

The Laser Microphone

Authors: Mark Chounlakone, Julian Alverio

Filtering Section: Justin Tunis

Abstract

The human voice can generate sound waves in the range of 300 Hz to 3400 Hz [8]. These sound waves vibrate nearby objects, making it possible for an analog electronic device to convert these vibrations into an audio signal[1]. One way to accomplish this conversion from movement to audio is to use a "laser microphone", which reflects a laser off the vibrating object and uses a receiver to capture the laser's reflection. The reflection of the laser gets deflected as vibrations shift the surface of the vibrating object. Therefore, if a receiver takes in the oscillating laser signal from a fixed location, the receiver will detect the laser deflections caused by the vibrations that were originally produced from an audio signal. The receiver can then filter and amplify this signal, and output it as audio. Through this process the laser microphone effectively reproduces the audio that induced the object's vibrations.

The laser microphone is able to reproduce audio detected from a vibrating surface with relatively high accuracy: less than 8% distortion. As an additional feature, the laser microphone is also able to transmit audio via amplitude-modulated laser signal, capture the laser signal, and output the audio. Thus by using a laser based system that captures oscillations in the position of the laser, the laser microphone is able to accurately reproduce both the audio that induced an object's vibrations and audio transmitted via laser communication.

Introduction

Human speech is composed of sound waves that vibrate the objects nearby. The goal of this project is to create an analog electronic device that can read in a signal from a vibrating object and reproduce the audio that originally caused it to vibrate. This project achieved this goal with the laser microphone, which is a device that shoots a laser at the vibrating surface and captures the reflection with a receiver. The vibrating surface translates the position of the laser's reflection, thus allowing the receiver to output the laser input signal through speakers as audio. Thus the laser microphone can reproduce the original audio signal that induced the object's vibrations through a speaker.

To function as desired, the device needed to be able to filter out all frequencies in the input laser signal outside of the human speech frequency range. Secondly, the device

needed to be able to output audio through a class D amplifier using a switching frequency of at least 400 KHz. Additionally, the device needed to be able to support direct laser-to-receiver audio communication. In order to meet these functional requirements, the electronic part of the system is composed of four stages. The first stage is laser transmission, where the laser will be amplitude-modulated if the device is transmitting directly from laser to receiver. The second stage is a laser reception stage, with a reverse-biased photodiode in a transimpedance amplifier. This transimpedance amplifier offers high common mode gain rejection and high gain. The third stage is a high order signal filter. The last stage is a class D audio amplifier, which switches at a frequency of at least 400 KHz and outputs audio at a maximum voltage swing of at least 8 Vpp.

While the main focus of the project is exploring the design of circuitry used to reproduce speech from a captured signal, some mechanical characteristics of the system are also key to the system's success. There are two key mechanical properties that the laser microphone depends on. The first is the acoustic properties of objects which generate vibrations from sound waves. The second is the reflectivity of the object's surface.




Figure 1
Depiction of how a laser can be used to detect mechanical deflections.

Figure 1 shows both of these mechanical properties at work. As the surface shown on the left shifts, the reflected beam of the laser changes position. Since this project was focused around the electronics, we chose to simplify the mechanical challenges by building a few setups with ideal mechanical properties. These setups are described in the system design.

System Design

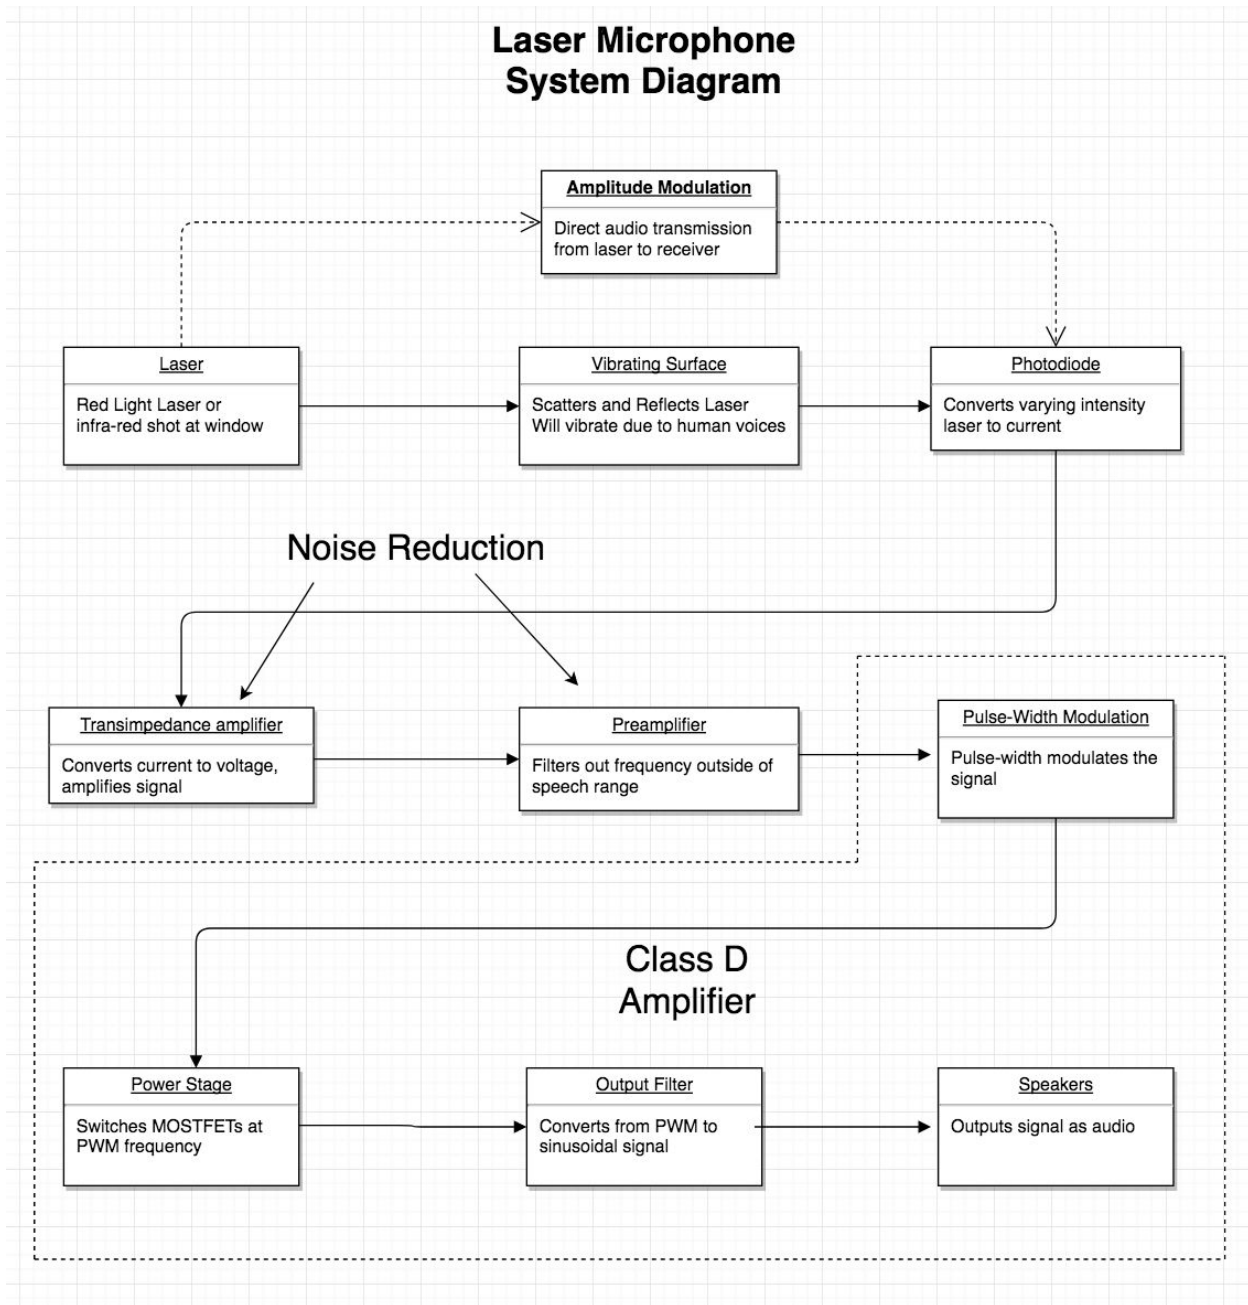


Figure 2
A flow diagram of audio reconstructed from the laser microphone

The laser microphone is a very sensitive device which must detect minute surface deflections caused by audio. The detection is done using a photodiode and a transimpedance amplifier. The signal observed by the photodiode is a noise injected version of the audio signal of interest. Noise is injected through mechanical vibrations, ambient light, and the photodiode itself. To produce a coherent reconstruction of speech, the laser microphone filters out the noise in the transimpedance amplifier and the preamplifier. Once the audio signal has been cleaned, it is sent through the class D amplifier which then plays back the signal on a speaker.

Laser Transmitter

The first stage of the laser microphone is the transmitter. For this project, there were two modes of transmission, each built separately. The first mode was reflected transmission. This is the mode that has been discussed the most so far. The laser is shot at a vibrating surface and the receiver captures the reflection. The circuit was simply a 5mW red laser pointer connected directly to a 5v supply which is not shown.

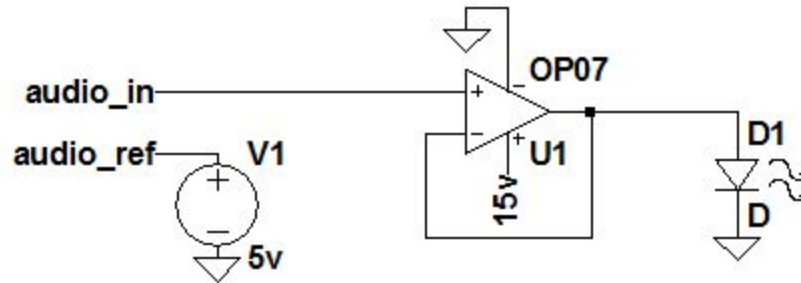


Figure 3

A depiction of the circuit used to directly communicate to the receiver with the transmitter

The second mode of laser transmission was a direct communication mode. In this regime, the laser is modulated with an audio signal and the receiver captures the beam directly. The circuit used for direct communication is shown in Figure 3. The laser diode is depicted as an LED, D1. The power output of the laser had to be modulated with the signal being sent so that the circuitry on the receiving end could reconstruct the signal. To do this an audio signal centered around 5V was sent through a simply op amp buffer which drove the laser diode. By driving the laser diode with a modulated voltage around it's operating voltage, the power is modulated. While this circuit works, it is not the most elegant way to modulate the laser. Laser diodes are fairly sensitive devices and will die without a proper current regulating circuit. This laser diode was not a laser pointer so it used a as resistor a current regulator rather than a proper laser driver. A smarter design would use a voltage controlled current regulator to modulate the current at a fixed voltage.

Laser Receiver

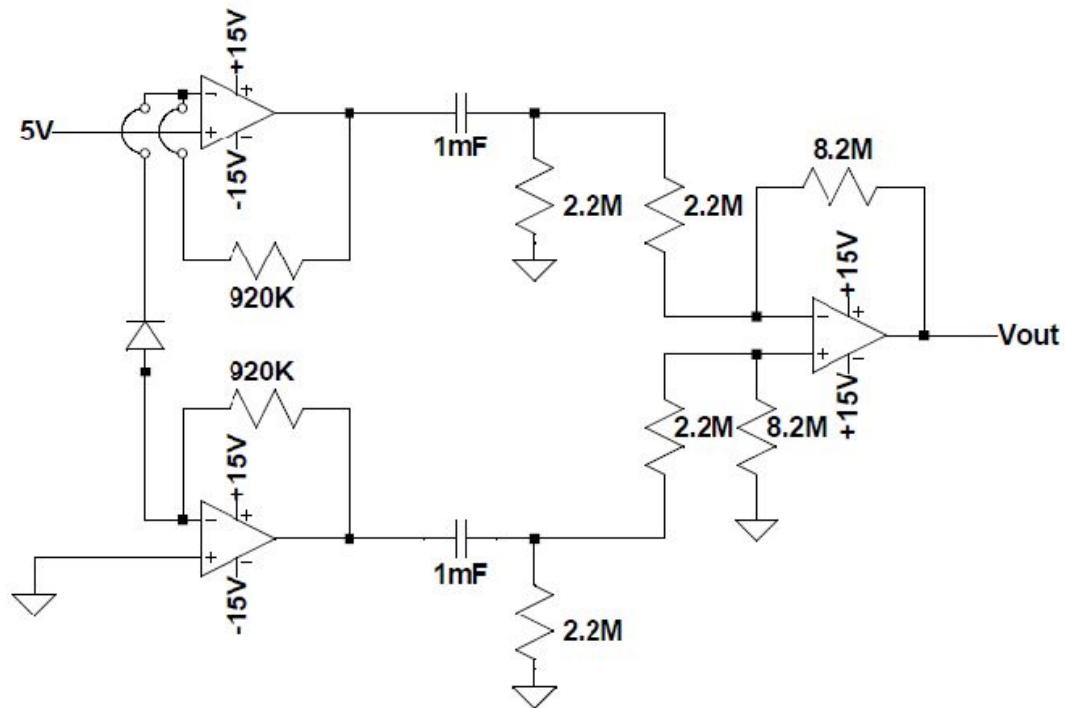


Figure 4

A depiction of the photodiode and transimpedance amplifier used to receive the signal.

The second stage of the device was the laser receiver. This stage was designed to capture the reflected laser signal and turn it into a voltage. The receiver was composed of a photodiode and a transimpedance amplifier (TIA). In Figure 4, the photodiode is depicted as a normal diode between the inverting inputs of two op amps. The photodiode is the sensor that captures the laser signal. In reverse biased mode, the photodiode has a reverse current which is linearly related to the light irradiance. Once a reverse current is generated, the TIA the current into a voltage. The TIA topology shown in Figure 4 incorporates a differential amplifier. The differential amplifier gets rid of interference caused by ambient light using common mode rejection. The reason the is necessary is explain later in this section.

This TIA topology can be described in two parts. The first part is composed of two op amps to create two inverted signals. The non-inverting inputs of the op amps are set to 5v and 0v such that the photodiode is reverse biased. When light generates a reverse current in the photodiode, the current must pass through both resistors. Ohm's law states that $V = IR$, so the current passing through the resistors generates proportional voltage signals. The photodiode's datasheet states that the reverse current is on the order of tens of microamps, so 920K resistors were selected to gain the current. The output of the bottom op amp is the inverse of the top op amp's output with some dc offset. These two signals are then ac coupled by the 1mF/2.2M capacitor resistor pair, which gets rid of the offset. The capacitor and resistor pairs are ac couplers which also

act as high pass filters. I just choose a frequency, .07 Hz, below the vocal range as a corner frequency. The second stage of the transimpedance amplifier is a differential amplifier which was previously described.

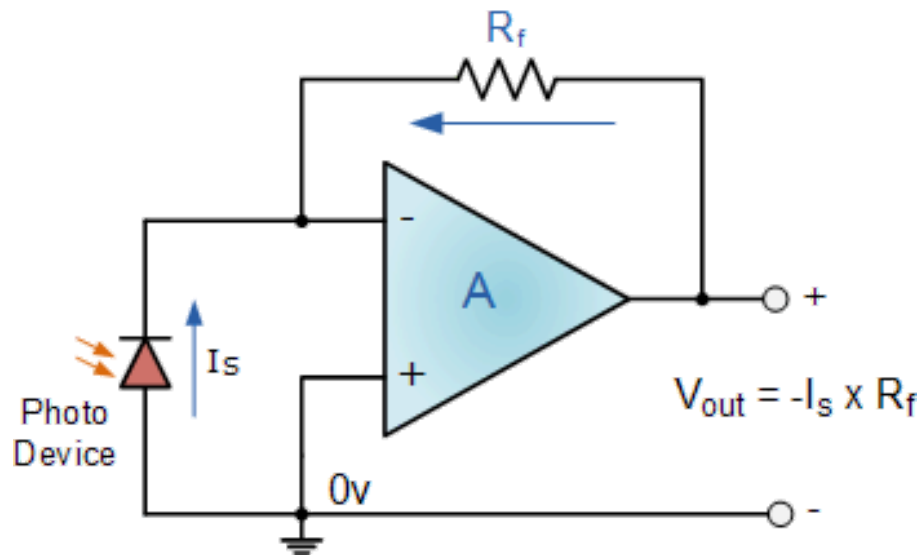


Figure 5

A depiction of the photodiode and the original transimpedance amplifier design.

So why was the differential topology chosen? Originally we had a different design. This one was more simple but without the second stage of the amplifier, ambient light was interfering with the signal. The difficulty with the ambient light did not come from the steady state offset, but from changing ambient light. Just moving around the photodiode would change the amount of ambient light detected, causing the entire signal to shift. Originally we built a box to isolate the photodiode from ambient light, but alignment became much more difficult. Professor Jacob White proposed this new topology as an elegant solution to the ambient light issue. So to reduce the amount of time spent fixing a mechanical issue, the design was expanded to compensate for the interference and the box was no longer necessary. Another nice attribute of the circuit above was that it doubled the signal we received since the current flowed through two resistors. This topology is also able to add more gain in the differential amplifying stage.

Signal Filtering

The signal from the transimpedance amplifier contains noise from sources like AC hum from equipment, low frequency rumbles, and mechanical vibrations. Amplifying the signal at this point would produce a distorted output and snooping would be infeasible. In order to cope with this we have included a voice frequency band-pass filter to attenuate noise outside of the frequency range of human speech, which is from 300 Hz to 3 kHz.

The circuit consists of a 3rd order high-pass filter cascaded with a 3rd order low-pass filter. Each of those components consisted of a first-order RC filter cascaded with a

Sallen-Key topology second-order filter. The actual circuit had measured cutoff frequencies of 315 Hz on the high-pass filter and 3050 Hz on the low-pass filter, which is right around the vocal-range we were aiming to isolate. Additionally, the circuit provided a measured -60 db/decade attenuation outside of the passband which translates to a -1000x gain, meaning noise significantly outside of the 300 Hz to 3 kHz range was attenuated to a negligible level.

We also experimented with higher-order filters as well as adding gain with the filters. A fifth-order bandpass filter that had been constructed was difficult to tune to the proper cutoff frequencies because a large number of components were involved. Most of the resistors and capacitors in lab have a 5% margin of error, which adds up quickly when building a module with many of these components. While it would certainly be possible to have the circuit operational, the time-to-benefit ratio we would have received from the extra work would not have paid off. Adding gain to the system was also a little tricky as it ended up injecting noise back into the system and created a distorted output. One work-around would have been to include a simple, non-inverting op-amp configuration with gain after the output of the filter; however, this became unnecessary as the signal from the transimpedance amplifier was large enough that no extra gain was needed from this part of the system.

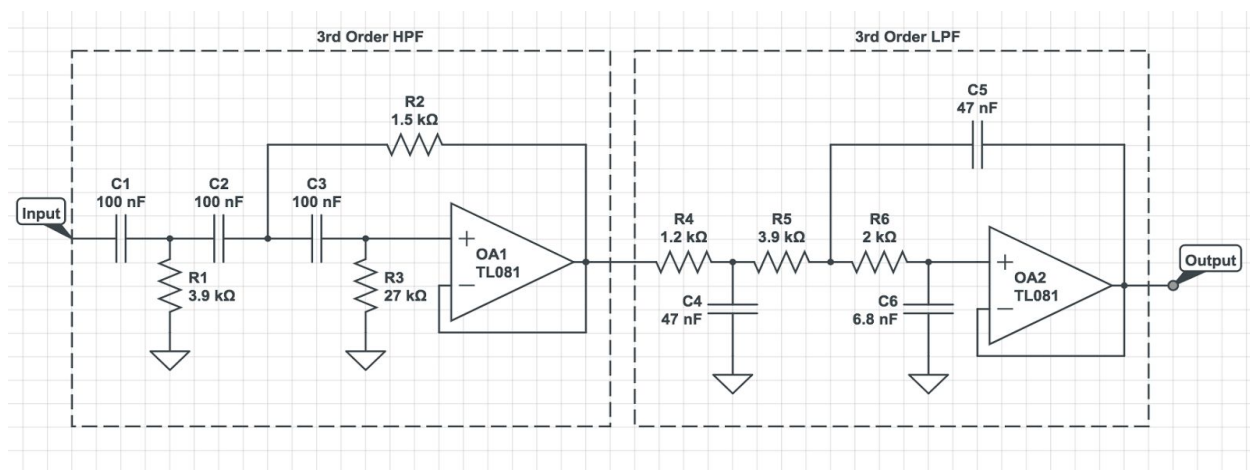


Fig 6. Vocal Band-Pass Filter

Audio Amplifier

In order to output the audio signal from the filter stage, the laser microphone uses a full bridge class D amplifier. Below is a diagram of the amplifier:

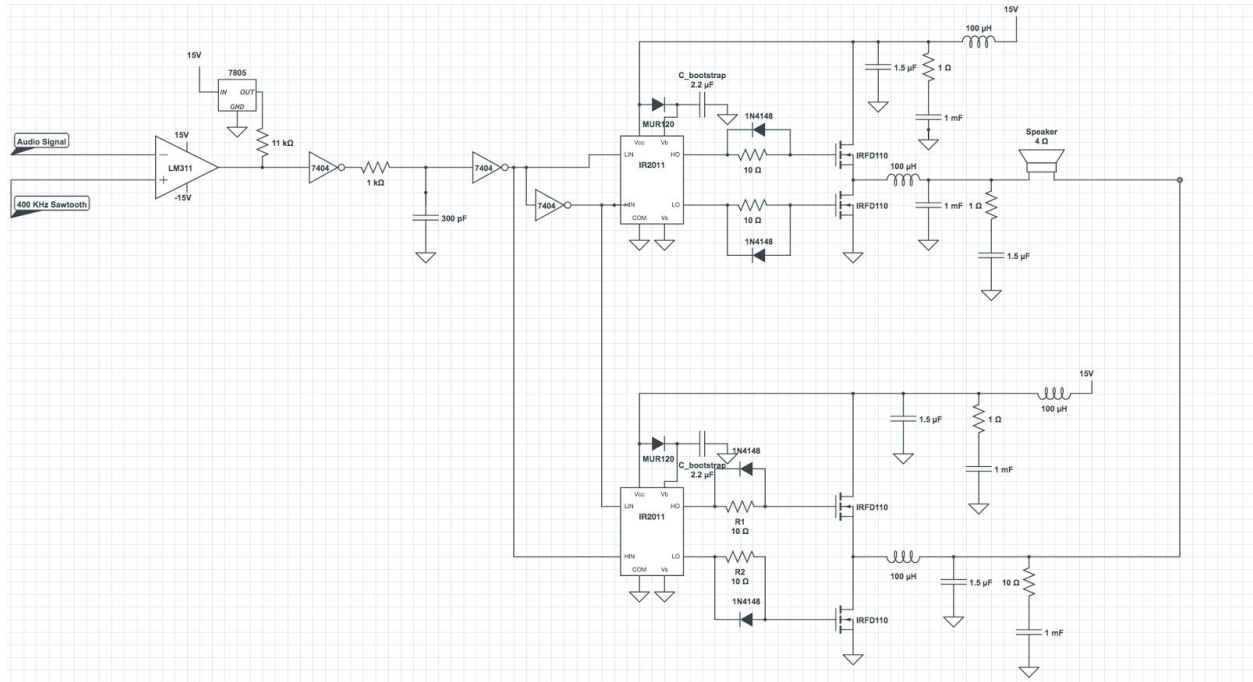


Figure 7: Full-Bridge Class-D Amplifier

The implementation operates in four stages: the pulse-width modulation stage, the gate driving stage, the power stage, and the output stage. Each stage presented a unique set of design challenges.

Pulse-Width Modulation Stage

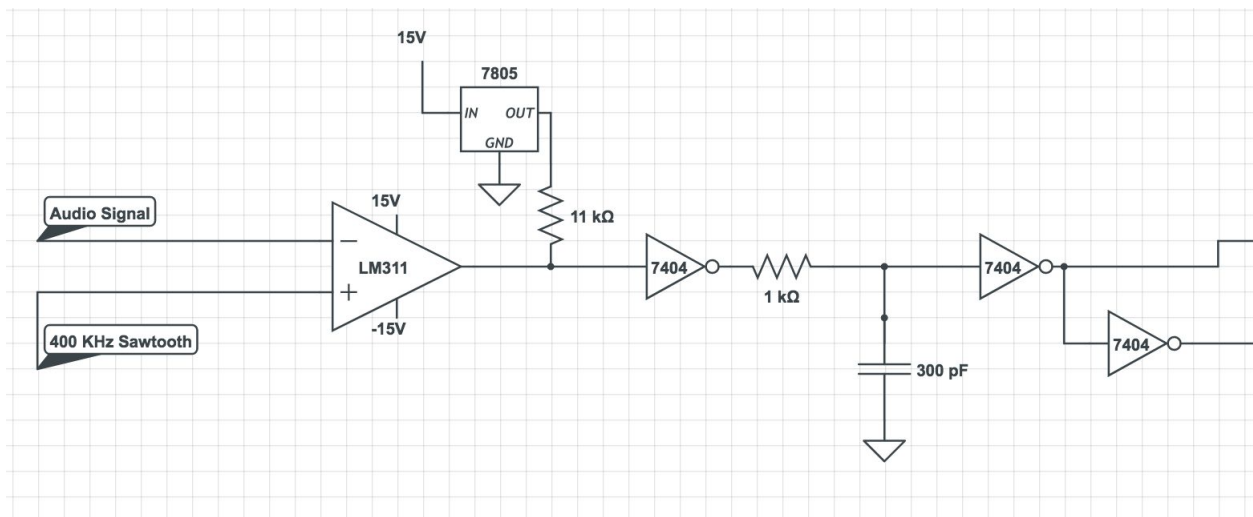


Figure 8: Pulse-Width Modulation Stage

The amplifier converts from a sinusoidal audio input signal to a pulse-width modulated signal by comparing the audio input with a sawtooth wave. The frequency of the sawtooth wave determines the "switching frequency", or the frequency with which the MOSFET's turn on and off to produce audio output. In order to meet the specifications,

the sawtooth had to have a frequency of at least 400 KHz. In the next section is discussion of the sawtooth generation design using a comparator.

In order for the pulse-width modulated signal to be able to activate the gate driving stage, the pulse-width modulation stage had to produce a modulated output, as well as its complement, and both of these signals had to be very clean.

By providing an LM311 comparator with a sawtooth and audio input, the output of the LM311 was the pulse-width modulated signal. To ensure the signal output signal was quite clean, laser microphone passes the comparator output through an inverter, then through a low pass filter, and then through another inverter. The low pass filter acts as a "glitch filter", to eliminate any high frequency noise present in the signal. This high frequency noise may come from unwanted noise in the comparator inputs causing unwanted, high-frequency switching. The inverter steps force the signal to be at a clean 5V or 0V at all times, as opposed to something in between. This, in addition to decoupling capacitors, allows the microphone to attain a very clean signal, as shown below in figures 9 and 10. The PWM stage used an additional inverter to produce the complement of the original PWM signal.

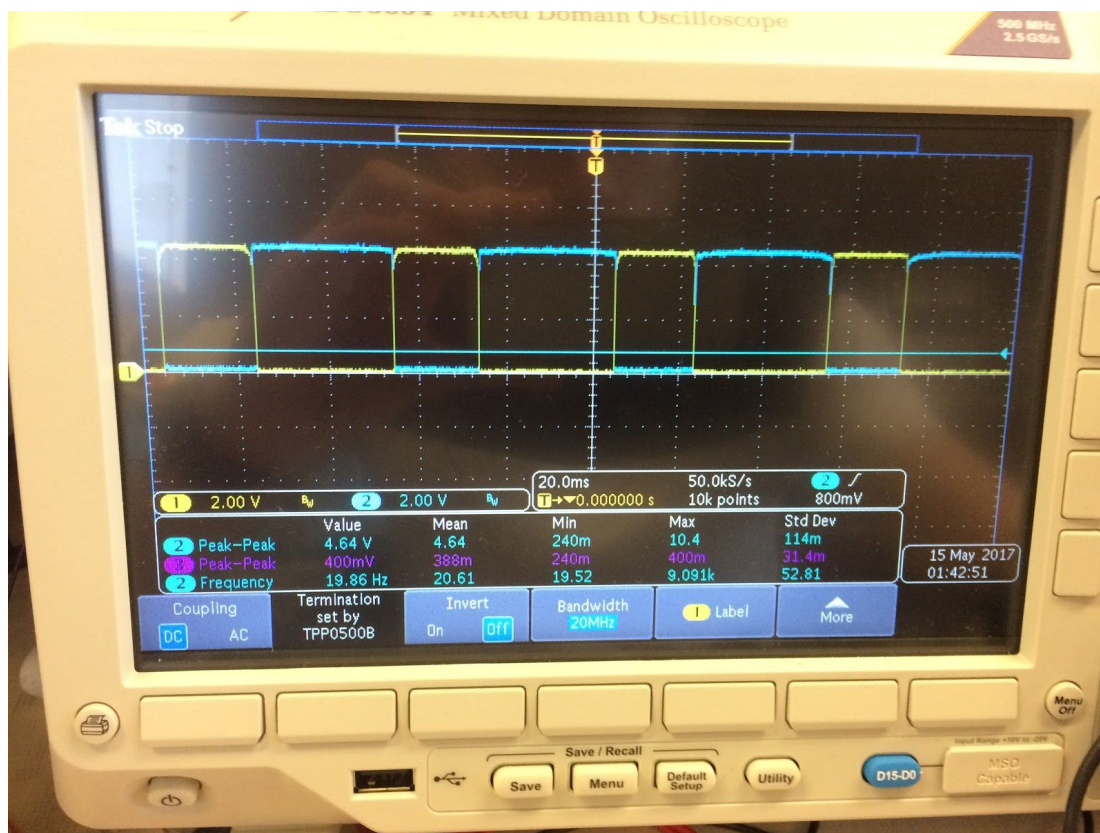


Figure 9: Oscilloscope trace showing the PWM output and its complement. The blue cursor shows the gate driver's "logic low" cutoff

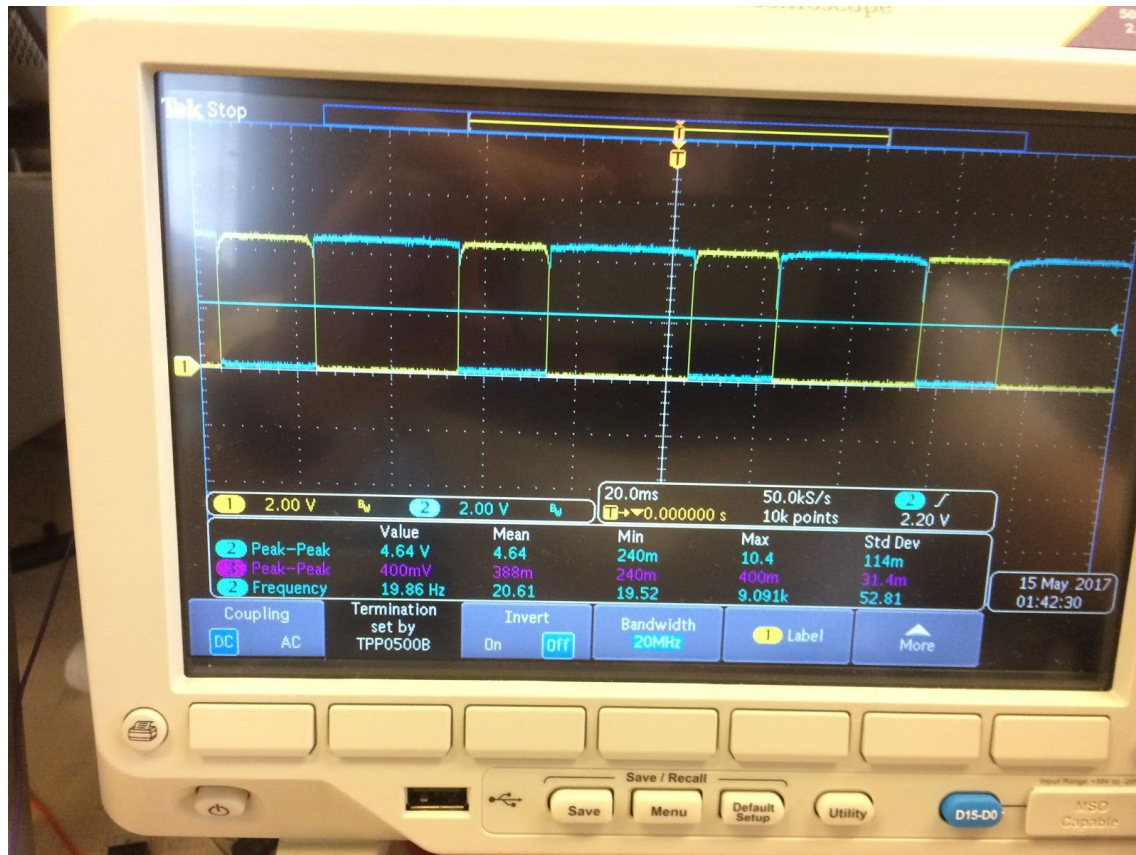


Figure 10: Oscilloscope trace showing the PWM output and its complement. The blue cursor shows the gate driver's "logic high" cutoff

Figures 9 and 10 demonstrate a very clean output of the pulse-width modulation stage, where the logic high is well above the logic high cutoff and the logic low is well below the logic low cutoff. This output then fed into the next stage: the gate driving stage.

This stage in particular presented a number of design challenges. One key objective was produce two clean, complementary PWM outputs without causing so much of a delay between the two signals that it causes shoot-through with the MOSFET's in the power stage. The original solution to this was to use a complementary comparator. The LT1016 produced a complementary output with minimal delay, and seemed like the ideal solution. However, the chip proved to be simply too sensitive to provide a clean output and the implementation was unable to use the chip altogether. It turned out that, even a small stray capacitance could cause the output to be very noisy. After the implementation no longer used the LT1016, the new designs were able to produce a very clean signal.

Gate Driver Stage

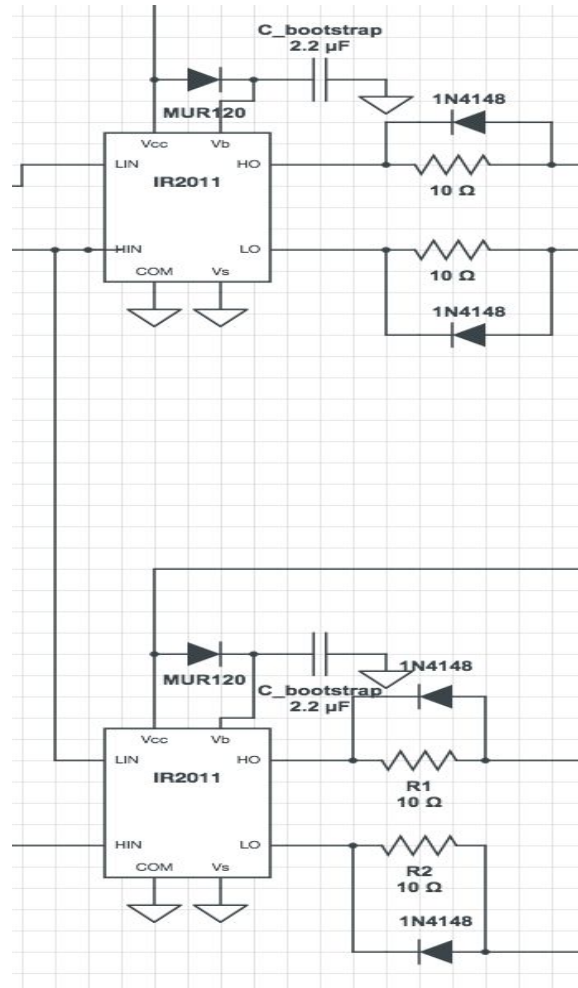


Figure 11: Gate Driver Stage

The gate driver stage takes inputs from the PWM stage and turns the MOSFET's on and off very rapidly. The gate driver is able to switch the MOSFET's quickly by pumping a large amount of current into the gate, in order to overcome the gate capacitance. One key design consideration is the generation of dead time. If the high and low side MOSFET's turn on and off without any delay in between, shoot-through can occur. Shoot-through is when both MOSFET's are on at the same time. This can damage the MOSFET's and will cause distortion in the signal output to the final stage. In order to avoid shoot-through, there must exist a dead time in between the two MOSFET's turning on; a time in which neither MOSFET is on. Though the gate driver does have some dead time generation, the MOSFET's will still experience shoot through without additional dead time generation. In order to generate additional dead time, the laser microphone utilizes a schottky diode in parallel with a small resistor. This pair causes the turn-on time to be slightly delayed due to the RC time constant the resistor creates, while the turn off time is not delayed. The resistors required particular attention -- too small of a resistor does not generate sufficient dead time and allows for shoot-through, while too large of a resistor generates too much dead time.

Since the gate driver switches very rapidly and draws a relatively large amount of current, the laser microphone utilizes a passive second order low pass filter in order avoid injecting high frequency noise into the power supply. The filter is the same used for the high side MOSFET and will be discussed in the next section.

The gate driver stage produces a very clean output to the MOSFET's. In figure 12 below you will see the gate driver input compared to the MOSFET gate voltage on the output.

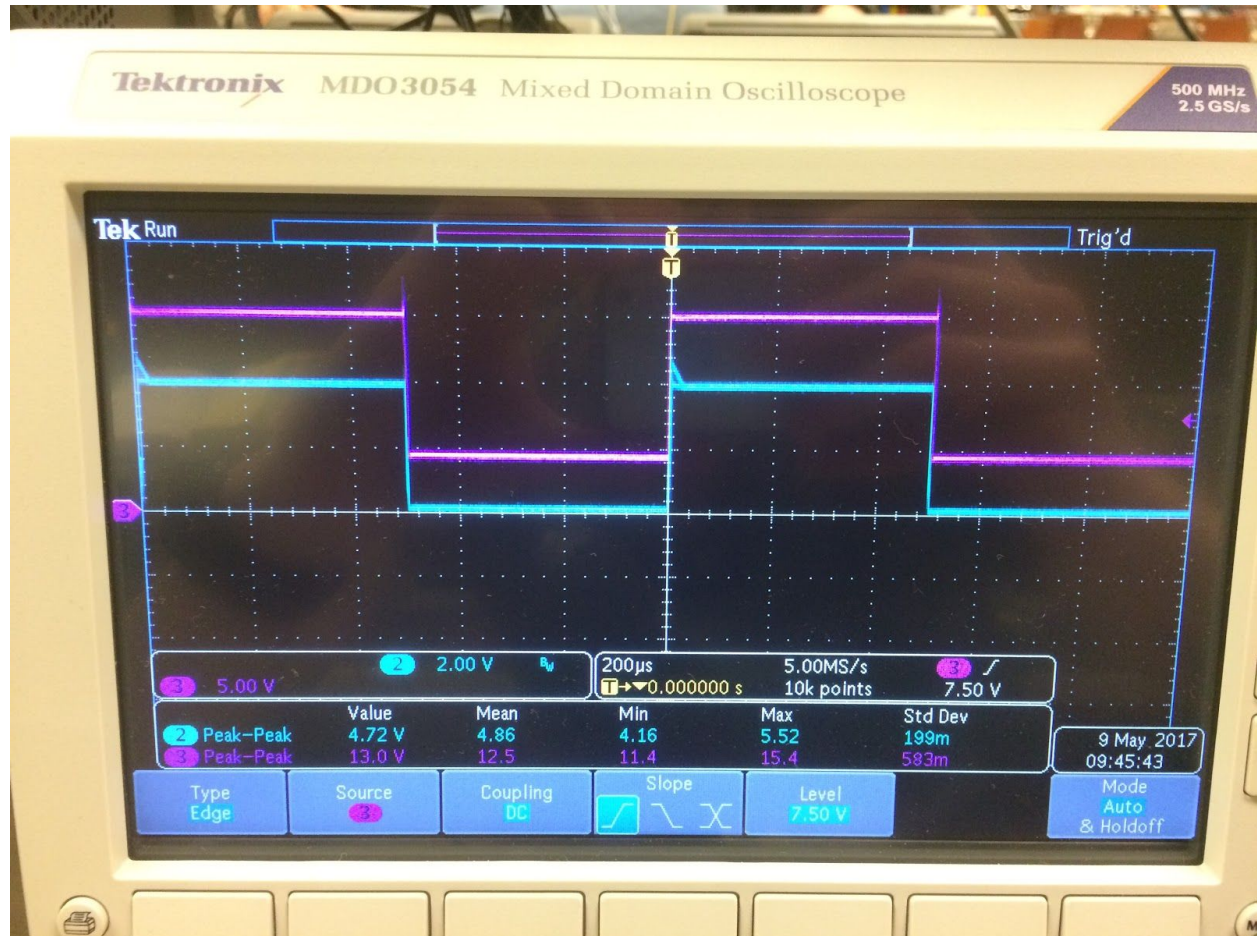


Figure 12: Oscilloscope Trace Showing the Low Side Gate Driver Input in Blue and the Low Side MOSFET Gate Voltage in Purple (Purple Trace Shifted Up for Visibility)

Power Stage

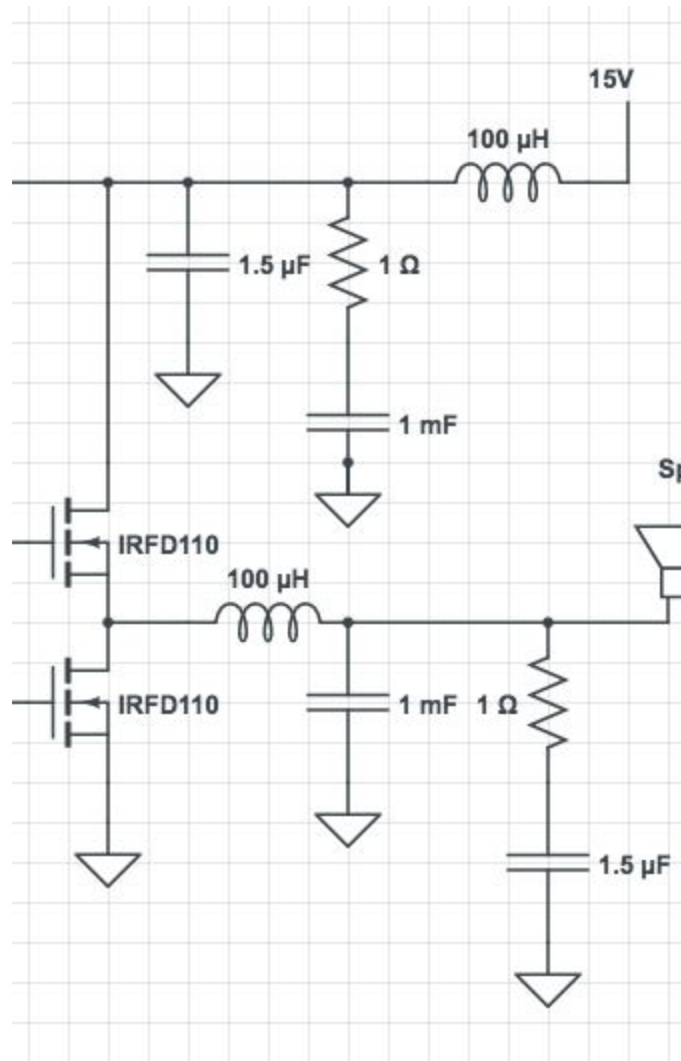


Figure 13: Half of the Power Stage

The power stage is composed to two symmetrical halves, like the one shown in figure 13. The power stage works by rapidly switching high side and low side MOSFET's (shown above) very rapidly to amplify the PWM signal. However, since this stage is drawing substantial amount of current and switching very rapidly, the laser microphone implements a passive second order low pass filter in order to keep the power supply noise-free.




Figure 14: Simulation of Current Demand on Power Supply, using a Passive Second Order Filter Without Damping

Designing this filter presented some interesting design challenges. A regular second order filter would have a very high resonant peak, as shown in figure 14. A filter with the laser snoop parameters but no damping would have a resonant peak of almost 4 dB. After adding damping, the resonant peak almost completely disappeared, as shown in figure 15.

Choosing values for the circuit was difficult. Choosing values was a matter of maximizing the size of the resistor so it was at least a 1 ohm resistor, minimizing the size of the inductor such to minimize its parasitic resistance, maximizing the size of the large capacitor such to minimize DC power dissipation through the resistor, and optimizing the size of the small capacitor to set the frequency cutoff at an appropriate value. The laser microphone's functional frequency cutoff is approximately 13 KHz, after taking into account parasitics. This is a near ideal frequency cutoff for the laser microphone.

Since 1.2 ohm resistors were not available, the laser microphone implements a 1 ohm resistor and utilizes the ~0.2 ohm parasitic resistance from the large damping capacitor.




Figure 15: Simulation of Current Demand on Power Supply, using a Passive Second Order Filter Without Damping

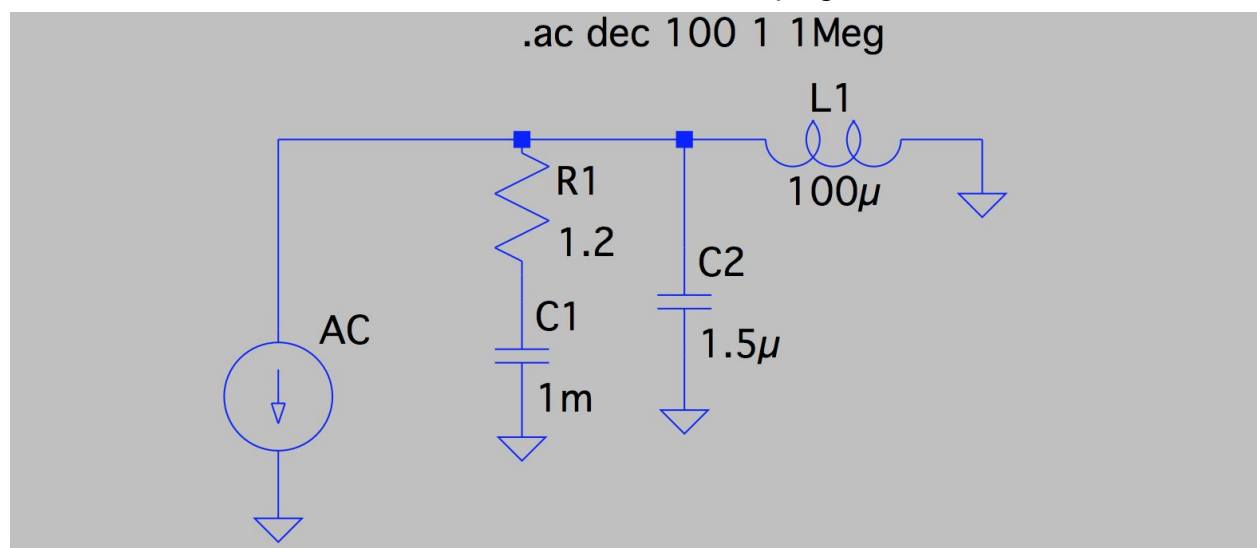


Figure 16: The Simulation Run to Produce Figure 15

Output Stage

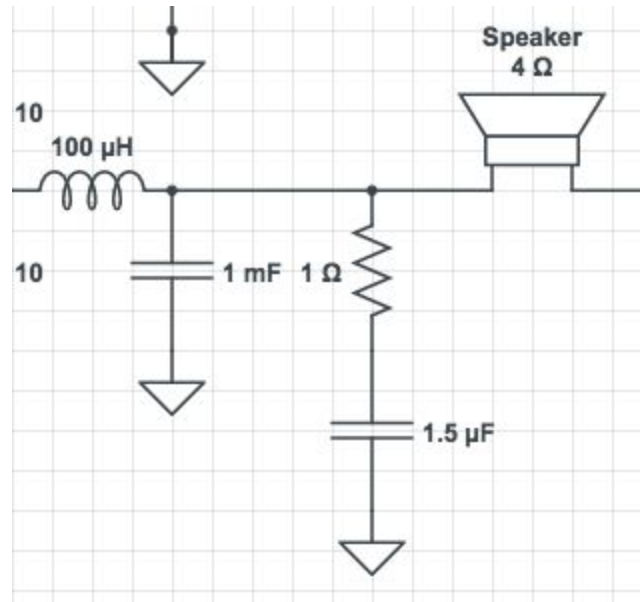


Figure 17: Half the Output Stage

The PWM from the output from the power stage then passes through a second order, low pass filter to the speaker. The filter is the same as the one used with the power supply filtering, only with the inductor in the first part of the filter to avoid rapidly charging and discharging the capacitors. This then produces an output to the speaker. The output was actually relatively clean at times, allow the class D amplifier to play music and test tones. Below is a scope trace of the outputs across load resistors.



Figure 18: Output From the Full Bridge Class D with a Sine Wave Input

The Ramp Generator

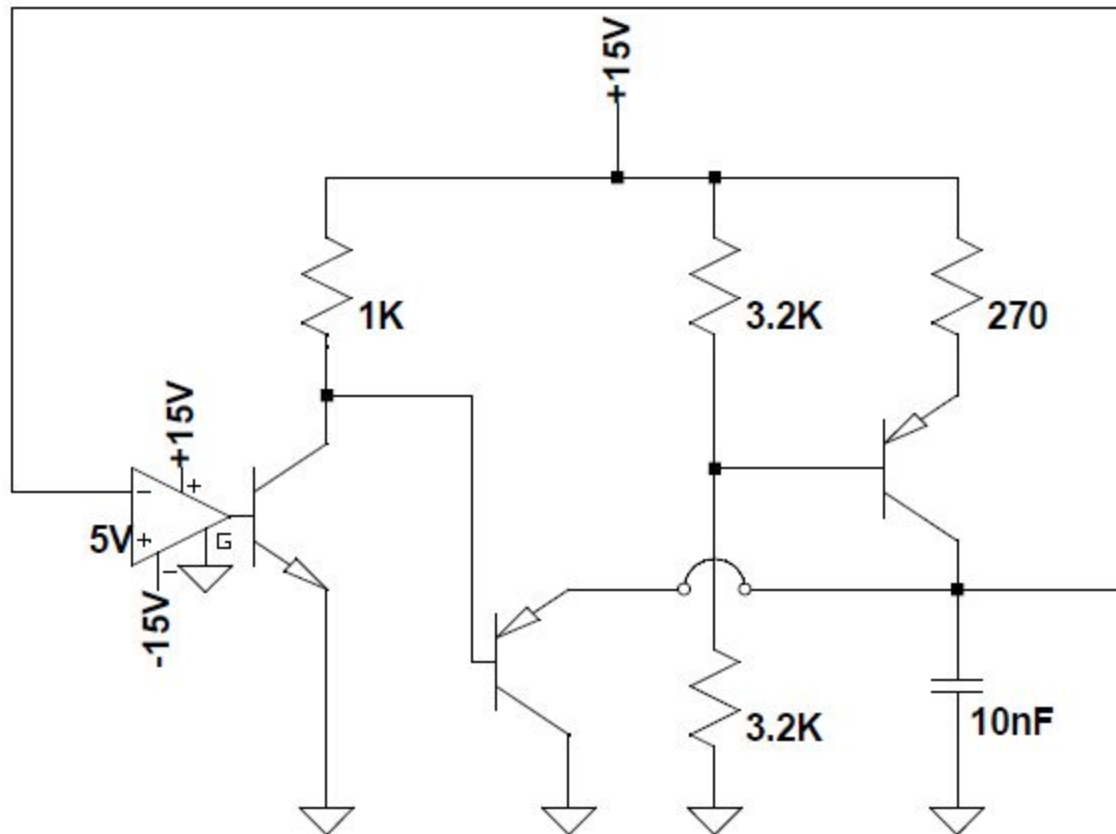



Figure 19
A depiction of the ramp generating circuit.

One of the requirements we had was to have the class D amplifier operate at 400kHz. This frequency is set by the ramp generator. Figure 19 shows an implementation of the ramp generator. This design uses a current source and an LM311 comparator charge and discharge a capacitor. The ramp can be observed by probing the capacitor. The pnp to cap setup on the right is a current source to linearize the rise of the ramp ($I = \Delta V / \Delta t \cdot C$). Since the fall time is relatively small compared to the rise time of the ramp, $1/\Delta t$ is approximately 400kHz. The current source is calibrated to about 23.11 microamps and the desired voltage swing is 5Vpp. So $C = I \cdot \Delta t / \Delta V$ which is about 10 nF. The non-inverting side of the LM311 was connected to the threshold voltage and the inverting side was connected to the capacitor. The LM311 has an npn type transistor internally connected to the output. Another npn was connected to the internal npn as an inverter and its collector was tied to the base of a pnp with a 1K pullup resistor. The allowed the pnp to act as a proper discharge path. So in operation if the capacitor charged up to the threshold voltage, the LM311's internal npn would turn off, allowing the external npn to turn on. This would then trigger the pnp to turn on, which discharged the capacitor.

Test Setup

There was a total of three test configurations. The first level was testing the laser's ability to pick up vibrations. This was done by attaching a mirror to a speaker and controlling the speaker with a signal generator. The laser was reflected off the mirror onto the receiver, the speaker was then driven at some known frequency, and the received signal was compared to the known frequency. The results confirmed that the received signal matched the known frequency. The process was iterated to tune gains in the TIA to obtain a voltage range large enough for audio. Then the receiver was connected to a class AB amplifier to test the audio quality. Once the filter was complete, it was connected between the receiver and the audio amplifier. The addition of the filter significantly reduced noise. Finally, we moved onto testing the system on a window. The window turned out to be too dirty, and our tools for alignment were mediocre. If we had been able to shoot the laser at the center of a clean window we believe the system would have performed well. So to emulate an ideal window we taped a mirror to a cardboard box and reflected the laser off of that. When comparing the output of the system to a pure sine wave input from the function generator, the mean percentage error was less than 3%. Qualitatively, this means that the output was clear and crisp audio.



Risk Assessment Chart

Functional Requirement	Risk	Countermeasure
Audio output must be human interpretable	The sound output might be too noisy for humans to easily understand	Adding modulation to our laser at the transmitted and demodulating the signal at the receiver. High order filtering

		and amplifying audio signal before outputting.
The receiver must receive the laser's reflection directly	The receiver may not be in a good position to receive the laser.	At first we will have to set up our testing apparatus very carefully to optimize the location and angle of the receiver. Eventually we hope to add a controls system to manipulate a base and optimize the location and angle.
Laser capturing audio signal	Medium - It's been documented and done. This will come down to calibrating and limitations of the components we choose.	We have the speaker test rig which should allow us to observe our systems behaviour in a controlled setting. We also will be touring a laser lab and checking our data sheets to better our understanding and intuition of our system's behaviour.
Window bouncing the laser	High - This is very dependent on the properties of the glass we bounce on and we cannot control this.	Again we have the speaker and we have the ability to try bouncing the laser off other objects. We could also purchase a small sheet of glass or acrylic (whichever is more ideal) and make a demoing mirror with ideal physical properties (should get a better idea during lab tour).
Audio waves strong enough to shake the window.	Low - This is dependent on the properties of the glass as well, but we can easily increase the volume to get a larger wave.	As mentioned, we can generate a louder sound against the glass to get it to vibrate.

Optics to Audio	Medium - This is something that none of us have experience with and it is possible that the conversion is nonlinear	The system has been built before so there must be either a work around to nonlinearity. We just need to familiarize ourselves with the properties of the phototransistor we select. There also is filtering and amplification involved, but we have had lab experience with this and so the risk there is relatively small in comparison
-----------------	---	--

Results and Conclusion

After testing and modifying the design of the laser microphone, it was calculated that the total harmonic distortion was less than 8%. This was calculated using a mean square percent error. Two signals were measured one from an input speaker and one from the output of the laser microphone. The two signals were captured on an oscilloscope, and exported as csv files. The percent error of each data point was calculated ($\text{percent error} = (\text{actual} - \text{expected}) / \text{actual}$), squared and averaged. The laser microphone was more sensitive when the sun was down. This could have been due to saturation of current in the photodiode during daylight. Overall at peak performance, the laser microphone was able to pick up music, and even people talking on the other side of the lab. There did seem to be some form of feedback which plagued the system. It was unclear whether it was audio feedback or electrical feedback. Jason suggested that isolating power across the system could reduce the feedback if it was electrical. Given the nature of the system it is plausible that the speaker was vibrating both the table that the receiver was on and the cardboard box. This would cause large “feedback” spikes as the speaker got louder. The project overall was a great exercising in analog circuit design.

References

1. https://www.academia.edu/13549827/Sound_Waves_Transmission_via_IR_Laser_The_Laser_Microphone_Device?auto=download

2. <http://www.lucidscience.com/pro-laser%20spy%20device-1.aspx>
3. <http://www.biophot.caltech.edu/publications/pdf/mcdowell-2008-pink-noise-fundamental-limit.pdf> (pink noise in photodiodes)
4. https://en.wikipedia.org/wiki/Amplitude_modulation
5. https://en.wikipedia.org/wiki/Optical_communication
6. <http://mathworld.wolfram.com/Circle-CircleIntersection.html>
7. <http://www.vishay.com/docs/81521/bpw34.pdf>
8. https://en.wikipedia.org/wiki/Voice_frequency
9. <http://www.joemario.com/blog/photodiode-amplifier-design-for-optical-vibrometry/>