The Design and Validation of a Wireless Bat-Mounted Sonar Recording System

A Thesis Presented for the
Master of Science
Degree
The University of Tennessee, Knoxville

Jeremy Joseph Langford

August 2016

To those who taught me the value of an education from a young age.

Acknowledgements

I would like to thank Dr. Robert Bodenheimer and Mr. Mike Crabtree for making the Bodenheimer Fellowship available which I received to aid my in my studies as a graduate student. I would also like to thank Dr. Jeremy Holleman and everyone in his research group who helped so often along the way and made this research project possible.

Abstract

Scientists studying the behavior of bats monitor their echolocation calls, as their calls are important for navigation and feeding, but scientist are typically restricted to ground-based recording. Recording bat calls used for echolocation from the back of the bat as opposed to the ground offers the opportunity to study bat echolocation from a vantage otherwise only offered to the bats themselves. However, designing a bat mounted in-flight audio recording system, (bat-tag), capable of recording the ultra-sound used in bat echolocation presents a unique set of challenges. Chiefly, the bat-tag must be sufficiently light weight as to not overburden the bat, thus altering the bats behavior or causing harm to the bat itself. This thesis details the design and validation for the hardware and firmware of a wireless bat mounted ultrasonic recorder implemented using only commercial off the shelf parts.

Table of Contents

Chapter 1

Introduction

Scientists studying the behavior of bats monitor bat's echolocation calls, as their calls are important for navigation and feeding. Experiments in the past have involved using either a single or an array of ground based microphones to capture echolocation calls [3] [21]. While ground based recording has yielded important data in the past (and will continue to into the future) it is limited in its ability to capture the outgoing echolocation call and the corresponding echos from the bat's perspective. This is particularly problematic for studying echolocation as the amount of delay caused by the time it takes for sound to propagate in air between the outgoing echolocation call and the returning echos contains important information about the distance from the bat to its target.

One solution is to place the microphone on the bat itself, that way all of that sounds that are recorded correspond to the same information that the bat has available to to make all of its own navigational decisions. However, designing a bat-mounted echolocation recording system poses many difficult engineering challenges, chiefly system weight. The entire bat-mounted recording system must be light enough as to not alter the bats natural behavior nor cause harm to the bat itself. This presented thesis details the design and validation of such a bat-mounted recording system referred to in this paper as the “bat-tag”. The bat-tag, was designed to allow researchers to place the bat-tag on the bat's back and then release the bat. Then, as the bat is naturally interacting in its environment the bat-tag records the bat's echolocation calls. After the bat returns to its roost, the bat-tag wirelessly transmits the recorded data to a base recording station allowing the researchers to collect the recorded data at their convenience.

To provide a similar system for reference, this work will be compared to a system used by neuro-ecologists at Tel Aviv University to perform in-flight recordings of echolocation calls correlated with GPS data [4]. The work described in the paper from Tel Aviv University weighed 3.8 g sampled the bat ultrasound at a rate of 94 kHz, had a radio tracking unit for physical recovery of the system, and had GPS data logging capabilities. The work presented in this thesis weighed 3.7 g, sampled ultrasound at a rate of 222 kHz, and featured a radio for wireless transmission of recorded data.

Chapter 2

Design Methodology

2.1 Tag Weight

The root challenge at creating an in-flight sonar recorder for a bat (bat-tag) is system weight. The weight of the target bat species dictates the maximum weight of the bat tag. The Guidelines of the American Society of Mammalogist for the Use of Wild Mammals in Research state that a radio sensor tag should weight less than 5 to 10 percent of the target mammals weight [10]. In this case, the species *Phyllostomus hastatus* was chosen as the target of investigation. *Phyllostomus hastatus* female weights range from (75 g - 85 g) while males range in weight from (90 g - 100 g) [18]. Following the conservative end of the guidelines, if the bat-tag is to weight no more than 5 percent of the body weight of the target bat, a sensor tag should weigh in a range of (3.75 g - 5 g). Being under the lower weight limit would allow for the bat-tag to be used on smaller female specimens, while being closer to the upper weight would limit application to larger male specimens. At the start of this project, the target weight was set to 5 g to ease the engineering challenges.[10]

The largest impact of the weight limit is on the amount of available energy for the bat-tag to use. As the only practical way of powering the bat-tag at this time is using batteries, there is a clear trade off between how much energy one wants to make available for the bat-tag to use and the weight of the bat-tag. The chosen energy source for this application was non-rechargeable lithium polymer batteries due to their high energy density, the specifics of which will be discussed in a later section.

The second largest attributing factor to bat-tag weight is the actual mass of the printed circuit board that all of the components are mounted on. To help predict how large of an impact that the mass of the printed circuit board would be in the final bat-tag weight, a MATLAB script was written. The MATLAB script used the known densities of FR-4 and copper, information about the number of layers in the board, the thickness of the copper on each layer, the copper area density per layer, the thickness of the FR-4 between each layer, and the coordinates of the board verticies. Using this MATLAB script, the printed circuit board mass was predicted to be 0.967 grams.

At this point, the issues of final bat-tag weight and battery capacity have been brought forward. Care must be taken in the selection of the remaining components used to implement the functional requirements of the system so that any one particular component does not draw an unreasonable amount of current, thus limiting battery life. Additionally, the components must be selected to reduce their raw weight contribution as well as their contribution to board area.

2.2 Analog Front End

In designing the analog front end, it was first necessary to consult the functional requirements of the bat-tag. Adult *Phyllostomus hastatus* are capable of producing a variety of calls. Screech calls are documented as having frequency content as low as 6 kHz while echolocation calls can range as high as 80 kHz.[3] This information set the bandwidth requirement for both the microphone and the analog filter. Additionally, as with the rest of the components on the board, the analog front end needed to operate at supply voltage of 1.8 V, require minimum board area, and operate with as small of a current draw as possible. Thus, the SPU1410LR5H was selected as the microphone as it could operate on supply voltages ranging from 1.5 V to 3.6 V, had a documented frequency response up to 100 kHz, had a nominal current draw of 120 μA and compact (3.00 mm X 3.76 mm, 11.3 mm^2) package .[14][15].

Phyllostomus hastatus when navigating towards a target will first start with high intensity calls and as it detects the target and closes in, it will begin to use lower intensity calls. An exponential decay model for outgoing call intensity has been suggested by Ulrik Nørum, Signe Brinkløv, and Annemarie Surlykke based off of their observations. Their observations show that when further than 3.2$\pm$.2 m away from an obstacle *Phyllostomus hastatus* will create echolocation calls of high intensity at 112.1$\pm$1.9 dB referenced 10 cm from the source. Once it locates and begins to close in on a obstacle, their model predicts that *Phyllostomus hastatus* will begin to lower the intensity of the outgoing calls by 10.7$\pm$1.0 dB per halving of distance to the target. They suggest that the reason for this is so that the returning echos do not overburden the ears of the individual bat. Following this model, an outgoing call at a target 0.8 m away could be expected to be 90.7 dB referenced 10 cm away from the source [21].

To utilize the dynamic range of the Analog to Digital Converter (ADC) on board the micro controller and to provide the flexibility required for gathering experimental data from bats in flight, it was necessary to have a means to provide adjustable gain to the output of the microphone. Setting the gain to high results in the outgoing calls clipping, while setting the gain to low could result in not being able to detect echos. It should be noted that the intent of providing adjustable gain was not to monitor the microphone's output during operation and adjust the gain to keep it rail to rail, but to allow the gain to be reprogrammed in between test runs. This is because adjusting the gain to always keep the output rail to rail during operation would remove from the recorded data information about how much echolocation calls are attenuated as they travel through air and return back to the bat as echos.

2.2.1 Analog Filter Design

Figure 2.1 shows the filter topology described subsequently in this section. Implementing the adjustable gain was also taken as an opportunity to implement an anti-aliasing filter to band-limit the microphone's output. To this end, it was necessary to design an adjustable gain amplifier circuit with a fixed cutoff lowpass filter. The LT6257 by Linear Technology was selected as the part to implement the filter circuit as it contained four op-amps in one (4.04mm X 4.9mm) 16 pin package. The LT6257 could run on a supply voltage ranging from 1.8V to 5.25V and each amplifier in the package needed only 65 μA when active [16].The first amplifier was used to provide adjustable gain while the other three amplifiers implemented a 6th order Chebyshev low pass filter with a roll off starting at 80 kHz.

The gain stage was implemented as a standard inverting amplifier circuit with a 1 kΩ fixed resistor plus the 400 Ω from the output impedance of the microphone on the inverting terminal and a digitally controlled MCP4014 rheostat by Microchip Technology in the feedback loop providing adjustable gain. The MCP4012 is quoted as being able to provide a minimum impedance of 75 Ω and a maximum impedance of 50 kΩ. Using the gain equation for a inverting amplifier $-\frac{\mathrm{R_f}}{\mathrm{R_i}}$, the adjustable gain should be variable from a factor of -25.4 dB up to a factor of 31.1 dB. [19].

The Chebyshev low pass filter was designed roll off at 80 kHz with a passband gain of 12 dB and permitting .1 dB of passband ripple using the Analog Filter Design (AFDT) Toolbox available for MATLAB and simulated in LTspiceIV [23]. The roll off was chosen to preserve most of the high frequency content of *Phyllostomus hastatus* s echo location calls, while still sufficiently band-limiting the microphone's output signal as to ease the sampling rate requirement of the ADC. After the passive component values for the low pass filter were specified by the AFDT, the filter was next simulated in LTspiceIV using the LTC6255 as the component model for the individual amplifiers. This was particularly convenient because the model already existed in the stock

LTspiceIV libraries and more importantly, it shared the same data sheet with the LTC6257.

Furthermore, it was necessary to AC couple the output of the microphone to the input of the gain stage of the amplifier circuit, and to add a 0.9V DC bias to the output of the last filter stage which fed into the ADC of the micro-controller. To this end, a 100 nF capacitor was added in series with the output of the microphone forming a high pass filter with a corner at 1.13 kHz. The output of the last stage used two 100 kΩ resistors to form a resistor divider to divide the 1.8V supply down to 0.9V to DC bias to the output of the filter. To ensure that there was no preexisting DC bias before the 0.9V bias was added, the output of the filter was first AC coupled using a 3nF capacitor. This formed an additional high pass filter with a cut off of 1.06 kHz. The particular capacitor values were chosen to provide a high pass component to the filter which would be helpful for filtering out low frequency noise when trying to record bat calls while still preserving the lower frequency screech calls. The cut offs for the high pass filters could have been chosen to be slightly higher or lower, but in the way that they were selected, at 6 kHz the filters frequency response is mostly flat preserving the screech calls. Additionally, the particular values were chosen because they proved to provide a stable circuit in simulations which had provided problematic in previous designs. The simulated frequency response of the filter at ti's maximum and minimum gain setting are are shown in figures 2.2 and 2.3.

Doing a DC operating point simulation at the maximum gain setting showed a current draw of 280μA from the passives and LTC6267. The MCP4012 is quoted as having a max static current draw of 1μA and the microphone is quoted as having a max current draw of 160μA meaning the described circuit should have a max current draw of 441 μA [19][16].

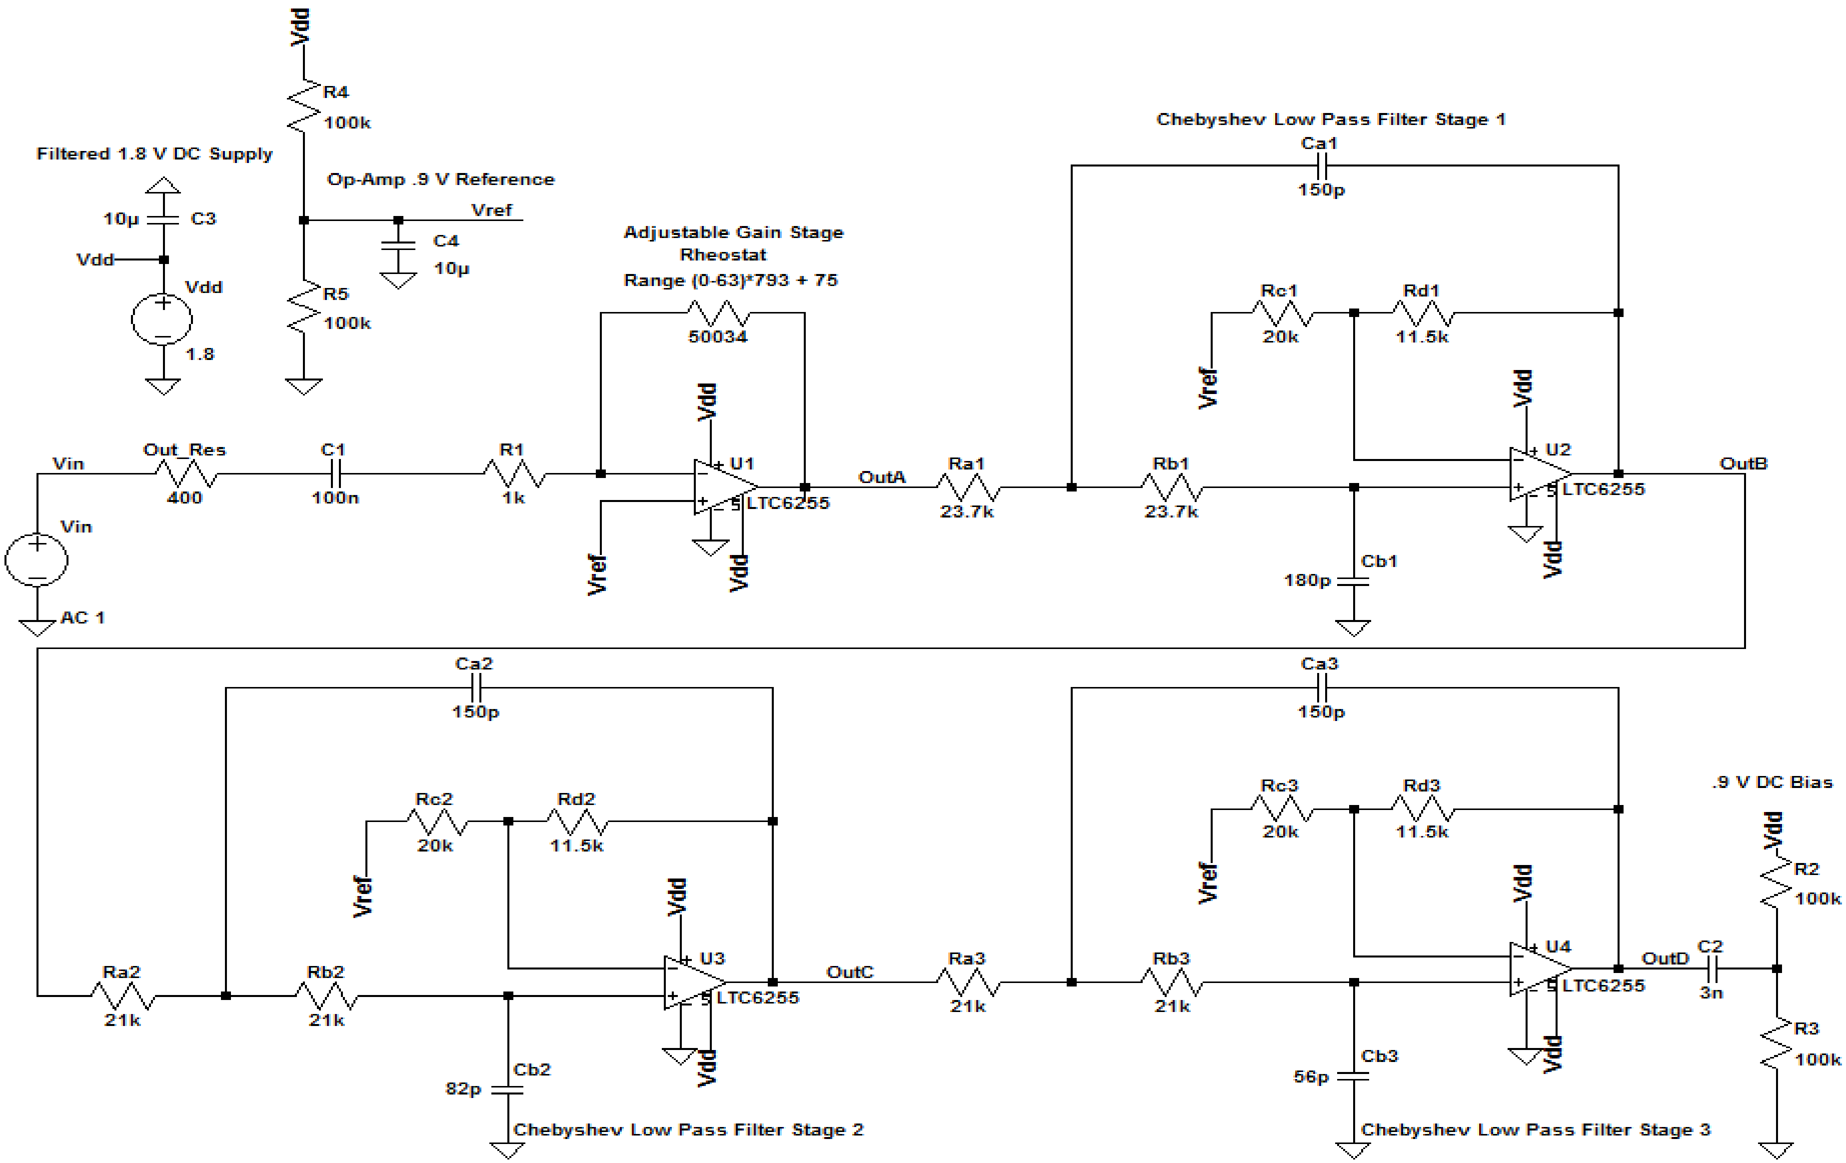

Figure 2.1: Schematic of Analog Front End

2.2.2 Analog Noise Calculations

In order to predict the combined noise performance of the analog filter and microphone an input referred noise spectral density simulation was performed on the analog filter in LTSpice, the results of which were imported into MATLAB. Subsequently, the noise spectral density information from the microphone's data sheet was also imported into MATLAB. A plot was produced from both sets of data showing that over the frequencies of interest (1 kHz - 100 kHz) showing that the microphone should be the largest contributor of noise 2.4. Both sets of data were then processed from 1 kHz to 100 kHz to compute a net input refereed RMS noise value for the analog front end of 23.2 μV.

2.3 Radio

The radio chosen to transmit the data off of the flash memory chip was the CC1101 by Texas Instruments. This particular radio was chosen because of it's ability to run on a supply voltage of 1.8 V, modest current draw of 17.2 mA when transmitting with 0 dBm output power at 915 MHz, ability to perform FSK-2 modulation at data rates as high as 500 kbps, ability to enter a low power wake up on radio mode, and small size of (4.15 mm X 4.15 mm) [24]. First, it should be noted, that this particular radio is not necessarily the best choice for all applications. This is mainly in reference to experimental data that will be discussed later, but in brief, it was found that to reduce the bit error rate to an acceptable level without using error correcting codes, it was necessary to lower the data rate of the radio significantly below it's quoted maximum FSK-2 data rate of 500 kbps. In the radio results section 4.3 the actual realized performance of the radio and other alternative solutions will be discussed in detail.

To minimize power loss caused by energy being reflected back from an antenna during a transmission, it is necessary to to ensure that the impedance of a radio

Figure 2.2: Simulation at Maximum Gain

Figure 2.3: Simulation at Minimum Gain

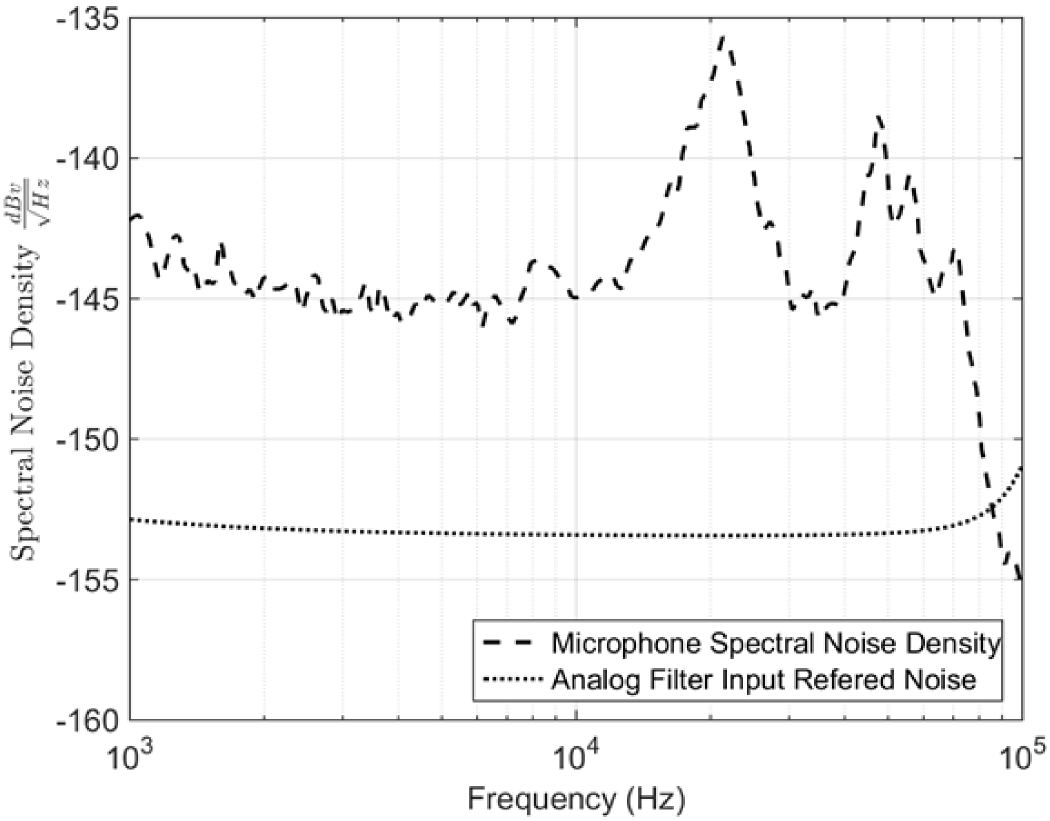

Figure 2.4: Analog Front End Estimated Noise Spectral Densities

matches the impedance of the antenna that is used. In this design, a single ended chip antenna, ANT-916-CHP-x, by Antenna Factor was used due to it's relatively compact size of (16 mm X 3 mm), omni-directionality, and it's 50 Ω impedance [2]. The differential impedance of the CC1101 is quoted as being (86.5 + j43) Ω when transmitting at 915 Mhz. That is why the CC1101's data sheet included the specifications for a circuit to convert the differential RF port's impedance to a single ended impedance of 50 Ω. However, the reference circuit required the use of a network of 14, 0402, passives to implement which would have occupied much board area. Additionally the circuit could have been sensitive to variances from the reference design in board layout. A solution to these problem came in the form a single 0805 balun by Johason Technology designed to convert the CC1101's RF port's differential impedance to a single ended impedance of 50 Ω [13].

2.4 Non-Volatile Memory

As the micro-controller only held 64kB of SRAM, and the ADC generated 12-bit samples encoded in 16 bit unsigned integers at a rate of 222 ksps, the micro-controller's SRAM would be filled in under 150 ms. Thus, it was necessary to have an external non-volatile memory capable of holding a useful amount of information. The MT29F8G08ABBCAH4-IT by Micron was selected due to it's 1.8V supply voltage, 8 Gb of storage, and high data throughput capabilities. Details regarding the specifications of the flash memory chip must be acquired on Micron's website as the contents of the data sheet are confidential and proprietary. However, in the Non-Volatile Memory Results section 4.4, the measured write latency and data throughput of the entire digital data path are discussed which includes the measured data throughput and write latency to the flash given this particular implementation of flash memory drivers running on this particular micro-controller at a particular clock speed.

2.5 Micro-Controller

The micro-controller selected for this design was the PSoC® 5LP (Programmable System on Chip) by Cypress Semiconductor. It was selected because of the flexibility lent by the platform's digitally controlled analog systems, FPGA-like programmable hardware which could be programmed in Verilog to synthesize many useful subsystems, multiple ADCs and comparators, and Digital Filter Block which could be used to easily synthesize digital signal filters [5]. Additionally, the PSoC® 5LP's flexible clocking system allowed for many of the components, including the Cortex M3 CPU, to run a range of different clock frequencies allowing for a more optimal balance between performance and power to be reached. For example, it would be possible to set the bus clock for the CPU at 12 Mhz while clocking a SPI interface synthesized using the digital blocks at 24 Mhz [17]. For the application of making a light weight bat-tag, a PSoC® 5LP variant was selected that came in a QFN package to reduce board area and weight. The QFN package occupied 64 mm^2 and weighed 178 mg compared to the 196 mm^2 and 638 mg of the TQFP package [7] [6].

2.5.1 Active Current Modes

Cypress Semiconductor provided a spread sheet for estimating the overall current consumption of it's PSoC® 5LP over a defined set of parameters [8]. To estimate the amount of current consumed by the PSoC® 5LP during recording, the following parameters were provided. VDDA: 1.8 V, VDDD: 1.8 V, IMO: 3 MHz, ILO: 100 kHz, PLL: 3 In / 63 Out, CPU Speed: 12 MHz SAR ADC Speed: 250 ksps, and all other components were selected as being turned off. This provided an estimated current of 6.04 mA. To estimate the amount of current used by the PSoC® 5LP during radio transmission, the SAR ADC was changed to off for a new estimated current consumption of 5.75 mA. It should be noted, that these numbers do not include any of the current used by the programmable digital blocks, such as would be used by the synthesized SPI interfaced during radio transmission.

2.5.2 Low Power Modes

Important for a low power design, the PSoC® 5LP is capable of entering Sleep and Hibernate power modes. The sleep power mode halts the CPU and shuts down most of the subsystems allowing the PSoC® 5LP to retain the state of its SRAM and draw around only 2 μA. During sleep mode, the PSoC® 5LP could be awoken in under 25 μs in a number of ways including but not limited to, a pre-set sleep timer expiring, a comparator event, Port Interface Control Unit (PICU) interrupt (any interrupt generated from a GPIO pin), an I^2C address match, or device reset. If the PSoC® 5LP is set to enter hibernate mode, it will retain the state of it's SRAM and draw around 300 nA. In this state the PSoC® 5LP can only by woken be a comparator event, PICU interrupt, or device reset. Additionally, waking up from hibernate mode takes longer than waking up from sleep mode, less than 125 μs compared to less than 25 μs [22].

2.5.3 Overview of Operation

The PSoC® 5LP served as the main control unit for the bat-tag. The overall data path is described in brief in this section and subsequently elaborated on. Analog data was generated from the microphone and sent through bat-pass filter with programmable gain. The PSoC® 5LP was responsible for switching on the power to the analog front end and setting the gain of the bat-pass filter by changing the impedance of a digitally controlled rheostat in the feedback loop of an inverting amplifier as described in 2.2. The analog signal would next be digitized by one of the on-board 12-bit SAR ADCs at a rate of 222 ksps. The digitized samples were then buffered in circular buffer in RAM. Samples from the circular buffer were then flushed out a flash page size at a time to the flash memory chip for non-volatile storage. After the desired number of samples were taken, the PSoC® 5LP would turn off the power to the analog front end, then enter a low power sleep mode waiting for an interrupt from the radio. Once the PSoC® 5LP was woken up by the radio telling it that it was in range of the base station and then begin reading the data off of the flash chip and transmitting it using the radio to the receiving base station until either the entire contents of the flash chip were sent, or the battery powering the bat-tag died.

2.5.4 Analog Front End Control

As mentioned in 2.2 the analog front end's power could be turned off and on by the using a P-MOSFET as a power switch to the analog supply that powered the microphone, rheostat, and operational amplifiers. From the perspective of the PSoC® 5LP, this was accomplished by calling a function that would drive one of the GPIO pins high/low depending on the desired state of the power switch. The main cost in terms of hardware in implementing this was having to reserve the usage of one of the PSoC® 5LP's GPIO pins. From a software perspective, once the power was switched off, the rheostat responsible for setting the gain of the bat-pass filter loses it's old state, so once the power was switched back on, not only would one have to wait for the output of the filter to become valid, but one must also reprogram the wiper setting on the rheostat. Early in the design process, it was decided that even though it would be possible to implement dynamic gain control on the bat-pass filter using the rheostat, that doing so would cause ambiguity in the relative amplitude of the recorded data, as all of the data would not be recorded with the same gain. One of the goals of the bat-tag was to be able to record outgoing sonar pulses from a bat, and catch the sonar returns from the bat's perspective. For example, one would expect the outgoing sonar pulse to be at a higher amplitude than the returning echo, and if the PSoC® 5LP had controlled the gain, so that the recorded signal always took up the entire dynamic range of the ADC information would be lost regarding the strength of the returning echo with respect to the strength of the outgoing sonar pulse. A possible solution to this problem would be to encode the gain of the bat-pass filter into the data whenever the gain was adjusted, which would allow for the amplitude information to be later recovered. However, for this design, it was decided to go with the simplest solution and allow for the gain to be adjusted at program time of the bat-tag. The disadvantage of this approach was that finding the proper gain setting that would allow weak echos to be resolved and at the same time not clip on strong outgoing sonar pulses would take some trial and error during field testing.

2.5.5 Sampling

In this design, the PSoC® 5LP's 12-bit SAR ADC's sampling frequency was chosen to be to 222 kHz. This frequency was actually chosen before the analog front end was designed and was chosen as a compromise between reducing the effects aliasing and reducing the amount of storage required to represent a bat recording. The original desired sampling frequency was 240 kHz to provide a sampling frequency of three times the highest frequency of interest, but at the time that this design choice was made, the clocking systems on the PSoC® 5LP were not as well understood by the author, and the closest sample rate that could be achieved was 222 kHz. This section explores the sufficiency of sampling at 222 kHz provided the simulated results from the bat-pass filter and the frequency response described in the microphone's datasheet.

When a continuous time signal is sampled by multiplying it with an infinitely long impulse train, each sample is separated in time by period of T, in this case, 4.5 μs as shown in figure 2.5.

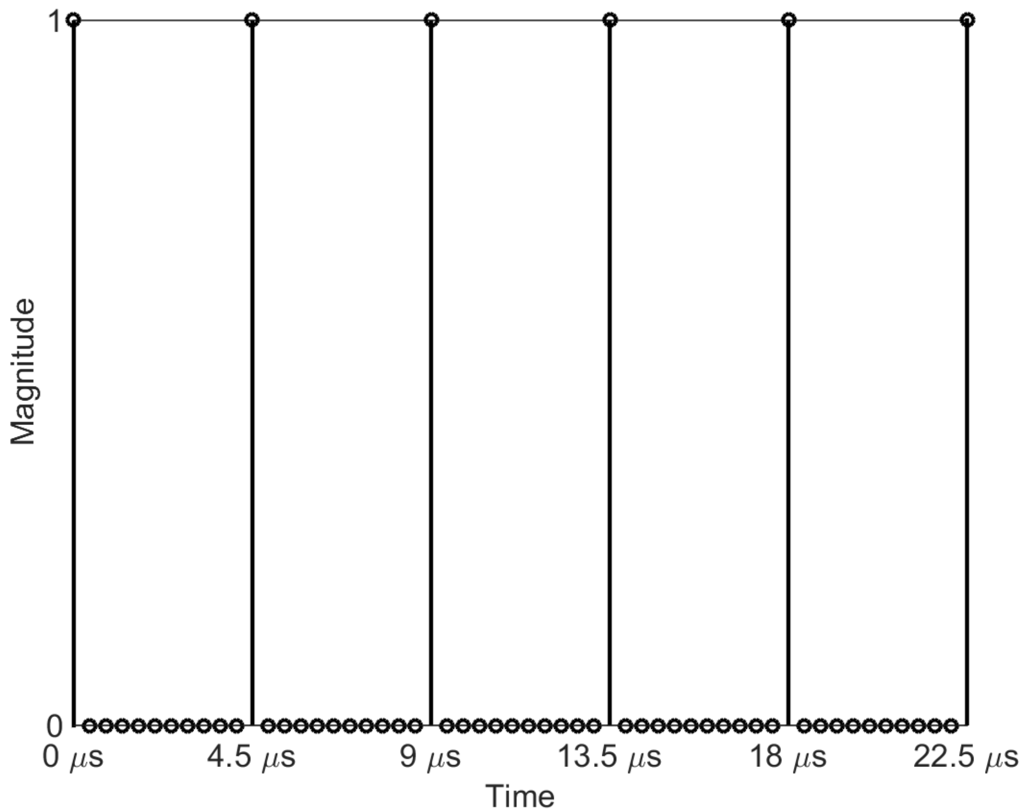

Figure 2.5: Time Domain Impulse Train

The frequency at which the samples are taken is equal to $\frac{1}{\mathrm{T}}$ which will be referred as $\mathrm{F_s}$. The frequency domain transform pair for an infinite impulse train in time, is a frequency domain infinite impulse train with each impulse separated by $\mathrm{F_s}$, for this particular application, that is 222 kHz as shown in figure 2.6.

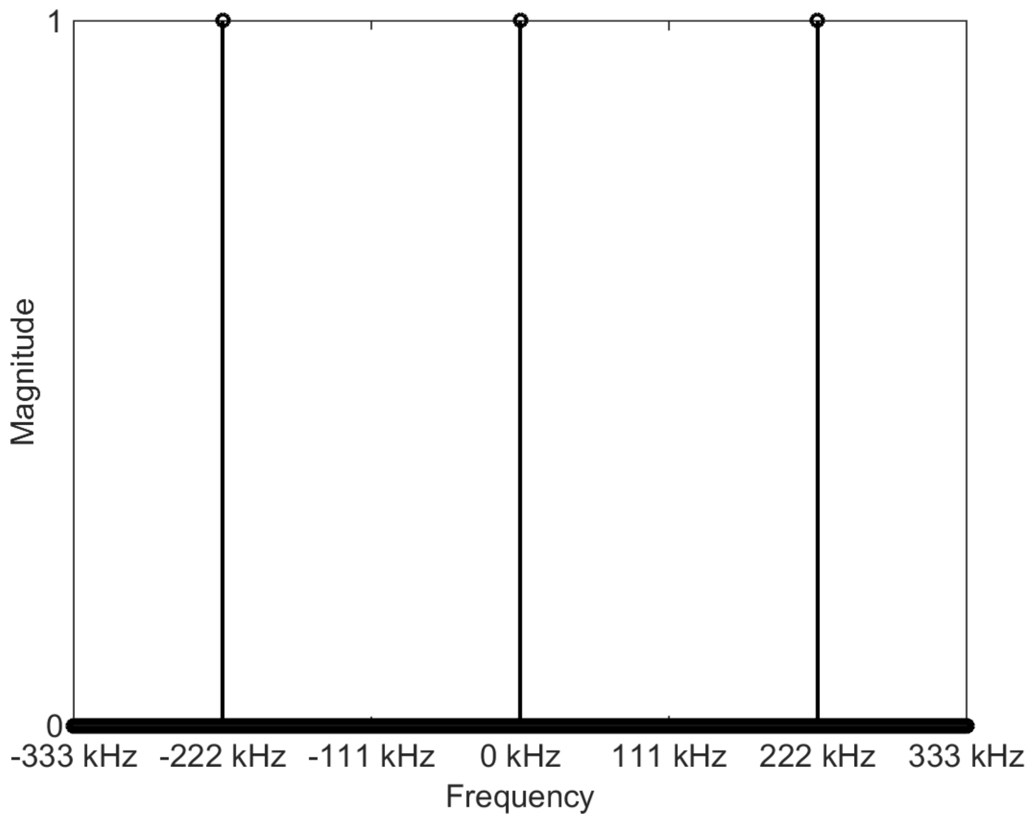

Figure 2.6: Frequency Domain Impulse Train

Since multiplication in time is equivalent to convolution in frequency, multiplying the time domain impulse train by the continuous time signal results in convolving the spectrum of the original continuous time signal with each impulse of the frequency domain impulse train. The impulse function has a special property known as the shifting property were convolving a function $f(x)$ with an impulse delayed by x_0 units from the origin will result in $f(x)$ being translated to $f(x - x_0)$. In the context of sampling a continuous signal with a sample frequency of F_s, the net result is that in frequency the original spectrum of the sampled continuous time signal is repeated every integer multiple of F_s. Since it is not possible to take an infinite number of samples, any sampled continuous time signal is said to be time-limited. However, a time-limited signal can not be band-limited, meaning that there will be a point in frequency in which the repeated multiples of the spectrum and the original spectrum of the continuous-time signal centered at the 0 frequency point will begin to overlap. This point at $\frac{F_s}{2}$ is were the phenomenon known as aliasing will begin to occur. Figure 2.7 shows the original spectrum and only the first repetitions at $-F_s$ and F_s, one can clearly see were the overlapping frequency components would begin to add together to cause aliasing.

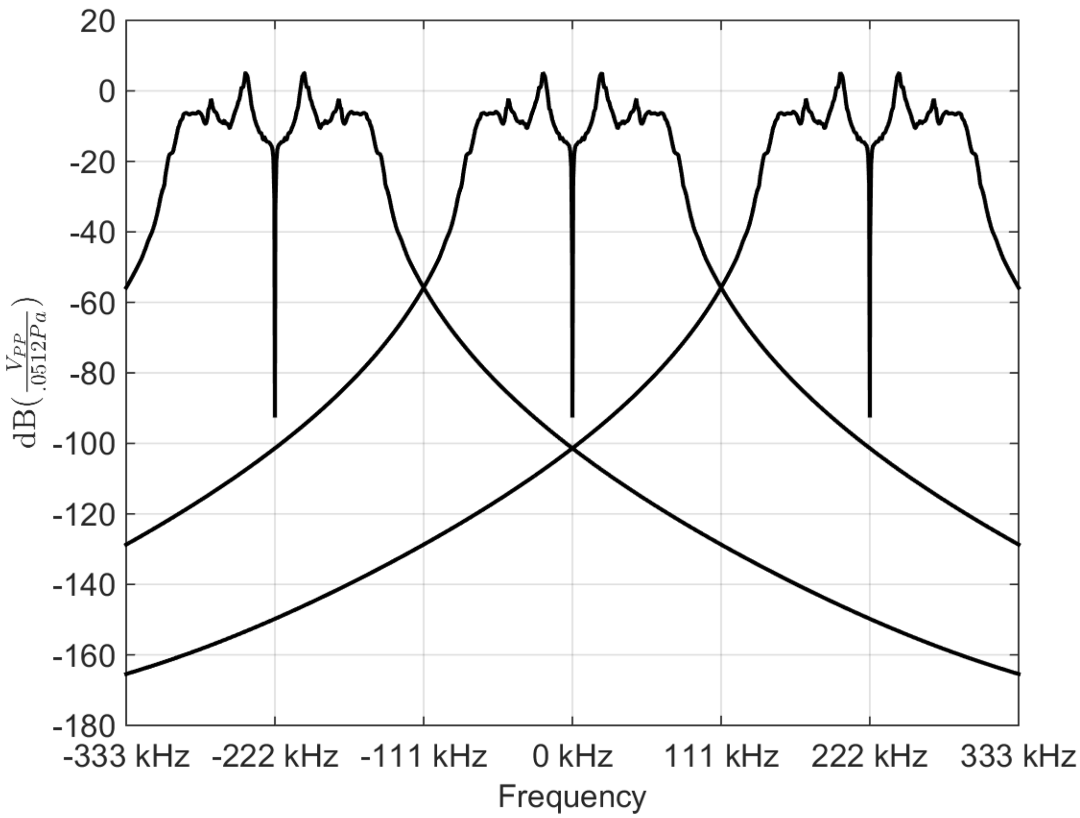

Figure 2.7: Frequency Domain Repeats

The extent of the severity of the aliasing is dependent on how much signal frequency content exists above $\frac{F_s}{2}$. To aid in this end, a low pass filter can be employed which will attenuate any part of a signal with frequency content above the filter's cut off frequency. The ideal frequency domain low pass filter can be represented by the rectangle function which preserves any frequency content within the rectangle, while completely removing any frequency content outside the rectangle. So long as the sampling rate is greater than the frequency at which the rectangle cuts off at, there will be no aliasing, and the parts of the time domain signal with frequency content contained within the bounds of the rectangle can be perfectly reconstructed in the time domain. However, constructing an ideal low pass filter would not be possible, as perfectly band-limiting a signal would imply that the signal would no longer be time-limited, therefore impossible to represent in any practical digital system. As such, one must settle for filters such as the ones described in section 2.2 which have frequency responses like the ones shown in figures 2.2 and 2.3. The low pass filter described in figures 2.2 and 2.3 are what approximately band-limit the frequency response of the analog front end. Which is why at the points were aliasing does occur as shown in figure 2.7, its effects are mitigated, but not completely removed. All of this discussion on sampling implies that aliasing is something that can not be completely removed; however, there are ways of quantifying how bad it will be and techniques to control it.

As described earlier, once a continuous-time signal is discretized, its spectrum will be repeated every integer multiple of the sampling rate. The filter shown in figures 2.2 and 2.3 described how much attenuation a signal would have at a given frequency. Real analog to digital converters have a finite resolution, and if one were to attenuate the frequency components of a signal that are at a frequency greater than or equal to $\frac{-\mathrm{fs}}{2}$ and $\frac{\mathrm{fs}}{2}$ to a point that their amplitude was less than half the resolution of the ADC, then aliasing would contribute no noise to the sampled signal. On the PSoC® 5LP using the 12-bit SAR ADC with (0 V - 1.8 V) as the valid input analog range yields a resolution of $\frac{1.8\ \mathrm{V}}{2^{12}}$ or 439 μV. This implies that in order for aliasing to not attribute any noise, all signals past $\frac{-F_s}{2}$ and $\frac{F_s}{2}$ need to be less than 220 μV or -73.2 dBV.

To help evaluate the performance of the bat-pass filter design, a model of the analog front end was created in MATLAB using the LTSPice simulation results of the analog filter described in 2.2.1 and information from the microphone's data sheet and application note. One assumption that had to be made was that of the units of the frequency response provided by the microphone's data sheet. The vertical axis of the magnitude of the frequency response quotes the units as $\frac{\text{dBV}}{.1\text{ Pa}}$; however, this would not have provided a linear response to signals at the same frequency with different amplitudes, as one would normally assume that an analog microphone would do, so it was instead assumed that the axis was poorly labeled and that the units were actually instead in $\text{dB}\frac{V_{RMS}}{.1\text{Pa}}$. The created model assumed that the highest gain setting on the filter was used, and simulated the effect of an input signal with uniform frequency content from DC to 10 MHz. Since the microphone's data sheet and application note did not provide information about its frequency response before 1 kHz nor past 100 kHz, its frequency response was assumed to be constant with respect to the last documented value on the edge of those ranges. Interpolation was used to give the model created from the microphone's data sheet the same number of points as the LTSpice simulation, in such a manner that both the LTSpice simulation results and the model of the microphone's frequency response would both share the same domain. This allowed for the frequency responses of the LTSpice simulation and the microphone's model to be multiplied together to form a model of the net frequency response of the analog front end to an impulse in time. The units of the model were in $\text{dB}(\frac{V_{PP}}{\text{Pa}})$ which would allow the effect of different strengths of inputs into the microphone to be evaluated. For the purpose of evaluating how much of an impact aliasing might have, the strength of the simulated input was selected so that the frequency component with the largest gain would have a rail to rail output of 1.8 V. In other words, the strongest input at the largest gain setting that would not result in clipping was used to simulate a 'realistic' worst case scenario. The simulated net frequency response of the analog front end is shown in figure 2.8.

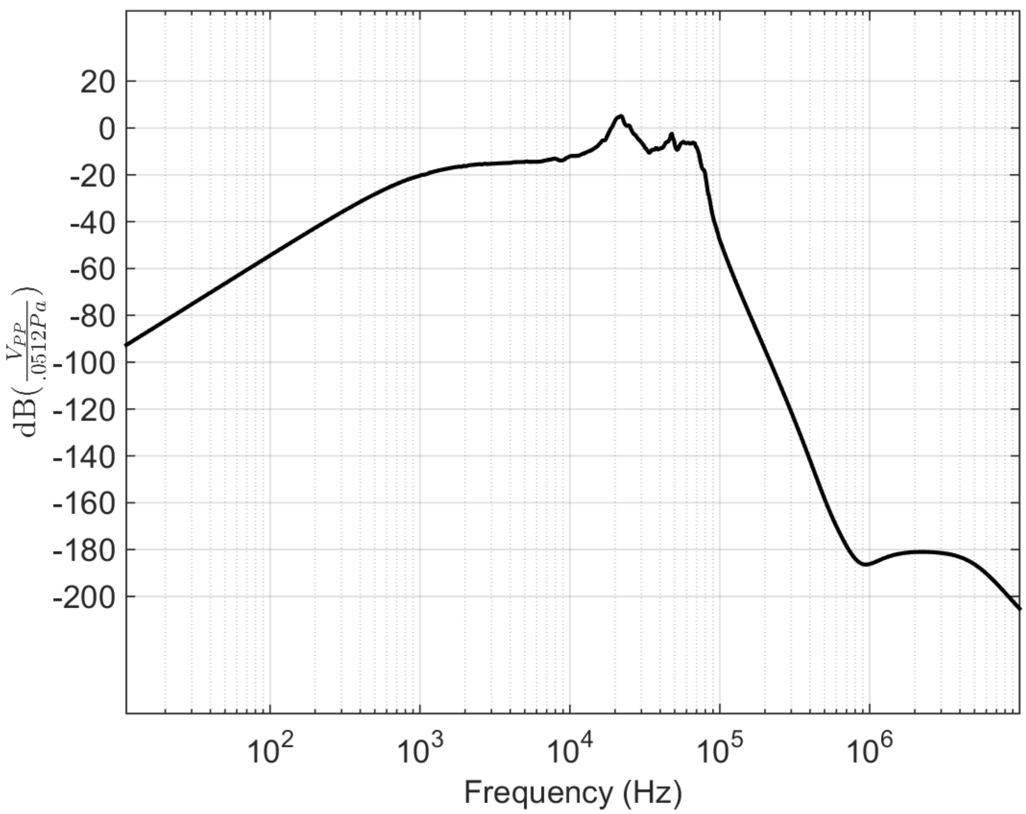

Figure 2.8: Net Analog Front End

In this simulation, the -73.2 dBV frequency was at 143 kHz which would require a sampling rate of at least 286 kHz to ensure that all aliased components would be an amplitude of less than half of the resolution of the ADC. However, as stated in 2.2, only frequencies up to 80 kHz are of interest, meaning that aliasing can be tolerated as long as the aliased frequency components aliased at frequencies between 80 Khz and $\frac{F_s}{2}$, as anything in that range could be removed in post-processing of the recorded data. Figures 2.9 and 2.10 respectively show how many ADC counts at a given frequency would be distorted by aliasing and the frequency at which a particular aliased component would show up at.

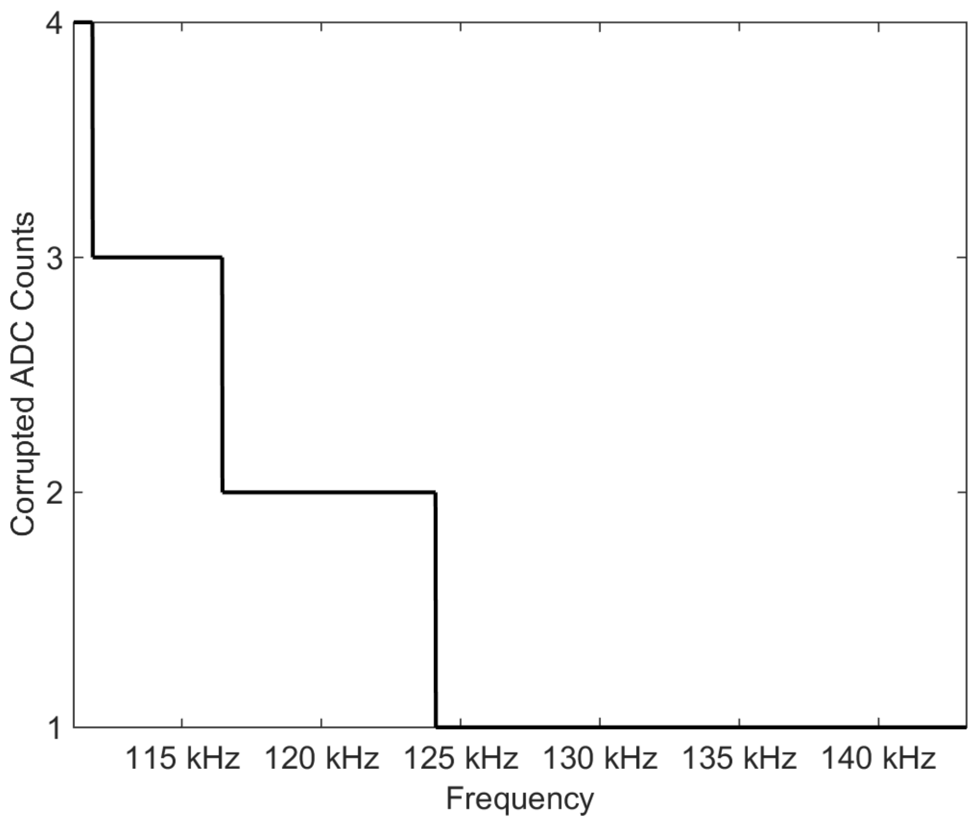

Figure 2.9: ADC Counts Corrupted by Aliasing

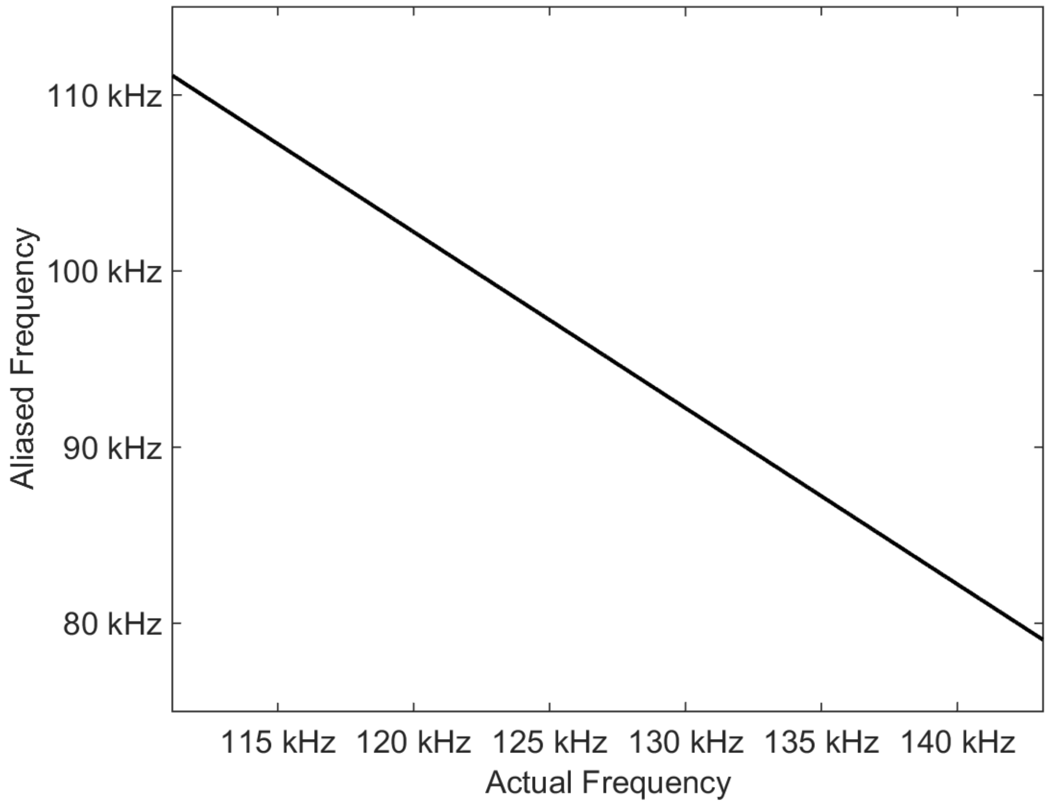

Figure 2.10: Frequencies Causing Aliasing

A sampling rate of only 223 kHz would be required to prevent aliasing from manifesting in the sampled data at 80 kHz or below in frequency, which puts sampling at 222 kHz well in range of only corrupting a single ADC count at 80 kHz. Furthermore, the point in which aliasing would be worst, at 111 kHz would result in only four ADC counts being corrupted. An additional MATLAB simulation was performed, to help further understand the effects of sampling the output of an analog filter at a specific sampling rate which was provided an input with equal amplitude across all frequency. The frequency response of the analog front end was first centered at 0 Hz so that it extended to 10 Mhz as described earlier in this section. Next, it was mirrored across the vertical axis so that the frequency response now extended from -10 MHz to 10 MHz to represent the fact that a real time domain signal is symmetrical in magnitude in the frequency domain. Next, copies of this frequency response were shifted every positive and negative multiple of the sampling rate, then multiplied by the magnitude of the un-shifted two sided frequency response at those particular shifted frequencies. The now attenuated copies were summed together along with the original unweighted/unshifted frequency response. The resulting frequency response was plotted from 0 Hz to 10 Mhz and showed how the frequency response of a real analog filter would be altered once it's output was discretized. This frequency response is shown in figure 2.11.

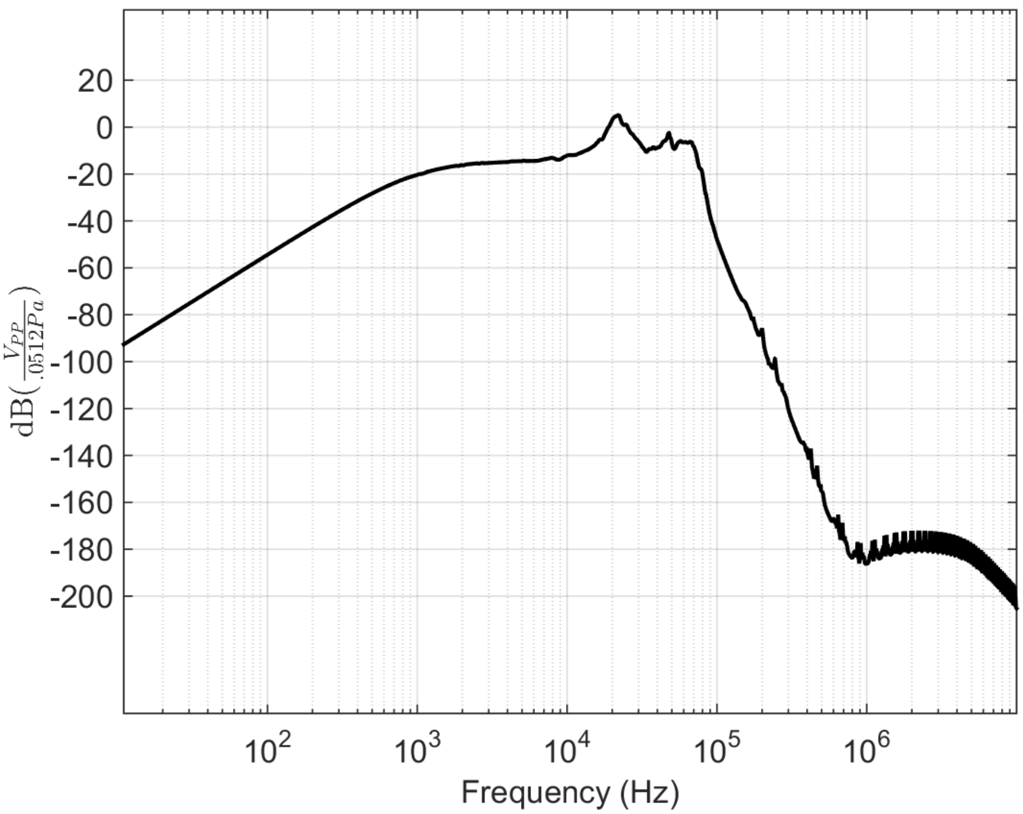

Figure 2.11: Net Analog Front End Discretized

Figure

The NAND flash had a standard Open NAND Flash Interface (ONFI), and writing a page required that first a page program command be sent, followed by 5 address bytes, then 2048 data bytes, concluded with a final command byte. This entire sequence was performed over a 16-bit bus, 8-bits for control logic, and 8-bits for commands, data and addresses. Writing a flash page required 2055 bus cycles, followed by the time for the flash chip to actually program the page to it's internal memory [12]. In order for the entire recording system to work as intended, the fundamental challenge that must be overcome is the time to send a page's worth data over the flash bus and for that data to be written to the flash must be less than the amount of time it takes for the ADC to generate a flash page's worth of data. This meant that the code responsible for controlling the flash bus had to be very carefully written. Function calls had to be either in-lined by the compiler, or in-lined by hand. Particular attention had to be paid to the loop which was responsible for writing the 2048 data bytes to the flash chip as it dominated the time to write the data over the flash bus. Improvements to the execution time of the source code were tested by observing flash-bus activity on an oscilloscope. If a change in the source code resulted in a page write bus transaction occurring faster, then the change was kept. An alternative approach could have been to increase the clock speed of the PSoC® 5LP during data recording; however, this would increase the amount of current required to operate the bat-tag, so it was advantageous to operate the PSoC® 5LP at the lowest clock speed possible, while steel meeting all real time requirements.

2.5.6 Radio Handshake

One of the other challenges to overcome in this project was determining when the bat-tag should wake up from its low power sleep mode and begin transmitting it's previously recorded data to the base station. The radio used, the CC1101, introduced in section 2.3 featured a low power, Wake up On Radio (WOR) mode which allowed

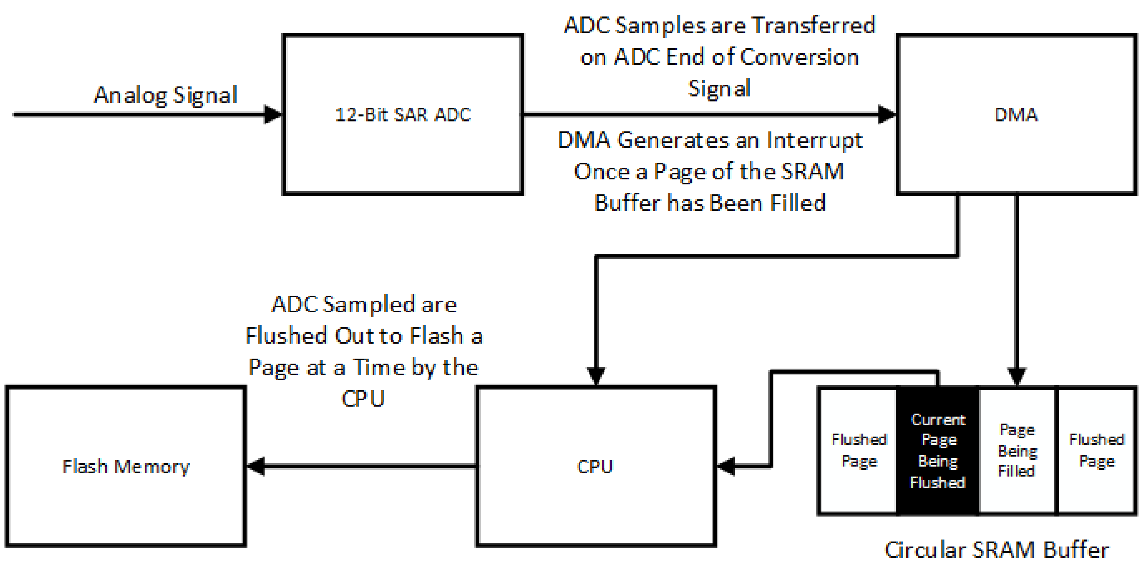

Figure 2.12: Sampling Data Path

the radio to alternate between low power sleep mode, and receive mode in a predetermined duty cycle. If the radio received a packet while it was in receive mode, it would stay in receive mode, and then generate an interrupt that would be sent to the PSoC® 5LP and wake it out of it's low power sleep mode, thus notifying the bat-tag that it was in rage of the base-station. To show how the WOR saved power over instead just staying in receive mode waiting for a packet from the base station, the sleep mode is documented as drawing 200 nA and the receive mode is documented as drawing around 18 mA [24]. So, if the radio was configured to sleep for 60 seconds, then sit in receive mode for 3 ms, the average current draw would then be 1.10 μA, low enough that if the radio was the only device powered by an ideal 1.8 V 150 mAh battery and powering the radio from the battery was perfect it could run for over 15 years. After the micro-controller has been woken up by the radio, a hand-shaking process between the base station and the bat-tag is initiated. The base station alternates between transmitting hailing packets and listening for a response from the bat-tag. Once the bat-tag has been awoken, it transmits a series of packets to let the base-station know that it was awoken then enters receive mode. Once the base-station receives one of these packets it transmits a series of packets to let it know that it's OK to begin data transmission, and then enters receive mode. Finally, the tag receives one of these "TX OK" packets, waits for a predetermined delay, and begins transmitting it's recorded data. Prior to this last step, if the base-station does not get a timely response, it reverts back to it's "Hailing TX" state and if the bat-tag doesn't get a timely response, it reverts back to it's WOR state.

2.5.7 Radio Transmission

Once the bat-tag had completed its radio handshake, it would begin the process of transmitting the stored data on the flash chip. It would do this by reading a flash page's worth of samples, 2048 bytes, loading the 64 byte transmit FIFO of the radio over a SPI interface, and then switching the radio to transmit mode by sending a command over the SPI interface to the radio. After 32 radio transmissions had been completed, the next flash page would be read, and the process would continue until either the battery powering the bat-tag died, or the entire contents of the flash chip had been transmitted. Refer to figure 2.13 for a diagram of the radio's data path. At the time of the writing of this thesis the data transmitted over the radio was 12-bit ADC samples encoded into 16-bit unsigned integers, leaving the upper 4 bits unused. There are plans for in the future to use the extra bits as space to implement a (15, 11) Hamming Code for single bit error correction of the 11 most significant bits of every ADC sample [9]. In this technique, a new 16 bit transmit word would be constructed by first setting the least significant bit of this new word to value of the least significant bit of the ADC sample. Next, 4 parity bits would be calculated from the 11 most significant bits of each ADC sample and would be combined together to form the upper 15 bits of the new word. To avoid the overhead of having to do this computation for every ADC sample that is transmitted over the radio, the finite range of the ADC output will be exploited. Since there are only 4096 possible outputs from the 12-bit SAR ADC, every possible transmit word could be stored in a table keyed on ADC values. This table would be precomputed in a program such as MATLAB and stored in an array of unsigned 16-bit integers in the SRAM of the PSoC® 5LP. This would mean that at run-time, the only necessary operation to 'compute' a new transmit word, would be to index an array. It is hoped that this way, the integrity of the transmitted data would be increased while having a minimum impact on radio throughput.

2.6 Printed Circuit Board

The first development prototype board had designed dimensions of (41.8 mm X 47.5 mm, 199 mm^2) and out of 5 bare boards had a measured average weight of 6.07 g. This made it clear from the beginning of the development of the deployment bat-tag that if left unchecked, PCB weight could overwhelm the entire system weight budget.

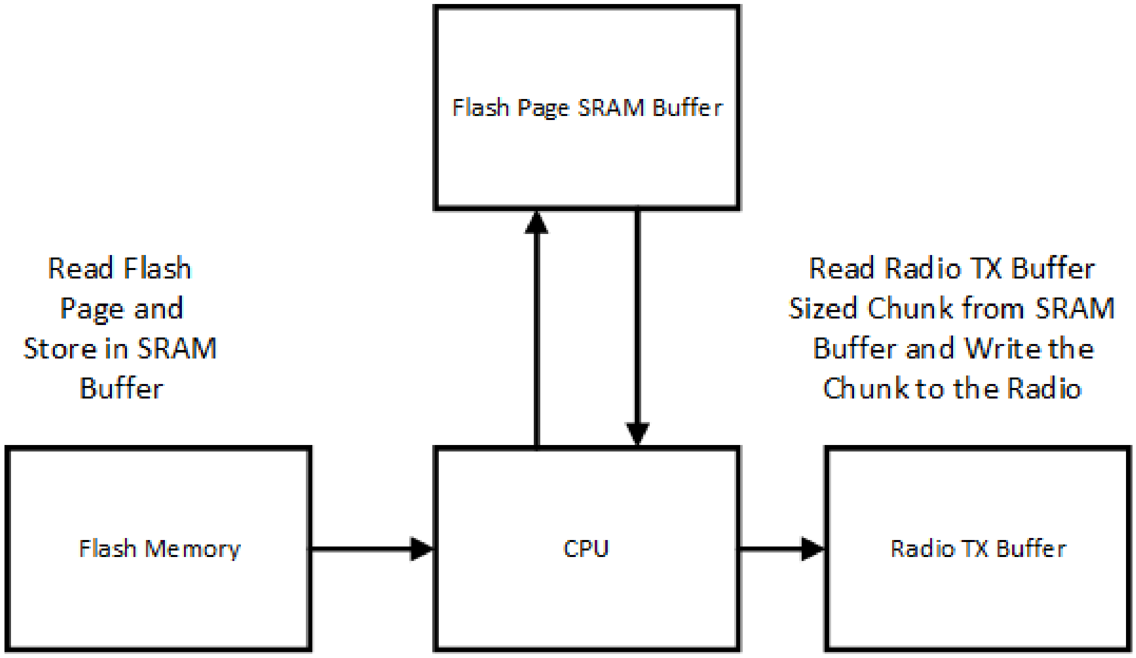

Figure 2.13: Radio Data Path

While this first prototype was designed more to evaluate the hardware rather than for minimum weight it did demonstrate how much care must be taken in the PCB design in order for the system to meet the design requirements. Minimizing the PCB weight was achieved in a few steps. The first step was getting rid of extra devices where possible. For example, the first prototype had a "gas gauge" circuit that could be used to let the micro-controller measure the voltage across the battery to allow the micro-controller to enter a low power state once the battery voltage had crossed a critical threshold. The firmware to implement this feature was never written, and from a functional standpoint, having a working gas gauge wouldn't be particularly useful since the data acquisition phase of device operation was expected to draw very little current compared to the data transmission phase of device operation, and in the intended use case, once data transmission began, data would be sent until either the entire memory was transmitted, or the battery died during transmission. This implied that it was unlikely that during data acquisition to prove advantageous to to stop gathering data based off of a battery event, to enter a low power sleep mode in preparation for a later data transmission. Another such optimization that was made to reduce board PCB weight was to replace the matching network composed of 14, 0402 passives described in the CC1101's data sheet with a 0805 sized single component matching network by Johason Technology [24][13]. The next step in minimizing board weight was to employ a non standard board shape. While many PCBs are quadrilaterals, having a PCB made in any number of nonstandard shapes is no issue for some board houses. During board layout, the components were first roughly grouped by digital and analog and then orientated in such a way as to ease routing difficulty. Just as importantly though, the components were placed so that the largest components were situated in the middle of the board, and the antenna and microphone were situated at opposite ends of the board. This allowed the board to be wider where more area was needed and to be tapered towards to ends to the extent that the antenna stuck out on it's own with just enough board underneath it to provide mechanical stability. The inspiration at the start of layout was the shape of a torpedo, narrow near the antenna, and narrow in the rear for the microphone so that the antenna would stick out like a tail for the bat, and the microphone would stick out closer to the head of the bat. Additionally, in this design, where possible both sides of the board were used for components. Care was taken to place nothing under the radio, crystal, nor antenna. Additionally, measures were taken to keep digital signals from being routed under or near sensitive analog signals. For example, the flash memory chip was actually placed on the side opposite of the micro-controller; however, it was placed in such a way that it did not overlap the micro-controllers analog portion. The final measure taken to reduce board weight was to reduce the thickness of the board. The first prototype had a measured average mass density of 3.05 $\frac{\text{g}}{\text{cm}^2}$ with a 61 mil board thickness. As mentioned in the tag weight section 2.1 a MATLAB script was written to help estimate the mass of a printed circuit board. The user would define the vertices of their printed circuit board in an array, moving in a clock wise fashion with the bottom left most vertices being referenced to (0, 0) in mils. Next the user would provide information about the board stack up, in particular, oz of copper in each signal layer, mils of FR-4 between each copper layer, and the estimated copper area density per copper layer. For example, in the bat-tag the following specifications were used, 1 oz copper for the outer signal layers, .5 oz copper on the internal copper layers, 3.5 mm of FR-4 between the signal and internal layers with 18 mm of FR-4 between internal layers. Next, it was specified that the top and bottom signal layers by area were composed of 10 percent copper and the internal copper layers being composed of 80 percent copper by area. The script first calculated and output the area of the printed circuit board and a figure showing the outline of the board specified by the vertices input by the user 2.14, next it calculated the volume of copper and FR-4 in the specified board, used the specific gravities of copper and FR-4 to estimate the weight of the specified board. Additionally, the information from this script showed that the weight of a 62 mil thick 4-layer board with 1 oz copper outer layers and .5 oz copper inner layers was dominated by FR-4 and not copper as one might intuitively expect. Using this knowledge it was decided to use a 31 mil thick board instead of the standard 62 mil thickness with 1 oz copper on the outer layers and .5 oz copper on the inner layers. The exact thickness of the finished board and it's comprising layers is subject to both manufacturing tolerances and the specific layer stack up used by the board house. Using information from Sierra Circuit and from this particular board layout, it was predicted that the bat-tag should have a finished board weight of 0.967 g with an area of 647 mm^2. This prediction was of course subject to manufacturing tolerances, and the accuracy of the parameters provided to the script.

2.7 Power

The board was supplied by a single CP251525 3 V Lithium/Manganese Dioxide primary battery. This particular battery was chosen because of it's 150 mAh capacity, 80 mA max pulse discharge current and 40 mA continuous discharge current, compact dimensions (25.5 mm X 15.5 mm X 2.60 mm), and being quoted at 2.0 g, light weight.

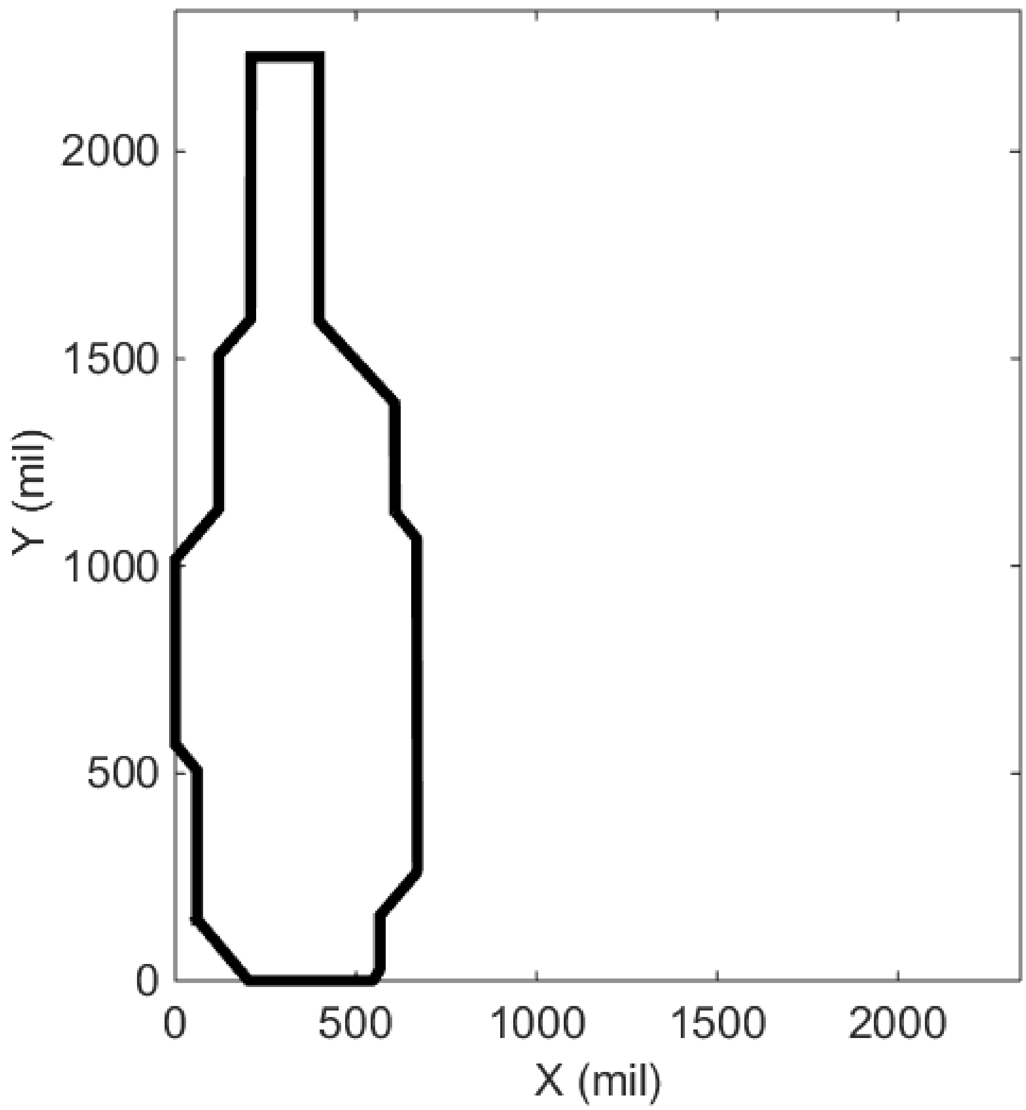

Figure 2.14: Bat-Tag Supplied Board Coordinates

Beneficially, the examples of this battery weighed in the lab weighed closer to 1.5 g [11].

The battery was used to power two 1.8V ADP172 linear regulators, under no load the regulators are quoted as drawing a 23 μA quiescent current, while under a 150 mA load they are quoted as drawing a 130 μA quiescent current. Additionally, under a 150 mA load, the regulators are complemented by their 25 mV dropout voltage. Most critically, the regulators themselves came in a wafer level chip scale package and were only (.95 mm X 1.025 mm) and required only two 0402 capacitors each in terms of passives [1].

As mentioned in section 2.6, the PCB was designed with separate analog and digital power domains. The power domains were created by breaking the ground and power layers of the PCB into separate planes. The ground layer was broken up into two planes, VSSD and VSSA, digital and analog ground respectively. However, the power layer was broken up into three planes, VDDD, VDDA, and Switched VDDA. VDDD and VDDA were digital and analog power as one might expect; Switched VDDA was created by breaking VDDA into two separate planes and connecting the planes together using a Si1013R/X P-Channel MOSFET (P-FET) and using one of the GPIO pins on the micro-controller to switch the MOSFET on and off as needed. VDDA, was used to power components that must always remain powered such as the analog system of the micro-controller, while the other plane, Switched VDDA, was used to power the analog front end. According to LTSpice simulations of the analog front end and the microphone's data sheet, by switching the supply to the analog front end off, it should be possible to reduce the total current draw of the system by 441μA [14]. While 441μA is a small amount of current compared to the 17.2 mA used by the radio while transmitting, it's important to be able to bring the sleep current of the entire system to a minimum so that the bat-tag can wait in a low power state for extended periods of time while waiting to come into range of the base station [24].

2.8 Post Fabrication Design Changes

2.8.1 Switched Analog Front End Power Supply

During testing of the micro-controller, it was discovered that switching the P-FET on that allowed the switched analog front end's power supply to be turned off and on would cause the analog supply voltage of the micro-controller to experience a transient drop and would then subsequently result in the micro-controller resetting it self. It was actually possible to for the bat-tag to enter what appeared to be an infinite cycle of turning the analog front end on and resetting it self. To get around this issue, the P-FET was removed and instead of using one of the micro-controller's GPIO pin to switch the boards main analog supply to the analog front end's supply, the GPIO pin was directly connected to the analog front ends voltage supply plane, allowing the analog front end to be powered directly by the GPIO pin instead. This allowed for the analog front end to continue to be switched off and on as needed, and resolved the issue of the bat-tag resetting itself. However, it should be noted that this solution is not optimal as the analog front end is now supplied by a digital pin while still using the analog ground plane as it's reference.

2.8.2 Analog Front End Gain Control

During testing of the analog front end, it was discovered that controlling the gain of the analog front end using the rheostat in the feedback loop of the first amplifier would result in unpredictable behavior. When the prescribed resistor of 1K was used on the inverting input and the resistance of the feedback loop was adjusted using the rheostat, the gain would stay constant. When the resistor on the inverting input was changed to a 5k resistor and the rheostat was adjusted, the gain would adjust over a small subset of the values it was supposed to take on, and would become unstable at lower gain settings. It was decided to replace the rheostat with a fixed resistor which final value would be determined after field testing. Once the rheostat was replaced with fixed value resistors, the analog front end's stability issues went away and the experimental gain values tracked the simulated gain values reasonably well. It is hypothesized that the rheostat caused parasitic interactions with the gain stage amplifier effecting the stability of the amplifier. As of the writing of this document, the observed experimental behavior has yet to be recreated in simulations using the provided models for the amplifiers.

Chapter 3

Experimental Methodology

3.1 Tag Weight

All weight measurements were taken using a AWS-100 digital scale calibrated with a 100 g weight. The scale was marketed as being capable of measuring weights as small as .01 grams. Measurements were taken of the bare printed circuit boards, assembled printed circuit boards, batteries, and assembled printed circuit boards with battery and wires wrapped in water resistant heat shrink film.

3.2 Analog Front End

The analog front end's frequency response was measured using an SR785 Dynamic Signal Analyzer. The microphone was removed from the bat-tags and a lead was soldered onto the pad were the microphone's output was routed. Another lead was soldered onto the pad of a surface mount resistor to access the output of the analog front end. The source output of the signal analyzer was attached to the microphone's output lead, the network analyzer's Channel One port, and to an oscilloscope. The analog front end's output lead was attached to the network analyzer's Channel Two port, and to another channel on the oscilloscope. Next a swept sinusoidal output was generated from the network analyzer's source port. As the frequency of the sinusoid was swept, the output was monitored on the oscilloscope to ensure that the analog front end's output was not distorted nor clipped. Meanwhile, the network analyzer compared the signal it received on it's Channel Two input to that of the signal received on it's Channel One input to calculate the gain at a particular frequency. Sweeping the source frequency from 10 Hz to 100 kHz resulted in the experimental measurement of the transfer function of the analog front end within that respective frequency range.

3.3 Radio

The bit error rate of the radio was measured experimentally out in an open grassy courtyard in the late evening. A psudeo-random sequence of 512 bits was transmitted by the bat-tag repeatedly until a total of 100 million bits had been sent to the base station. Once the data was recorded, it was stored on a laptop so that bit error rates could computed at a later time in MATLAB. This test was first conducted with both the bat-tag and base station sitting close to the earth at the same elevation from the ground at distances of 1 m, 2 m, 3 m, 4 m, and 5 m away from the base station. Next, the bat-tag was attached to the top of a 2 m long pole and the experiment was repeated at distances of 1 m , 2 m , 3 m, 4 m, 5 m, 10 m, 15 m, and 20 m away from the base station.

In other in-lab experiments, the effective throughput of the radio was estimated by observing the rate at which the base station was storing data to it's non-volatile storage. While this not a precise measurement of the radio's throughput, it should be close because the radio is by far has the slowest specified throughput of any device in the datapath (350 $\frac{\text{kbit}}{\text{s}}$). Thus, the storage data rate of the base station should be close to the effective data rate of the transmitting radio.

3.4 Non-Volatile Memory

During the recording of audio data to the flash, the flash's chip enable pin was monitored on an oscilloscope. As the chip enable was dropped low at the start of a page write and not driven high until after an interrupt had been serviced notifying the micro-controller that the flash page had been written, the amount of time the chip enable pin was low was indicative of how long it took to write a page. Knowing the number of bytes that were written and how long it took to write the number of bytes, it was then possible to calculate an effective net flash write throughput and write latency.

3.5 Current Measurements

Current measurements of the bat-tag were taken by supplying power to the bat-tag through a Keithley source meter. The bat-tag was programmed to perform the tasks that it would spend a majority of it's operational time performing such as recording data and using the radio to transmit data. Current measurements of the bat-tag performing operations such as erasing the flash memory or calibrating the radio were not taken as a significant amount of time nor energy would used performing those operations in deployment. It should be noted that the measurements taken by the Keithly source meter were indicative of average current consumption, this was verified previously by monitoring currents of an older version of the bat-tag using an oscilloscope and comparing the average value of the current observed on the oscilloscope with that which was displayed by the Keithly source meter.

3.6 Full Data Path Tests

The following three tests were designed to test the performance of the bat-tag's entire datapath specific particular conditions. In all cases signals from the microphone were first filtered and provided gain through the analog filter circuit. Next, the signals were discretized into samples by the micro-controller's 12-bit SAR ADC and buffered in SRAM before being stored in the external flash memory. After all of the desired samples had been stored, they were then read from the flash memory and transmitted over the radio to the base station. It should be noted that to expedite the gathering of experimental data, the radio handshake between the base station and tag were disabled. In all of these experiments, a 49.9 kΩ resistor was used in the feedback loop of the gain stage of the analog filter

3.6.1 System Noise Test

This experiment was designed to help gain insight into how much noise could be expected to be recorded in a quite environment. In this experiment, the bat-tag was placed inside of a quiet room and left to record 20 seconds of data before being carried in range of the base station. In this case, the quiet room was actually an anechoic chamber built inside of a Faraday cage designed to block 2 GHz rf signals. Qualitatively speaking, the anechoic chamber seemed to attenuate outside sound which is why it was chosen as the location for this experiment.

3.6.2 Sonar Data Acquisition

This test was designed to evaluate the ability of the bat-tag to resolve both high and low amplitude signals. Outgoing bat echolocation calls can be quite loud; however, the returning echos can be very faint, so this experiment was particularly important.

A commercially available sonar range finder was used in this experiment to simulate bat calls. The range finder produced 40 kHz sonar pulses every 100ms. To test the ability to capture an echo, the range finder was placed 4.25 m away from a wall with the bat-tag being placed between the wall and sonar range finder. Specifically, the bat-tag was 1.75 m away from the wall, and 2.5 m away from the sonar range finder. This would allow allow for a sonar pulse to pass by the bat-tag, reflect off the wall, and pass again by the bat tag in an approximately 10 ms interval.

3.6.3 Effect of Sampled Signal from Wrapping in Film

During deployment of the bat-tag, some method of protecting the bat-tag from moisture and rain is required. To this end, the bat-tag will be sealed inside of heat shrink film while it is in field use. However, it is necessary to understand what effect that this will have on the attenuating the output of the microphone. This was tested by using the same experimental setup as described in section 3.6.2, except for one recording the bat-tag was wrapped in heat shrink film. The sonar range finder produced very strong outgoing sonar pulses that would cause the output of the analog font end to clip if the sonar range finder was placed to close to the bat-tag. For this reason, when determining how much signal attenuation was caused by adding heat shrink film to the bat-tag, it was the sonar echos that were analyzed.

Chapter 4

Experimental Results

4.1 Tag Weight

Table 4.1 shows the measured weights of 10 batteries, 10 bare boards, 10 assembled boards, and 10 assembled boards wrapped in heat shrink film along with the corresponding mean and standard deviation for each set of measurements. Each column of data should be interpreted independently from every other columns of data. For example, the "Assembled Board" in sample three is not intended to be the same "Assembled Board" used in sample three for "Assembled With Battery and Film" entry.

Table 4.2 shows measurements taken of some of the more massive individual components. Only a single sample of each of the components were weighed due to the limited availability of spare bat-tag parts in the lab. The flash chip weighed was of a 4 Gb varient of the 8 Gb flash chip used on the bat-tag, and the mass quoted for the PSoC® 5LP came from the vendor's website [7].

The bare printed circuit boards on which all of the bat-tag's components were mounted weighed approximately 1.09 g each. The MATLAB based bare printed circuit board weight simulation discussed in section 2.6 predicted that the bare boards should weigh approximately 0.97 g, thus the measured result validates the MATLAB simulation. Most importantly however, in section 2.1 it was discussed that an adult examples of *Phyllostomus hastatus* could be expected to weigh anywhere between 75 g for a smaller female and 100 g for a larger male and that a sensor attached to the bat should not weigh more that 5 percent of the bat's own weight. This gave a target bat-tag weight range of 3.75 g to 5 g for the fully deployment ready bat-tag. With the bat-tag weighing in at 3.65 g, .1 g less than the smallest target weight, it is suitable for deployment even on smaller female examples of *Phyllostomus hastatus*.

Figure 4.1 shows the bat-tag and battery placed next to a penny for size reference.

4.2 Analog Front End

Figure 4.4 shows the approximate maximum non-clipping output voltage of the analog filter with the microphone removed. The yellow trace is the analog filter's output voltage and the green trace is the output from a function generator that was divided down using a resistor divider and fed into the input of the analog filter. The yellow trace shows that the analog filter is capable of generating an approximately 1.8 VPP output, which would allow for most of the dynamic range of the ADC to be used. It must be noted that this measurements was taken with the microphone not as the input into the system and that the microphone will begin to generate a distorted output in the presence of excessively high sound pressures. The microphone's data sheet quotes it's total harmonic distortion at 100 db SPL as being 1% and at 115 db SPL as being 10% [14]. The quoted total harmonic distortion of the microphone could be a concern when capturing the outgoing sonar pulses of *Phyllostomus hastatus* as the pulses have been estimated to be as loud as 112.3±1.6 dB SPL 10 cm away from the mouth of the bat [21]. So, it is very possible that field test will show that the effective output range of the entire analog front end is less than that of the analog filter.

Figures 4.5, 4.6, and 4.7 show the measured and simulated frequency responses of the analog front with varying feedback resistor values adjusting the gain. The simulation curves shown here are different than the ones shown in figures 2.2 and 2.3 as they used different resistor values for the feedback loop and those simulations included the resistance that the microphone itself contributed to the node on the inverting input of the gain amplifier.

At the lowest tested gain setting with a 1 kΩ feedback resistor, the filter frequency response peaks significantly higher than the simulation at the low pass cut off frequency. This 13 dB peak suggests that this configuration may be potentially unstable and thus it should not be used in deployment. The tests with the 11.5 kΩ and 49.9 kΩ feedback resistors track their respective simulations well and do not suffer from the same excessive peaking as the configuration with the 1 kΩ feedback resistor. This suggests that feedback resistors in the entire range of (11.5 - 49.9) kΩ could be used to provide a range of gain configurations with a reasonable expectation of predictability and success. The limiting factor on setting the gain yet higher is the fact that bandwidth of the amplifier is dependent on how much gain is provided, so it is possible that even higher gains could be provided, so long as it was understood that desired high frequency content could be lost to the bandwidth limitations of the amplifier.

4.3 Radio

The results from the radio's bit error rate test are shown in figure 4.8. The vertical axis is the number of bit errors found in the received data while the horizontal axis shows how far the bat-tag was from the base station in meters in terms of 1-D distance. The series in dotted lines shows the results of leaving the bat-tag and base station both roughly the same distance from the earth as distance was varied. The series in the dotted lines shows the results of leaving the base station on the ground and attaching the bat-tag to the top of an approximately 2 m tall pole as distance was varied. Series with circular markers counted every bit in a dropped radio packet as a bit error, while the series with triangular markers only counted bit errors in the received data.

While the raw radio data rate was set to $\frac{350 \text{ kbits}}{\text{s}}$ the observed data rate of the radio when there was a low rate of dropped packets was observed to be around $\frac{250 \text{ kbits}}{\text{s}}$. This is because of several factors. First of all, one must consider the time involved with writing a packet's worth of data to the radio's TX FIFO over the SPI interface. Secondly, and most importantly, the raw $\frac{350 \text{ kbits}}{\text{s}}$ dis not reflect the fact that in the radio's current configuration it must transmit 8 preamble bytes and a 4 byte sync word for 64 bytes of data transmitted allowing for a max theoretical throughput of $\frac{295 \text{ kbits}}{\text{s}}$ if the FIFO was always kept full. While the radio's raw data rate could be increased, and the number of preamble bytes could be reduced to increase the realized data throughput, it was found that these settings provided a good balance between bit error rate and data throughput given the provided radio.

The most important thing to take from the data throughput result is that it would take approximately 8.87 hours of transmission time to send the entire contents of the flash memory chip to the base station which is according to table 4.3 would consume more energy than what the entire battery can provide. A radio that could transmit data at a higher throughput using a similar amount of current would perform far better in this application.

4.4 Non-Volatile Memory

Figures 4.2 and 4.3 shows the chip enable pin on the flash chip being monitored while ADC samples are being moved in real time from the SRAM buffer on the microcontroller to the flash chip. The time while the chip enable pin is low shows the time when the ADC samples are being written over the bus and the time required for the flash chip to actually store those samples. At a CPU clock frequency of 12 MHz writing a page to the flash chip takes 2.76 ms. The time between the microcontroller attempting to write pages to the flash chip is 4.52 ms, meaning that there are 1.76 ms of spare time between flash page writes at this CPU clock frequency. This information can be used to help validate that the real time constraints of the digital datapath are being met. The DMA controller is programmed to generate an interrupt whenever 1024 samples have been written into the circular SRAM buffer alerting the micro-controller that it needs to write a page from the buffer into the flash memory chip. At a sampling rate of $\frac{222{,}222 \text{ samples}}{\text{s}}$ raw data should be generated at a rate of $\frac{444{,}444 \text{ bytes}}{\text{s}}$, a rate that should be enough data to fill a 2048 byte flash page in 4.61 ms. The fact that flash pages begin being written every 4.52 ms suggest that data is being moved from the ADC to the SRAM buffer and that the interrupt notifying the micro-controller to write the flash are all functioning in a timely manner. Just as importantly, the fact that there are 1.76 ms left over between flash page writes shows that once samples are gathered in the SRAM buffer, that are able to be flushed out to the flash memory in a timely manner.

4.5 Current Measurements

Table 4.3 shows the current drawn by the bat-tag in various states. "Idle, Analog Front End Off" and "Idle, Analog Front End On" were recorded with the micro-controller spinning in a while(1) loop with the analog front end switched off and on respectively. The purpose of having the micro-controller spin in a while(1) loop was to provide a set of simple instructions for the micro-controller's processor to execute that would provide a set of test conditions to represent an "idle state" with the micro-controller not actually sleeping. This allowed for the current being drawn by the analog front to be estimated as approximately 360 μA which compares reasonably to the simulated 441 μA described in section 2.2. Additionally, this suggests that the micro-controller itself appears to be drawing more current than expected. Section 2.5 describes that the current used by the micro-controller while recording data should be 5.75 mA, while the measurement shows that the bat-tag drawing 9.44 mA, which should be mostly contributed by the micro-controller. As of the writing of this document, the reason for this is under investigation.

The "Recording Data" state consisted of the analog front end being switched on, the micro-controller sampling the analog front end's output signal and buffering the samples while writing them to the flash memory. This state was intended to be representative of the "normal operational mode" of the bat-tag with the micro-controller operating in a mode were all real time constraints were being met and raw ADC samples were being moved to the flash. In this state, it was predicted that it would take 0.67 hrs to fill the entire flash memory and that 7.98 mAh of charge would be used which is a modest portion of the entire available battery charge of 135 mAh.

The "Reading and Transmitting Data" state was measured when the analog front end was switched off, and the micro-controller was reading samples off of the flash memory and transmitting them over the radio. At 114 mAh this state was shown to use a majority of the battery's available charge. This could be reduced by using a radio that could modulate data at higher rate while using a similar amount of current, but as the system presently functions, the amount of charge required to transmit the contents of the entire flash memory chip exceeds the capacity of the battery.

Finally, the "Wake Up On Radio Sleep" state was measured with the analog front end turned off and the micro-controller in sleep mode waiting for the radio to detect the base station. The amount of current used by the bat-tag "Wake Up On Radio Sleep" is higher than expected. Sections 2.5 and 2.7 describe that the regulators were expected to draw a combined 46 μA, the micro-controller 2 μA, and the radio 1.10 μA in this state. These currents were expected to comprise most of current used in the "Wake Up On Radio Sleep" state. However, from a functional standpoint drawing 1.09 mA in the "Wake Up On Radio Sleep" does not compromise the functionality of the bat-tag unless that bat is out for several days before returning back to its roost.

4.6 System Noise Test

Figure 4.9 has three curves. The top curve is the one sided FFT of the audio recording taken in the anechoic chamber. The middle curve is the top curve passed through a 64th order 1-D median filter to suppress that outliers in the data that make it difficult to see the actual trend of the data. The bottom curve represents the spectral noise density of the microphone as documented in its datasheet [14].

Observing these curves serve as a qualitative method to determine how well the amplifier circuit was performing. In a well designed instrument amplifier circuit, the instrument itself should dominate the noise measurements and the amplifier circuit itself should contribute minimal noise. Figure 4.9 shows that the spectrum of the audio data for noise test is shaped very much so like the documented spectral noise density for the microphone suggesting the microphone dominates the analog noise, which is suggestive of a good design in terms of noise.

To gain a quantitative measure of the system noise performance, the time domain data from the system noise test had it's DC component removed and the RMS value of the portion of the recording now showing significant amounts of radio bit errors was calculated as being 4.36 mV. To compare this to the input referred noise calculations in 2.2.2 this value was divided by the mid band gain of analog filter, approximately 130, to give a measured input referred noise of 33.5 μV compared to the 23.2 μV shown in simulations. Considering that the room the test was performed in was not perfectly quiet and that the simulations did take into account electrical noise created by writing the flash memory and from other sources, this value compares favorably well to the simulation.

The input referred noise floor of 33.5 μV RMS determines what the quietest sound that can be detected with the microphone. Any sound that does not create a signal with a greater amplitude than 33.5 μV RMS will not be distinguishable from noise created by the microphone and the analog filters. An application note for the microphone states that 40 kHz the sensitivity of the microphone is $\frac{-55 \text{ dBV}}{.1 \text{ Pa}}$ meaning that in order to generate an output signal greater than 33.5 μV RMS that the sound at the microphone would have to be greater than 0.00188 Pa or 39.5 dB SPL [15]. If it was desired for there to be 10 times more signal amplitude than noise, then the sound would need to be greater than 0.0188 Pa or 59.5 dB SPL. The loudest sound that the microphone can accurately represent is determined by the Total Harmonic Distortion of the microphone. According to the microphone's datasheet, with a 1 kHz input tone, if no more than 1% Total Harmonic Distortion is desired then the input sound should be less than 100 db SPL [14]. This information aggregated together means that the microphone and analog filter together should be able to operate with input sound levels in a range from 39.5 dB SPL to 100 dB SPL.

4.7 Sonar Data Acquisitions

Figure 4.10 shows a sonar and echo recording from the bat-tag with no film wrapping and figure 4.11 shows the same, except for the bat-tag wrapped in heat shrink film. At the start of both figures a large outgoing sonar pulse is visible. It appears that the recorded outgoing sonar pulse in figure 4.10 is reaching ground rail and could be a sign of clipping. The outgoing pulse in figure 4.11 is clearly attenuated by the heat shrink film and there is no evidence of it clipping. In both figures 4.10 and 4.11, 10 ms after the main outgoing pulse there is a recorded echo of the main pulse as predicted in section 3.6.2. Figures 4.12 and 4.13 show these echos in greater detail. If the largest peak to peak voltage in figure 4.12 is taken to be 325 mV (1.025 V - .7 V), and the largest peak to peak voltage in figure 4.12 is taken to be 130 mV (.93 V - .80 V), then the heat shrink film appears to attenuate the echos by approximately 8 dB. In this experiment the echos from the bat-tag wrapped in heat shrink film had enough amplitude to be well resolved by the ADC, field testing will reveal if this result extends to measurements taken on foraging bats.

There appears to be number of other echos present in the recording as well. The best explanation for these echos is that the outgoing sonar pulse reflected off other surfaces closer than the wall that the sonar range finder was aimed at. This experiment does show that the bat-tag is capable of recording outgoing sonar pulses and the returning echos in a laboratory test. Field testing will show how effective the bat-tag is at recording bat echolocation call echos from more challenging targets such as small insects.

4.8 Experimental Results Figures

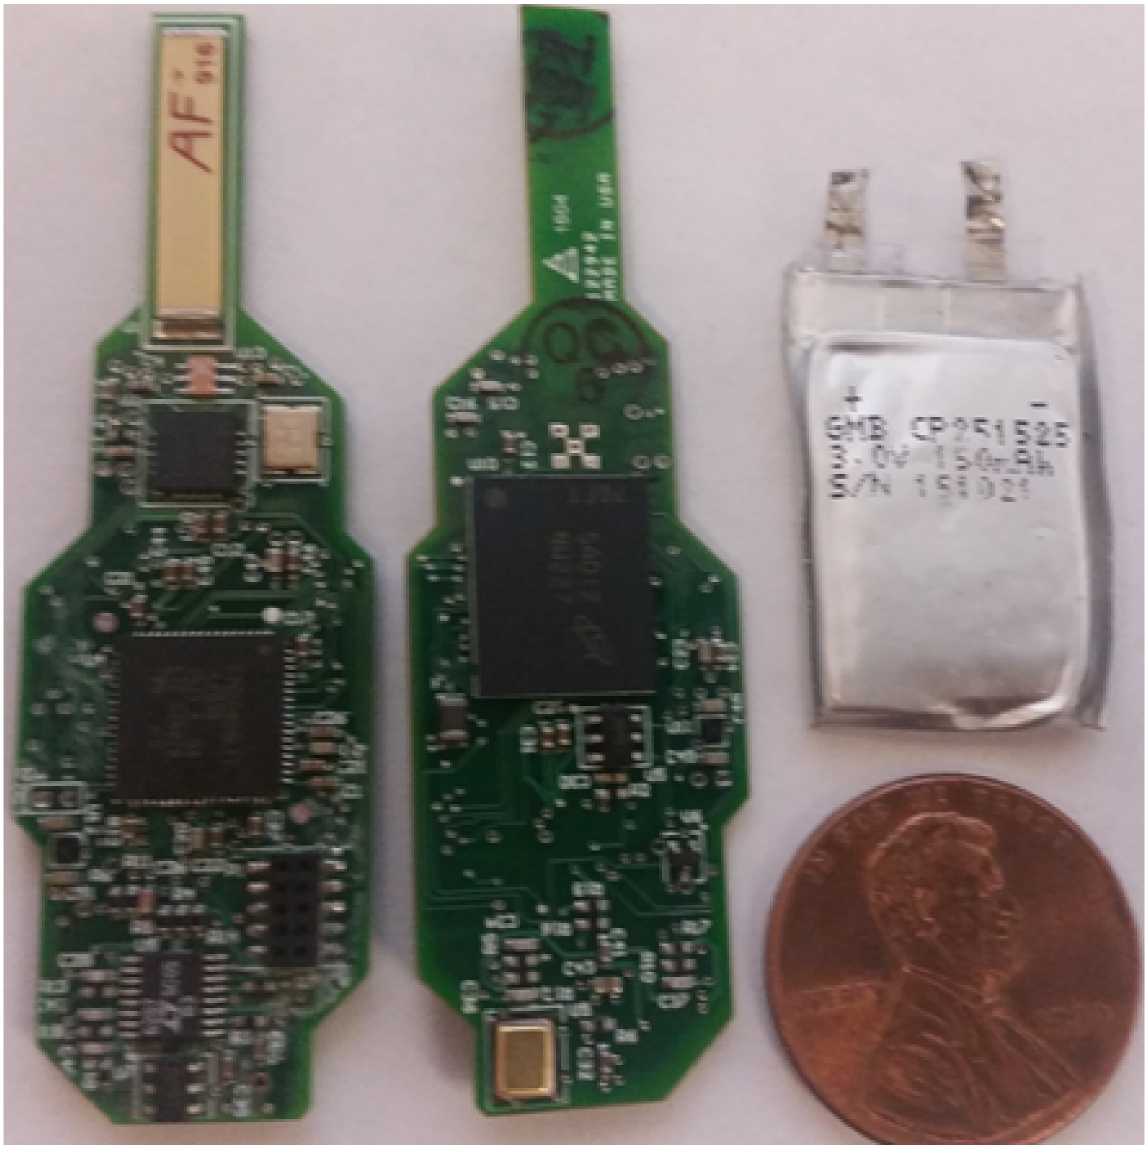

Figure 4.1: Bat-Tag With Battery

Table 4.1: Bat Tag Weight Measurements

Sample	Battery	Bare Boards	Assembled Boards	Assembled With Battery and Film
1	1.44 g	1.12 g	1.96 g	3.63 g
2	1.48 g	1.06 g	1.93 g	3.66 g
3	1.52 g	1.08 g	1.98 g	3.63 g
4	1.53 g	1.13 g	1.95 g	3.67 g
5	1.48 g	1.09 g	1.98 g	3.65 g
6	1.55 g	1.11 g	1.96 g	3.65 g
7	1.52 g	1.04 g	1.98 g	3.65 g
8	1.48 g	1.06 g	1.95 g	3.69 g
9	1.51 g	1.13 g	1.98 g	3.62 g
10	1.48 g	1.10 g	1.94 g	3.65 g
mean	1.50 g	1.09 g	1.96 g	3.65 g
std dev	0.03 g	0.03 g	0.02 g	0.02 g

Table 4.2: Bat Tag Component Weights

Part	Weight
PSoC® 5LP	0.18 g (vendor quoted) [7]
MT29F4G08ABBDAH4-IT:D (Flash 4 Gb)	0.16 g
Antenna	0.17 g
CC1101	0.17 g
Mean Total Part Weight 4.1	0.86 g
Accounted Part Weight	0.68 g
Unaccounted Part Weight	0.19 g

Table 4.3: Bat Tag Current Measurements

State	Current	Hours	Charge
Idle, Analog Front End Off	9.44 mA	0 h	0 mAh
Idle, Analog Front End On	9.80 mA	0 h	0 mAh
Recording Data	11.9 mA	0.67 h	7.98 mAh
Reading and Transmitting Data	22.8 mA	5 h	114 mAh
Wake Up On Radio Sleep	1.09 mA	12 h	13.1 mAh
Total Charge			135 mAh
Battery Charge			~135 mAh @ 20 mA

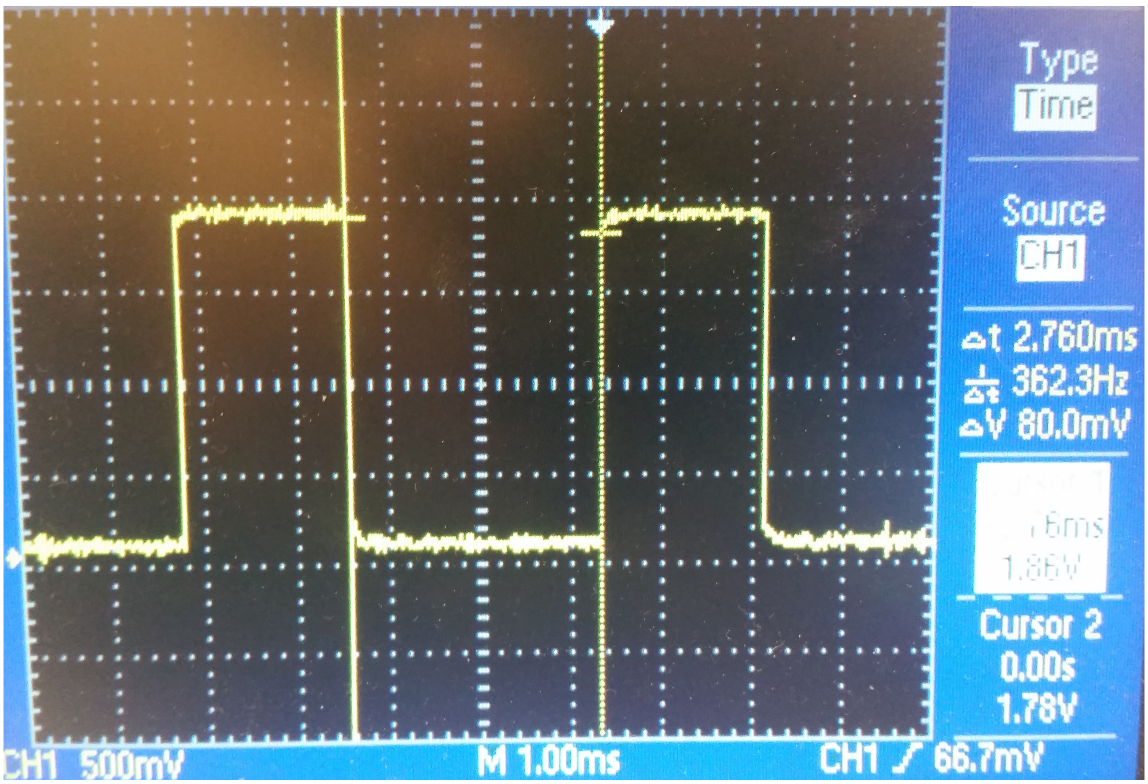

Figure 4.2: Oscilloscope Picture Showing Time the Flash Chips Chip Enable is Held Low Indicating Flash Page Write Time

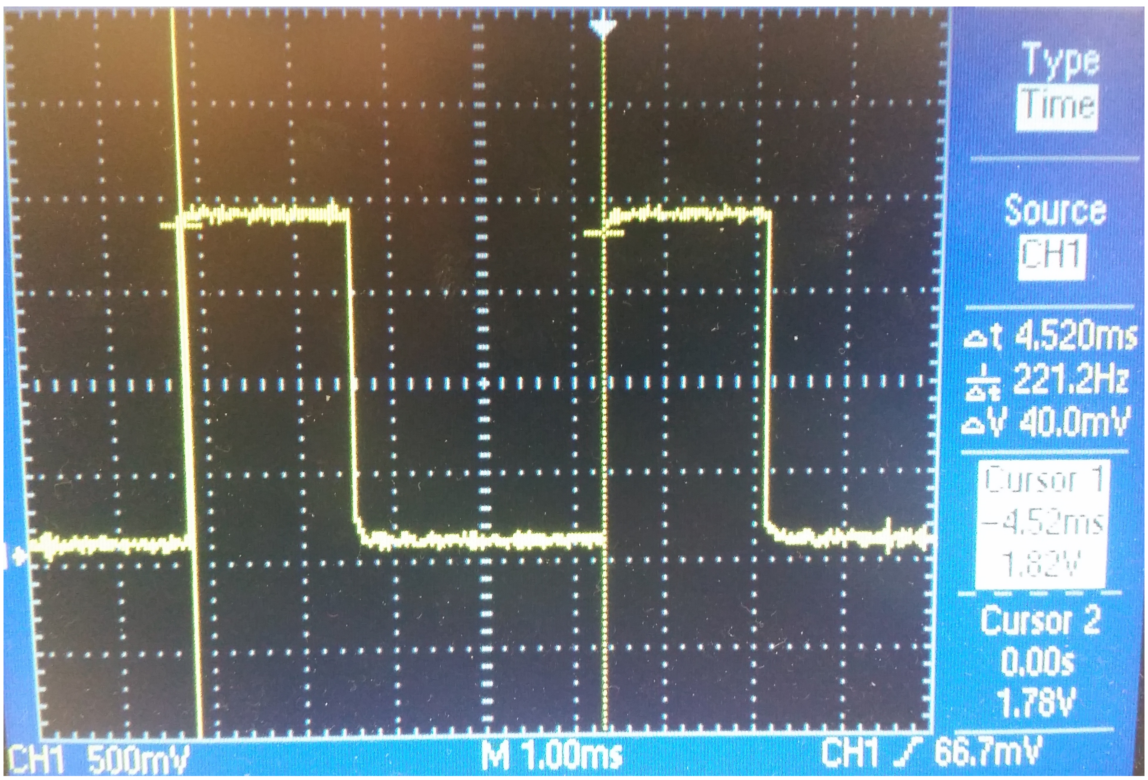

Figure 4.3: Oscilloscope Picture Showing the Time Between Flash Page Writes

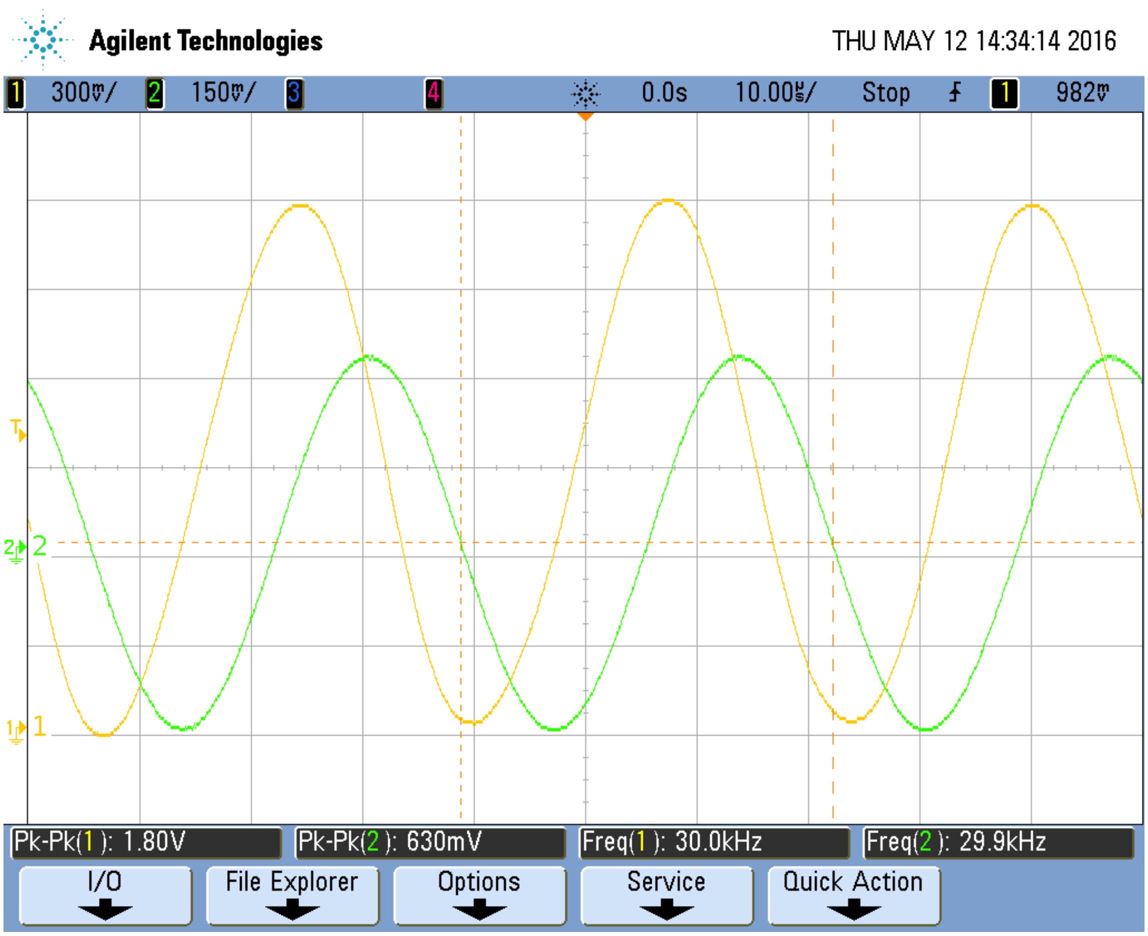

Figure 4.4: Analog Filter Max Output Voltage

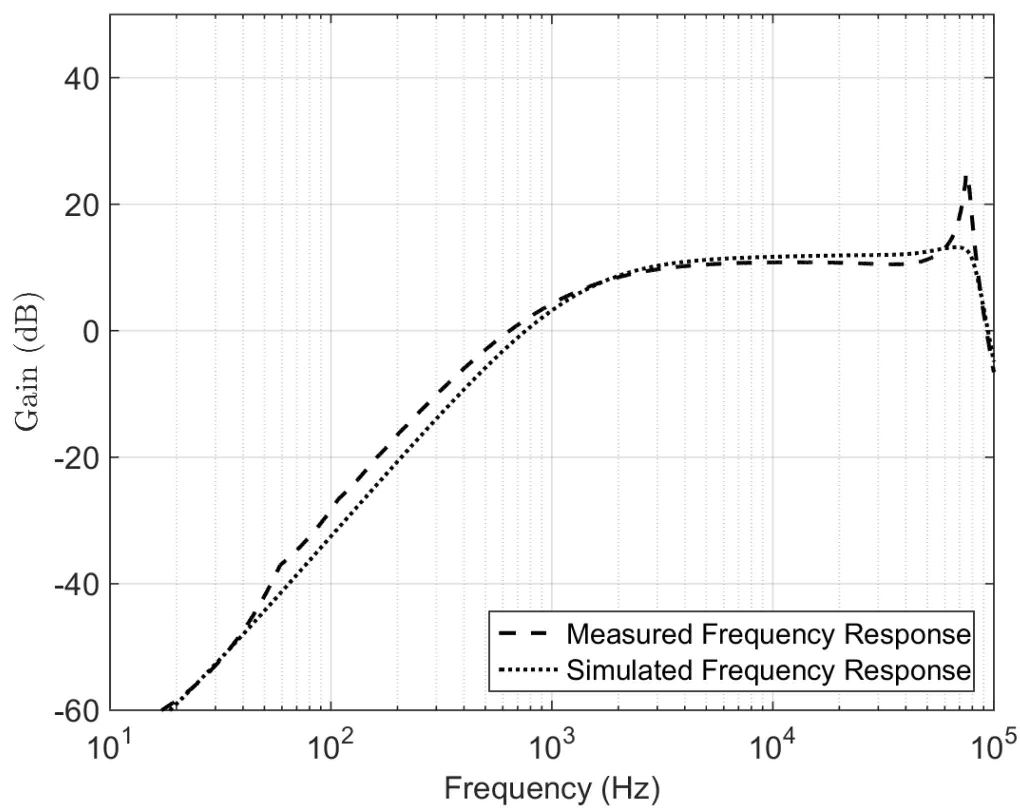

Figure 4.5: Measured Versus Simulated Analog Front End Frequency Response with a 1 kΩ Feedback Resistor

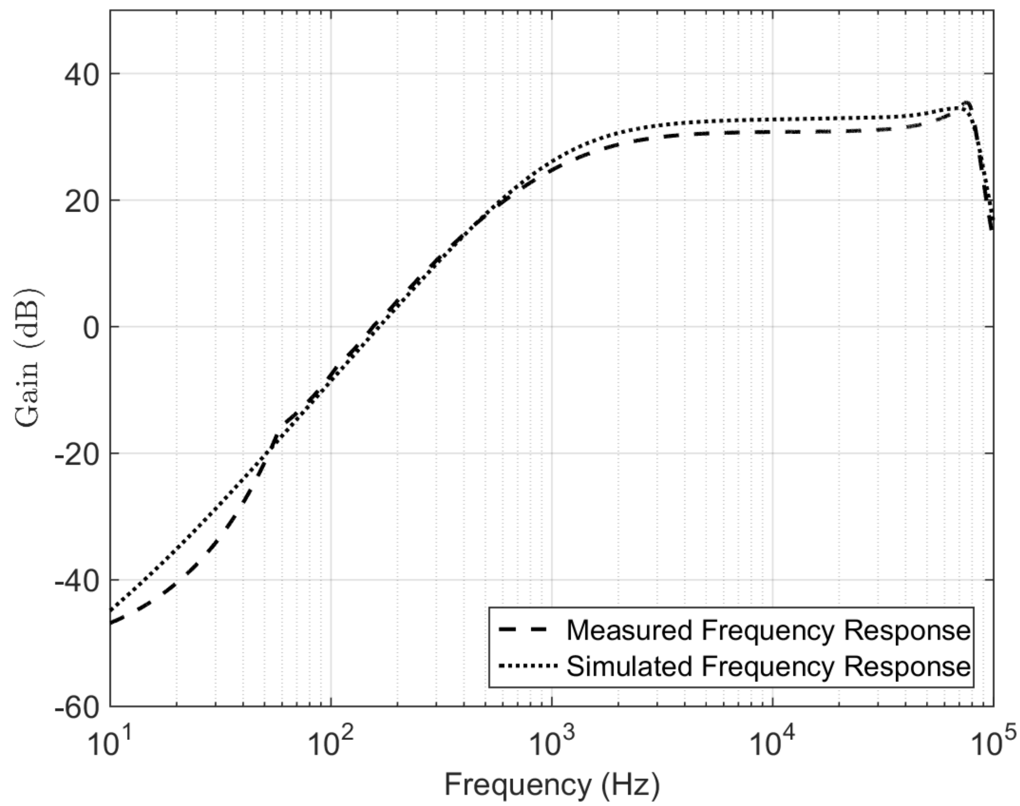

Figure 4.6: Measured Versus Simulated Analog Front End Frequency Response with a 11.5 kΩ Feedback Resistor

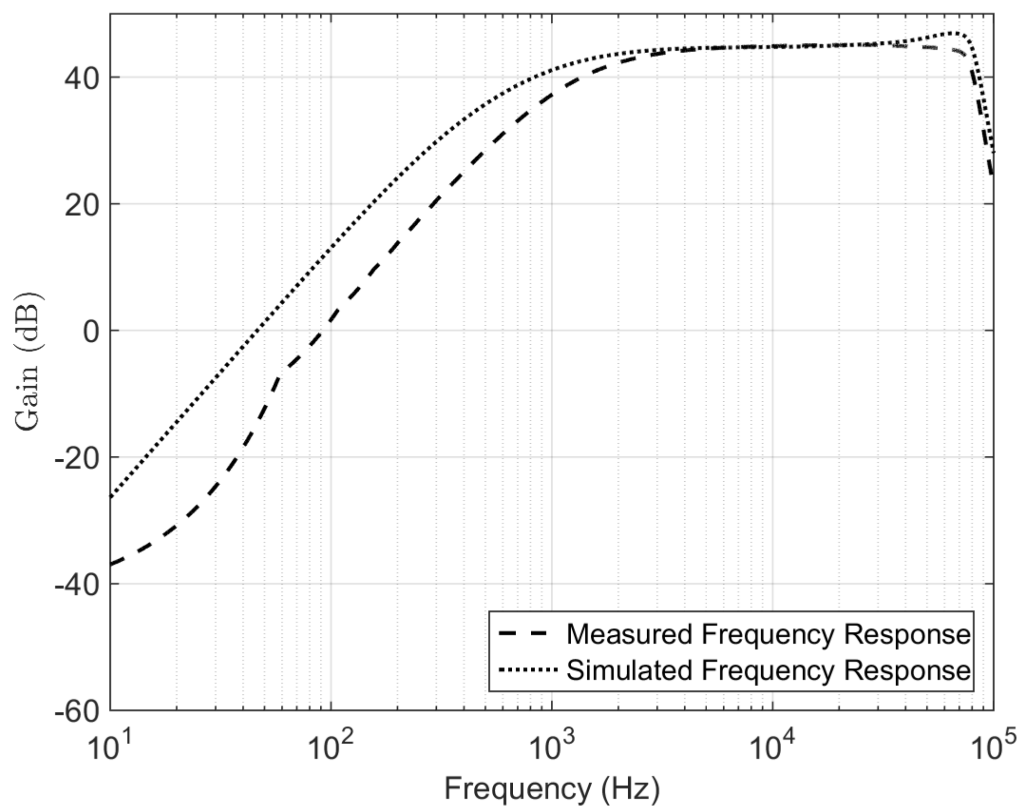

Figure 4.7: Measured Versus Simulated Analog Front End Frequency Response with a 49.9 kΩ Feedback Resistor

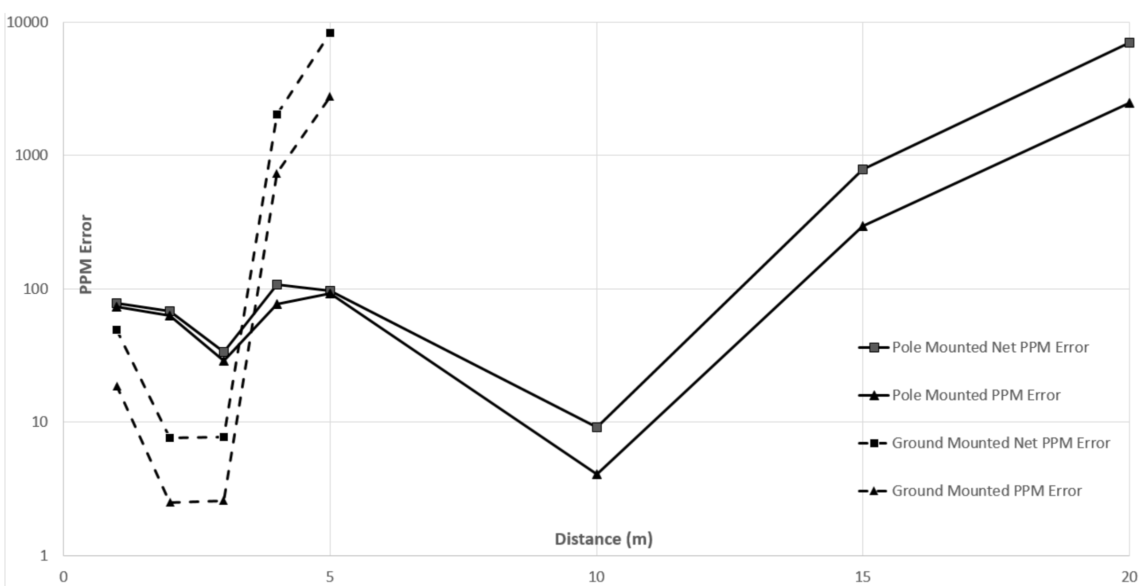

Figure 4.8: Radio Bit Error Rate as a Function of Distance

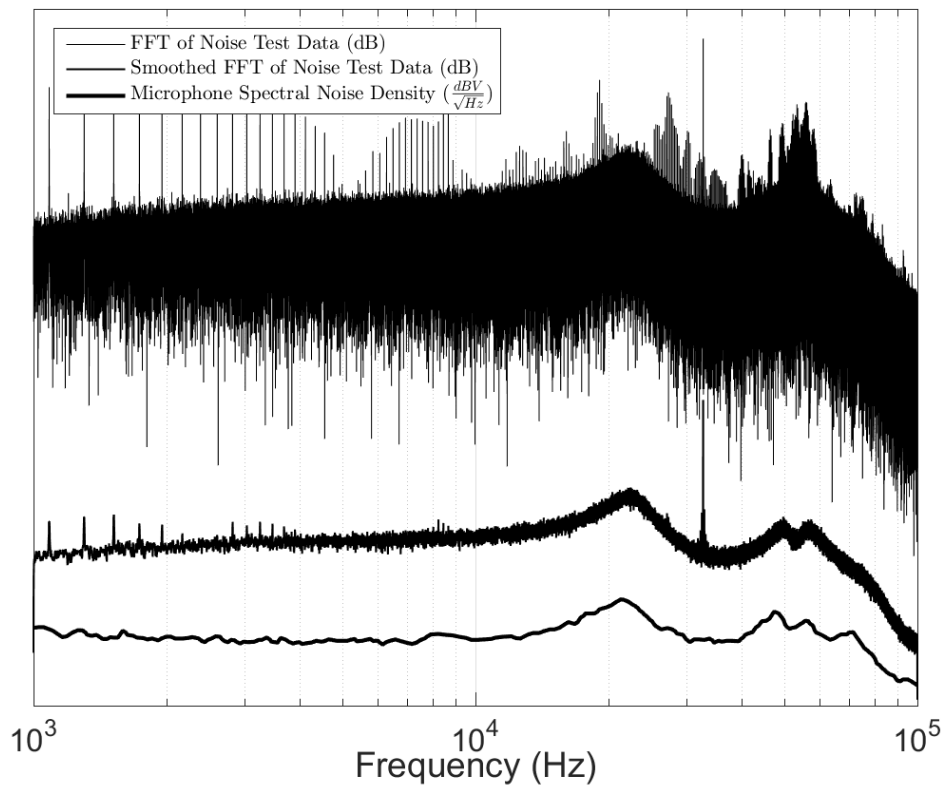

Figure 4.9: Documented Microphone Spectral Noise Density and FFTs of System Noise Test

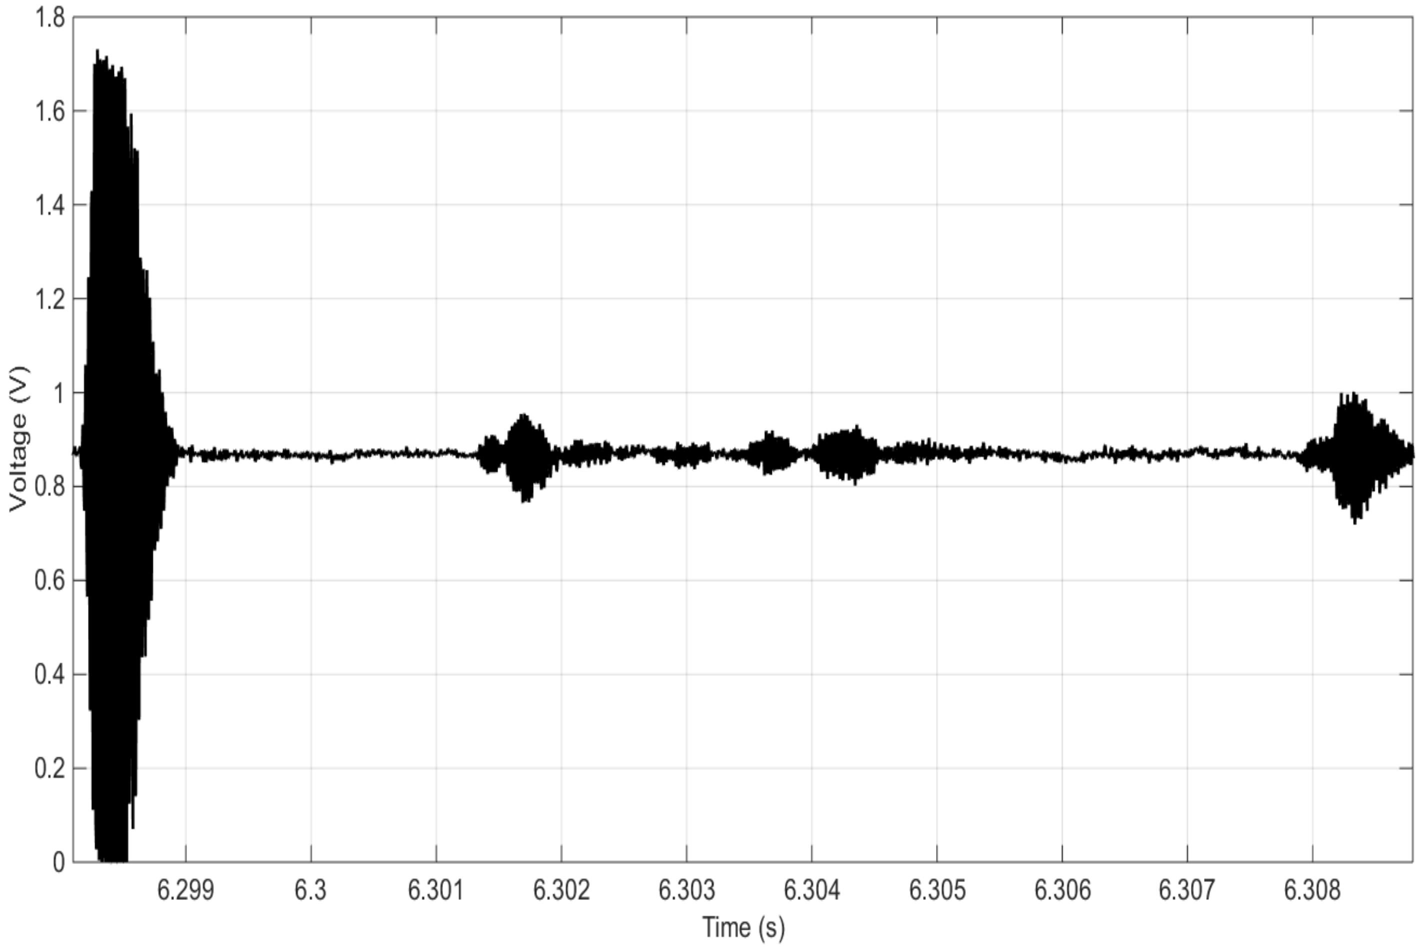

Figure 4.10: Sonar Test With No Film Around that Bat-Tag

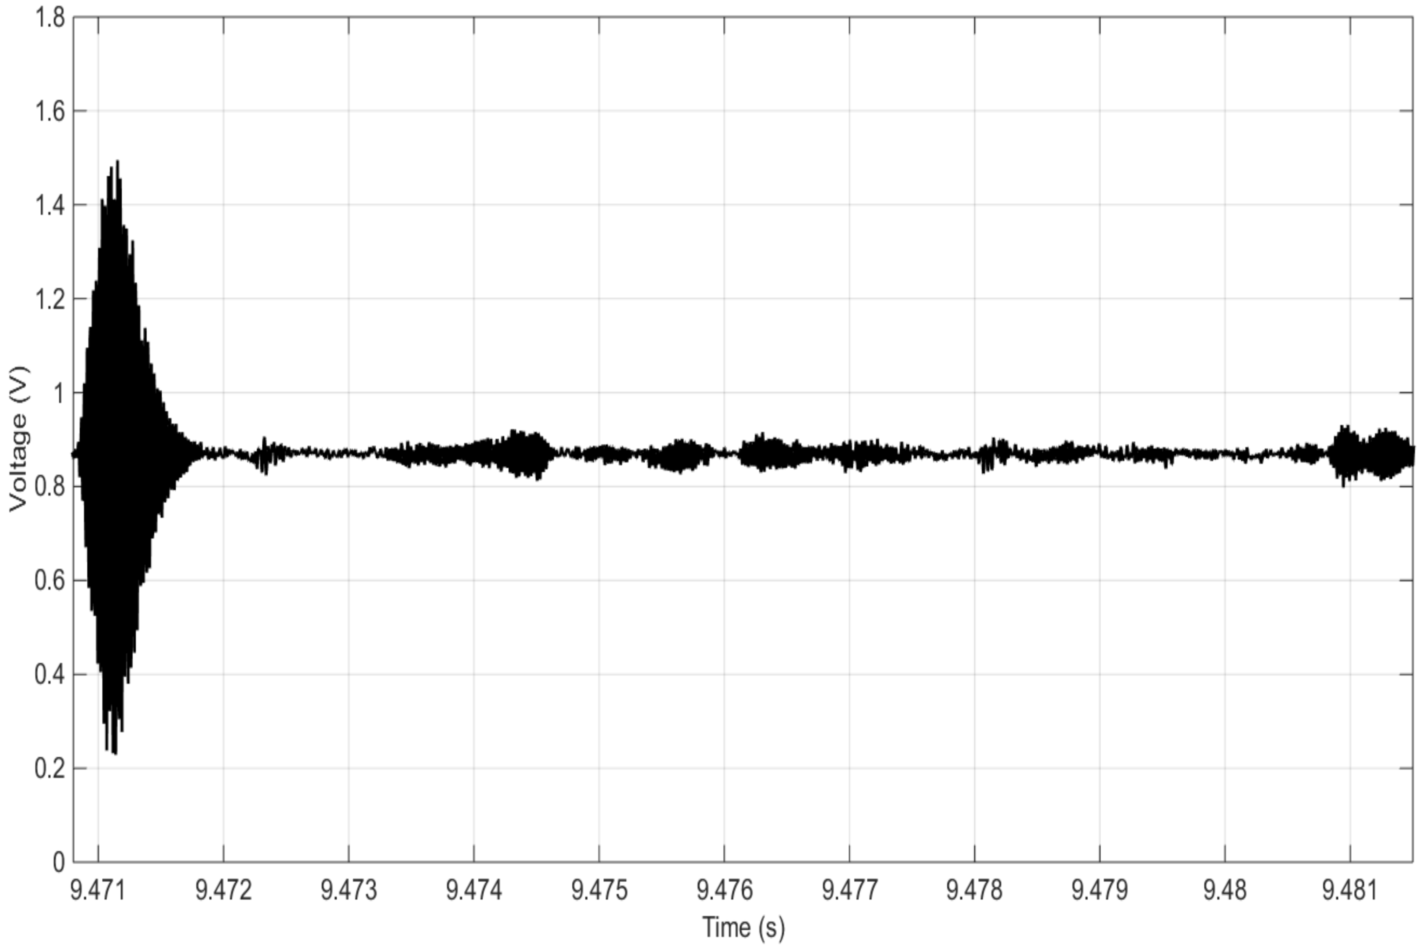

Figure 4.11: Sonar Test With Film Around that Bat-Tag

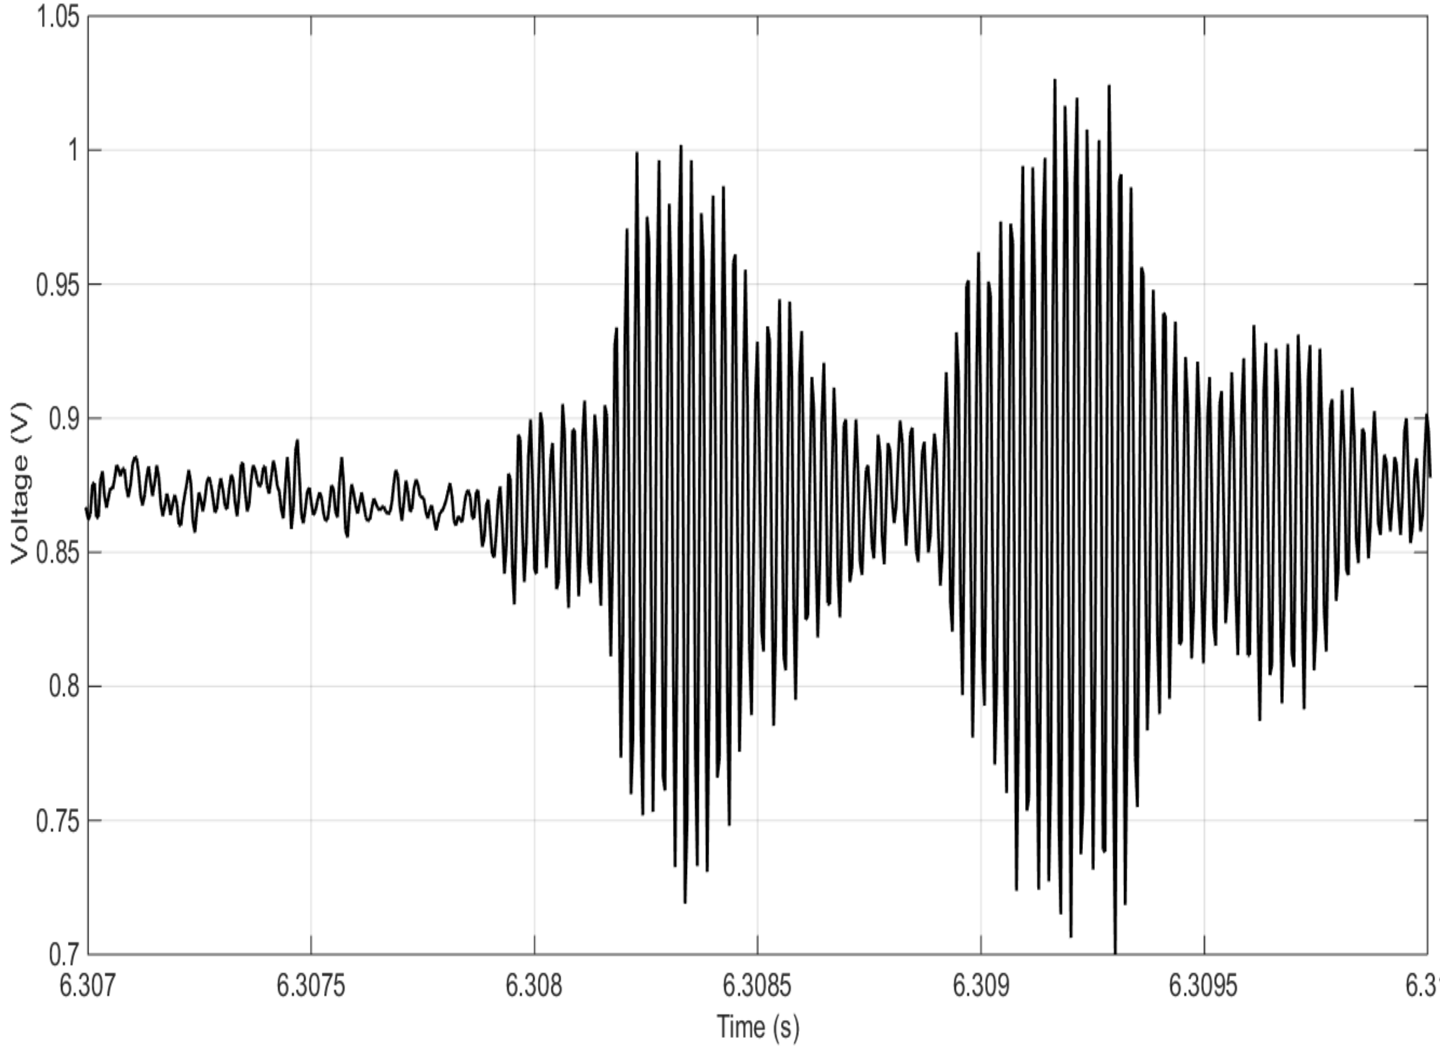

Figure 4.12: Detail of Echo Captured 10 ms After Sonar Pulse With No Film Around that Bat-Tag

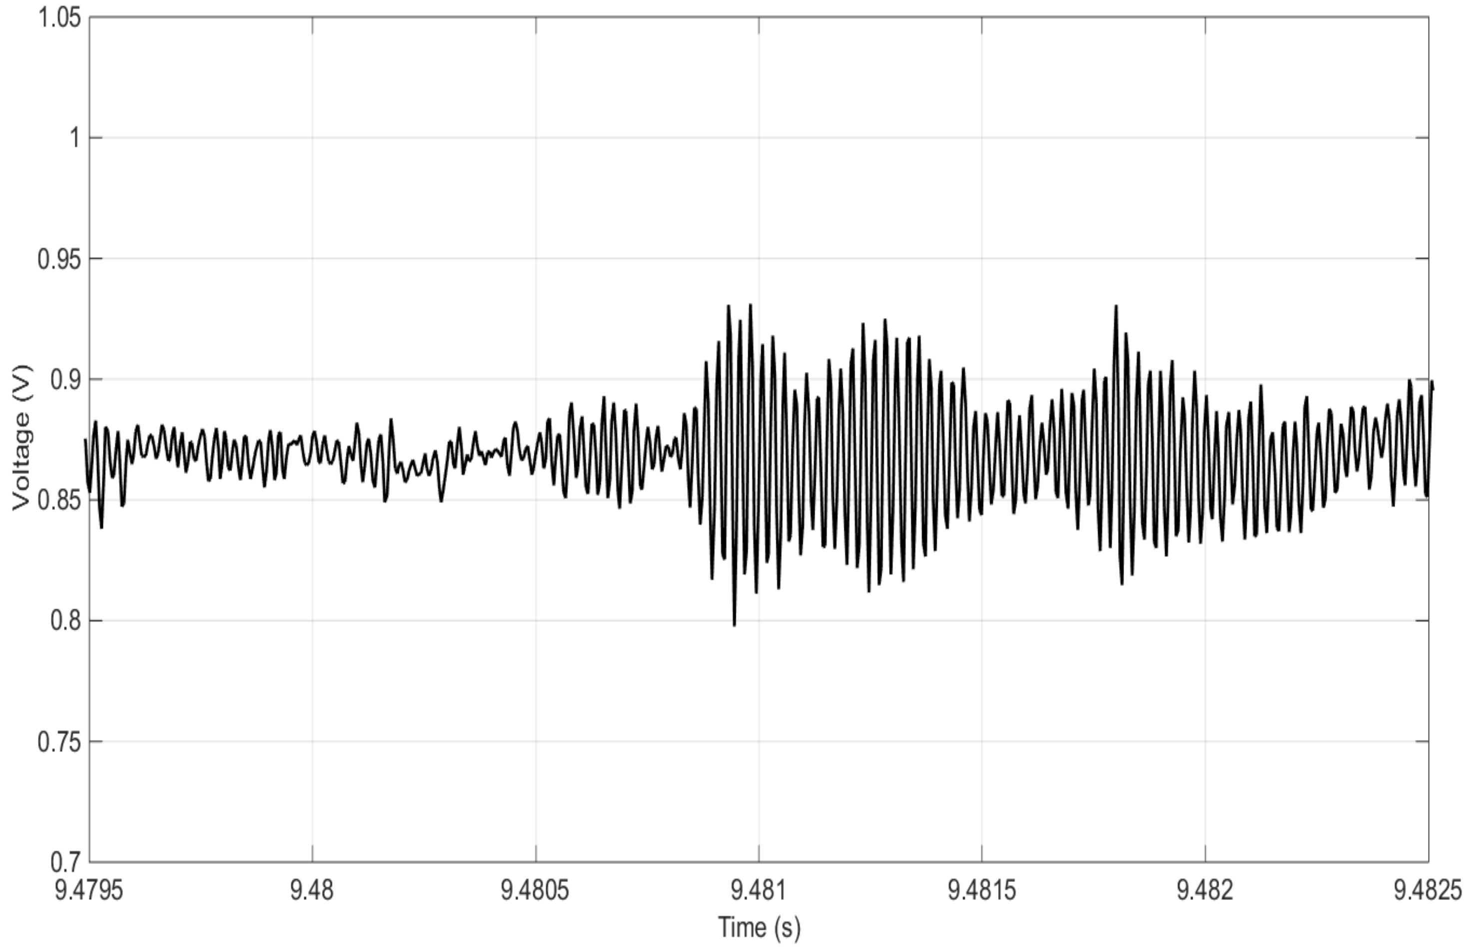

Figure 4.13: Detail of Echo Captured 10 ms After Sonar Pulse With Film Around the Bat-Tag

Chapter 5

Discussion

5.1 Radio Operation

The data shown in figure 4.8 shows three important things about the radio. First, as distance between the transmitter and receiver increases, the bit error rate tends to rise. Second, it appears that the bit error rate is significantly lower if the transmitter is not at the same elevation above the earth as the receiver. Third, it appears that it is possible to position the transmitter and receiver in such a way that the bit error rate at that position would represent a local minimum regardless of distance between the transmitter and receiver. Taking all of this information together it appears that for a practical application that the optimum placement for the receiver would be no more than 10 m away from where the bat would be expected to sleep. Most importantly however, it appears that performing a bit error rate site survey at the bat's roosting area would be essential to ensure reasonably predictable bit error rate performance. While assuming that a bat will sleep exactly where a bit error rate measurement was taken is very strong assumption to make, one can take measures to place the receiver in a spot were there is reasonably good bit error rate performance from most of the places were it would be expected for the bat to sleep.

5.2 Comparison To Recording System Used by Tel Aviv University Researchers

The bat-tag used by the neuro-ecologists from Tel Aviv University featured a GPS receiver that allowed position information to be recorded in addition to echolocation calls which provided information about bat flight patterns. The bat-tag described in this thesis did not feature a GPS receiver, but instead had a radio transceiver which would allow for the recorded bat echolocation calls to be gathered without researchers having to be present while the bat was in its roost. Additionally, the bat-tag described in this thesis sampled its microphone's output at 222 kHz compared to the 94 kHz used by the Tel Aviv researchers allowing for higher frequency content in echolocation calls to be resolved. The bat-tag used by the Tel Aviv researchers featured an external RF tracker to aid in tag recovery. While the work presented in this thesis presently does not have the ability to be covered with an RF detector, this feature could be added through software controlling the communications radio to act as an RF tracker. Finally, the work presented in this thesis weighed .15 g less than bat-tag used by the Tel Aviv researchers [4]. Directly comparing the merits of these bat-tags directly is difficult. For example selecting GPS over a radio transceiver is dependent on the nature of the experiment and deployment conditions, and using a higher sampling rate, while it allows for more high frequency content to be resolved, it also results in more storage being required in order to record for the same amount of time. However, it should suffice to say that both of these bat-tags should be useful in their own experiments.

Table 5.1: Comparison Between University of Tennessee and Tel Aviv University Bat-Tags

Characteristic	University of Tennessee	Tel Aviv University [4]
GPS	No	Yes
Radio Communications	Yes	No
Ultra Sound Sampling Rate	222 kHz	94 kHz
RF Tracker	Could be Implemented With Existing Radio	Implemented With .22 g External Module
Weight	3.65 g	3.8 g

5.3 Future Work

5.3.1 Power System

A way to extend the life of the battery further would be to regulate the output voltage of the battery using a buck converter before powering the linear regulators. While the buck converter would require additional board area, it would be more efficent to step the battery down from 3 V to 1.85 V then powering the linear regulators with that 1.85 V which would power the rest of the board. While buck converters are not perfectly efficient when stepping down a voltage, they are far more efficient than linear regulators, which generally waste all of the excess energy when stepping down a voltage.

5.3.2 Radio

Improving the radio communications of the bat-tag offers the single largest opportunity to improve the overall system performance. Currently transmitting the entire contents of the flash memory chip would take over 9 hours and would require approximately 90 mAh more battery charge. Finding a way to reduce the overall transmission time could also be the solution to the battery charge problem as well as most of the energy used in most radio communications is used to create the actual radio waves and not in the data modulation. One possible radio replacement would be the nRF24l01+ by Nordic Semiconductor. It operates at 2.4 GHz, has a maximum raw modulation rate of 2 Mbps, comes in a small 4 mm X 4 mm QFN package, and there appear to be several software libraries that could potentially ease the integration work with the chip and provide rich features such as acknowledgement of sent data and retransmission. The only apparent disadvantages to this radio is that it can not operate at 1.8 V which would require some rework on the bat-tag's power supply system. More importantly however, there does not seem to be a low power wake up on radio mode for the nRF24L01+ comparable to that of the CC1101, so some other method of deciding when to begin transmission would be needed [20].

5.3.3 Data Compression

Another step that could improve the ability of the bat-tag to transmit it's stored data in a timely fashion could be data compression. Some methods of data compression are considered lossless meaning that performing decompression on the compressed data should yield the exact original data. Other methods are considered lossy in that one is not guaranteed to be able to recover the original data after compression. Choosing an appropriate algorithm will require considering the desired compression ratio, how much of a quality loss is acceptable, and how the quality of the data compressed data is impacted by being transmitted over a noisy channel.

5.3.4 Analog Front End

Understanding the root cause of the gain control issues mentioned in 2.8.2 caused by the digitally controlled rheostat in the gain stage of the analog front end is an important next step. This would allow for either modifications to the current bat-tag done in the lab or a future iteration of the bat-tag to resolve the issue were the gain does not adjust as expected allowing the implementation of features such as dynamic gain control to be implemented.

Chapter 6

Conclusion

This thesis demonstrates the design and validation of a wireless bat-mounted sonar recording system. Designing the bat-tag required attention to detail in many areas that would be easy to overlook when first describing the system. Every component on the bat-tag had to be chosen with care to ensure that it would not violate the weight nor the power budget of the system. Meeting this requirement and engineering a system capable of recording bat sonar required careful attention to detail in the design of the analog front end to ensure that the bandwidth requirements for recording the bat's sonar was met, and that the data would be sufficiently low pass filtered to ease the sampling requirements on the digital portion of the design. Designing the digital datapath of the recording system required meeting real time constraints set by the sampling rate of the ADC and the performance of the flash memory chip. Meeting these real time constraints required extensive research into the architecture of the micro-controller and a solid grasp of how to write efficient source code. Implementing the radio communication capabilities required extensive trial and error in the lab to strike performance balance between transmit speed and bit error rate that required testing outside the lab to validate.

Experimental results showed that the presented bat-tag could sample higher frequency signals than the device used by the Tel Aviv researchers and offered the ability to offload recorded data wirelessly, while sacrificing the ability to record in flight GPS data. In terms of weight, both devices were of a similar weight. Table 6.1 summarizes the performance characteristics of the presented system. It is capable of effectively sampling bat echolocation calls up to 80 kHz. By adjusting the gain of the analog front end, components of echolocation calls and echos with intensities in the range of 39.5 dB SPL - 100 dB SPL can be sampled. The system is capable of storing up to 37.5 minutes of 12-bit ADC samples recorded at 222 kHz in its 8 Gb flash memory, but it only capable of transmitting roughly 57% of that data using the radio before the battery dies. Weighing in at 3.65 g, this device is suitable for deployment on female examples of *Phyllostomus hastatus* without exceeding the recommended 5% tag weight for a bat. The presented data shows that this device is competitive with the current state of the art in bat echolocation recording devices.

Table 6.1: Bat-Tag Performance Summary

Characteristic	Performance
Effective Input Bandwidth	80 kHz
Sound Intensity Input Range	39.5 dB SPL - 100 dB SPL
Recording Time	37.5 m
Transmission Speed	$\frac{250\ \text{kbits}}{\text{s}}$
Transmission Range	10 m
Required Transmission Time	8.8 h
Actual Battery Limited Transmission Time	5 h
Weight Including Battery and Protective Film	3.65 g

References

[1] Analog Devices (2014). 300 mA, Low Quiescent Current, CMOS Linear Regulator. 37

[2] ANTENNA FACTOR (2011). ULTRA COMPACT CHIP ANTENNA DATA GUIDE. 13

[3] Bohn, K. M., Boughman, J. W., Wilkinson, G. S., and Moss, C. F. (2004). Auditory sensitivity and frequency selectivity in greater spear-nosed bats suggest specializations for acoustic communication. *Journal of Comparative Physiology A*, 190(3):185–192. 1, 4

[4] Cvikel, N., Levin, E., Hurme, E., Borissov, I., Boonman, A., Amichai, E., and Yovel, Y. (2014). On-board recordings reveal no jamming avoidance in wild bats. *Proceedings of the Royal Society of London B: Biological Sciences*, 282(1798). 2, 68, 69

[5] Cypress Semiconductor Corporation (2015). PSoC®5LP: CY8C58LP Family Datasheet. 14

[6] Cypress Semiconductor Corporation (2016a). Cy8c5868axi-lp035. http://www.cypress.com/part/cy8c5868axi-lp035. Accessed: 2016-03-29. 14

[7] Cypress Semiconductor Corporation (2016b). Cy8c5868lti-lp038. http://www.cypress.com/part/cy8c5868lti-lp038. Accessed: 2016-03-29. 14, 45, 55

[8] Cypress Semiconductor Corporation (2016c). Psoc 3/5 power estimator spreadsheet. www.cypress.com/file/45771/download. Accessed: 2016-03-29. 14

[9] Department of Electrical and Computer Engineering - University of New Brunswick, Fredericton, NB, Canada (2016). Error correction and the hamming code. http://www.ece.unb.ca/tervo/ee4253/hamming.shtml. Accessed: 2016-03-27. 32

[10] Gannon, W. L., Sikes, R. S., Care, T. A., and of the American Society of Mammalogists, U. C. (2007). Guidelines of the american society of mammalogists for the use of wild mammals in research. *Journal of Mammalogy*, 88(3):809–823. 3

[11] Huizhou Markyn New Energy Co., LTD. (2015). GMBJSYG150310. 37

[12] Hynix Semiconductor, Intel Corporation, Micron Technology, Inc., Phison Electronics Corp., Sony Corporation, STMicroelectronics (2006). Open NAND Flash Interface Specification. 29

[13] Johanson Technology (2014). 915MHz Impedance Matched/Balun/BPF Integrated Component for T.I. CC110X, CC111X, CC113X and CC115X, CC110L, CC113L, CC115L, CC430 and RF430. 13, 34

[14] Knowles Acoustics (2012). Zero Height "Ultra-Mini" SiSonicTM Microphone Specification With MaxRF Protection and Extended Low Frequency Performance. 5, 37, 46, 51, 52

[15] Knowles Electronics (2013). Ultrasonic Applications for Knowles Electret and MEMS Microphones. 5, 52

[16] Linear Technology (2011). 6.5MHz, 65μA Power Efficient Rail-to-Rail I/O Op Amps. 6, 7

[17] Max Kingsbury (2015). PSoC®3 and PSoC5LP Clocking Resources. 14

[18] McCracken, G. F. and Bradbury, J. W. (1981). Social organization and kinship in the polygynous bat phyllostomus hastatus. *Behavioral Ecology and Sociobiology*, 8(1):11–34. 3

[19] Microchip (2006). Low-Cost 64-Step Volatile Digital POT. 6, 7

[20] Nordic Semiconductor (2008). nRF24L01+ Product Specification. 70

[21] Nørum, U., Brinkløv, S., and Surlykke, A. (2012). New model for gain control of signal intensity to object distance in echolocating bats. *Journal of Experimental Biology*, 215(17):3045–3054. 1, 5, 46

[22] Rodolfo Lossio (2015). PSoC®3 and PSoC5LP Low-Power Modes and Power Reduction Techniques. 15

[23] Squire, J. (2012). Analog filter design toolbox. http://www.mathworks.com/matlabcentral/fileexchange/9458-analog-filter-design-toolbox. Accessed: 2016-02-25. 6

[24] Texas Instruments (2013). Low-Power Sub-1 GHz RF Transceiver. 9, 31, 34, 37

Vita

Jeremy Langford enrolled in Computer Engineering at the University of Tennessee in 2010. During his time as an undergraduate he interned with Emerson Process Management in Knoxville, TN and Garmin in Olathe, KS. In December, 2014 he graduated magna cum laude and in January, 2015 enrolled into the Computer Engineering M.S. program. During his time as a master's student he received the Bodenheimer Fellowship and interned once more with Garmin. He will graduate with a M.S. in Computer Engineering in May, 2016. After graduation, he plans on moving to Olathe, KS were he will work full time on the hardware integration team with the marine product group at Garmin.