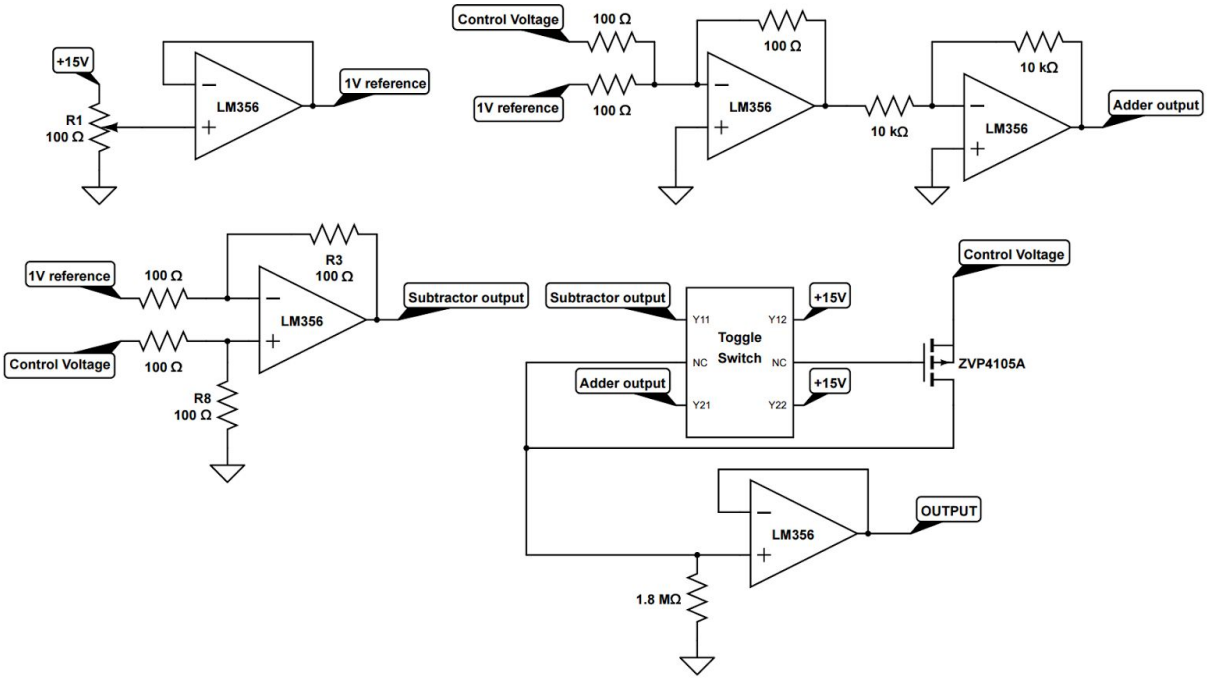

Figure 4. Octave switching circuit, with adder and subtractor op-amp configurations with a physical switch for the output selection

The octave switching feature of the laser harp was implemented with the above circuit, where a physical three-state switch connected either the lower octave control voltage, the higher octave control voltage, or the original control voltage. The toggle switch was built such that flipping the switch up connected the NC nodes to its respective sides' Y_{11} or Y_{12}, and vice versa with nodes Y_{21} and Y_{22} on the down switch. The resting middle position of the switch is on the NC nodes. The two sides of the switch are not electrically connected to each other.

The control voltages will add a set voltage either to increase or decrease the control voltage, in either an adder or subtractor op-amp configuration respectively. The linearly increased/decreased control voltage signal will be fed into the linear-to-exponential circuit, which will generate the correct exponential increase/decrease for the required control voltage for an octave transition.

The circuit diagram in Fig. 4 shows a simple op-amp subtractor circuit, with a 1V reference voltage being created via a potentiometer and a voltage buffer on the inverting input of an op-amp and the control voltage on the non-inverting input of the op-amp. The resistors were all chosen to be as close in measured value to each other to hold the direct subtraction relationship of the control voltage minus the 1V reference. The op-amp adder circuit uses the same 1V reference and the control voltage, with the output fed into another inverting unity gain op-amp, in order to account for the inverting result of the op-amp adder circuit.

These two outputs with either 1V subtracted or added to the control voltage and connected to the Y_{11} and Y_{21} pins respectively, and on the Y_{12} and Y_{22} pins, +15V rails are connected. The original control voltage is connected to the NC on the left side of the toggle switch through a P-FET which is turned off whenever the toggle switch is flipped up or down. As such, flipping the switch up will turn off connection of the control voltage and route the subtractor output to the input of the VCO, and likewise for the adder output.

The output of the switch is once again buffered with a LM356 voltage follower configuration, and the output of the switch is tied to ground through a 1.8MΩ resistor to provide a discharge path for the voltage at the node.

D. Volume Control

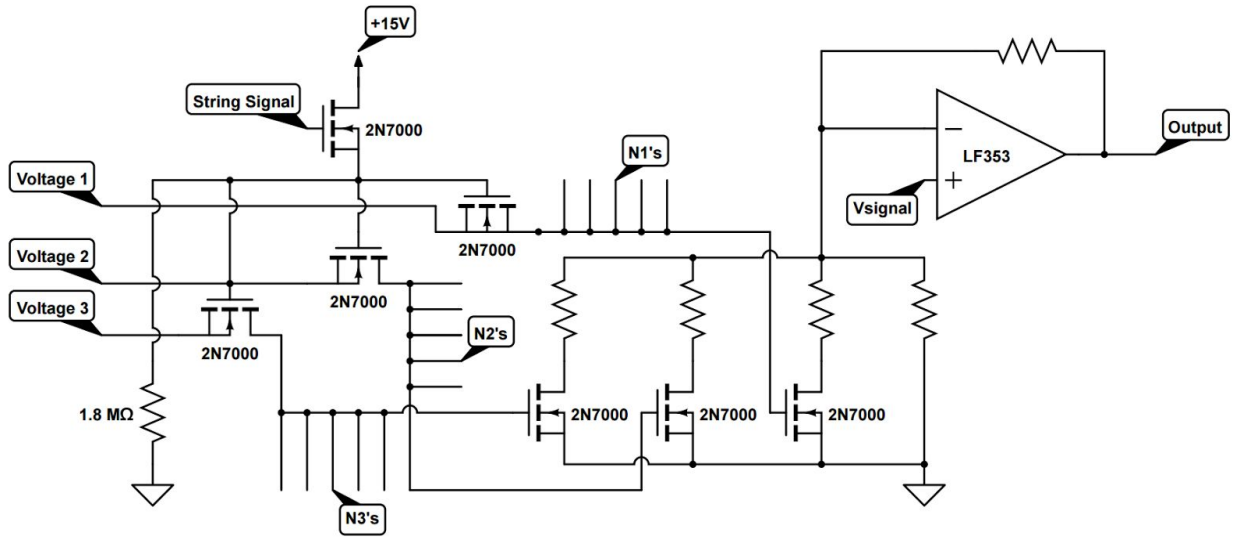

Figure 5. Volume control circuit, with the left half of the circuit representing a single string's three voltage dividers, and the right a basic inverting-gain op-amp configuration

The above volume control circuit was designed with a basic inverting-gain op-amp configuration, in the right of the diagram in Fig. 5. Four resistors are in parallel to ground, with three paths being controlled by the channel of the 2N7000 N-FETs. When a string is triggered, the MOSFET tied to +15V at the top left of the diagram opens, charging the gates of three N-FETs (with a 1.8MΩ resistor to ground for a discharge path), allowing the signal of each a string's three voltage dividers of photoresistors to pass through, joining a single node with the input of the five other strings' N-FETs, tied to each laser beam's own respective NFET to ground.

When a string is hit, then the signals from that string's voltage dividers are allowed to pass through and to turn on its respective MOSFET, thus decreasing the parallel combination resistance to ground in the inverting-feedback gain op-amp, and increasing the gain of the amplifier. Hitting three beams will then amplify the signal output by the synthesizer more and correspond to a louder tone, versus only two beams or just one beam. Since the signals are only allowed to pass by a string's total signal, only the current string's voltage dividers that is being played can turn on the MOSFETs pulling resistors to ground in the parallel combination.

Unfortunately, the circuit above only works using an ideal model of the MOSFET. Due to non-ideal leakage current, when an AC signal flows through the circuit, the MOSFETs' drains will be pulled to lower potentials than ground, and so a non-uniform amplification is presented in the positive sweep of the AC signal. The precise issue presented itself upon construction of the circuit and was not able to rectified given the time. However, a solution to the non-linear amplification is to apply a negative voltage to the gate of the MOSFET, instead of simply grounding the gate of the MOSFET. This will ensure the drain of the MOSFET is not pulled

negative potential and the current will route fully through an open branch in the parallel combination.

E. Final Thoughts

The interesting challenge in the circuits designed above was the originality of the circuits, due to the fact that there is no standardized way to try to perform the operations that were done above in a more discrete, analog method, when a quick digital solution could have sufficed. In this vein, I decided to attempt to design the most compact system, with as few components as possible. It would be easy to tackle the problem with simply an overwhelming number of comparators and the like, but lack any interesting solutions. It was a challenge to come up with a more analog circuit while knowing that with fewer constraints, one would most often opt for a much more digital solution, perhaps with a small microprocessor performing all the operations with ease.

IV. Synthesizer (Briana)

At the heart of the problem is the challenge of processing a signal from any given string and generating a waveform with the predetermined frequency, corresponding to the string's pitch. This was accomplished by implementing three main systems: a linear to exponential converter, a voltage controlled oscillator, and a voltage controlled filter. Additional systems were designed such as a low frequency oscillator and an envelope generator which will be discussed later.

A. The Voltage Controlled Oscillator

Designing and building a synthesizer is a classic analog project and as such, many tried and true circuit designs already exist and are largely variations on similar design concepts. As such, design for the voltage controlled oscillator was based on a common circuit topology as is described in the Horowitz and Hill's *Art of Electronics Third Edition* (p.267). The final implementation of the VCO is shown in **Figure 6**.

This circuit consists of an integrator, a Schmitt trigger, and a reset to cause the oscillation. The 10kΩ voltage divider sets the positive input to half the control voltage which the op amp sets to be equal to the negative input. As a result, current is driven across the 3.9kΩ resistor and into the capacitor, charging it and causing a steady rising output.

The second op-amp is a Schmitt trigger that takes the integrator signal as input. When the input voltage rises above the threshold of 2/3 Vref, it outputs 5 V and the threshold voltage falls to 1/3 Vref. When the input voltage falls below that, the output goes to 0 V and the threshold moves back up and thus, the output is a square wave.

This wave is connected to the gate of a 2N7000 N-channel MOSFET, which when turned on, will ground the negative input of the integrator and cause the capacitor to discharge. This

creates the falling edge of the integrator output and results in a square wave. In this way, the MOSFET serves as the reset for the circuit.

This version of a VCO relies on the relationship between Vin and Vref in order to set its wide range of achievable frequencies. By setting Vref to a constant 5 V, a larger Vin will charge the capacitor faster, allowing the upper threshold voltage to be surpassed faster than if Vin were smaller and driving less current across the 3.9kΩ resistor. Thus the speed at which the Schmitt trigger thresholds are exceeded at are variable and so are the corresponding frequencies.

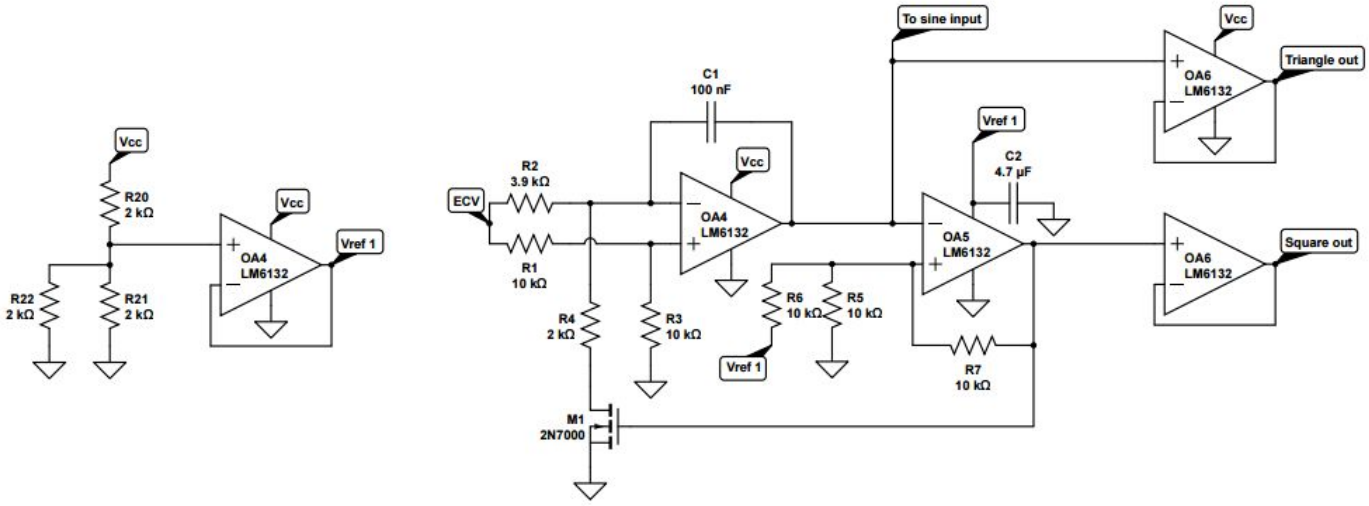

Figure 6. Final VCO design

The triangle wave and the square wave outputs are voltage buffered and sent as inputs to the voltage controlled filter to eliminate any unwanted loading effects in . These waves produce slightly different sounds as they are made up of different harmonics of the fundamental frequency. For an additional tonal option, the triangle wave can be converted into a sine wave using a differential pair of matched NPN BJT's (**Figure 7**). This circuit makes use of the BJT's hyperbolic tangential current relationship to produce a smooth sine wave. The primary source for the calculation and theory for this circuit is from the 1976 IEEE Journal of Solid-State Circuits (Meyer, Sansen, Lui, Peeters, 418-420). The addition of a unity gain op amp and the 50kΩ potentiometer are added to address issues in amplitude attenuation and rectification of the full synthesized sine wave. Even after these additions, the sine wave output was very approximated and only truly became sinusoidal after filtering.

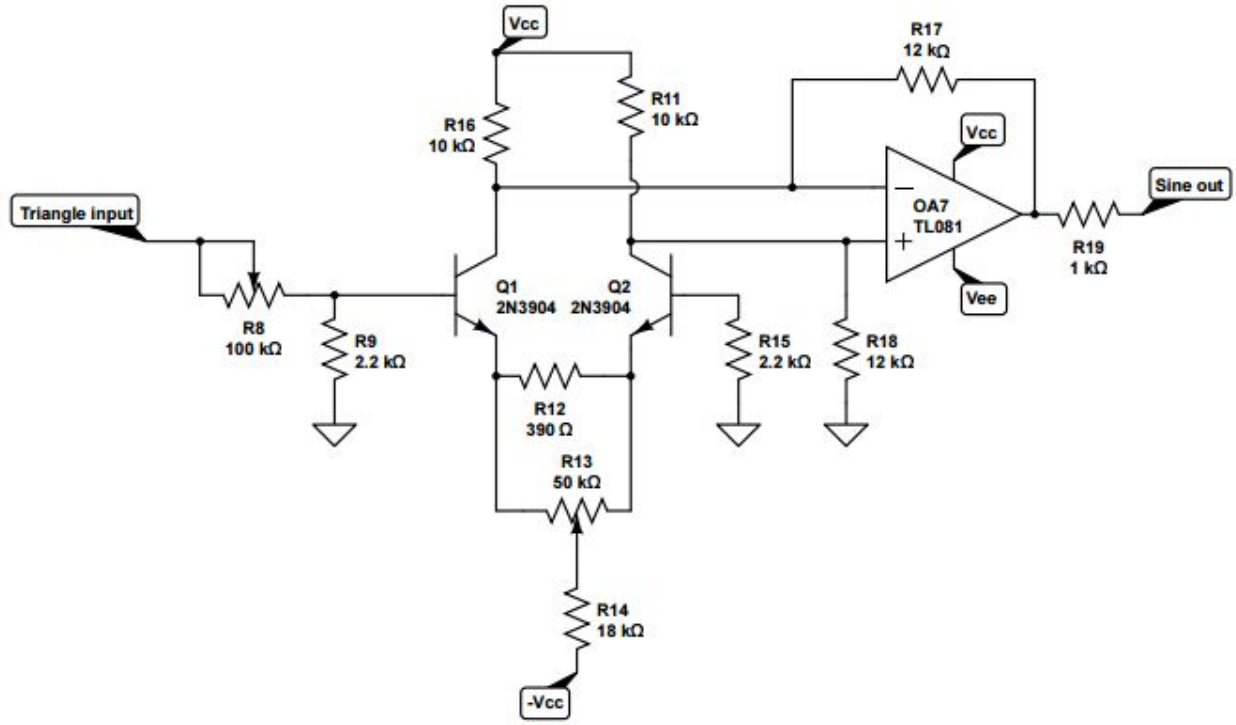

Figure 7. Triangle wave to sine wave converter circuit

B. Linear to Exponential Circuit

As a standalone module, the VCO is completely equipped to take in any control input in the range of 0-10V and produce over five octaves of pitches. All that it requires is a carefully tuned input voltage from the harp interface that matches the voltage required for the VCO to produce the desired frequency. However, it one of the most interesting objectives for this laser harp's operation is the capability to switch to a different octaves or keys while automatically remaining in tune.

We decided to focus on creating an instrument that could produce a pentatonic scale starting on C2 and switch to the C2 or the C1 scales in the Western musical style. However a critical challenge of tuned tone generation is the fact that as the pitch of a tone increases, the fundamental frequency of that note increases exponentially. This means that the VCO alone has no singular change switching mechanism in control voltage that would ensure that all the notes in the scale would be accurately adjusted to a different octave or key. A conversion circuit between linear to exponential control voltages is needed.

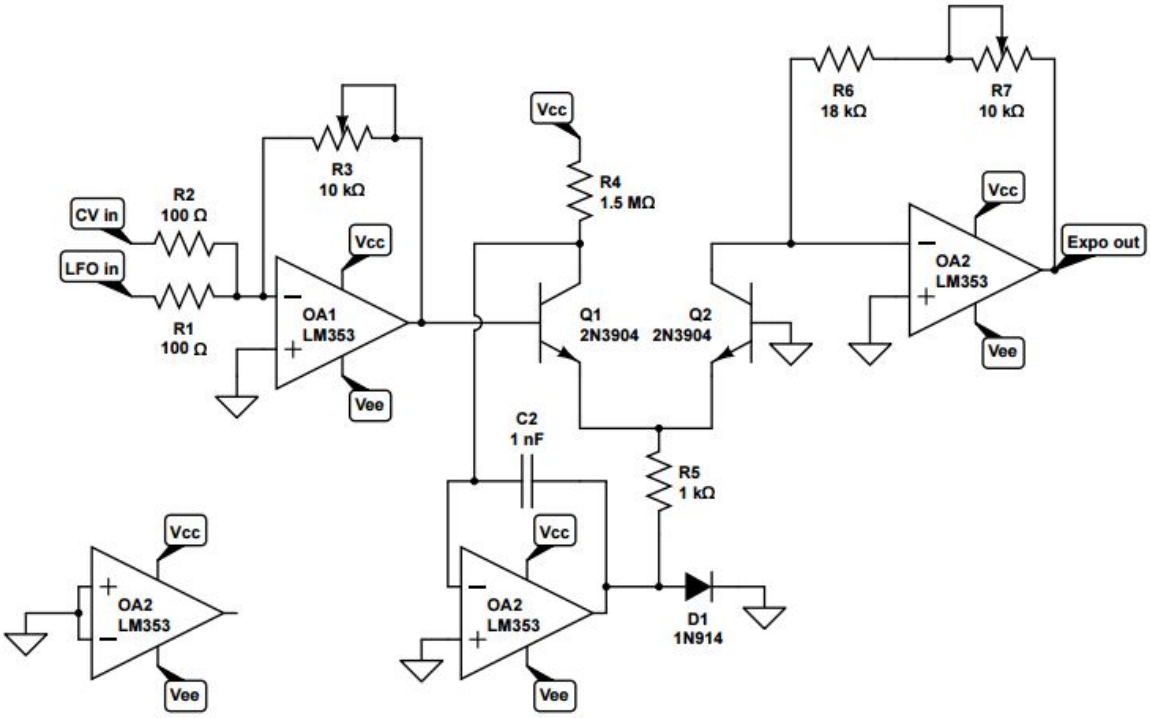

Figure 8. Linear to exponential converter circuit

The most important aspect of this circuit's design as shown in **Figure 8** is the matched NPN BJTs implementation as current mirrors. The current of a single BJT is described as,

$$I_C = I_S(e^{(\frac{qV_{be}}{kT})} - 1), \qquad \text{(IV.1)}$$

in which the collector current is exponentially related to the voltage drop across the base and emitter. Using this relationship, it is simple to set the input as the base voltage to Q1 and get the appropriate current to then convert back into an exponential control voltage. However, there is temperature dependency that factors into the equation and threatens to cause drift in the output as the device heats up over time. By using a second very closely matched BJT as a current mirror, the saturation current of both Q1 and Q2 are nearly identical. As such, the collector current of Q1 is equal in magnitude to that of Q2, with opposite signs. Similarly, Vbe of Q1 will equal -Vbe of Q2 and the temperature is no longer a factor.

The collector current of approximately 10 microamps of Q1 is generated by the voltage drop across R4 from the 15V power supply. In order not to saturate the transistor, the input from the harp and any addition effects circuitry is scaled down so that when the input is 1V, the base voltage of Q1 is near 10 mV. The op amp across the C and E terminals of Q1 and the 1N914 diode are added precautions to stabilize any drift in current driving R4. The output exponential current is converted back to voltage using a 10kΩ variable resistor that was essential to finding a reasonable range of input voltages for the desired exponential control voltages.

When it came time to integrate the full linear to exponential circuit and the voltage controlled oscillator, there was a bias output voltage of around 200 mV at all times. This means

that even when all laser beams were uninterrupted, there was enough voltage being sourced to the VCO to generate a constant tone. In order to cancel out this floating voltage, I first attempted to match the impedances at the non-inverting nodes of all the op amps of the linear to exponential converter circuit. I added 200nF bypass capacitors to the power rails and grounded any unused op amps in dual packages, to no observable improvement.

When attempts to generate -200mV passively and then using an adder op amp configuration (**Figure 9**) failed to eliminate what was likely a problem of biasing in the op amps, the circuit that successfully addressed the constant output voltage is shown in **Figure 10**.

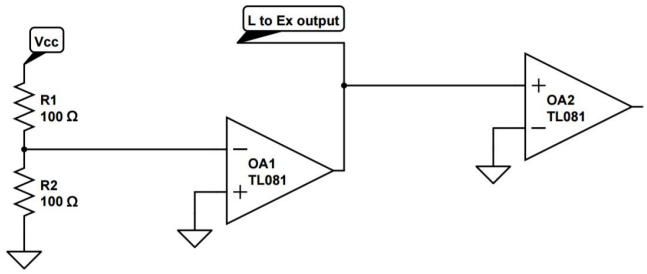

Figure 9. Failed voltage compensator circuit

Firstly, recognizing that the lowest needed control voltage needed to produce the lower C1 frequency provided an approximate lower threshold for which the control voltage circuit is expected to drive the oscillator. Testing the voltage controlled oscillator showed that that it was tuned to produce this frequency with around 300 mV. Thus the objective became to automatically ground the linear to exponential output when the input was zero and no laser beam was interrupted.

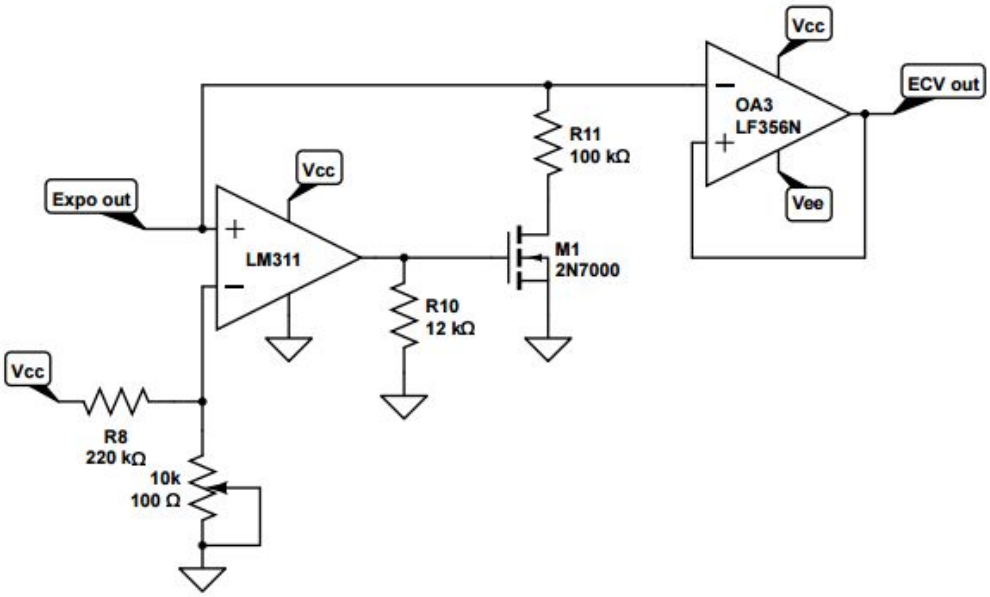

Figure 10. Final voltage limiting circuit

By creating a tunable reference voltage around 250 mV, the output of an LM311 turns on or off an N-channel transistor, using the same theory for the grounding or reset MOSFET

implemented in the VCO. If the non-inverting input is greater than the inverting input, the output of the comparator rails to the +15V rail, turning on the MOSFET. The 100k pull down resistor ensures that the exponential voltage is preserved, accurate and free of attenuation. If the reference voltage is higher than the input, the gate voltage of the FET rails to zero, pulling the voltage at the non-inverting input to the voltage buffer to ground. So for all voltages below the threshold of the reference voltage, the input to the voltage controlled oscillator is zero.

C. The VCF

The initial design for the VCF was based on the need for an adjustable cutoff frequency for different pitches. If correctly implemented, we decided that passing only the first three harmonics present in the square and triangle waveforms would result in the most pleasing tonality. Throughout this course, much of our understanding of using active and passive filtering techniques relied primarily on an resistor and capacitor relationship. Following this reasoning, designing a resistor with an impedance that responded to a control voltage seemed the best way to achieve a sharp and variable cutoff frequency. The VCF design is based on using a pair of operational transimpedance amplifiers (OTAs) to create a two-pole, voltage controlled, Butterworth filter (**Figure 11**). This design is based on the LM13700 application section of the datasheet.

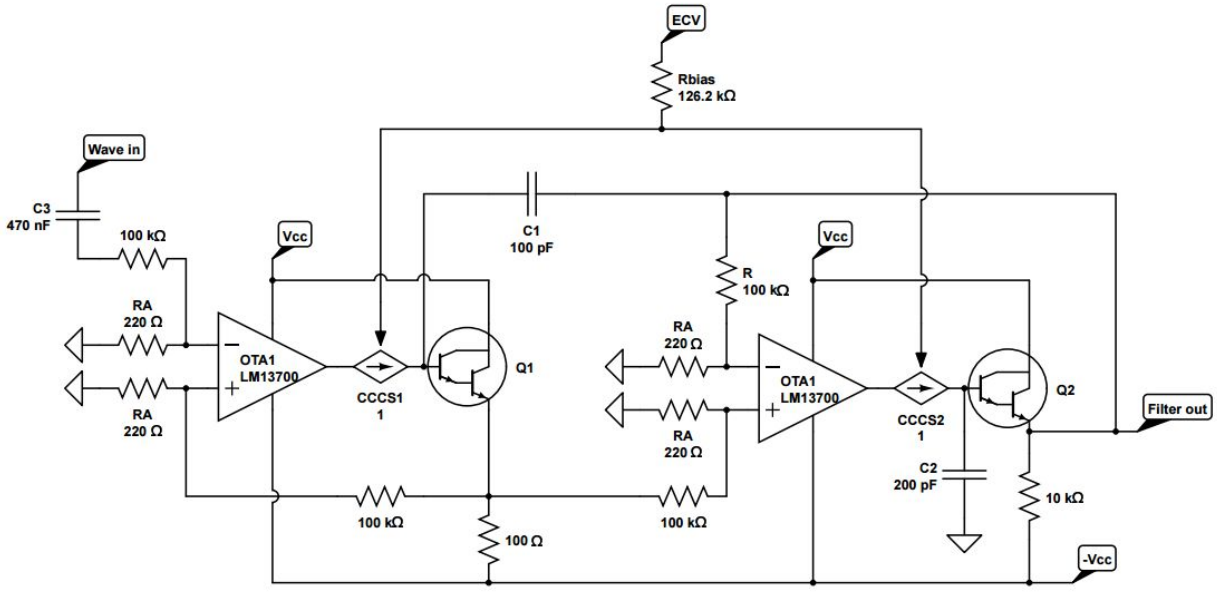

Figure 11. LM13700 OTA Voltage controlled filter. Note that the CCCS1 and CCCS2 current sources correspond to the current bias input. Additionally, the Darlington pairs shown in this figure are built into the OTA device on pins 6-11.

Essentially, the OTA is a voltage controlled current source with bias current that is used to charge and discharge a 1 microFarad capacitor. At high frequencies, the capacitor will behave

more like an open circuit that draws more current than the bias current can source. As a result, the signal begins to attenuate. The point at which this occurs is given by the equation IV.2,

$$F = \frac{RA * gm}{(R+RA)2\pi C} \quad \text{(IV.2)}$$

and $gm = 19.2 * Ibias$. (IV.3)

Using these equations, I determined the appropriate resistor value needed to set Ibias to a current that would pass the desired harmonics for the mid-range tones (about 350-450 Hz) we were generating. After iterating through a calculated and physically implemented Rbias values, the best resistor that filtered across the entire range of frequencies that was ultimately chosen was 126.2kΩ. This circuit was fairly straightforward in implementation and the filtered frequencies did prove to resemble more of a sine wave shape. This eliminated the harsh, buzzy overtones that are fairly easily aurally detected from the original waves.

Integration of the linear to exponential converter, voltage controlled oscillator, and voltage controlled filter was largely seamless once the functionality of each module was guaranteed. The control voltage for the VCF is the same as the non-linear control voltage that the VCO receives as input.

D. Final Thoughts on the Synthesizer (Briana)

A consistent challenge in this module was selecting the appropriate operational amplifiers and their appropriate power configurations. The VCO in particular was not functioning as simulation and calculation suggested it should and after testing combinations of single and dual package amps with either single or double supply rails, the single supply LM6132 proved the strongest choice. Physically debugging this problem took time as neither simulation nor calculation took the effects of the power rails into account.

Designing the oscillator circuits began with a somewhat misguided search for originality and simplicity that was challenging to recognize in some of the preexisting synthesizer schematics available in literature and on the internet. A good deal of time was spent designing and simulating individual oscillators per string, many of which came to resemble the topology for the LFO. This ultimately was constricting in that it only ever producing a single waveform with unchanging frequency and would require building 18 individual circuits in order to reach the range that the final design was capable of.

Through this period of trial and error, the power of a voltage controlled circuit became apparent. The fact that a single circuit has such a robust response for a variety of input signals is possibly the most impressive characteristic of this system. In retrospect, more of the synthesizer's stretch goals would have been accomplished by establishing a trust in the synthesizer techniques and best practices earlier on in the process.

V. Additional Sound Shaping Effects

A. The Envelope Generator (Henry)

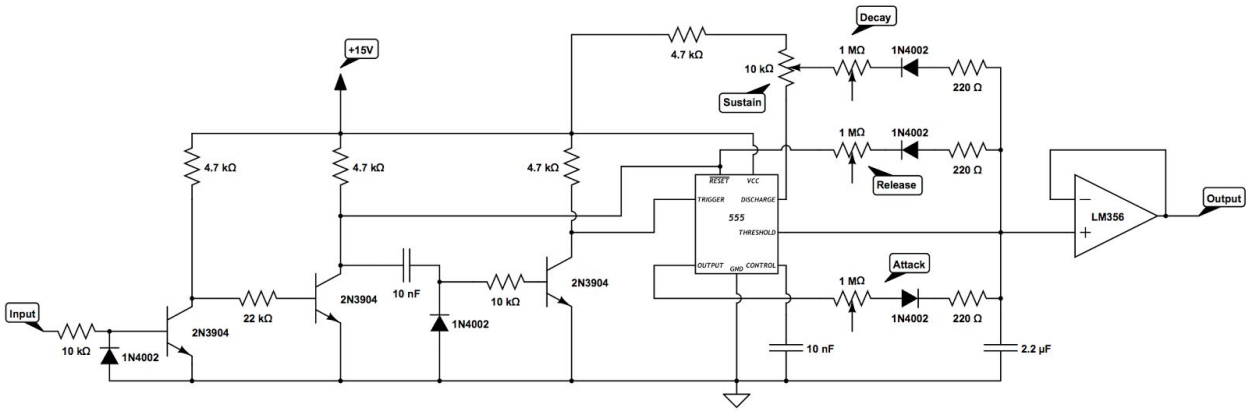

Figure 12A. Envelope Generator

The envelope generator creates an envelope which has a varying amplitude, which is used to control an amplitude modulator, described in the following section, in order to integrate the envelope shape with signal output of the voltage-controlled amplifier. The figure (12B) below depicts the ideal shape of the envelope generated, with an ADSR form, standing for an attack, or initial pluck of a string, the decay of the note, the ringing sustain of a note, and an eventual sustain, and release of the note into silence. The envelope would add a more realistic sound to each note being generated by the voltage-controlled oscillator.

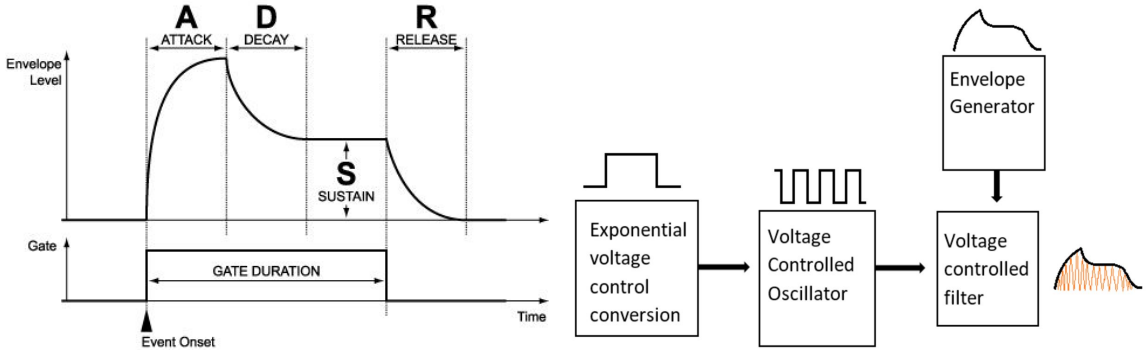

Figure 12B. Description of the format of an ADSR envelope, and the integration of the module

The envelope generator shown in figure 12A first buffers the input with a 2N3904 BJT, and generates a step pulse into the trigger pin of the 555 timer. The 555 timer will then charge up the 2.2 uF capacitor through the output pin, through a potentiometer, which can be used to tune the time constant of the attack charge time. The diode will restrict the back voltage flow, and so discharge through the decay potentiometer to the level of the sustain voltage, set by the sustain

voltage divider of the potentiometer, and then discharge fully to ground through the release potentiometer to the reset pin.

Another add-on that would be added, time permitting, would be a pedal, which would simply increase the sustain voltage and the resistance of the release resistor, increasing the end ringing of the note. It would most likely be a button that could trigger a pulse that opens a FET that performs the previous operations, and upon release, the time constants revert back to the original decay times.

As described in the following section, the amplitude modulator was crucial for integrating the envelope shape generated but we lacked the time to fully integrate the envelope generator into the rest of the project.

B. The Amplitude Modulator (Briana)

The recommended method of integrating the envelope shape with the signal output was voltage controlled amplifier that worked as an amplitude modulator. This would take both the signal and the envelope and scale the signal by the magnitude of the envelope. We didn't have time to build and debug this type of circuit but design for this module was based on a current mirroring BJT pair in a configuration much like the one shown in **Figure 13**.

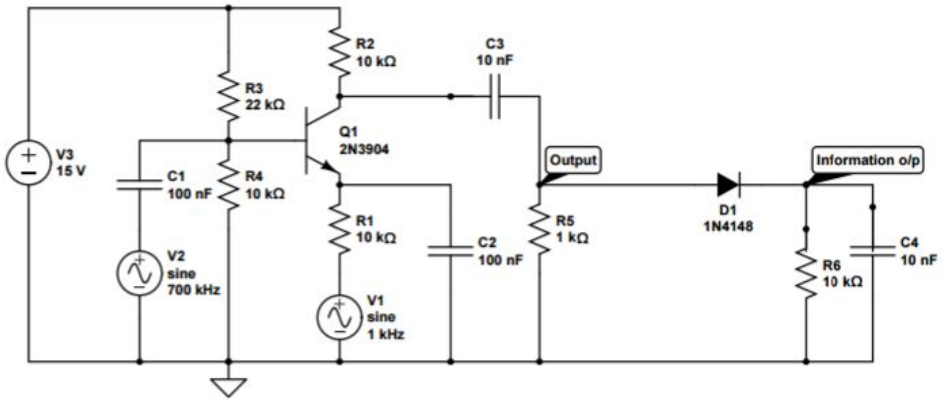

Figure 13. Proposed design for an amplitude modulator

This type of modulation could have also been accomplished with an OTA.

C. The Low Frequency Oscillator (Briana)

An added effect that we were able to integrate was the addition of a low frequency oscillator. Our LFO produces a 10 Hz approximated sine wave that is added to the linear control voltage. This second signal attenuates the control voltage up and down which the VCO then generates into a tone with small oscillations in frequency. To the ear, this sounds like a tremolo or vibrato that moves around the fundamental frequency of the note.

An elegant approached to this is a single 555 timer circuit that produces a 10 Hz sine wave. By configuring the timer in astable mode as shown in **Figure 14,** the timer becomes a fast switching oscillator that produces a square wave as long as V1 is on. The frequency of this oscillation is set by the RC pair present at the threshold pin using the equation,

$$F = (2\pi * RC)^{-1} . \qquad \text{(V.1)}$$

The final implementation shown in **Figure 14** shows the appropriate RC relationship that will achieve around 10 Hz, as well as some additional passive filtering to approximate the waveshape as more of a sine wave.

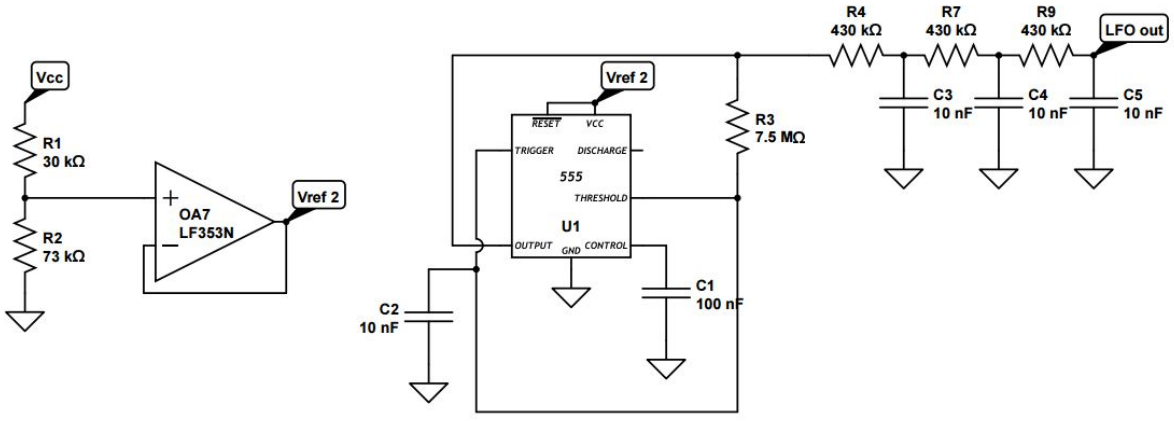

Figure 14. A 10 Hz 555 Time sinusoid waveform generator with three stages of passive signal filtering.

VI. Audio Amplifier (Mira)

A. Overview

Because the output signal from the synthesizer does not have sufficient current to drive a speaker at high volume, an audio amplifier is required to hear the laser harp's music. For this project, we chose to construct a class-D amplifier. The reason for this selection is that the class-D amplifier, as a switching amplifier, has the highest efficiency of all amplifier classes, with an achievable efficiency of 85-95%. Other benefits of a class-D amplifier include a 0 dB power supply rejection ratio (indicating that the gain is proportional to the power supply voltage, as compared to a class AB amplifier which has constant gain.) Class-D amplifiers are also much more efficient than class AB amplifiers for driving low-impedance loads, such as a speaker. They are thermally stable, and avoid crossover distortion in the zero crossing region. The minimization of loss in a class-D amplifier design is explained mathematically in **Figure 15**, shown below.

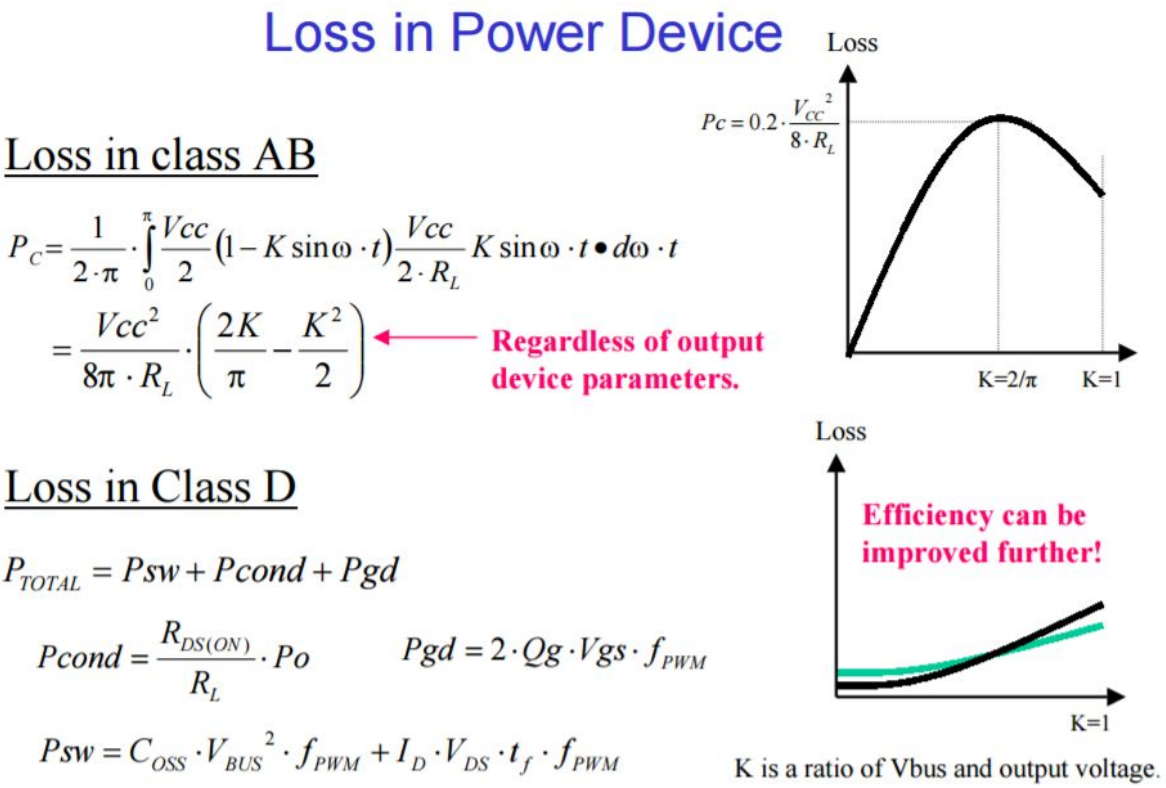

Figure 15. Comparison of loss in the power devices in a class AB amplifier versus a class D amplifier. The class D amplifier clearly wins!
(Source: www.irf.com/product-info/audio/classdtutorial.pdf)

The class-D relies on pulse-width modulation (PWM) to convert the input signal into a pulse train before amplification. The pulse train is amplified by a switching controller and output stage, after which a low-pass filter is applied to create the output – an amplified version of the original input signal.

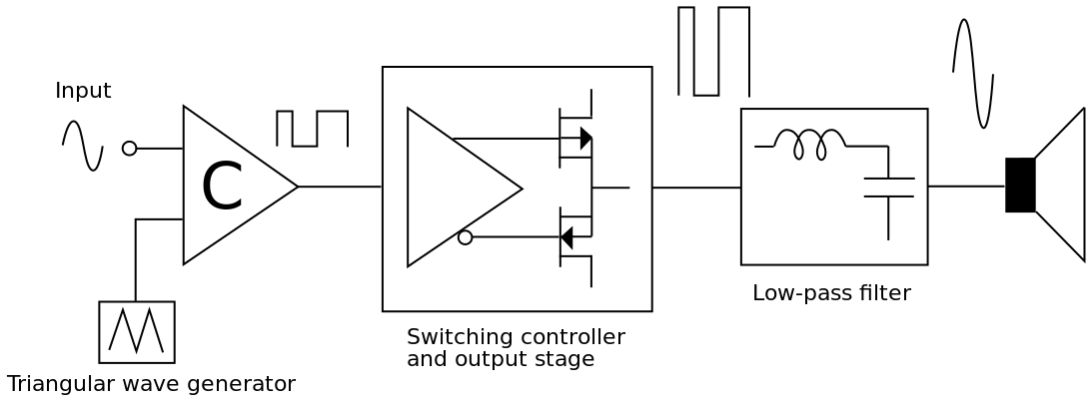

Figure 16. Simplified block diagram of a Class-D amplifier.
(Source: www.en.wikipedia.org/wiki/Class-D_amplifier)

Creating the theoretical design of the class-D amplifier was relatively straightforward. However, the circuit was extremely sensitive to noise due to the use of high frequency signals and highly sensitive devices. Thus, significant difficulties were encountered that required extensively reconsidering choices of components for the circuit, as well as devising creative methods for mitigating noise generated along the signal pathway.

B. Theoretical Design

Prior to breadboarding, an initial circuit design was drafted consisting of six submodules: input filtering, triangle wave generation, PWM via a comparator, gate driver circuitry, the MOSFET power stage, and output filtering. The circuit was originally designed to take as input an audio signal between 0-5V, and would amplify it to be 30V peak-to-peak. The circuit would operate using a $\pm$15V power supply obtained from the power supplies available on the lab benches. Successive iterations of testing and debugging saw many significant changes to the original schematic. The final, successful design was indeed able to successfully operate with +5 and $\pm$15V power rails. However, due to the specifications of the particular gate driver being used, and difficulties in level-shifting the input signals to the gate driver, the power stage topology was altered. As a result, the final circuit design amplified the audio signal to only 15V peak-to-peak; but as we discovered, this was sufficiently loud coming from our speaker. In addition, due to noise, overloading of the power rails, and other problems that were encountered during the testing process, additional circuitry was added to filter from the power supply, and to clean up the signal between the comparator and gate driver. Various choices for a comparator and for buffering op-amps were tested before the optimal devices were settled upon. The final schematic for our working class-D amplifier is shown below.

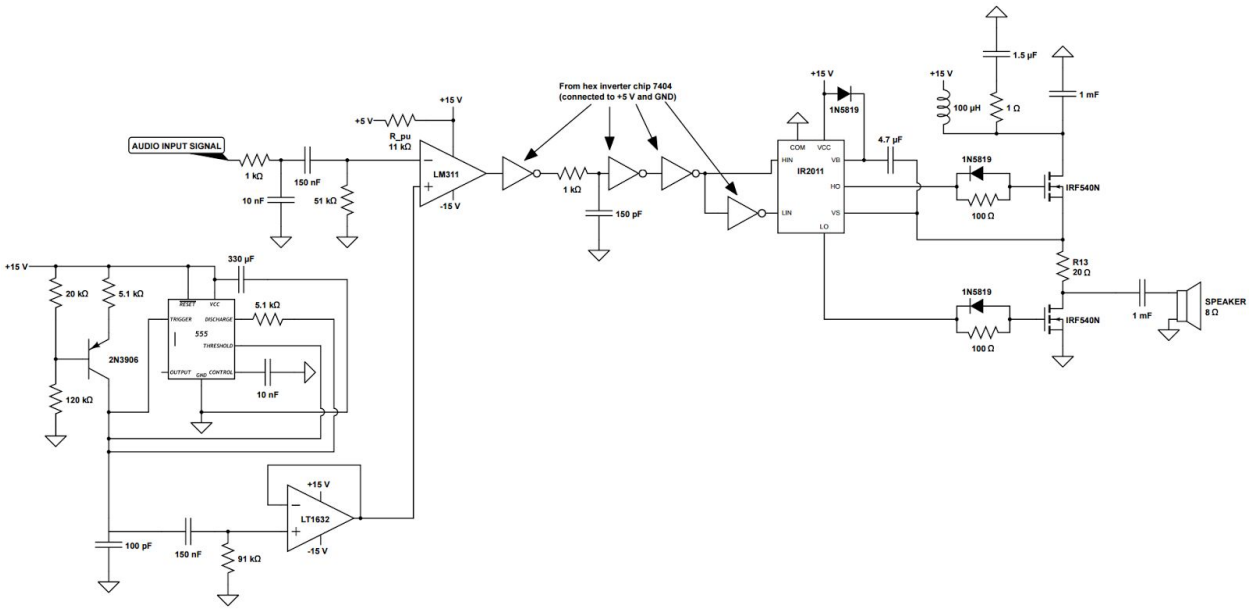

Figure 17. Complete circuit diagram of the Class-D amplifier.

Additional measures to mitigate noise include the liberal use of bypass capacitors of varying capacitances. The amplifier was also built on two separate breadboards – one for input filtering, triangle wave generation, PWM, and noise removal; and another for the gate driver and power stage. To avoid injection of noise into the power rails, which presented a significant problem for this amplifier, separate wires connected to the lab bench power supply were used to provide the ±15 and 5V rails at various points in the circuit; and star grounds were used wherever possible. After all these modifications, the amplified output audio signal was impressively clean. Each sub-module, including descriptions of the design process and ensuing complexities, will be discussed individually in the sections below.

C. Input filtering

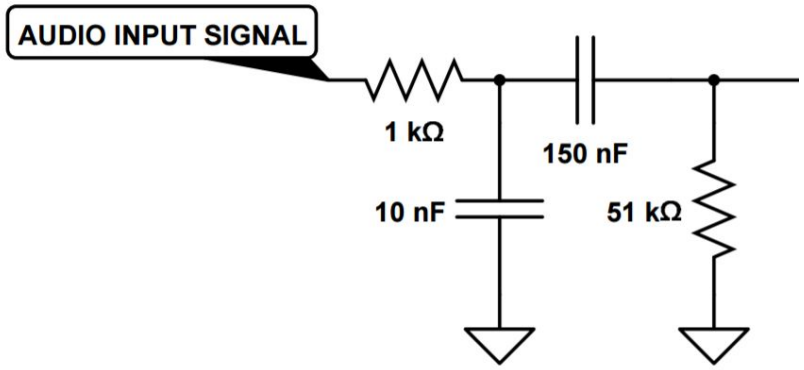

Figure 18. Input filtering of the input audio signal.

The output signal from the synthesizer serves as the amplifier's input signal, which has an amplitude between 0 and 5 V. It is passed through a cascade of two filters, first a low pass filter with a cutoff frequency of roughly 16 kHz, followed by a high pass filter with a cutoff frequency of roughly 21 Hz (all our tones have audio frequencies within this band). Since there is no active buffering between the high pass and low pass filters, the overall transfer function is no longer truly that of a bandpass filter. However, measurements of the cutoff frequencies and the gain across the passband demonstrated that this cascaded filter adequately reduces noise on the input audio signal. The purpose of the input filtering is to clean up the signal so that no noise is amplified. Given that a class-D amplifier relies on the accurate transmission of high frequency signals between extremely sensitive devices, it is particularly important for this application that noise is minimized at every stage of the amplifier.

D. Triangle Wave Generation

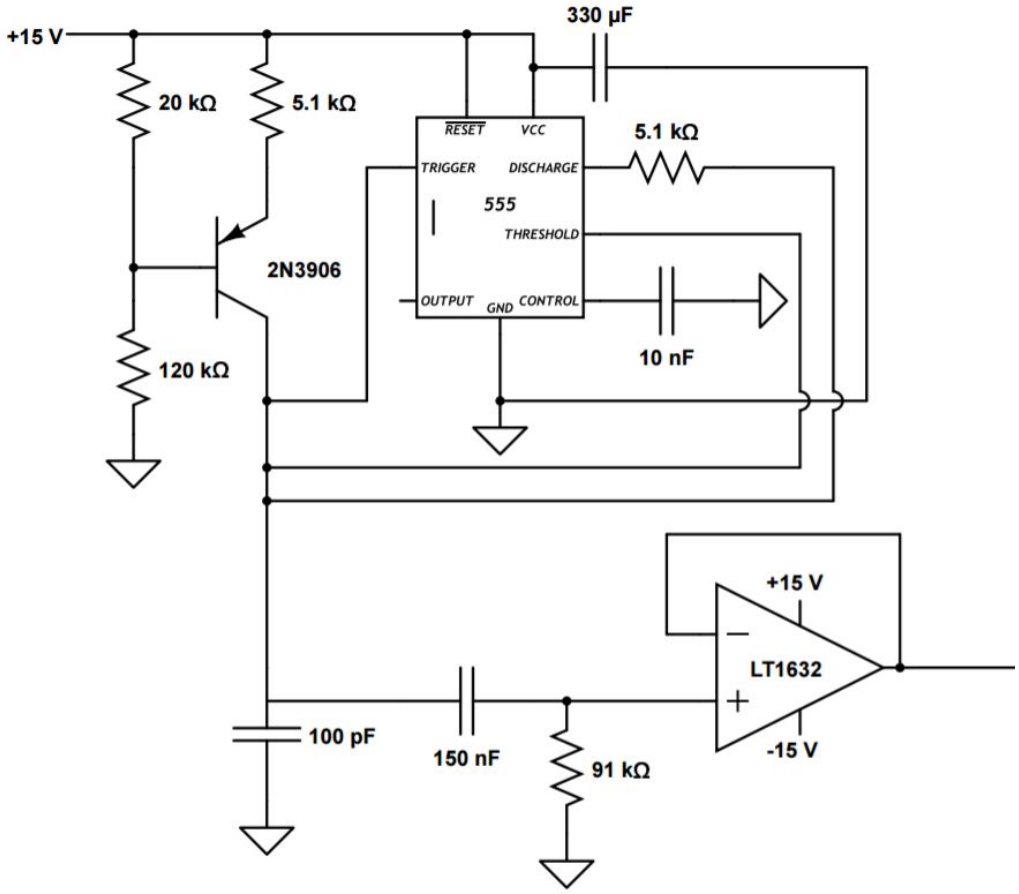

Figure 19. Generation of a 200 kHz triangle wave using a 555 timer.

For pulse-width modulation, a high-frequency triangle or sawtooth wave is required to compare the input signal against. Here, the sawtooth oscillation is generated using a 555 timer. A PNP transistor biased using a resistive divider acts as a current source to charge and discharge a 100pF capacitor. The voltage across the capacitor is a 200 kHz sawtooth wave. The switching frequency of the pulse-width modulation was chosen to be roughly 200 kHz as this frequency is more than ten times the upper bound of the human hearing range, but not too high of a frequency for the MOSFETs' switching capabilities. Initially when compared against the audio input signal, the comparator output appeared to be railing. This turned out to be due to a significant DC offset on the sawtooth wave. Accordingly, a high pass filter with a cutoff frequency of roughly 12Hz was added to center the signal around zero volts. It was also noticed that the sawtooth signal would appear significantly altered when power was switched on and the comparator was connected; to compensate for this, active buffering before the comparator was added to prevent loading the output impedance. The buffering is done using an LT1632 op-amp, which is particularly well-suited for high frequency applications due to its 45 MHz gain-bandwidth product.

E. Failed Attempts with the Comparator and Noise Removal

Choosing the right comparator proved to be one of the most significant challenges and hindrances to progress in building a class-D amplifier. When designing the circuit, out of consideration for the high switching frequency planned for use in the pulse-width modulation, the LT1016, an "ultra-fast, precision comparator" with 10ns response time, was proposed for use. The dual-output nature of the comparator was desired in order to produce two outputs that are exactly complementary (identical but precisely out of phase), to be able to feed the gate driver. However, numerous issues were encountered when using the LT1016. Its high speed and precision, previously thought to be an asset given the high switching frequency of the PWM, proved problematic, as the device was able to catch all the noise from the power supply and in the input signals. The addition of bypass capacitors along the power rails had minimal effect. As a result, the output pulse train was extremely noisy, appearing as in the pictures shown below.

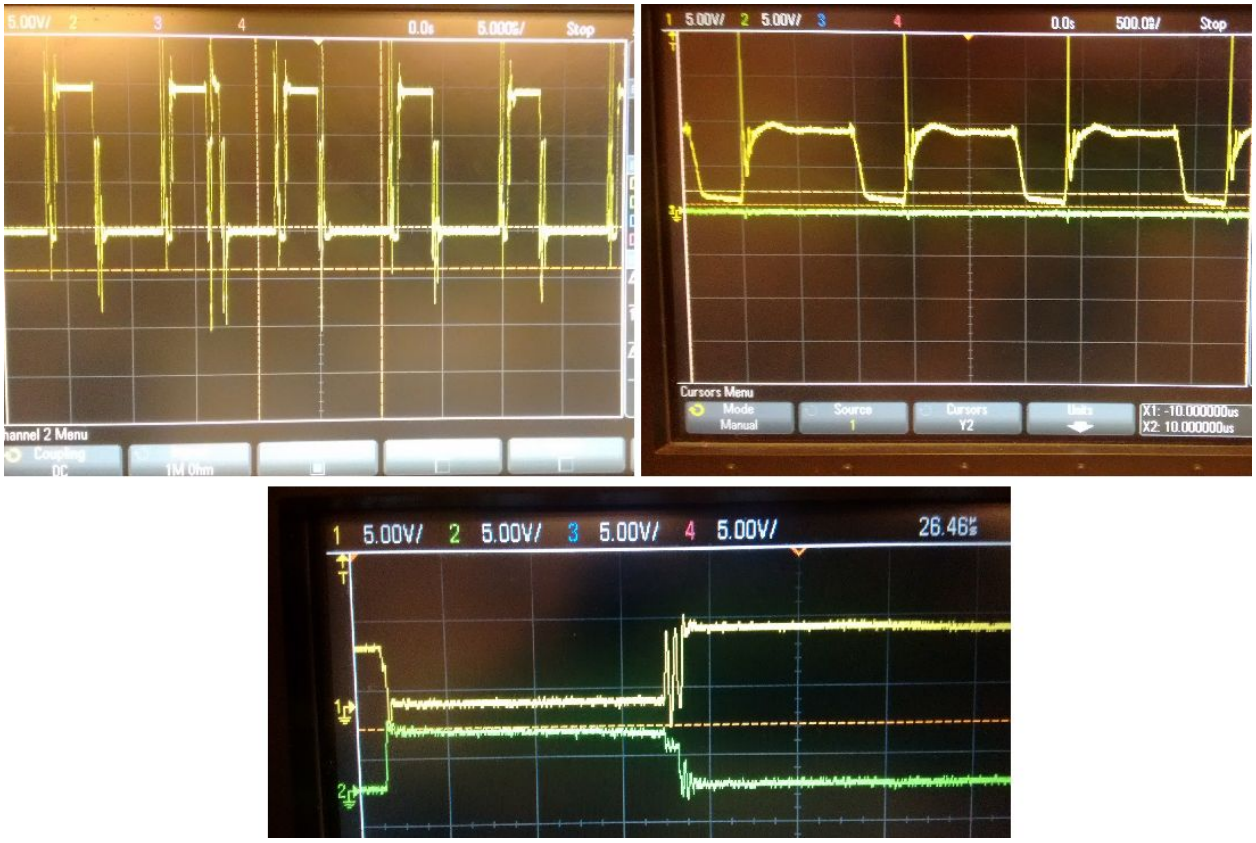

Figure 20. Noise observed on comparator output with differing amounts of noise mitigation across the circuit.

Before ensuring that the comparator output was sufficiently clean, the gate driver and power stage were connected to the output of the comparator. The noise in the signal being input to the gate driver resulted in shoot-through across the MOSFETs, as the dead-time matching from the gate driver was incapable of handling the ringing on the input signals. Attempts were made to compensate for this in the power stage of the circuit, by adding RC snubber circuits

across the MOSFETs as well as other measures to avoid shoot-through; however, with the comparator output remaining so noisy, these efforts proved futile.

Returning to the comparator, hysteresis was added in attempt to remove the high frequency oscillations on the comparator output. Designing hysteresis for this purpose was somewhat complicated because both the input signal and the reference (sawtooth wave) were AC waveforms. In the end, adding hysteresis on the LT1016 failed to produce any notable improvement. It was concluded that the LT1016 was too fast for such a noisy circuit, and no amount of noise mitigation would be enough to compensate for the LT1016's precision

Later examinations of the LT1016 datasheet provided numerous clues that this attempt was doomed from the start. ("An inch of wire between the capacitor and the LT1016 can cause problems." "Do not use sockets." "… oscillation is more stubborn and persists well after the output has gone low. This condition is due to stray capacitive feedback from the outputs to the inputs. A 3kΩ input source impedance and 3pF of stray feedback allowed this oscillation." "The villain of this situation is a large output load capacitance. This could be caused by cable driving, excessive output lead length or the input characteristics of the circuit being driven." " The output transitions are initially correct but end in a ringing condition. The key to the solution here is the ringing. What is happening is caused by an output lead that is too long." - excerpts from the application notes of the LT1016 datasheet.) Attempting to use the LT1016 turned out to be a costly waste of time.

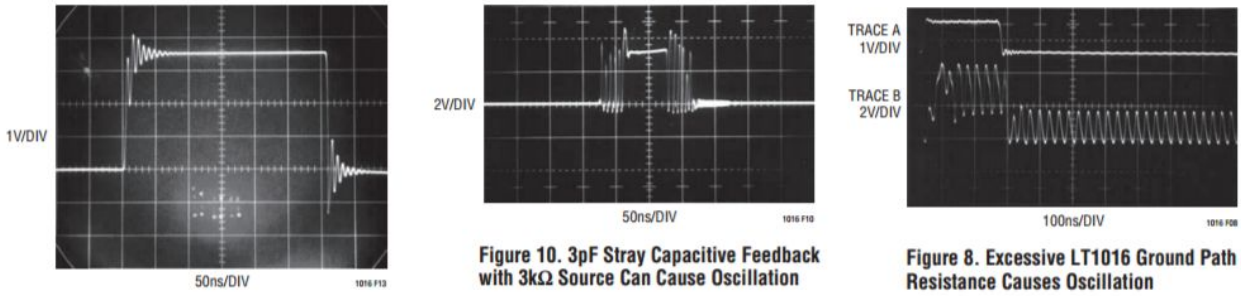

Figure 13. Lengthy, Unterminated Output Lines Ring from Reflections

Figure 10. 3pF Stray Capacitive Feedback with 3kΩ Source Can Cause Oscillation

Figure 8. Excessive LT1016 Ground Path Resistance Causes Oscillation

Figure 21. Figures from the LT1016 datasheet describing various kinds of noise on the comparator output. These resemble some of the signal outputs seen in **Figure 20**.

F. (Successful) PWM and Noise Filtering

The LM311, with a response time substantially slower than that of the LT1016 but still relatively fast (165ns as compared to 10ns), proved the ideal choice of comparator for this circuit. The comparator takes as inputs the filtered audio signal and the 200 kHz sawtooth wave. It is powered with the ±15V supply, with an 11 kΩ pull-up resistor connected to +5 V. The output of the LM311 is a much cleaner output pulse train from 0-5 V.

Some noise compensation remains necessary, so the signal is passed through multiple inversions using a 7404 hex inverter chip. The purpose of digitizing the signal is to sharpen the

edges of the pulse train. More than one inversion is included because of the extreme sensitivity of the circuit to noise. A low-pass filter with a frequency cutoff of roughly 1 MHz is placed in between two of the inversions in order to remove more of the high-frequency noise. The resulting signals, taken from the input and output of the last inversion, are characterized by duty cycles proportional to the instantaneous value of the input audio signal at any given moment. Though the signals are not exactly in phase, as there is an 8-12 ns propagation delay for each inversion on the 7404 hex inverter chip, the gate-driver dead time of 20ns sufficiently compensates for the lag. The final schematic for this portion of the circuit is shown below:

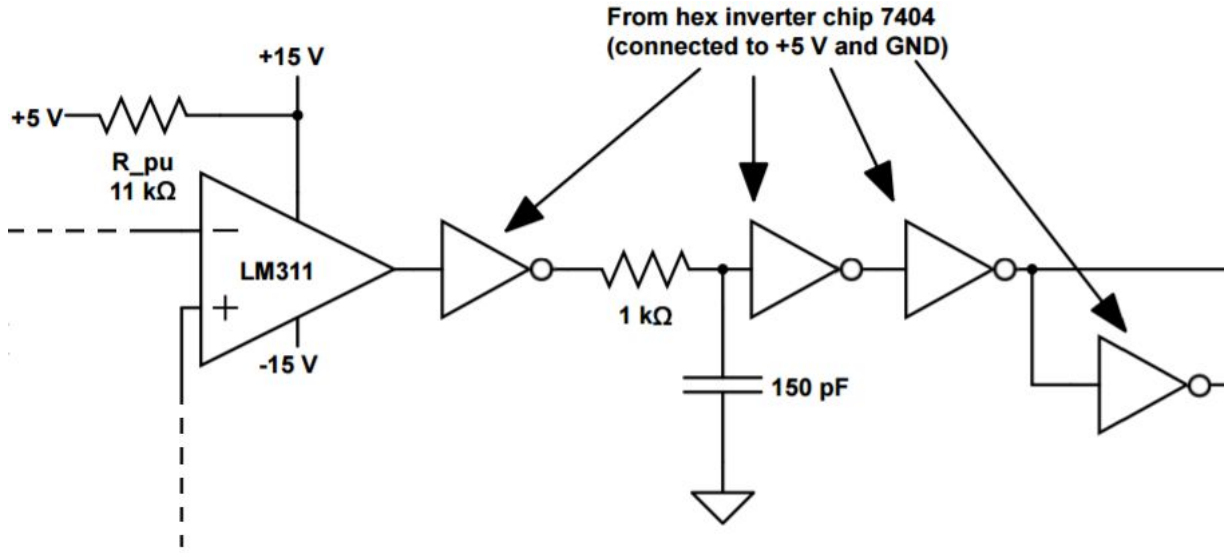

Figure 22. PWM using LM311 comparator, plus noise removal using a 7404 hex inverter.

Figure 23. The outputs of the comparator and noise reduction phase, now looking substantially cleaner.

G. Gate Driver Chip and Power Stage Topology

The two pulse train outputs from the comparator and noise removal circuitry are passed as the high and low inputs to a specialized MOSFET driver, the IR2011. This driver performs several important functions which include level-shifting the high side MOSFET drive such that its gate voltage is referenced to its source, ensuring that the gate voltage is 10-15V higher than the drain voltage, controlling the high side and low side MOSFETs independently to inject dead time, and providing under-voltage lockout. Dead time is time intentionally inserted between the ON states of the high and low side MOSFETs, necessary to avoid creating a low impedance connection between the power rails if both MOSFETs are ON simultaneously; such a low impedance connection, called “shoot-through,” risks thermal failure of the MOSFETs. The gate driver ensures that dead time is low enough to avoid distortion of the output, and avoid operation of the MOSFETs in linear mode (which would lower efficiency). The under-voltage lockout function of the driver is to prevent operation of the MOSFETs in linear mode rather than in saturation or cutoff. This avoids excessive power loss in thermal dissipation which could lead to MOSFET failure.

The IR2011 driver provides the gate voltages for both the high and low side MOSFETs, ensuring that only one MOSFET is ON at a time. For the high side MOSFET, a floating power supply referenced to the switching node is used to drive the gate. When the low side MOSFET is ON and V_S is pulled down to ground, the bootstrap capacitor charges through the bootstrap diode from the V_{CC} supply, thus providing a supply to V_B. A bias voltage referring to the negative bus voltage drives the gate of the low side MOSFET.

Because shoot-through is often encountered during the testing process as a result of noise in the gate driver inputs or incorrect connections to the gate driver, power MOSFETs are necessary. In this circuit, the IR540N power N-channel MOSFET (from Infineon) is utilized. It survived numerous iterations of shoot-through throughout the testing and debugging process, demonstrating its ideality for this application.

The finickiness of the driver and the complexity of its connections, which were not considered carefully while initially designing the circuit, proved to be the next most significant challenge encountered in building the class-D amplifier. Initially, the MOSFETs were connected between +15 V and -15 V, and the COM pin on the driver was connected to -15 V. However, the input signal to the driver was between 0 and 5V with respect to ground, rather than with respect to the COM connection. This meant all inputs to the driver were considered high, which caused the driver output to rail and led to shoot-through on the MOSFETs. After it was noted that the input signal needed to be referenced to COM, level-shifting circuitry was added between the comparator and gate driver. An LT1632 op-amp was used in an adder configuration, adding the comparator output and -15V. The output of this adder was inverted, so another LT1632 was included to invert the signal. This partially worked; however, the signal (noisy as it was) couldn't get quite close enough to the -15V rail to meet the high and low thresholds (<0.7 V above COM for a low, >2.2 V above COM for a high) required by the gate driver.

Attempts to level-shift the input signal to the gate driver were abandoned, and the COM connection was shifted from -15V to ground. As the MOSFETs were now connected between +15 V and ground, it seemed necessary to create a stable +7.5 V supply between the source of the high side MOSFET and the drain of the low side MOSFET. A variety of attempts were made to create such a node, involving potentiometers and voltage regulators (among others); however, none of these could reliably sink *and* source a considerable amount of current (nearly 2A). Thus, during the tests, large amounts of current were fed back into the IR2011 driver chip. Multiple chips were burnt in this fashion.

Upon recognition of the futility of this endeavor, a slight modification on this topology was implemented, in consultation with the typical connection diagram shown below:

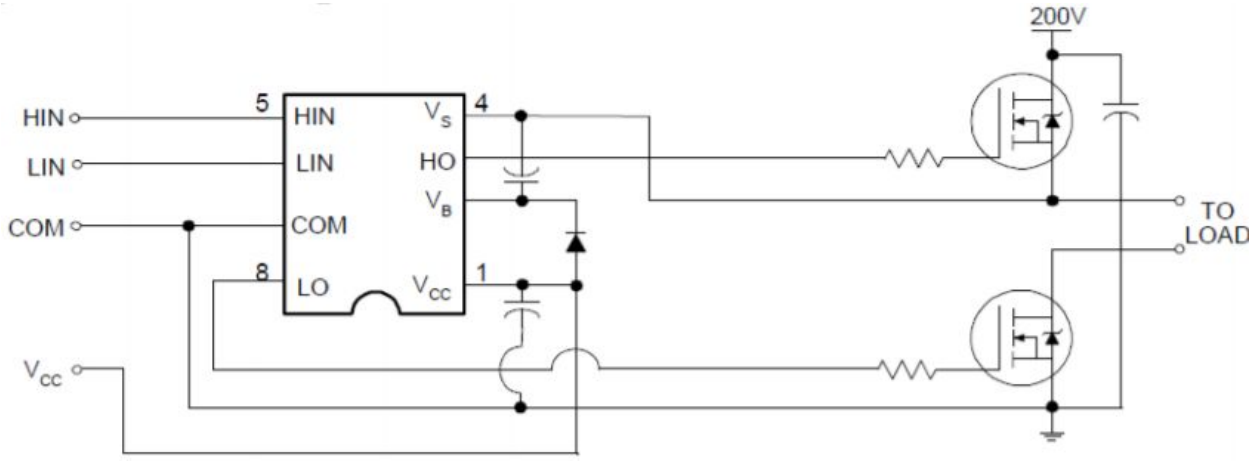

Figure 24. Typical connection of the IR2011 MOSFET driver.

In this topology, the COM pin is connected to ground, and the MOSFETs remain connected between +15V and ground. A 20 Ω resistor (with sufficiently high power rating) is placed between the source of the high side MOSFET and the drain of the low side MOSFET, rather than connecting these pins directly. After replacing the driver chip, this topology was successful. With no burning of the driver chip or shoot-through across the MOSFETs, a pulse train of nearly 15V peak-to-peak that retains its original frequency spectrum can be observed across the resistor.

As a half-bridge topology leads to a sizable amount of DC current in its output, a full-bridge topology was considered and attempted. However, the attempt failed (likely because it was made with a burn IR2011 driver still in place). Following this, a half-bridge topology was once again reinstated, and is used in the final, successful circuit shown below:

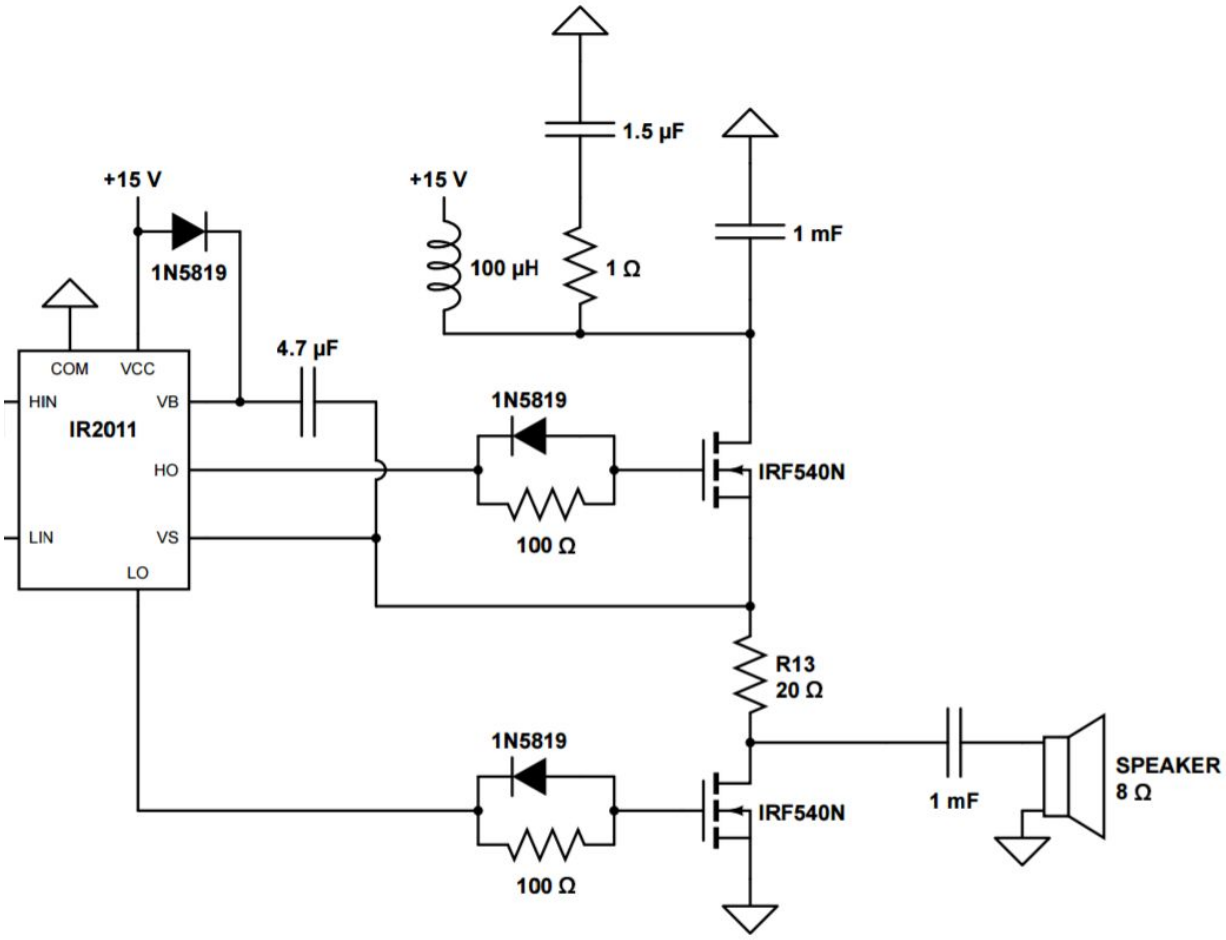

Figure 25. Complete circuit schematic of the driver, power stage, and output to the speaker.

The 20 Ω resistor remains included in the final design as it serves the purpose of lessening the current fed to the speaker. Connected at the drain of the low-side MOSFET is a 1000 µF DC blocking capacitor, which serves to remove the DC component of the output signal (this is necessary to drive a speaker). In place of the designed output filter, the built-in high-frequency filtering of the speaker itself is relied upon to retrieve the desired audio output from the amplified pulse train. Tests with a speaker demonstrated that this worked amply.

A power supply filter consisting of the parallel combination of an inductor, capacitor, and a series resistor and capacitor is included between the +15V supply and the drain of the high side MOSFET. This is to ensure that, as large amounts of current are intermittently drawn through the MOSFETs, no noise is injected back into the power rails to disrupt other stages of the circuit.

Some harmonic distortion was observed on the output, suggesting that there was still some noise left to be removed from the circuit.

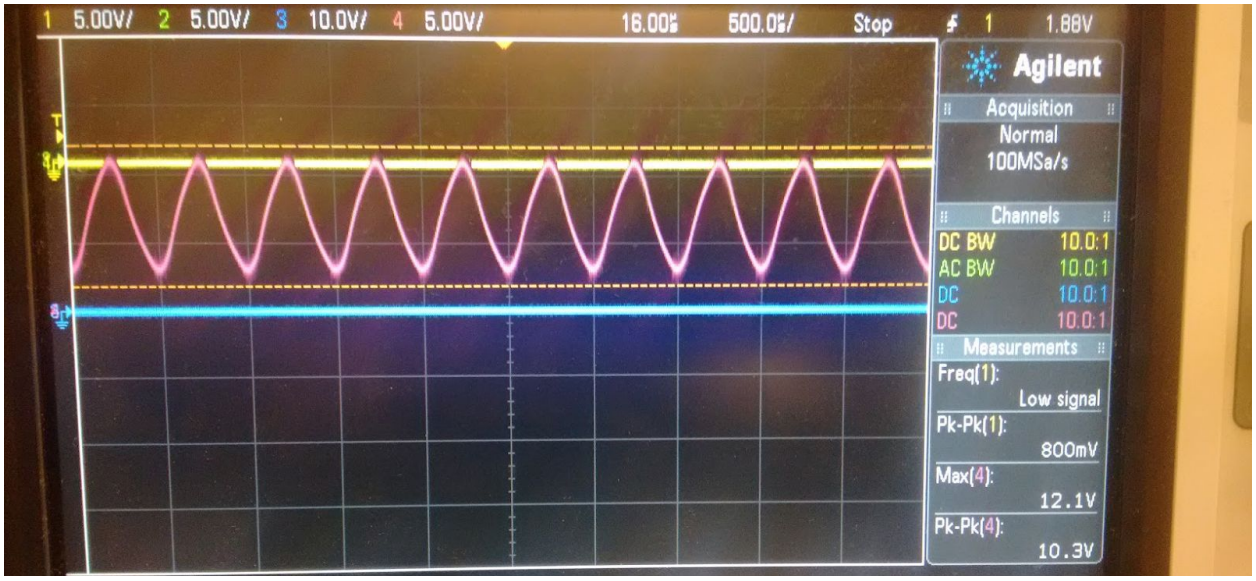

Figure #26. High harmonic distortion in output.

More bypass capacitors were added, and wiring for the power rails was reworked such that multiple wires were connected to the bench power supply itself. Star grounds were also used effectively. The combination of these noise mitigating measures successfully removed the harmonic distortion from the output. The output of the final circuit is a clean replica of the input audio signal, amplified to about 15V peak-to-peak.

With the class D amplifier connected to the remainder of the laser harp system, audio signals of any waveshape (triangle, square, or sinusoidal) and frequency are successfully amplified to drive the speaker. The sound quality of the laser harp's music is excellent.

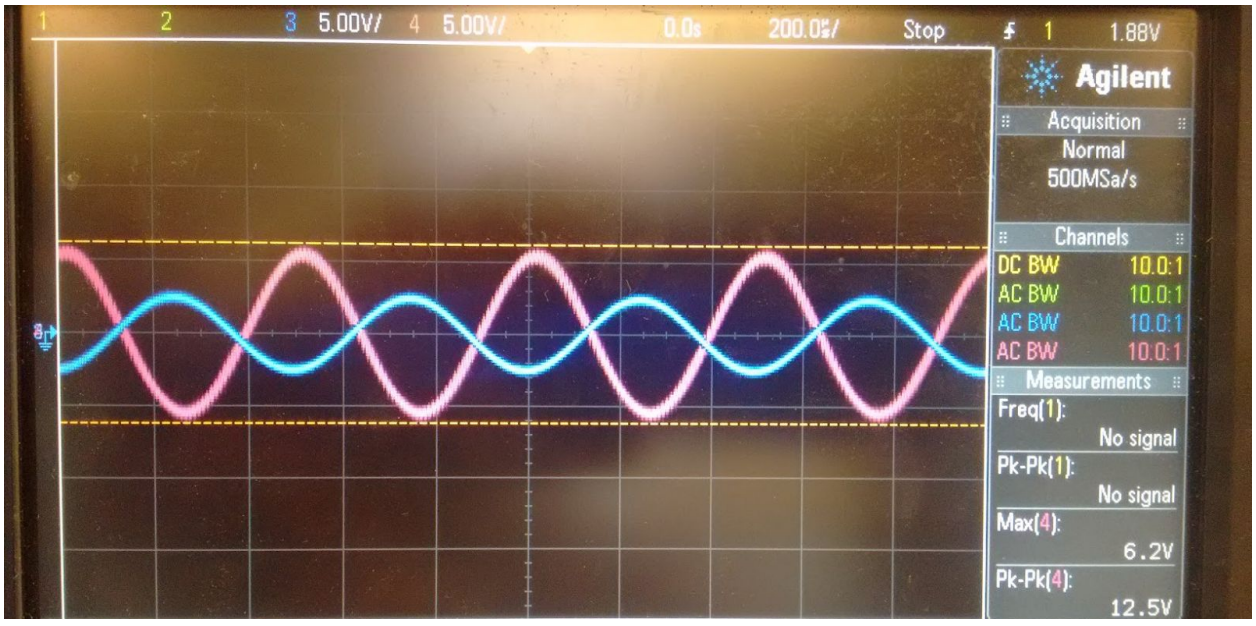

Figure 27. Oscilloscope shot taken while probing the input audio signal (blue) and the output audio signal (pink). Note the cleanness and lack of distortion in the output!

H. Class AB Amplifier

Because of the significant complexity involved in building a working class D amplifier, a class AB amplifier was also constructed as a backup. For this class AB amplifier, an LF356 opamp is used to amplify the audio input signal (which is filtered similarly as in the class-D amplifier). A push-pull output stage is used to provide enough current to drive the speaker load. The 1N914 diodes and 1.1 kΩ resistors are present to bias the output transistors such that crossover distortion is eliminated. The 5 Ω emitter resistors are there to stabilize bias current with changes in operating temperature. This design is nearly identical to the class AB amplifier design from lab 5 of 6.101. The class AB amplifier was indeed able to drive the speaker at high volume with minimal distortion of the audio signal.

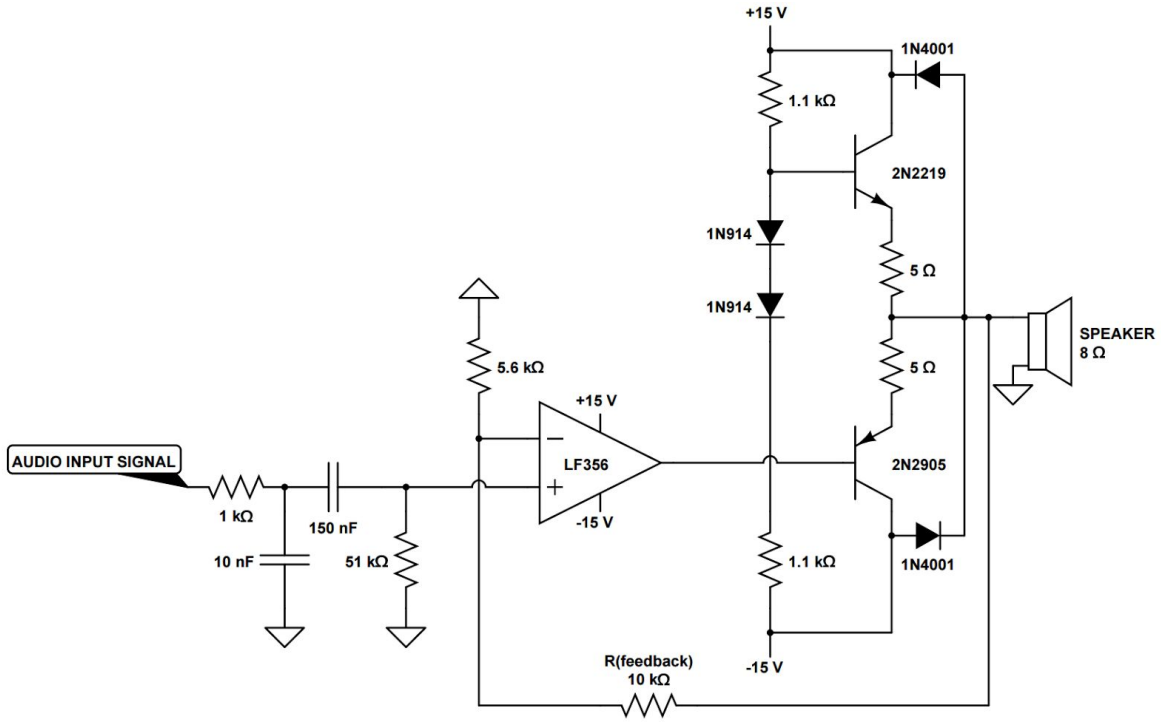

Figure 28. Circuit schematic of the class AB amplifier.

I. Final Thoughts on the Class D Amplifier

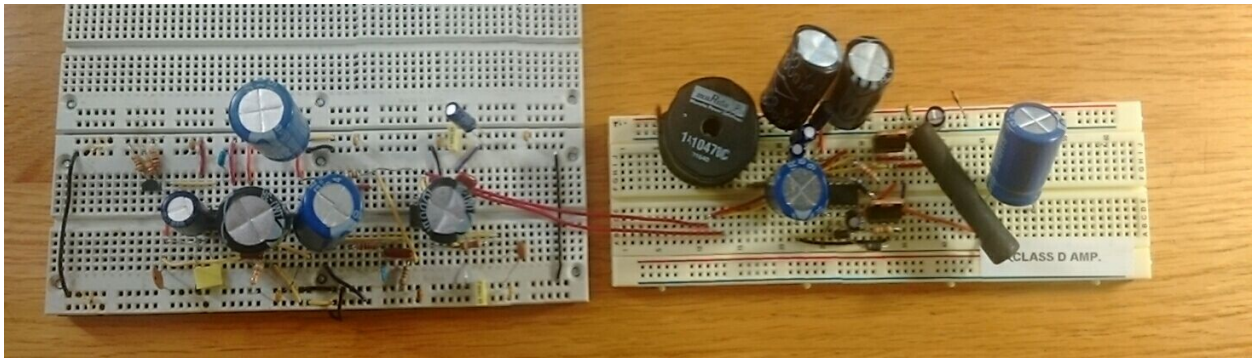

Figure 29. The completed, working Class D amplifier.

Though it took longer than anticipated to complete, the final class D amplifier provided excellent sound fidelity. Though limited to 15V peak-to-peak because of the power supplies used, this was sufficiently loud on the chosen speaker. The class D amplifier displayed a notable increase in efficiency compared to the class AB amplifier (demonstrably, the transistors in the class AB would warm up with high gain, whereas the MOSFETs in the class D amplifier remained cool.) Unlike with the class AB amplifier, no distortion was noted regardless of input volume or frequency, which corroborates the fact that the class D amplifier has a much wider frequency power bandwidth.

While the final result was impressively successful, the process to construct a working class D amplifier was beleaguered by numerous sources of error. This circuit was exceedingly sensitive to a variety of noise sources, including (but not limited to) the switching speed of the MOSFETs, the MOSFET packaging, length of wire leads, stray capacitance on the breadboards, the locations and quality of bypass capacitors, noise on the input audio signal, and power supply noise. Dead-time from the driver also makes a significant difference, as smaller dead-time leads to lower distortion, but too narrow dead-time causes shoot-through. The three most important steps to building a working class D amplifier are to minimize sources of noise in the switching stage, maximize noise immunity in the analog stage, and minimize noise coupling from the switching stage to the analog stage (source: www.irf.com). In hindsight, choosing to create the amplifier on a PCB rather than a breadboard would be a wiser decision.

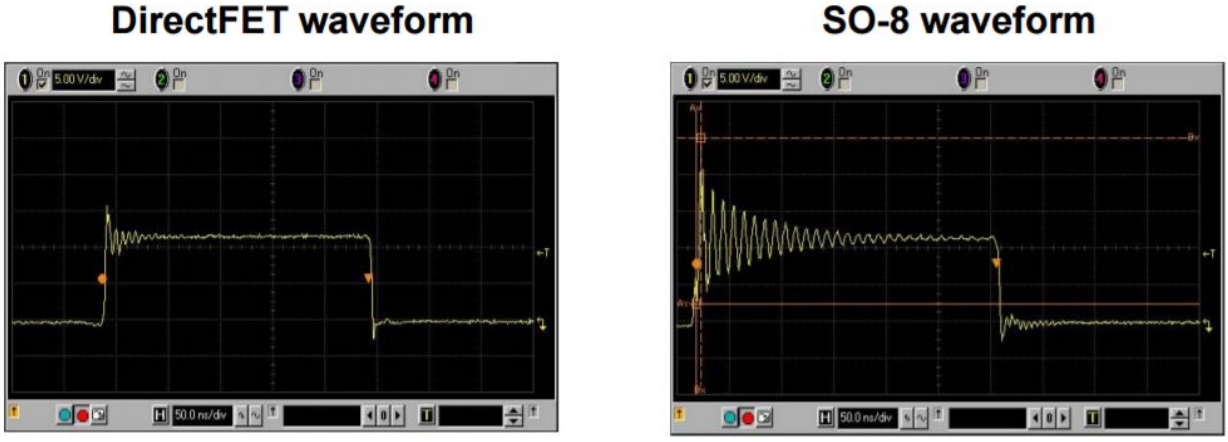

Figure 30. Ringing caused by inductance from the MOSFET packaging. Even this can make a sizable difference!
(Source: http://www.irf.com/product-info/audio/classdtutorial2.pdf)

Other lessons learned include avoiding overly precise components because of their ability to pick up on high frequency noise, awareness of extreme noise on the lab power supplies, the need to buffer between stages to prevent loading on input and output impedances, and various tricks to compensate for noise in different parts of the circuit. The importance of careful breadboarding was also highlighted by the fact that splitting the circuit onto two breadboards,

adding different wires from the power supply itself, and using star grounds made a significant difference in output quality.

Overall, the theory of a class-D amplifier is itself simple. But there is a tremendous difference between theoretical design, and implementation of this circuit in the real world. As a resource from International Rectifier mentioned, the circuit works well if "key components are carefully selected and the layout takes into account the subtle, yet significant impact due to parasitic components." Through the process of building this amplifier, I learned how to deal with non-idealities, and how to identify sources of error and devise inventive means of addressing these issues – important skills needed for real-life engineering.

VII. Conclusion

Integration of the subsystems proved fairly simple due to the fact that the system is designed to be highly modular. Once each subsystem was shown to function well in isolation, the quality of the signal processing end to end was surprisingly accurate and each effect was clearly audible. A majority of the challenges faced were the result of physical non-idealities in the devices being used. As a result, we all experienced valuable technical growth by frequently being challenged to understand theoretical expectations for our signals, identify causes of distortion, and devise robust solutions.

It is our goal over the summer to put all of our circuitry on PCBs and to rebuild a harp that is more aesthetically pleasing. We are excited by the range of possibilities that this particular analog circuit is capable of achieving, and we believe it would be incredibly rewarding to see our complete vision come to fruition. As engineers, it is vital to recognize what is feasible to achieve within the time allotted, but seeing as how we were very close to implementing a number of additional effects, it seems worthwhile to revisit this project in the future.

VIII. Acknowledgements

We would like to extend our sincerest gratitude for the guidance and support of Professor Gim Hom, without whom this project would not have been possible. The instruction we received over the course of the semester and as we struggled to begin this project provided us with the insight and confidence that were vital to producing a successful device in a limited timeframe. We would also like to thank Jason Yang for his invaluable help in building the Class D amplifier. He and Yanni Coroneos were incredibly generous with their time, knowledge, and encouragement throughout the entire process. In addition, we would like to express our appreciation for Dave Custer, who provided thorough feedback on our communication, and who spent considerable time and effort to document this project on video. The growth we have seen in ourselves as a result of 6.101 and this project in particular has been remarkable, and we are deeply grateful to the entire class staff.

IX. Works Cited

For the synthesizer:

1. Horowitz, Paul, and Winfield Hill. "The Voltage Controlled Oscillator." *The Art of Electronics*. 3rd ed. New York, NY: Cambridge UP, 2016. 267-68. Print.
2. https://www.allaboutcircuits.com/projects/diy-synth-series-vco/
3. http://schmitzbits.de/expo_tutorial/index.html
4. http://www.schmitzbits.de/adsr.html
5. http://www.birthofasynth.com/Thomas_Henry/Pages/VCO-1.html
6. http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=1050748

Data sheets for the LM13700 and LM6132 and all other devices were taken from the websites of Texas Instruments and the 6.101 class website.

For the audio amplifier:

1. https://en.wikipedia.org/wiki/Class-D_amplifier
2. https://www.allaboutcircuits.com/projects/how-to-build-a-class-d-power-amplifier/
3. http://www.irf.com/product-info/audio/classdtutorial.pdf
4. http://www.irf.com/product-info/audio/classdtutorial2.pdf
5. http://www.irf.com/technical-info/refdesigns/iraudamp1.pdf
6. 6.101 Lab 5

Datasheets for the LT1016, LM311, LT1632, SN74S04 Hex Inverter, and the IR2011 were taken from the websites of Texas Instruments, Linear Technology, International Rectifier, and the 6.101 class website.

Reference materials from International Rectifier (links 3-5 above) were particularly full of relevant information, as a class-D amplifier is the most commonly discussed application of the IR2011 gate driver.