Development of an FMCW Sonar System

Bernard Jin, Jack Bolte

Abstract (Bernard)—Frequency-modulated continuous wave (FMCW) radar is often used in industrial, automotive, and defense to detect and locate specific targets. Though this technology is powerful, cost and materials often limit the applications of this system. This report proposes an FMCW sonar system to allow for detection, bearing estimation, and rangefinding of targets. This system was able to be constructed with inexpensive components with similar capabilities to an FMCW radar system. The final device demonstrated object tracking capabilities up to a distance of 1 meter with an accuracy of 5.5 cm.

Motivation and Background (Bernard and Jack)

Ever since the dawn of radar, pulsed (Doppler) radar has been the premier method of rangefinding. This technology, which involves sending out a pulse of energy and timing how long it takes to return, is simple to implement and intuitive. However, pulsed radar has significant downsides; for example it is unable to detect targets at close range and it must transmit at very high powers [1]. This is why many rangefinding systems, especially those that operate in closer ranges, such as car proximity sensors, utilize a technique known as frequency modulated continuous wave (FMCW).

FMCW radar is a powerful technology that allows devices to quickly and precisely obtain speed and distance information without many of the downsides of other techniques (such as long minimum ranges) [1]. However, such devices are often expensive to create and develop, requiring specialized equipment and materials and leading to a high barrier to entry. Furthermore, RF attenuates quickly in some materials (such as water) [2], making radar unsuitable for such environments. In these cases, sonar is often able to achieve significantly better performance at a significantly lower cost. This project aims to create an FMCW sonar system utilizing sound to achieve similar direction-finding and rangefinding capabilities.

FMCW's principle of operation is not entirely dissimilar from that of pulsed systems, as it still measures the distance to a target using a time delay. The main difference is in the exact calculation of distance - while pulsed systems directly measure the time delay by timing when the signal is received, FMCW uses a frequency sweep and measures the difference in transmitted and received frequency (aka the "beat frequency") (Eq. 1). By transmitting a continuous wave, FMCW systems are able to operate at much lower power and at a higher update rate compared to pulsed systems [1]. Beat frequency is calculated by multiplying the incoming and outgoing wave (the identity $sin(a)sin(b) = (cos(a - b) - cos(a + b))/2$) to produce a wave with a frequency equal to the difference in frequencies. The higher frequency wave can then easily be discarded via filtering to extract only the beat frequency.

$$R = \frac{v|\Delta f|}{2\frac{df}{dt}}$$

Eq. 1: FMCW range equation [2]

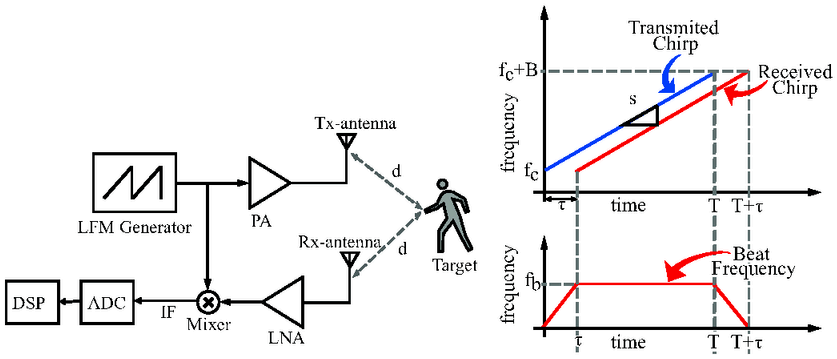

Fig. 1: A graphical representation of how FMCW works. [3]

Methods and Materials

The design of the FMCW sonar system is much like an FMCW radar (Fig. 2). The transmit signal is generated by a voltage-controlled oscillator (VCO), which is controlled by the MCU to generate a frequency sweep. The signal generated by the VCO is amplified to a voltage where the speaker is able to generate a fairly powerful transmitted wave. The wave generated by the speaker then reflects off of our target and is picked up by two microphones positioned on either side of the speaker. The tiny (10s of mV) signal generated by the microphone is then filtered to only include our sweep frequencies, and then normalized and amplified to allow the phase detector and MCU to receive a consistent signal regardless of the power received by the microphones (up to a point, as we don't want to amplify random noise). The transmitted and received signal is then used to calculate the beat frequency to determine distance and speed, and the phase detector is used to determine the angle of the received signal to determine the target's location.

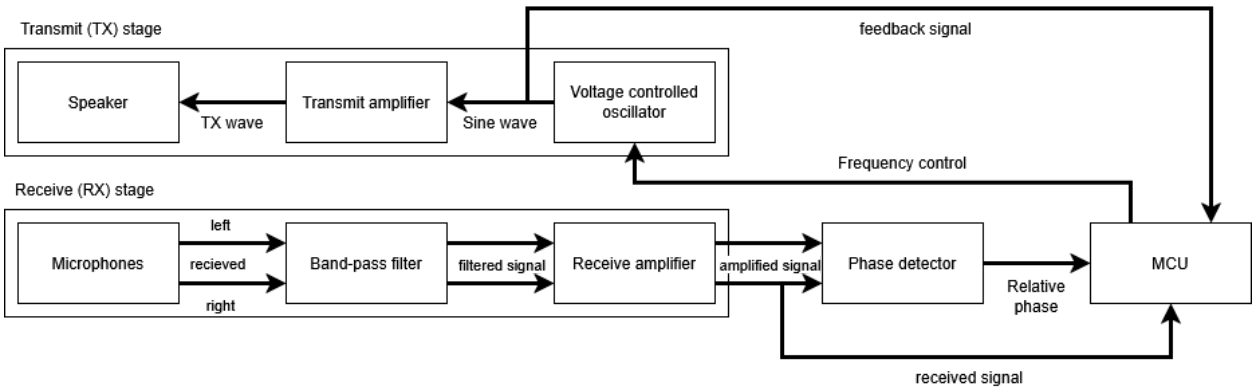

Fig. 2: FMCW Sonar block diagram

Transmit (TX) stage (Bernard)

The desired signal is generated by the VCO using a square wave generator and a filter. It is then amplified and transmitted through a transducer (speaker). A frequency sweep of 1-2 kHz was chosen to allow for a reasonable resolution of $v/(2 * BW) = 343\ m/s\ /\ (2 * 1\ kHz) = 17\ cm$ [2] while still

being able to be filtered to produce a sine wave and providing a reasonable spacing between RX microphones.

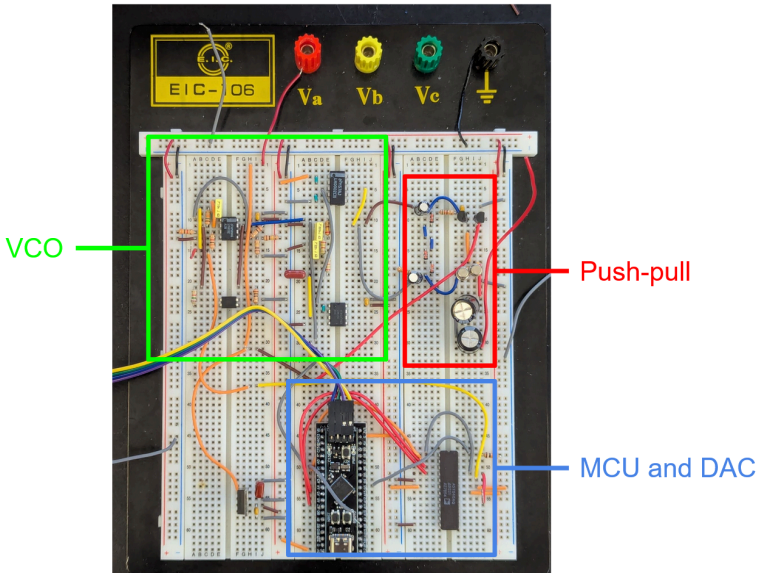

Fig. 3: constructed TX stage

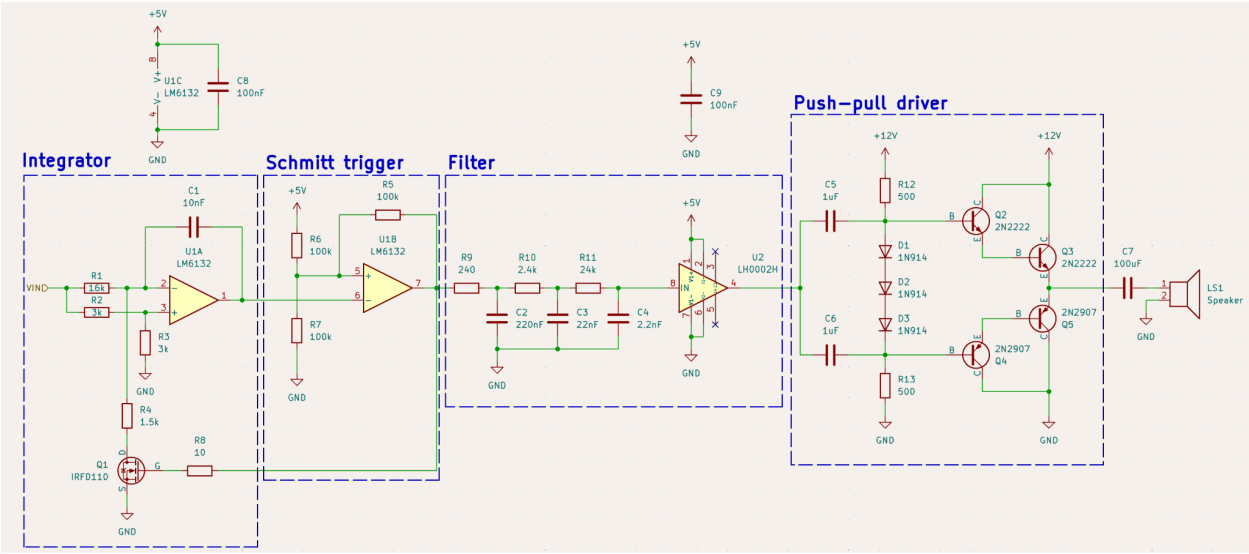

Fig. 4: TX stage schematic

Voltage Controlled Oscillator (VCO)

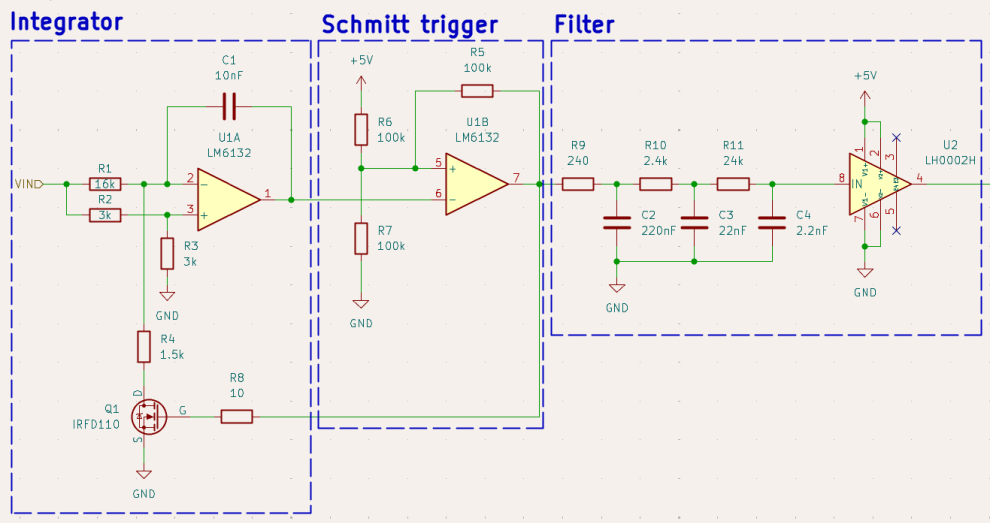

Fig. 5: VCO circuit

FMCW calculates distance by transmitting a frequency sweep and continuously comparing it to the current transmitted frequency, necessitating the inclusion of a VCO. This design uses a VCO instead of the MCU's DAC for a smoother signal, phase matching, and to offload the math required to digitally generate a sine wave.

The VCO is composed of three parts - an integrator, a Schmitt trigger, and a filter (Fig. 5). When the capacitor is uncharged, the control voltage initially causes the integrator's voltage to increase at a rate of $V_{in}/(2RC)$. Once the output voltage reaches $2*V_{CC}/3$ V, the Schmitt trigger charges the gate of Q1, causing the capacitor to discharge through R4. Once the output voltage reaches $V_{CC}/3$ V, the Schmitt trigger discharges the gate of Q1, letting the capacitor charge again. By tuning the ratio between R1 and R4, we were able to achieve a 50% duty cycle, resulting in an output period of $(4V_{CC}RC)/(3V_{in})$ and thus a frequency of $(3V_{in})/(4V_{cc}RC)$. R1, R4, and C1 were experimentally chosen to prove a 1-2 kHz range of approximately 1V to 2V.

A low-pass filter was added on the output to turn the square wave into a sine wave, which is fed into our transmit amplifier. The square wave output was chosen instead of the triangle wave due to the size of our bandwidth compared to our frequency. If the triangle wave output was chosen, creating a sine wave would be complicated by the fact that a triangle wave has a significant second harmonic, resulting in distortion. The square wave only exhibits odd harmonics, so by placing a filter at the ~3 kHz range, the VCO is able to effectively filter for only the first harmonic. The final filter was a 60 dB/dec filter with a cutoff frequency at 3 kHz, which was able to effectively turn the square wave into a fairly sinusoidal shape (Fig. 6). The output was also buffered to prevent the MCU and amplifier from loading the filter.

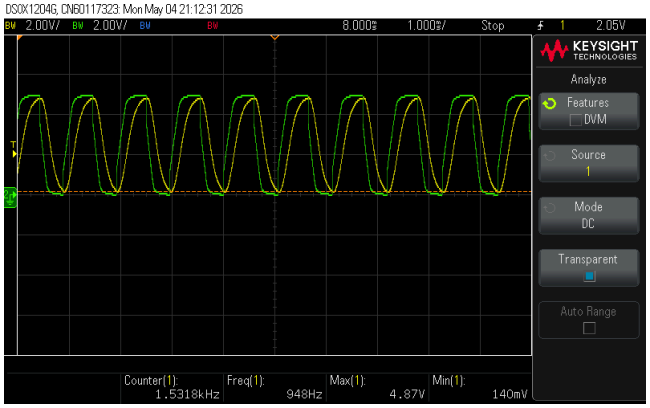

Fig. 6: VCO output. Green is the square wave from the Schmitt trigger, yellow is the filtered output.

Transmit (TX) amplifier

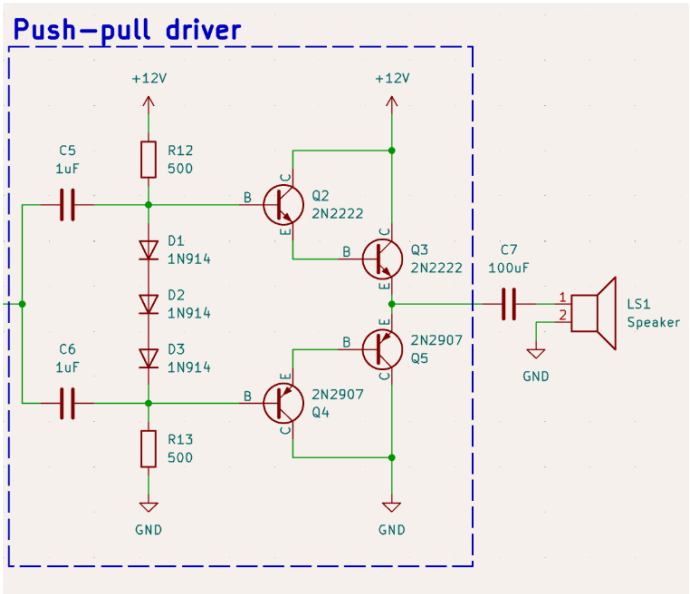

Fig. 7: TX amplifier schematic.

The TX speaker was a 3W 8 ohm speaker, capable of being driven by up to about 600 mA of current (about 5V). Since the output driving capabilities of our buffer is limited to 20 mA, a push-pull output stage was chosen to drive the speaker. A x300 gain requires a small signal gain β of about 300, which is not able to be consistently obtained by COTS BJTs. As such, a Darlington pair was chosen to be able to amplify with an effective β of over 1000. Three cascaded 1N914 signal diodes were placed to offset the base voltage of the Darlington pairs to overcome the saturation voltage and prevent crossover distortion. Due to difficulties in finding a 100uF nonpolarized capacitor, two 220uF polarized electrolytic capacitors were placed back-to-back to achieve a similar result.

Speaker

To allow the device to be made at a reasonable scale (i.e. prevent microphones from having to sit within each other), a 1-2 kHz sweep was chosen. This is well within the audible range, allowing ordinary

speakers and microphones to be used. A 1 kHz sweep allows us to have a 17 cm range resolution, which is comparable to some FMCW radar systems. As such, this device uses a commercial 3W 8 ohm mini-loudspeaker.

Receive (RX) stage (Jack)

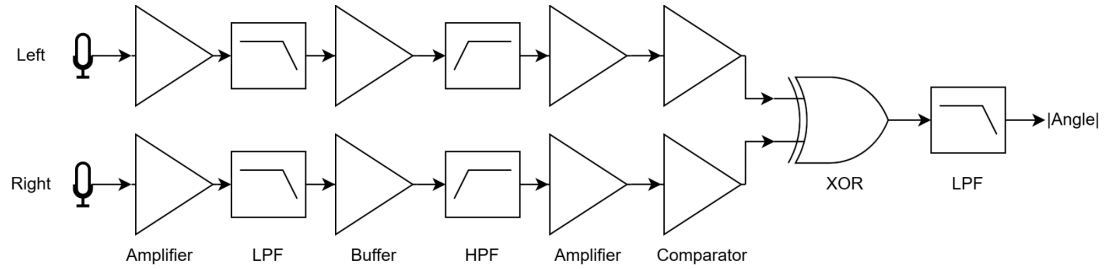

Fig. 8: Rx Stage Signal Chain

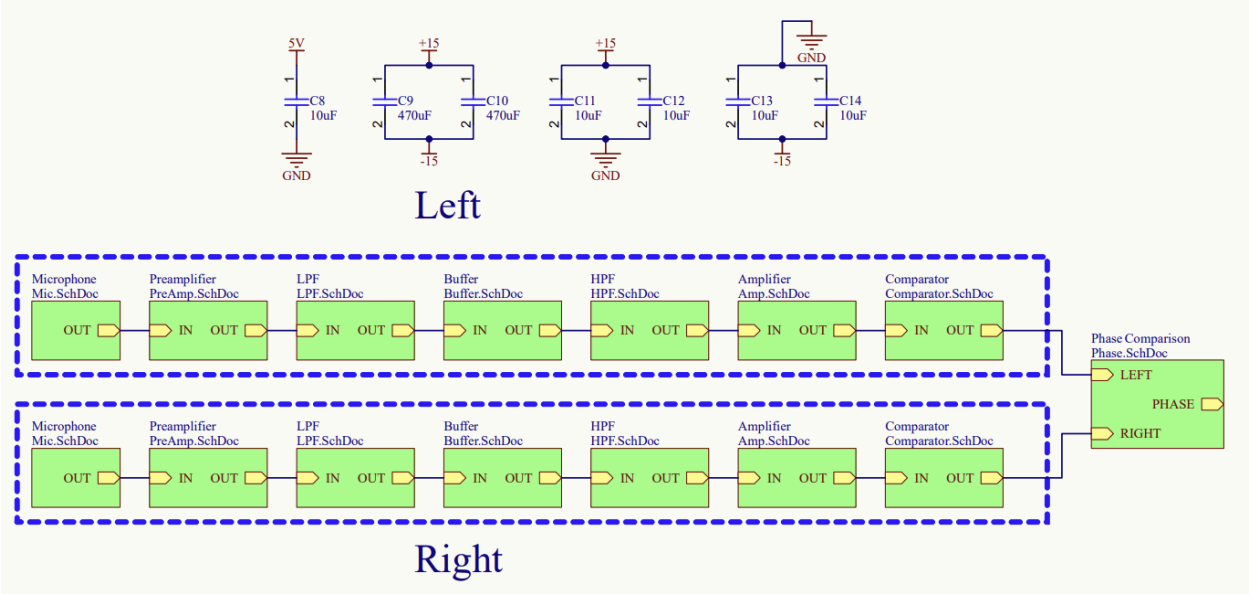

Fig. 9: RX Stage Schematics

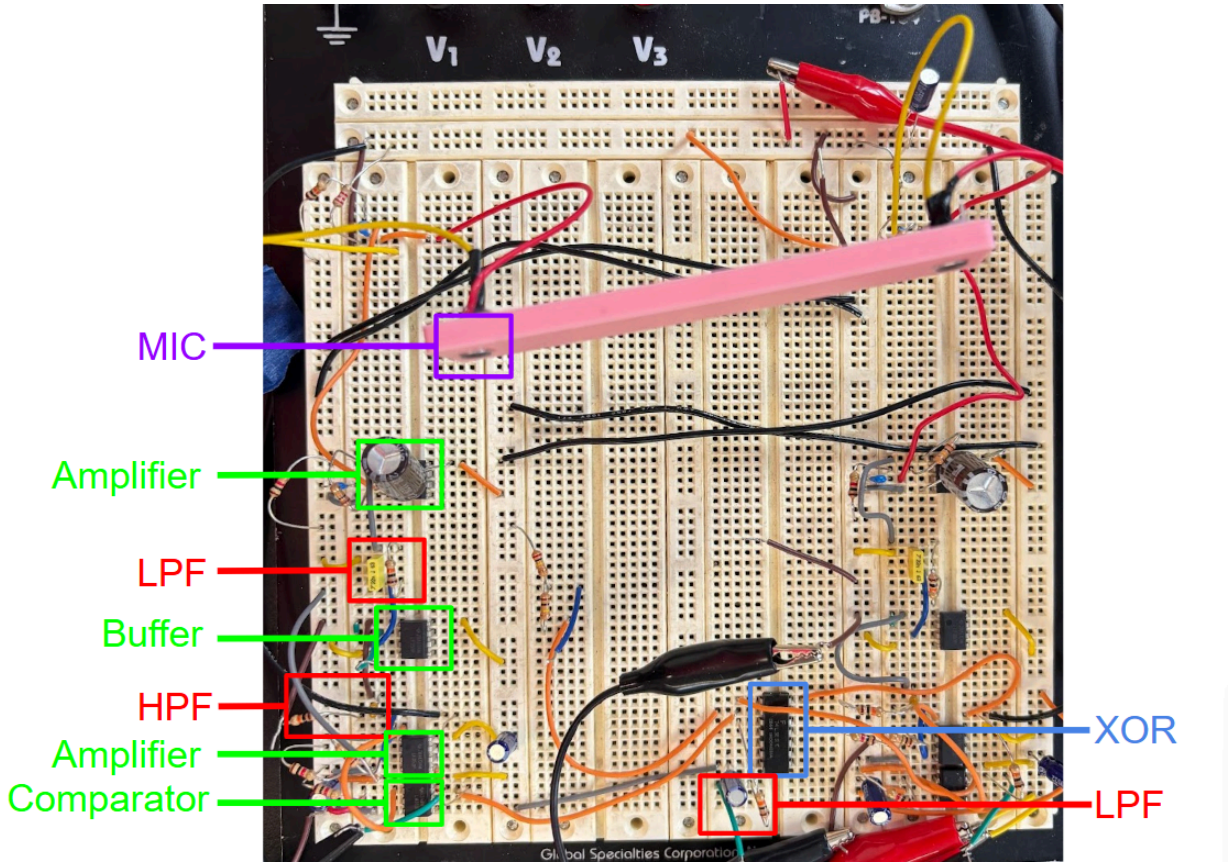

Fig. 10: Constructed RX Stage

The audio receiving front end consists of microphones (Fig. 11), amplification (Fig. 12), filtering (Fig. 13), comparison (Fig. 14), and phase measurement (Fig. 15). The output of this stage is a voltage proportional to the phase difference of the audio signal received by the two microphones.

Microphone

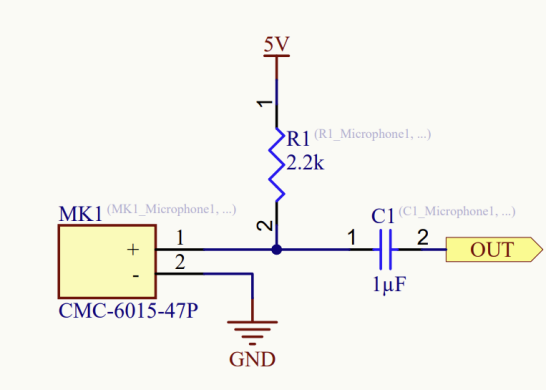

Fig. 11: Microphone schematic

The microphone chosen was a CMC-6015-47P. This microphone has a 56dB S/N Ratio and a sensitivity of -47dB. Choosing microphones with a higher S/N Ratio and better sensitivity would have significantly helped with signal quality issues we experienced in the RX chain. We used two of these microphones,

placed at 8.75cm apart, which is λ/2 apart at around 2 kHz, avoiding angle ambiguity towards the extremes of the received range.

Preamplifier

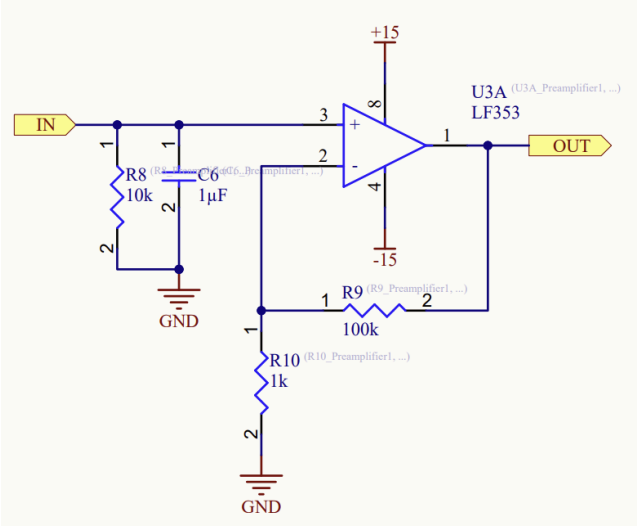

Fig. 12: Preamplifier schematic

The preamplifier boosts the high impedance, low voltage output from the microphone into a higher voltage, lower impedance output. In our case, we used a LF353 op amp, and set the gain to 101x. R8 prevents DC voltage from accumulating on C1 in the microphone circuit, and C6 is a decoupling capacitor. While the LF353 offered significantly better noise performance than the original choice of a LM358, a dedicated low-noise preamplifier would greatly improve the performance of our circuit.

Filtering

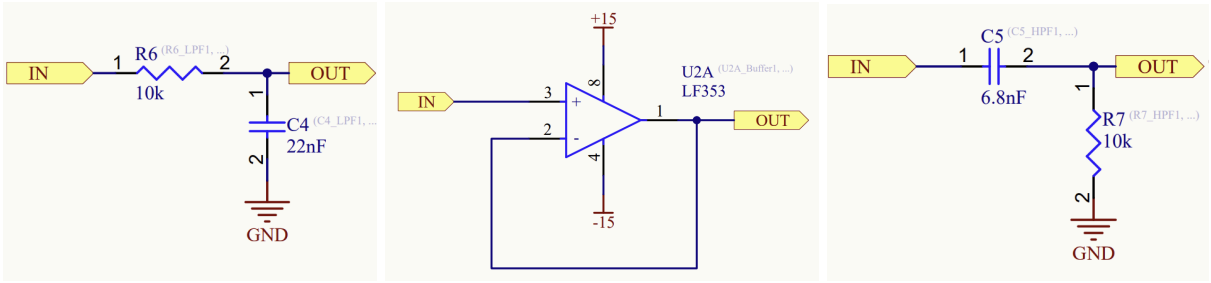

Fig. 13: Filtering schematic. Left: low-pass filter. Middle: buffer. Right: high-pass filter.

To isolate our frequencies of interest (1-2 kHz), we used a low pass filter, a buffer, and a high pass filter. Both the low pass and high pass filters are simple RC networks, with R = 10k to avoid loading of previous stages. For the low pass filter, a C of 22nF was chosen to achieve a cutoff frequency of 723Hz, and for the high pass filter, a C of 6.8nF was chosen to achieve a cutoff frequency of 2340Hz. A buffer, made from a LF353, was added between the low pass filter and high pass filter to prevent the high pass filter from loading the low pass filter.

Amplifier and Comparator

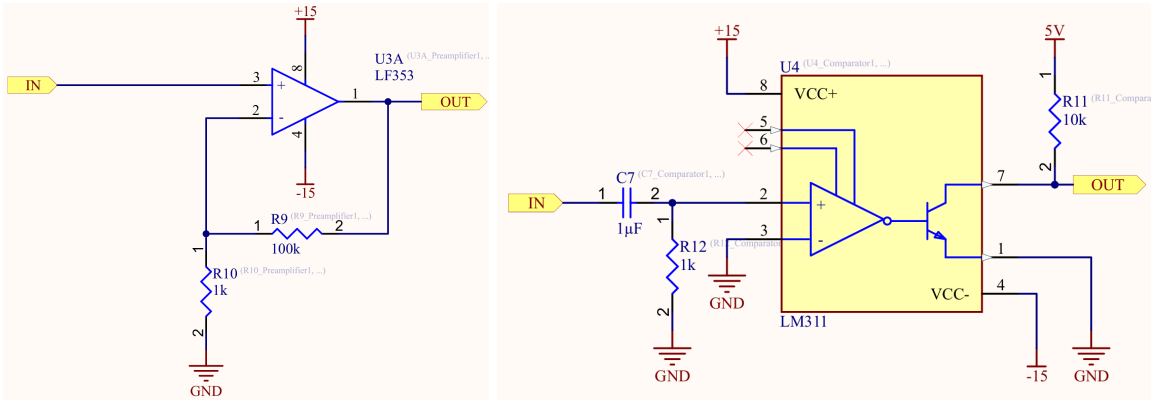

Fig. 14: Left: amplifier schematic. Right: comparator schematic.

After the filtering stage, we obtain a relatively clean sine wave at the frequency of interest. To convert this sine wave to a square wave at a digital logic level to feed into the phase comparison stage, we used a LF353 comparator, followed by a L311 comparator. The op amp stage, configured with a gain of 1001x, serves two purposes. First, the voltage gain means crossing 0V happens very quickly, and we increase our signal to noise ratio with regard to noise produced by the LM311. The second purpose is to provide a low impedance output that is more resilient to noise from the LM311 comparator. After the op amp stage, the signal is DC blocked using C7, and then compared to 0V. R12 was added to prevent a DC voltage from accumulating on C7. The collector terminal of the LM311 is pulled up to 5V with a 10k resistor, which provides a balance between signal rise time and switching currents and noise.

Phase Comparison

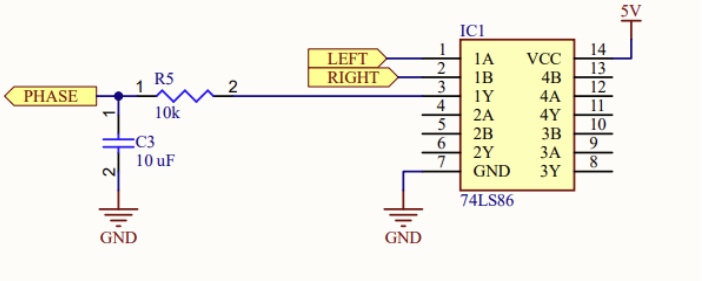

Fig. 15: Phase comparison schematic with output filter.

In addition to the ability to sense target distance and velocity, we also want to determine the direction of the target. We opted to calculate direction through phase difference of the received signal, as amplitude based methods would not be resilient due to the amount of background audible noise. From the comparators, we receive a square wave for each of the microphones. To measure the phase, we used an XOR gate IC, the 74LS86, to XOR these two square waves. The output of the XOR gate is related to the phase difference by the equation: $\Delta\phi = \pi D$ [4], where $\Delta\phi$ is the phase difference in radians, and D is the duty cycle of the output waveform. We can then put the output of the XOR gate through a RC filter to get an analog voltage proportional to the phase difference between the two signals, which can be measured by the microcontroller. Combining the earlier equation with $\Delta\phi = \frac{2\pi d}{\lambda} sin\,\theta$, where λ is the

wavelength of the sound wave, d is the microphone spacing, and θ is the angle the audio is arriving at, we get $\theta = sin^{-1}(\frac{D\lambda}{2d})$. In the case that the microphones are spaced $\lambda/2$ apart, which occurs around 2kHz in our design, the equation simplifies to $\theta = sin^{-1}(D)$.

MCU (Bernard)

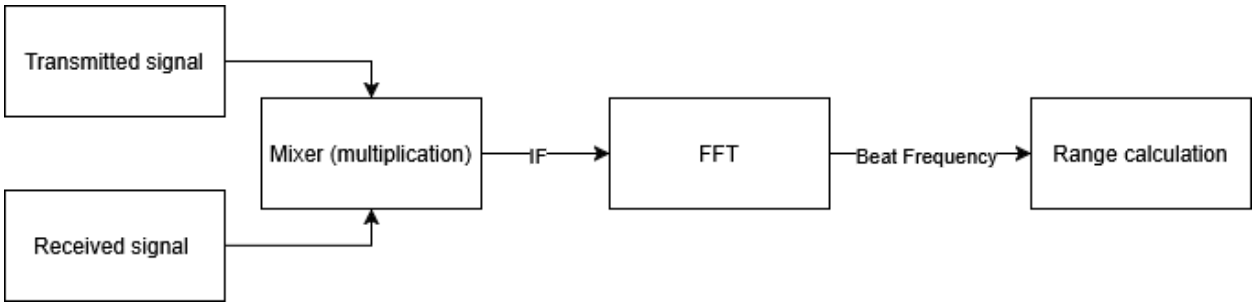

Fig. 16: Signal processing black diagram.

For signal processing and control, we used a black pill STM32 dev board, featuring an STM32F411. This was chosen for its availability, computing power, and high sample rate and resolution ADCs, while still remaining in a small form factor. Furthermore, the STM32 ecosystem is mature and both team members have experience working with it. This MCU calculated the beat frequency by digitally mixing the VCO's waveform and the incoming wave and combined the beat frequency with the phase information provided by the phase detector to give the location of a given target.

A 2.5 Hz triangle wave was generated by the MCU using an external AD7840 DAC. This triangle wave was connected to the input voltage of the VCO to generate our frequency sweep. A discrete DAC was chosen instead of a purely analog triangle wave generator to simplify testing multiple sweep frequencies and ranges.

Samples were taken at 10.24 kHz to avoid aliasing while allowing the MCU to complete the fast Fourier transform (FFT) between processing batches. Because of the fairly low frequency of operation, we are able to do digital mixing of signals instead of analog mixing, reducing development time. A 1024-point FFT of the signal is then taken for a frequency resolution of ~5 Hz (~17 cm), and the smaller frequency of the two frequencies with the largest amplitude is taken as the beat frequency.

Results (Bernard and Jack)

Due to the high level of noise in the environment, the device was tested with the TX stage transmitting directly towards the RX stage to maximize our signal-to-noise ratio. This means that the distance taken to the target is half as much, resulting in a halving of our resolution. Our final setup was able to obtain a raw range resolution of ~17 cm. Two rangefinding tests were conducted: a constant velocity sweep from 0.3 m to 1 m (Fig. 16), and a constant velocity sweep from 0.8 m to 0.15 m (Fig. 17). The p-value for the tests were calculated from a two-sided test with Pearson's correlation coefficient.

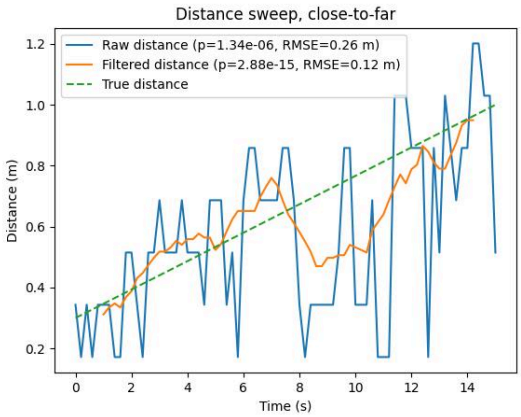

Fig. 17: Close-to-far sweep

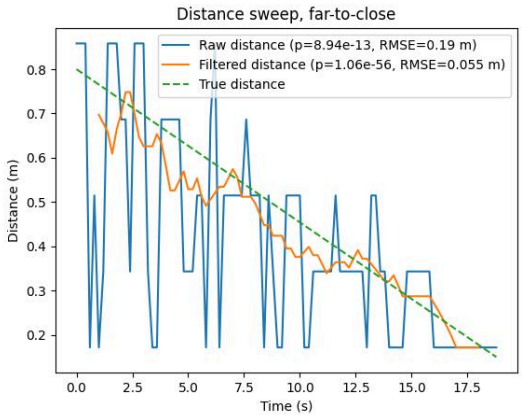

Fig. 18: Far-to-close sweep

Both tests demonstrate significant correlation between the true distance of the target and the distance estimated by the device. Performance was enhanced by convolving the raw data with a moving average filter of length 10, though this did result in an inability to get an accurate range estimate at the edges of the data.

Though the data shows a general trend towards the correct distance, the raw data itself still shows a significant amount of noise. Part of this is due to the coarse quantization of distance, as the device was only allowed to report six discrete values over the entire range of the test. Acoustic noise was also introduced during testing, particularly while readjusting the breadboards to allow the speaker to move through the entire range. This is the likely cause of the dip observed in the 8s-12s range in the close-to-far test, as the breadboard movement was done in a single large step, while the breadboard was slowly continuously moved in the far-to-close test. In total, our system was unable to achieve range resolution on-par with those of FMCW radar systems.

In the limited range that we tested in, direction finding functioned as expected.

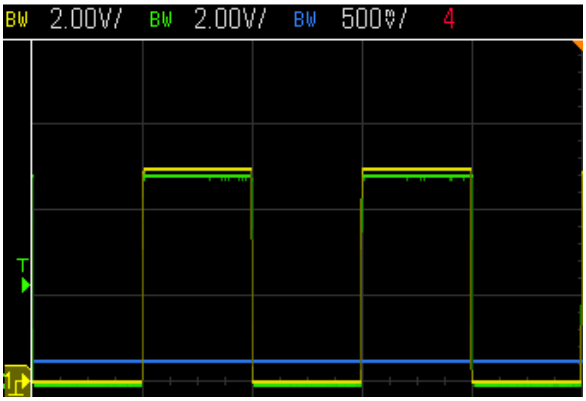

Fig. 19: In Phase

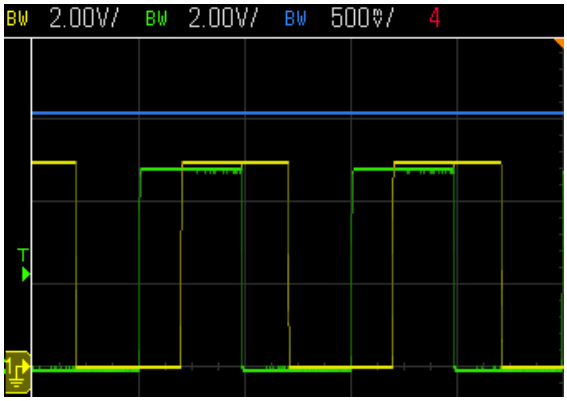

Fig. 20: Out of Phase

These scope images show the output of the comparator for the left (green) and right (yellow) audio, as well as the output from the phase comparison stage (blue). In figure 19, we see a scope image from when the audio was received in phase on both microphones, meaning the source was not at an angle. The output of the phase comparison is near the expected 0V. In figure 20, we see a scope image from

when the audio was received out of phase. The output of the phase comparison is 2V, which at an audio frequency of 1kHz corresponds to a calculated angle of 51.6 degrees.

Challenges and Lessons Learned (Bernard and Jack)

In the process of designing this system, we vastly underestimated the amount of noise in the 1-2 kHz range. This resulted in the signal having an extremely poor signal-to-noise ratio even at a range of 1 meter, limiting the effective range of our system. A second iteration of this device should likely utilize a much less noisy band, such as the ultrasonic range, allowing the device to get a much better range even without increasing transmit power. Additionally, more sensitive microphones with a higher SNR, as well as lower noise preamps would significantly improve the quality of the received signal.

Another difficulty in development was the creation of a sine wave at the output of the VCO. This was due to our sweep bandwidth being equal to the lowest frequency of the sweep. This forced the device to use the square wave output and heavy filtering to create the resultant sine wave. While this was able to create a usable wave, the amount of filtering used meant that a significant amount of power was being lost at the output of the VCO. This would also be helped by moving to a higher frequency range - by keeping the same bandwidth but making the frequency higher, the ratio of bandwidth to frequency decreases, allowing for easier filtering while maintaining the same range resolution.

Implementation of the DAC turned out to be particularly difficult. Due to the age of the DAC, the only methods to communicate with it were a semi-nonstandard SPI interface and a 12-bit parallel interface. Limitations on the number of pins on our MCU meant that we were forced to use the SPI interface, which took a significant amount of time to implement. The internal voltage reference of the DAC also proved to be not entirely accurate, and a significant amount of testing was required to find which values corresponded to our desired outputs. In the future, to simplify implementation and reduce development time, and possibly cost, an analog waveform generator should be used, or a significantly more modern DAC.

A recurring problem on the RX stage was waveform distortion. Due to the small magnitude of the signals we were working with, noise, crossover distortion, and device input currents were all concerns. In the future, we can pay more attention to properly decoupling all ICs, keeping a light load on op amp outputs at all times, and using lower noise devices.

Conclusion (Bernard)

We have demonstrated that FMCW sonar is a possible surrogate for similar radar systems in short-range, low power applications in open-air, and have demonstrated single-target tracking capabilities with inexpensive components. Our sonar system was unable to achieve similar range and resolution performance compared to commercial radars, likely due to our chosen frequency range.

Many improvements could be made to the system to improve performance. In particular, using a transmit band in the ultrasonic range would allow for easier filtering of the VCO’s output signal, as well as a wider bandwidth to allow for more frequency resolution and a faster FFT for improved temporal resolution.

References

[1] "What is FMCW vs Pulse Radar?," Counter UAV Radar — Low-Altitude Surveillance Radar. Accessed: May 11, 2026. [Online]. Available: https://www.counteruavradar.com/en/knowledge-base/what-is-fmcw-vs-pulse-radar/

[2] P. Saini, R. P. Singh, and A. Sinha, "Path loss analysis of RF waves for underwater wireless sensor networks," in *2017 International Conference on Computing and Communication Technologies for Smart Nation (IC3TSN)*, Oct. 2017, pp. 104–108. doi: 10.1109/IC3TSN.2017.8284460.

[3] A. Kumar, E. Easha, D. Sarkar, and G. Banerjee, *A Compact Quasi-Yagi Antenna for FMCW Radar-on-Chip based Through-Wall Imaging*. 2022. doi: 10.48550/arXiv.2208.11034.

[4] "Understanding the Exclusive-OR Phase Detector - Technical Articles." Accessed: May 11, 2026. [Online]. Available: https://www.allaboutcircuits.com/technical-articles/understanding-the-exclusive-or-phase-detector/