

Low-Power Spectrum Sensing for Cognitive Radio Applications

Xiao Xiao

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2012-146

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2012/EECS-2012-146.html>

May 31, 2012



Copyright © 2012, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Low-Power Spectrum Sensing for Cognitive Radio Applications

by Xiao Xiao

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, in partial satisfaction of the requirements for the degree of **Master of Science, Plan II**.

Approval for the Report and Comprehensive Examination:

Committee:

Professor Borivoje Nikolić
Research Advisor

(Date)

* * * * *

Professor Ali Niknejad
Second Reader

(Date)

Table of Contents

Chapter 1 - Introduction.....	1
1.1 – Scope of Research	3
1.2 – Organization of Report.....	3
Chapter 2 – Spectrum Sensing in the TV Band	4
2.1 – Characteristics of the UHF TV band.....	4
2.2 – IEEE 802.22 Requirements	5
2.3 – State of the Art and Related Works.....	6
Chapter 3 –Detection Modes	8
3.1 – Detection Techniques	8
3.1.1 – Energy Detection	8
3.1.2 – Pilot Detection	9
3.1.3 – Feature Detection.....	9
3.1.4 – Autocorrelation.....	10
3.2 – Proposed System	11
3.3 – Simulation Setup	12
3.2.1 – Detected SNR	12
3.2.2 – Signal Modeling	13
3.4 – Energy and Pilot Detection	15
3.5 – Autocorrelation.....	18
3.6 – System Sensitivity	20
Chapter 4 – System Architecture	21
4.1 – Subsampling	21
4.2 – Autocorrelation.....	24
4.2.1 - Single-Bit Autocorrelation	25
4.3 – System Architecture	27
Chapter 5 – System Implementation.....	30
5.1 - LNA	30
5.1.1 Bias	32
5.2 – RF Tracking Filter	33

Chapter 6 – Conclusion and Future Work	37
Chapter 7 – References	38

List of Figures

Figure 1.1 – Projected mobile data traffic in North America [1].	1
Figure 1.2 - Spectrum utilization in downtown Berkeley [4].	2
Figure 1.3 – Cognitive radio operating in time and frequency domains [4].	2
Figure 2.1 - Spectrum of a DTV channel: (a) ideal (pilot not shown) [8], and (b) measured [9].	4
Figure 2.2 - ATSC blocker profile.	5
Figure 3.1 - Energy detection on a 4 MHz-wide QPSK signal [16].	8
Figure 3.2 - ATSC pilot tone.	9
Figure 3.3 - SCF of (a) white noise, and (b) a QPSK signal [4].	10
Figure 3.4 - Time-domain waveform and autocorrelation function of (a) a sinusoid, and (b) white noise.	11
Figure 3.5 - High level block diagram of proposed system.	12
Figure 3.6 - Equivalent block diagram when (a) only noise is present, and (b) signal and noise are present.	13
Figure 3.7 - Transmit signal generation in simulation.	14
Figure 3.8 - Worst case sensing scenario, when target channel is (a) idle, and (b) occupied.	14
Figure 3.9 - Receiver modeling in simulation.	15
Figure 3.10 - Simulated (a) energy, and (b) pilot detection at baseband.	15
Figure 3.11 - Detected SNR as a function of input SNR for energy detection.	16
Figure 3.12 - Probability density functions for signal and noise.	16
Figure 3.13 - (a) Pfa, and (b) Pd for ideal channel energy detection.	17
Figure 3.14 – Normalized PDF of white noise and its autocorrelation with (a) 10^6 samples, and (b) 10^3 samples.	18
Figure 3.15 - Spectrum of detected pilot and noise for (a) energy detection, and (b) autocorrelation detection.	19
Figure 3.16 - Detected SNR vs. input power for pilot and autocorrelation detection.	19
Figure 3.17 - Detected SNR vs. input power for all detection modes.	20
Figure 4.1 - Illustration of subsampling with (a) subsampling ratio of 1, and (b) subsampling ratio of 2.	21
Figure 4.2 – Noise folding from subsampling.	22
Figure 4.3 - Conversion to baseband with sampling frequency set to (a) carrier frequency, and (b) channel center frequency.	23
Figure 4.4 - Time-domain waveform of (a) a sinusoid, and (b) its autocorrelation.	24
Figure 4.5 - Simulated time-domain waveforms of (a) a noisy pilot, and (b) its autocorrelation.	24
Figure 4.6 - Subsampled response to phase shifts in the sampling clock.	25
Figure 4.7 - Block diagram of a sigma-delta converter.	25
Figure 4.8 - Sigma-delta conversion of a sinusoid.	26
Figure 4.9 - Frequency spectrum of (a) sigma-delta input, and (b) sigma-delta output.	26
Figure 4.10 - Comparison of various autocorrelation implementations.	27

Figure 4.11 – Block diagram of proposed system.	28
Figure 5.1 - Schematic of LNA.....	30
Figure 5.2 - (a) NF, and (b) S11 as a function of R_{in}	32
Figure 5.3 – Frequency response of LNA (a) NF, and (b) voltage gain for high- and low-gain modes.	32
Figure 5.4 - Schematic of LNA bias.	33
Figure 5.5 - High level schematic of RF tracking filter.....	33
Figure 5.6 - Schematic of a Gm-cell.	35
Figure 5.7 - Frequency responses of filter (a) NF, and (b) voltage gain.....	35
Figure 5.8 - Schematic of one filter capacitor bank.....	36
Figure 5.9 - Frequency-tuned response of RF tracking filter.....	36

Chapter 1 - Introduction

In recent years, with the growth of smartphones, there has been exponential growth in mobile data usage. As seen in Figure 1.1, this trend is expected to continue for the foreseeable future as internet-connected mobile devices become ever more ubiquitous, and as consumers increasingly use these devices for streaming video and other data-intensive applications [1].

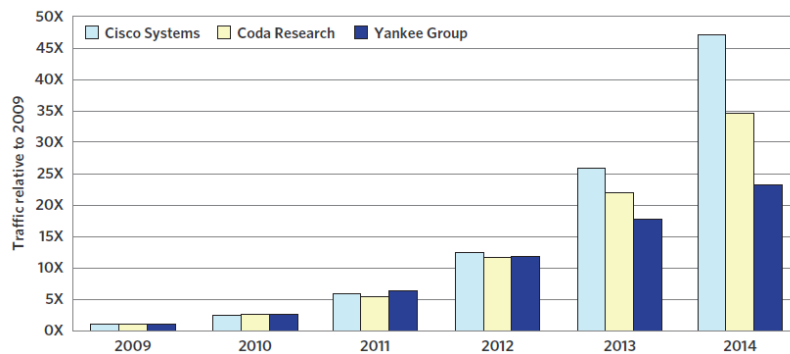


Figure 1.1 – Projected mobile data traffic in North America [1].

Accompanying this growth is a problem of spectrum scarcity, as users crowd the existing cellular bands. Opening up new bands mitigates the problem somewhat – the Federal Communications Commission’s (FCC) recent re-allocation efforts have increased commercial spectrum threefold. However, re-allocation alone cannot keep up with the rate of growth in mobile data traffic, which requires at least an order of magnitude increase [2].

A potential solution to this problem is the concept of a cognitive radio, conceived by Joseph Mitola in 1999 [3]. The solution takes advantage of the fact that no band is being used everywhere at all times. For example, Figure 1.2 shows a sample spectrum from downtown Berkeley, where utilization is about 30% below 3 GHz and virtually nil above 3 GHz [4]. The proposed solution is to allow devices to communicate on bands that are licensed to others – primary users – but are currently idle. If every band can be filled to near capacity in this way, that effectively grants us the order of magnitude increase in spectrum.

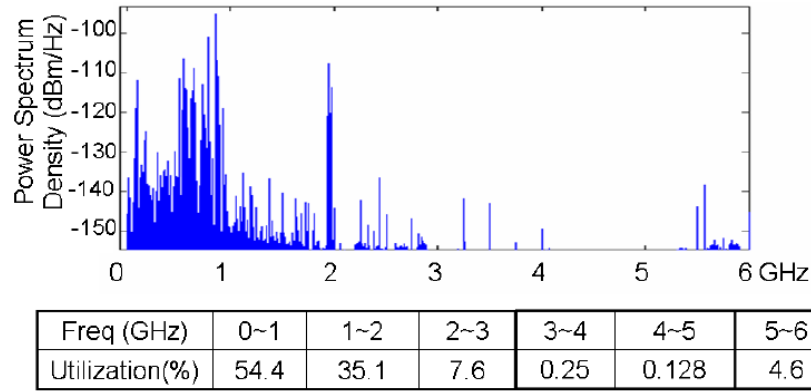


Figure 1.2 - Spectrum utilization in downtown Berkeley [4].

How a cognitive radio would operate in time and frequency domain is shown in Figure 1.3. At fixed time intervals, shown in yellow, the cognitive radio would scan the spectrum and identify idle bands, shown in white. It would then operate on those bands, shown in blue, until the return of a primary user has been detected. When this happens, the cognitive radio would vacate the re-occupied band for other bands that are now idle.

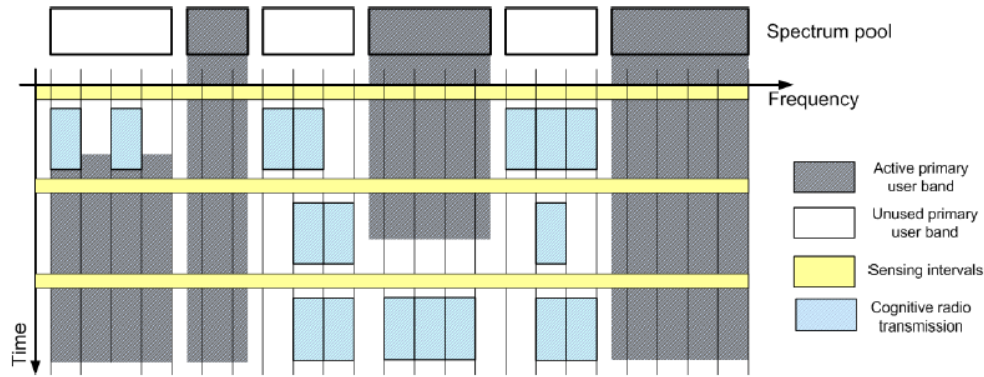


Figure 1.3 – Cognitive radio operating in time and frequency domains [4].

There are two main aspects of a cognitive radio that differ from a conventional radio: its reconfigurability and its spectrum sensing capabilities. The former allows the cognitive radio to reconfigure itself to receive and transmit on a variety of frequencies with a variety of bandwidths and standards. The latter, the focus of this research, allows the cognitive radio to robustly sense and identify idle bands. Since cognitive radios are unlicensed devices operating in licensed bands, spectrum sensing is essential in guaranteeing non-interference with the activities of the bands' primary users.

In 2008, the FCC approved the operation of unlicensed cognitive radio devices in the TV bands, with unlicensed mobile devices approved in the UHF TV band [5]. In 2010, the FCC dropped the spectrum sensing requirement for unlicensed TV band devices (TVBDs) due to concerns over its viability, and instead ordered TVBDs to identify idle bands by accessing a geolocation database [6]. However, because spectrum sensing offers much greater flexibility and lower overhead than

the database method, it remains a relevant area of research, and is expected to play a role in the future of cognitive radio systems.

1.1 – Scope of Research

One major contribution to the FCC’s decision to drop spectrum sensing requirements for TVBDs is the failure thus far of any system to meet the detection specifications set by the FCC. During a demonstration of industry prototypes conducted in 2008, no system achieved the necessary detection sensitivity of -114 dBm in real world scenarios [7]. In academic research, while there has been much work done on theory topics like detection algorithms and cooperative sensing, there has been a lack of work on physical front-end designs for spectrum sensing applications. Some works may embed sensing functionality into receiver designs, but since the sensitivity required for sensing is much higher than that of conventional receivers, these works all fail to achieve the necessary detection sensitivity as well. Thus, a system targeted specifically for TVBD sensing – achieving the required -114 dBm detection sensitivity with good dynamic range and low power – remains elusive.

The focus of this work, then, will be spectrum sensing receiver design in the UHF TV band for mobile applications. Specifically, this work seeks to produce a receiver front-end that is capable of meeting the FCC’s guidelines in [5] given realistic scenarios, and that does so with a low power overhead appropriate for mobile applications. This design is targeted solely for sensing functionality, and a fully-functional receiver design (i.e. capable of decoding data) is outside the scope of this work. Detection algorithms and other backend digital signal processing required for setting decision thresholds are also outside the scope of this work.

1.2 – Organization of Report

The rest of this report is organized as follows:

Chapter 2 reviews the characteristics of the UHF TV band and the requirements for performing spectrum sensing in this band. An overview of existing state-of-the-art is then presented.

Chapter 3 reviews known theoretical techniques for performing detection, and uses the theoretical background to propose our spectrum sensing system. We then describe how we simulated and evaluated the performance of the various detection modes in our proposed system.

Chapter 4 expands on the spectrum sensing system proposed in Chapter 3 into a realizable system architecture.

Chapter 5 describes the currently ongoing circuit implementation of the proposed spectrum sensing system.

Chapter 6 concludes this report and proposes several potential areas for future research.

Chapter 2 – Spectrum Sensing in the TV Band

In this chapter, we first look at the characteristic of the UHF TV band and its resident signals. Next, we discuss the requirements for spectrum sensing in the IEEE 802.22 standard and their challenges. Finally, we survey existing work that has been done on spectrum sensing for mobile applications.

2.1 – Characteristics of the UHF TV band

In the United States, the UHF TV band currently ranges from 470 MHz to 698 MHz, and it is composed of 38 6-MHz-wide channels. After the full transition to digital TV (DTV) was completed in 2009, all channels became purely digital, adhering to the Advanced Television Systems Committee (ATSC) standard.

According to the ATSC standard, terrestrial DTV uses 8-level vestigial sideband (8VSB) modulation. Within each 6 MHz channel, data power is spread over the center 5.38 MHz bandwidth, giving a 10.76 MSymbol/s datarate. Each channel has two 610-kHz-wide transition regions, one on each edge, from the response of the root raised cosine filters [8].

In addition, each channel also has a pilot tone at the frequency of the buried carrier, 310 kHz away from the lower channel edge. The pilot, within its narrow 10 kHz bandwidth, has a power that is 11.3 dB lower than the total data power of the 6 MHz channel [8]. Figure 2.1 illustrates the ideal and measured spectrums of a DTV channel.

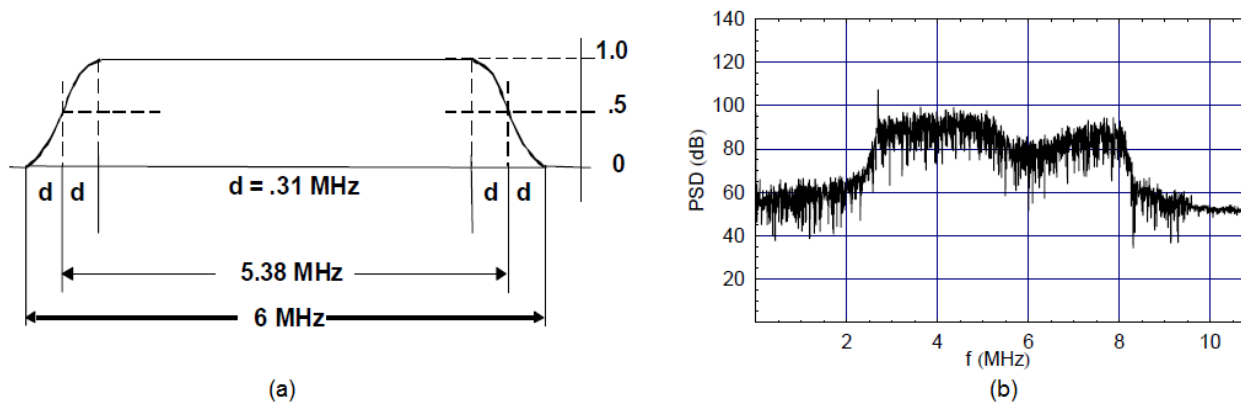


Figure 2.1 - Spectrum of a DTV channel: (a) ideal (pilot not shown) [8], and (b) measured [9].

The minimum signal-to-noise ratio (SNR) required for decoding an ATSC signal is set at 15.2 dB. With a maximum adjacent channel splatter of -46.5 dBc on the transmit side and a margin of about 4 dB, the adjacent channel blocker power is limited to a maximum of 27 dB higher than the wanted signal power in order to maintain sufficient SNR in the wanted channel. The maximum blocker power grows to 44 dB higher for the $N \pm 2$ channel, and increases by 4 dB

with each additional channel spacing until the $N \pm 6$ channel [8]. The overall ATSC blocker profile for weak signals is illustrated below in Figure 2.2.

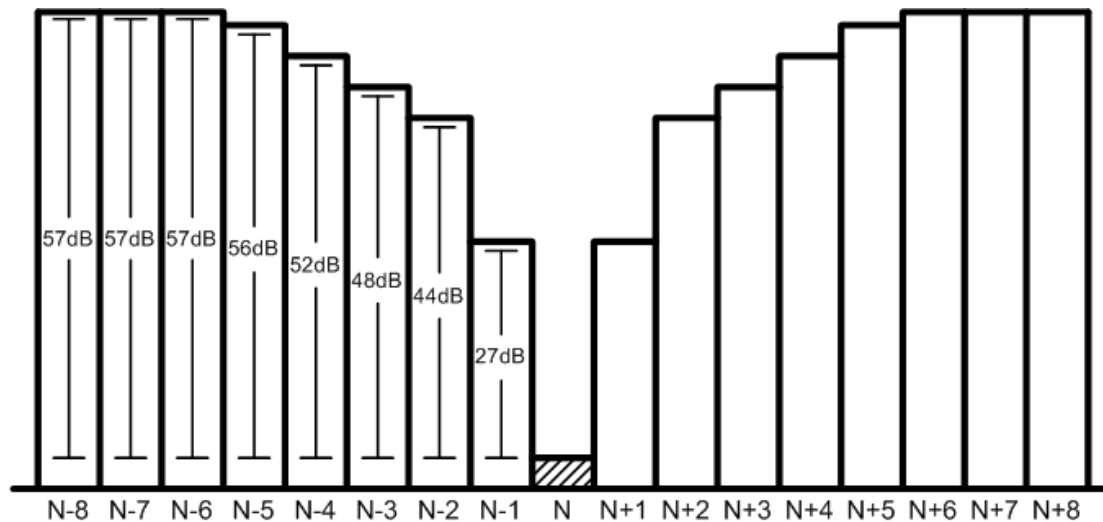


Figure 2.2 - ATSC blocker profile.

The thermal noise in a 6 MHz channel is about -105 dBm. Given the required decode SNR of 15.2 dB and a margin for receiver noise figure, the minimum sensitivity for DTV receivers, or the minimum input energy of a decodable signal, is set at -83 dBm [10]. However, since the sensitivity requirement must be met under worst case test conditions, many TV tuners in practice have higher sensitivity under common, realistic operating conditions.

In addition to DTV itself, wireless microphones also operate in the UHF band in vacant TV channels. There is no homogenized standard for wireless microphones, but they are usually narrowband, frequency-modulated (FM) signals with a bandwidth of 200 kHz [9]. Although wireless microphones are also unlicensed devices, they have priority over TVBDs and are considered primary users.

2.2 – IEEE 802.22 Requirements

The IEEE 802.22 standard governs Wireless Regional Area Networks (WRAN), which were developed to operate cognitively in the TV bands. The 802.22 work group, tasked with developing the standard, studied the geographical characteristics of base stations and primary receivers, and the interference and fading characteristics of wireless TV signals. The group then set sensing requirements for TVBDs to ensure an acceptably low level of interference with primary receivers.

In its preliminary standard published in 2008, the 802.22 work group set required sensing sensitivity level to -114 dBm for DTV in its 6 MHz channel bandwidth, for analog TV in its 100 kHz carrier bandwidth, and for wireless microphone in a 200 kHz bandwidth [5]. The sensitivity levels assume 0 dBi antenna gain, and that the receiver antenna is outside and at least 10 m

above ground, relatively free of obstructions. Since analog TV transmissions no longer exist in the U.S. and wireless microphones are under consideration to be moved into its own dedicated channels, we will focus only on DTV sensing in this work.

The challenge of spectrum sensing arises mainly from the -114 dBm sensitivity requirement, which is much lower than that of a standard receiver and is significantly below the -105 dBm thermal noise floor in a 6 MHz channel. Furthermore, sensing must occur with a low power overhead for mobile applications, and within a reasonably short sensing time. IEEE 802.22 set the sensing time at 2 seconds [11]; while this is a lengthy period, the actual sensing time should only be a small fraction since any time spent sensing cannot be spent transmitting and receiving on the detected idle channels. Additionally, IEEE 802.22 set both the probabilities of missed detection and false alarm to 0.1 [11]. Again, while these are relatively lax standards, it is in the interest of a cognitive radio system to minimize these probabilities for operational efficiency.

2.3 – State of the Art and Related Works

During the development of its WRAN policy, the FCC solicited TVBD prototypes for laboratory and field testing. Five such devices, submitted by Adaptrum, Institute for Infocomm Research, Microsoft, Motorola, and Philips Electronics, were tested in 2008. All prototypes reliably detected -114 dBm clean DTV signals in laboratory conditions; however, the sensitivities of the prototypes severely degraded, some up to 60-70 dB and others to the point of malfunction, when given real world conditions and the presence of interference [7]. The failure of the prototypes contributed to the FCC's ultimate decision to drop spectrum sensing requirements for TVBDs.

Aside from their less than desirable performance, the prototypes were designed as bulky and high-powered stationary devices analogous to traditional TV tuners. Detection was performed with power-hungry and time-consuming digital signal processing that put the shortest detection time at 0.1 seconds per channel, and the longer ones at tens of seconds per channel [7]. All of these factors make it evident that alternative approaches are necessary.

In the space of spectrum sensing for mobile devices, the first major work eliminated the need for complex digital signal processing as well as the preceding high-bandwidth, high-resolution analog-to-digital converter (ADC) by using analog correlation [12]. The work cross-correlated the downconverted receive signal with a window waveform, which effectively acted as an easily-tunable narrowband filter in the analog domain. However, since only energy detection was performed on the filtered signal, the sensitivity was limited by the noise floor to -74 dBm, which is no better than the sensitivity of standard TV band receivers. Furthermore, the receiver in [12] consumed 180 mW during both sensing and receive modes, whereas ideally, the sensing overhead should be much less than the power used for standard receiver operations.

The receiver front-end in a more recent work [13] was able to improve upon the performance of [12], achieving -84 dBm detection sensitivity while consuming only 30-44 mW. However, while [13] had a state-of-the-art, well-designed RF front-end that achieved good linearity, noise figure,

and harmonic rejection with low power, nothing separate or different was dedicated to spectrum sensing functionality save the design of the received signal strength indicator (RSSI) block itself. Thus, the detection sensitivity, while an improvement, was still constrained by the noise floor and no better than that of standard receivers.

The work in [14] proposed an architecture that cross-correlates the received signal from two identical, parallel front-end paths. This results in uncorrelated front-end electronic noise that could ideally average out to zero, effectively lowering the receiver noise figure without harming linearity or dynamic range. The receiver achieved a 4 dB noise figure (corresponding to an estimated detection sensitivity of -110 dBm) and 89 dB of spurious-free dynamic range (SFDR) in a 1 MHz bandwidth. This performance, however, came at the cost of high power, due to the use of two full analog front-ends as well as complex digital signal processing performing fast Fourier transforms (FFTs) and correlation. Not even including the power of two 14-bit ADCs, the receiver consumed 191 mW, which is much too high for mobile applications.

The work in [15], on the other hand, sought to demonstrate that spectrum sensing using digital signal processing could be performed in a fast and low power manner. The digital baseband processor in [15] achieved a detection sensitivity of -115 dBm with a detection period of 50 ms while consuming only 7.4 mW. However, in order to obtain this sensitivity, the processor required digital inputs with 20 bits of precision. Thus, while the processing power itself is low, the power that would be consumed by the RF front-end to provide this 20-bit high dynamic range input would certainly be prohibitively high.

Chapter 3 –Detection Modes

To construct a spectrum sensing system, we must first choose the method of detection – the method by which we evaluate the receiver input and make a decision on whether a signal is present. In this chapter, we review existing methods for signal detection, and evaluate their feasibility in our application of mobile spectrum sensing. Then, we propose a sensing system to implement the chosen detection modes. Finally, the theoretical performance of the proposed system and its various detection modes are evaluated using a Matlab simulation platform.

3.1 – Detection Techniques

There have been several proposed techniques to detect the presence of a signal or lack thereof. These methods have been analyzed extensively in theoretical literature, but we now evaluate them for their feasibility and applicability to our specific goal of sensing in the TV band.

3.1.1 – Energy Detection

The fastest and simplest detection method is energy detection, where we merely measure the received power. This can be done non-coherently, with knowledge of only the center frequency and the bandwidth of the target signal, and without any complex algorithms or processing. Simple energy detection, however, is incapable of distinguishing signal power from noise power, and hence it fails to detect in highly negative SNR regimes where signal is buried far under noise. For example, [16] evaluated the effectiveness of energy detection on a 4-MHz-wide QPSK signal, and the results are shown below in Figure 3.1.

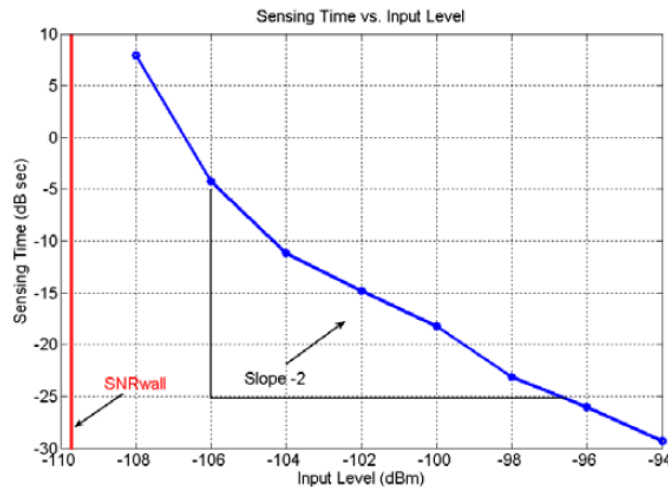


Figure 3.1 - Energy detection on a 4 MHz-wide QPSK signal [16].

The authors in [16] looked at sensing time versus input signal power with a constant false alarm rate of 5% and a detection rate of 60%. We see in Figure 3.1 that even using this extremely lax probability of detection, there exists an SNR wall below which detection is impossible given infinite sensing time. Furthermore, the sensing time remains prohibitively long for at least 10 dB

above the SNR wall. Thus, alternative detection methods are necessary if we seek to sense down to the -114 dBm level.

3.1.2 – Pilot Detection

Since we are targeting sensing specifically at ATSC signals, we can take advantage of certain characteristics specific to ATSC signals. One such characteristic is the sinusoidal pilot tone. As shown in Figure 3.2, the pilot in its narrow bandwidth has a higher SNR than the signal of the entire channel. Since the pilot is at a known position within the channel, we could simply narrow the bandwidth of detection to just that of the pilot, and gain higher sensitivity with energy detection.

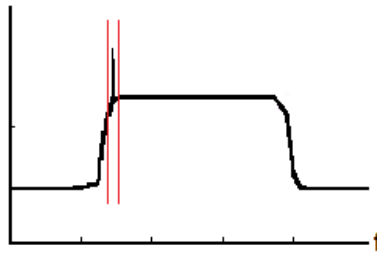


Figure 3.2 - ATSC pilot tone.

However, while pilot energy detection can achieve a higher sensitivity, it still has the same fundamental noise-floor limitations as channel energy detection. In addition, when dealing with such a narrow bandwidth, frequency offsets can be problematic and sharp filtering that approximates a brick-wall response is difficult to realize. These issues will all degrade the sensitivity further from the theoretical limit, making other, more robust approaches necessary.

3.1.3 – Feature Detection

A more robust detection method, cyclostationary feature detection, takes advantage of the fact that man-made, modulated signals are fundamentally different from noise in that they usually exhibit some kind of periodic behavior, either from carrier tones, cyclic prefixes, or other features. This periodicity can be extracted using a spectral correlation function (SCF), which measures the density of correlation between all the spectral components in a signal [4]. The SCFs for white noise and for a QPSK signal are shown below.

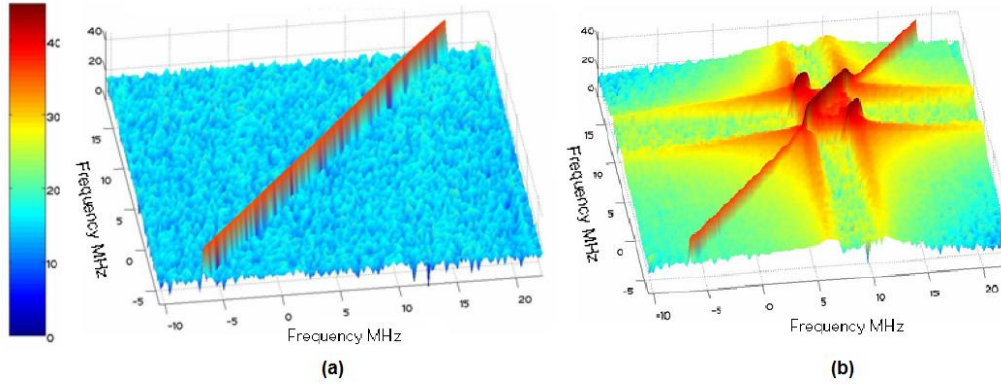


Figure 3.3 - SCF of (a) white noise, and (b) a QPSK signal [4].

As shown in Figure 3.3, white noise is only correlated at identical frequencies, while the spectral components at different frequencies of a modulated signal are also correlated, creating a distinguishable pattern in its SCF. However, feature detection presents many implementation challenges, the foremost of which is the complex processing required to generate the two-dimensional FFT and calculate the correlation functions. Complex processing indicates high power and long sensing time, both of which makes this technique unattractive for spectrum sensing in the mobile space.

3.1.4 – Autocorrelation

Another simpler approach that can be used to extract periodicity is to autocorrelate the signal. Autocorrelation – correlating a signal with a time-shifted version of itself – is a method that has long been used in spectrum analyzers, radio astronomy, and other applications that require high-sensitivity signal detection. Equation (1) below presents the autocorrelation function $R(\tau)$, where $X(t)$ is the signal being correlated, and τ is the time shift.

$$R(\tau) = E[X(t) * X(t + \tau)] \quad (1)$$

Since white noise is a stationary process that is uncorrelated with itself at different points in time, $R(\tau)$ for white noise would be 0 everywhere except for when $\tau = 0$. For modulated signals which are cyclostationary processes, the signal time-shifted by an integer number of periods should have the same statistical properties as the non-time-shifted version of the signal. Therefore, $R(\tau)$ should exhibit periodic behavior, with peaks occurring at τ equaling 0 as well as each integer period. Figure 3.4 illustrates this property.

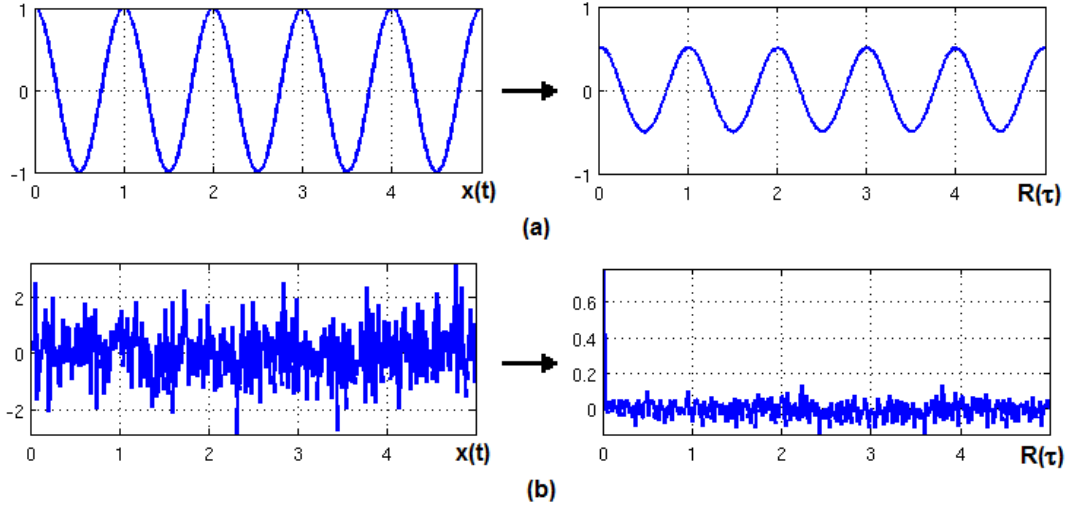


Figure 3.4 - Time-domain waveform and autocorrelation function of (a) a sinusoid, and (b) white noise.

If we look at $R(\tau)$ with τ equaling a non-zero integer period, the additive noise power would ideally average out to zero given enough time, and the resulting SNR of the autocorrelation function would be an improvement over the SNR of the signal itself. However, a digital implementation of autocorrelation still requires a relatively high amount of processing due to the multiplication operations needed to calculate the correlation function. As mentioned in Section 2.3, [14] demonstrated the power inefficiency of this approach. On the other hand, [12] demonstrated that correlation can be done in a lower power manner in the analog domain. While the system in [12] was high powered, much of it was used for generating the window waveforms, and the actual correlation operation itself consumed only 8 mW. We thus see that autocorrelation in the analog domain could potentially be a low-power, high-sensitivity solution to spectrum sensing.

3.2 – Proposed System

We propose a spectrum sensing system using a combination of channel energy detection, pilot energy detection, and autocorrelation. The pilot tone, as a pure sinusoid, provides an ideal candidate to extract the periodicity from using autocorrelation. The channel energy detection mode would be used for a fast, coarse scan of all the channels to eliminate ones that have strong signals residing within. Pilot detection and autocorrelation, being more time consuming, would only be used on a few likely-vacant channels where the signal, if present, is too weak to be sensed with energy detection.

The work in [17] proposed an efficient implementation of autocorrelation in the time domain using equivalent-time sampling. The input would be captured in parallel by two samplers, with the sampling clocks offset from each other by a time delay τ . By correlating the two sampled signals, and by varying τ in discrete steps, we effectively capture the autocorrelation of the input as a function of τ . Using this method, the resolution of the sampled signal is determined not by the actual sampling frequency but by the minimum τ -step, with effective sampling frequency

equaling $1/\tau_{step}$. Thus, this method enables us to save power by using a sampling frequency below the Nyquist rate.

Figure 3.5 below illustrates a high level block diagram of our proposed system. We adopt the method from [17] and sample the input signal, at frequency f_{RF} , onto two parallel paths. To target low power, we use equivalent-time sampling with $f_s < f_{RF}$. Since subsampling also acts as a mixer and downconverts our desired signal to baseband, the signal is then immediately channel-filtered following the sampler. In the first path, the filtered channel signal undergoes energy detection and pilot detection. If autocorrelation mode is required, then the signals from the two paths are multiplied together in the analog domain as we vary τ , the time delay between the two paths.

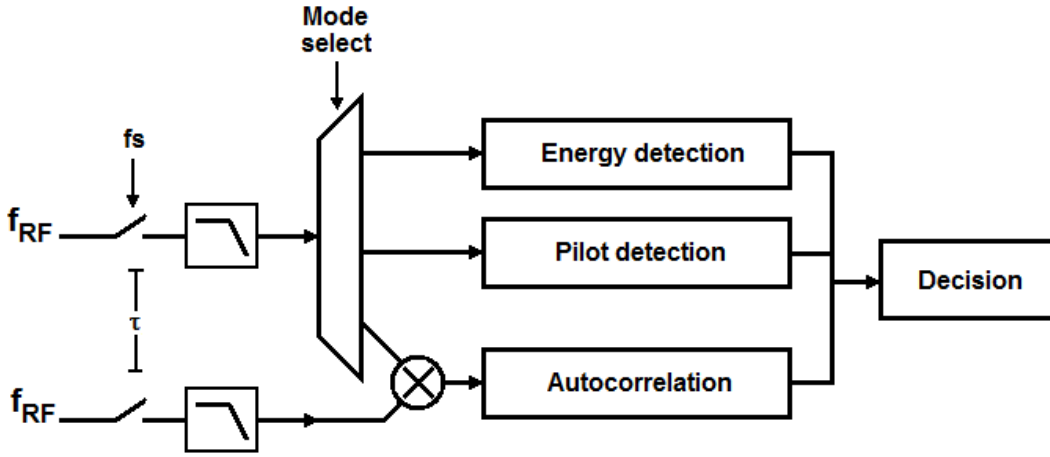


Figure 3.5 - High level block diagram of proposed system.

3.3 – Simulation Setup

We simulated our proposed system in Matlab in order to evaluate its different detection modes. In this section, we first describe the metric we used to evaluate the theoretical performance and sensitivity of each detection mode. Then, we look at how we modeled the DTV spectrum and the proposed receiver in order to test this metric.

3.2.1 – Detected SNR

To establish a metric to evaluate the performance of each detection strategy, we look at the two scenarios in Figure 3.6. In the first scenario, when no signal is present, the receiver will detect only the power of the noise from the environment; we denote this quantity N_{det} . For simplicity, we assume pure thermal noise and ignore non-white interference at this junction. In the second scenario, when a signal is present, the receiver picks up that signal along with the environmental noise, and we denote the total detected power under this condition S_{det} .

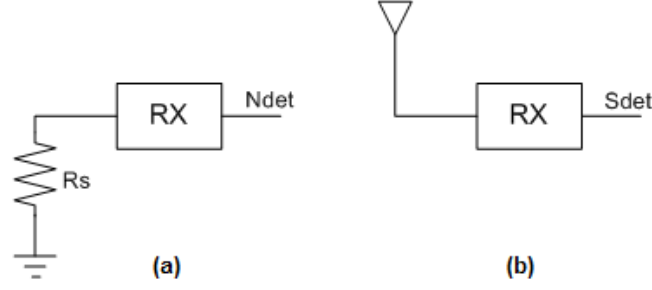


Figure 3.6 - Equivalent block diagram when (a) only noise is present, and (b) signal and noise are present.

If we denote the power of the input signal S_{in} , power of the environmental thermal noise N_{therm} , power of the receiver electronic noise at its output $N_{RX,out}$, and the total gain through the receiver A_{RX} , then for naïve energy detection:

$$N_{det} = A_{RX} * N_{therm} + N_{RX,out} \quad (2)$$

$$S_{det} = A_{RX} * (S_{in} + N_{therm}) + N_{RX,out} \quad (3)$$

Given a receiver noise figure, we can define $N_{RX,out}$ in terms of the input thermal noise and the noise factor F_{RX} :

$$N_{RX,out} = A_{RX} * N_{therm} * (F_{RX} - 1) \quad (4)$$

The SNR of an input signal is the ratio of the signal power to the thermal noise power. We can similarly define a term *detected SNR* to be the ratio of S_{det} to N_{det} , or the ratio of power detected when a signal is present to power detected when only noise is present. Using Equations (2)-(4), we find:

$$Detected\ SNR = \frac{S_{det}}{N_{det}} = \frac{S_{in}/N_{therm} + F_{RX}}{F_{RX}} = \frac{Input\ SNR + F_{RX}}{F_{RX}} \quad (5)$$

From Equation (5), we see that when the receiver adds infinite noise, the detected SNR reaches its lower bound of 1, or 0 dB. This indicates that inputs with the signal present and without the signal present are indistinguishable from each other at the point of detection. Conversely, when the receiver noise figure is very small, the detected SNR tracks with the input SNR.

3.2.2 – Signal Modeling

Figure 3.7 illustrates how transmit signals are modeled in simulation. We first generate a random stream of 8-level symbols in baseband. This data is then filtered to obtain the VSB frequency response, and a pilot is added as a DC level. Next, for every channel where we desire a signal to exist, we scale the signal power as desired and up-mix it to its appropriate RF channel position. Finally, additive white Gaussian noise (AWGN) is applied to model environmental thermal noise.

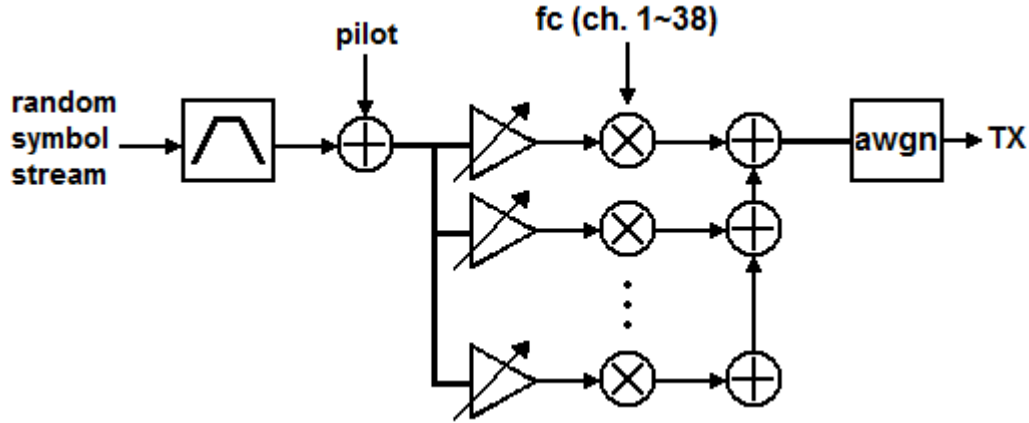


Figure 3.7 - Transmit signal generation in simulation.

For example, a simulated model of the worst case sensing scenario is shown below in Figure 3.8. The signal power (with additive thermal noise) is shown in blue and pure noise is shown in green. Figure 3.8(b) shows a 7-channel worst case blocker profile with a weak -114 dBm signal in the center channel. On the other hand, when a channel is idle like the center channel of Figure 3.8(a) is, the power of its surrounding channels can be arbitrarily high. We seek to differentiate between the two scenarios by detecting the presence or lack thereof of the weak signal in the center channel.

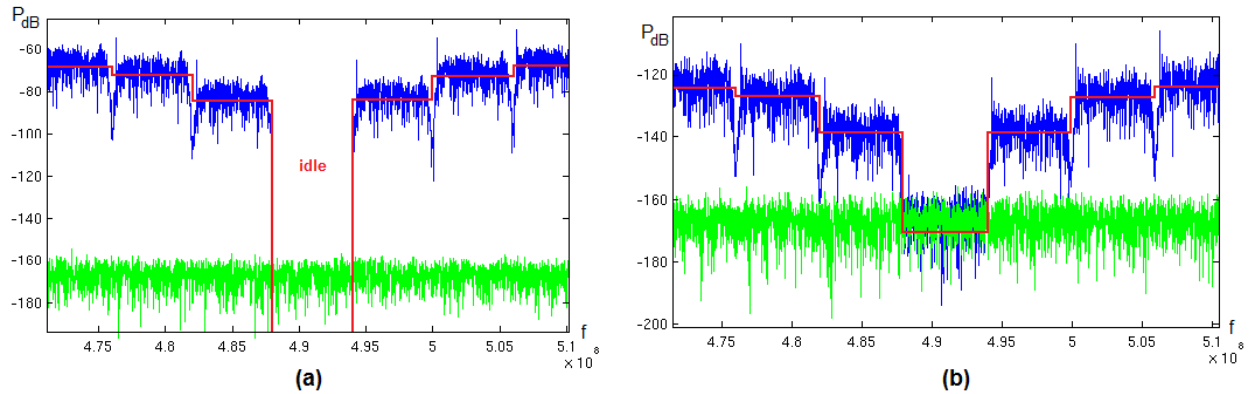


Figure 3.8 - Worst case sensing scenario, when target channel is (a) idle, and (b) occupied.

Figure 3.9 illustrates how the receiver is modeled in simulation for the measurement of detected power. First, a coarse bandpass filter is applied to the entire UHF band to attenuate out-of-band noise and thus reduce noise-folding during subsampling. We again ignore out-of-band interference and assume they have been sufficiently attenuated at this point in the receive path. We then subsample the signal at the desired channel with a sampling frequency equal to the channel frequency. The downconverted signal is filtered by either a channel filter or an ideal narrowband pilot filter, and the resulting power of the signal is measured using square law energy detection. For autocorrelation, we apply ideal mathematical autocorrelation to the filtered pilot, and then again apply square law energy detection to measure the power of the autocorrelated signal.

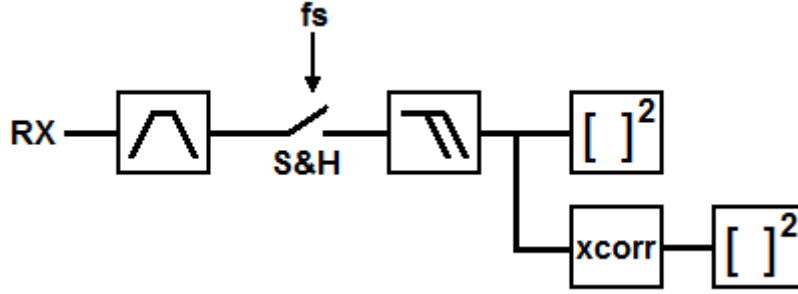


Figure 3.9 - Receiver modeling in simulation.

3.4 – Energy and Pilot Detection

Simulated energy and pilot detection are shown below in Figure 3.10, where the power spectrum densities of signal and noise are shown in blue and green, respectively. The red outlines illustrate example baseband filter responses.

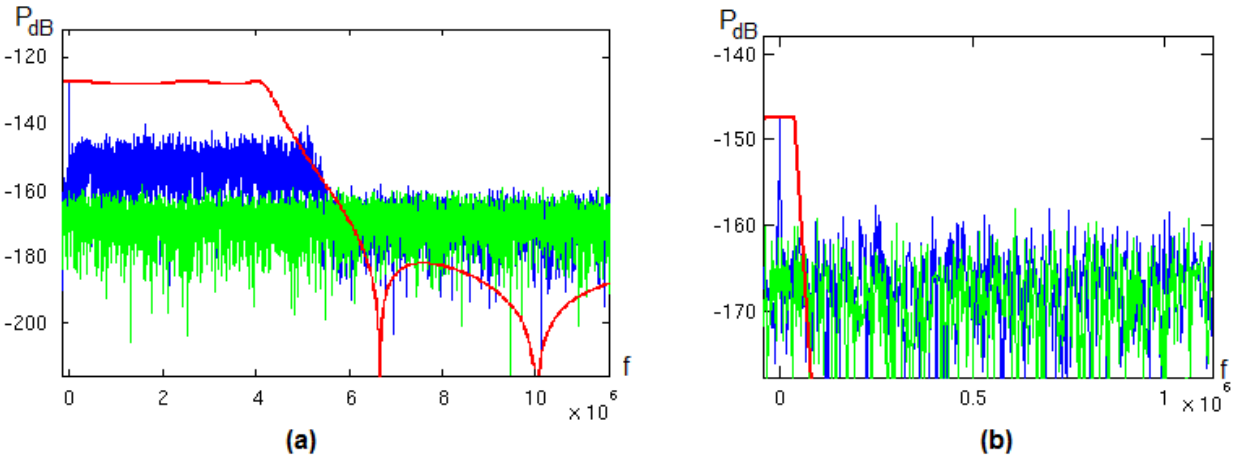


Figure 3.10 - Simulated (a) energy, and (b) pilot detection at baseband.

In Figure 3.10 (a), where signal power is distinctively higher than noise power, we perform energy detection by filtering the entire channel and measuring the resultant power. On the other hand, in Figure 3.10 (b), we see that for weak inputs, the signal power is completely buried in noise and indistinguishable from it. However, the narrowband pilot tone still rises above the noise floor. In this case, we perform pilot detection by narrowing the filter response to just the bandwidth of the pilot, and then, again, measuring the resultant power.

To evaluate the performance and limitations of energy and pilot detection, we first note that thermal noise power is given by

$$N_{therm} = k * T * \Delta f \quad (6)$$

where k is Boltzmann's constant equaling $1.38 \times 10^{-23} \text{ JK}^{-1}$, T is temperature in Kelvin, and Δf is the noise bandwidth under consideration.

At room temperature, $k * T$ gives -203 dB, or -173 dBm, of noise power per Hertz. In a 6 MHz bandwidth, the noise power becomes -105 dBm, and an input signal at -114 dBm therefore has a -9 dB input SNR. In a 10 kHz bandwidth, the noise power lowers to -133 dBm. However, the power of the pilot in its narrow bandwidth is 11.3 dB below the total channel power. Thus, the minimum pilot power is -125.3 dBm, making the minimum input SNR of the pilot 7.7 dB.

Using Equation (5), we can plot the detected SNR as a function of input SNR for various receiver noise figures. Shown below in Figure 3.11, for an input SNR of -9 dB for channel energy detection, the detected SNR in the ideal case (with no added receiver noise) is about 0.5 dB. Similarly, the detected SNR in the ideal case for pilot energy detection is about 8.4 dB.

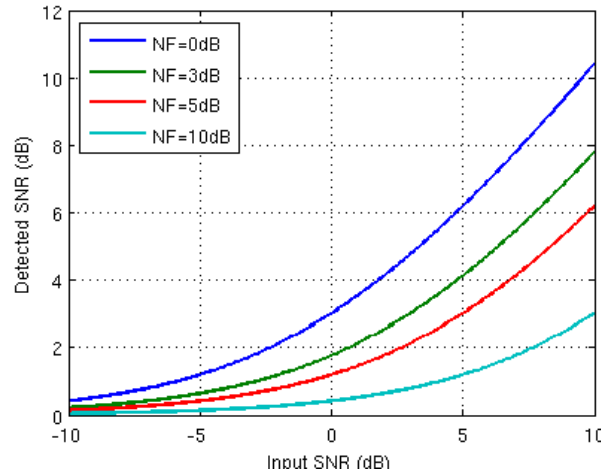


Figure 3.11 - Detected SNR as a function of input SNR for energy detection.

Although decision statistics and algorithms are outside the scope of this report, we will briefly discuss at this juncture the factors that determine the detected SNR threshold for sensing and their implications on system design. In theory, the probabilities of false alarm (P_{fa}) and detection (P_d) are functions of where the decision threshold is set and how many samples are used. More samples allow more averaging, resulting in a more accurate measurement that is less sensitive to instantaneous variances in noise and signal power. To find P_{fa} and P_d , we can model the probability density function (PDF) of noise with a Gaussian distribution and use Neyman-Pearson hypothesis testing, shown in Figure 3.12 below [16].

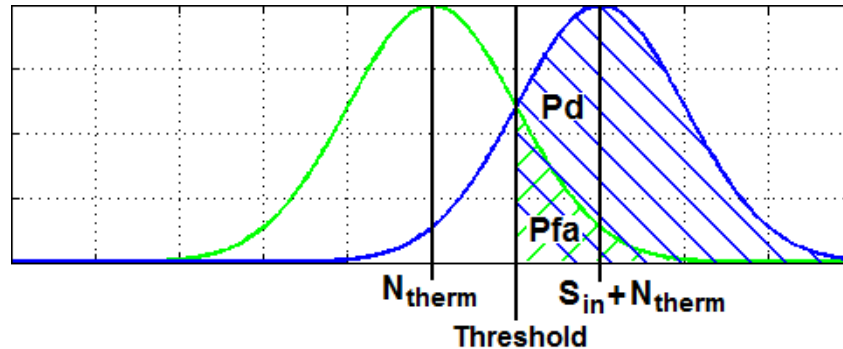


Figure 3.12 - Probability density functions for signal and noise.

In Figure 3.12, the green shows the PDF for pure noise, and the blue shows the PDF for signal with noise, where the mean power has shifted but the shape remains Gaussian. The decision threshold should be set somewhere between the two means, and anything above threshold in the noise distribution results in a false alarm, while anything above threshold in the signal distribution indicates a correct, desired detection. The shaded areas of P_{fa} and P_d can be quantified as:

$$P_{fa} = Q\left(\frac{S_{thresh} + N * N_{therm}}{\sqrt{2 * N * (N_{therm})^2}}\right) \quad (7)$$

$$P_d = Q\left(\frac{S_{thresh} + N(S_{in} + N_{therm})}{\sqrt{2 * N * (S_{in} + N_{therm})^2}}\right) \quad (8)$$

where S_{thresh} is the decision threshold power, N is the number of samples, and $Q(x)$ is the tail function for a normal distribution. For an ideal channel detection scenario, where the detected SNR is about 0.5 dB, we gain the following P_{fa} and P_d as a function of N shown in Figure 3.13. As expected, a higher threshold results in a longer detection time, but a lower false alarm rate. And when the decision threshold (in the form of detected SNR) is set to 0.7 dB > 0.5 dB, the system fails to detect completely.

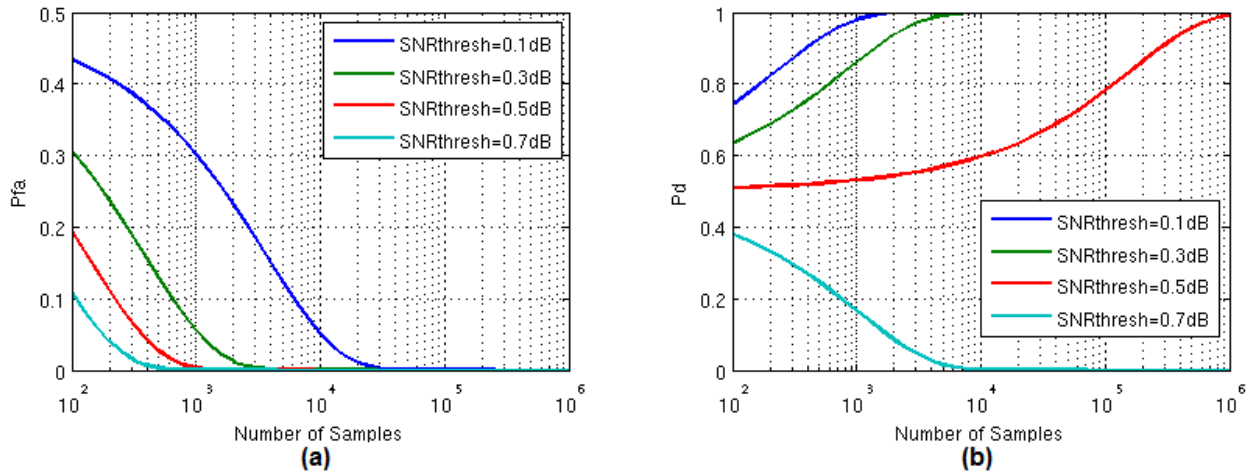


Figure 3.13 - (a) Pfa, and (b) Pd for ideal channel energy detection.

While theory seems to have indicated that we can set a detection threshold that will sense the 0.5 dB detected SNR with $P_{fa} < 0.1$ and $P_d > 0.9$, there are numerous problems in a realistic implementation that precludes this. The main limitation is the problem of noise uncertainty, where the actual noise power can differ greatly, more than several dB, from the expected, theoretical value due to process, temperature, and environmental variations. Periodically estimating noise using the detection system such as in Figure 3.6 (a) can alleviate the problem somewhat, ensuring that the system's own noise figure variations as well as gradually changing environmental conditions are accounted for. However, there would always be some residual estimation error, and other factors, such as interference profiles, can be rapidly time-varying.

These variations in noise, and consequently in detected SNR, make pure energy detection impossible at -114 dBm.

On the other hand, going back to Figure 3.11, we see for pilot detection's 7.7 dB input SNR, the detected SNR is about 8.4 dB, which should provide enough margin for noise uncertainty. However, there are additional factors to take into account. Firstly, up to this point, we have only looked at the ideal case without added receiver noise. Furthermore, a real system has imperfect filtering, and all residual out-of-channel noise *and* signals become additive noise from the perspective of our sensing system. So for example, if we add in a standard receiver noise figure of 5 dB, and make the conservative estimate that effective noise bandwidth is 2 times the desired signal bandwidth, then the input SNR lowers to 4.4 dB and the detected SNR to 2.9 dB. The detection margin has now become severely eroded. With possible additional erosions from non-linearities, interference, and wider effective noise bandwidths, a more robust detection method is necessary.

3.5 – Autocorrelation

When we autocorrelate thermal noise, since every value in the autocorrelation function is an average of numerous samples of noise, the variance of autocorrelated noise power should ideally be zero. This, in effect, makes the PDF of noise power approximate a Dirac delta function rather than a Gaussian. This behavior is shown below in Figure 3.14.

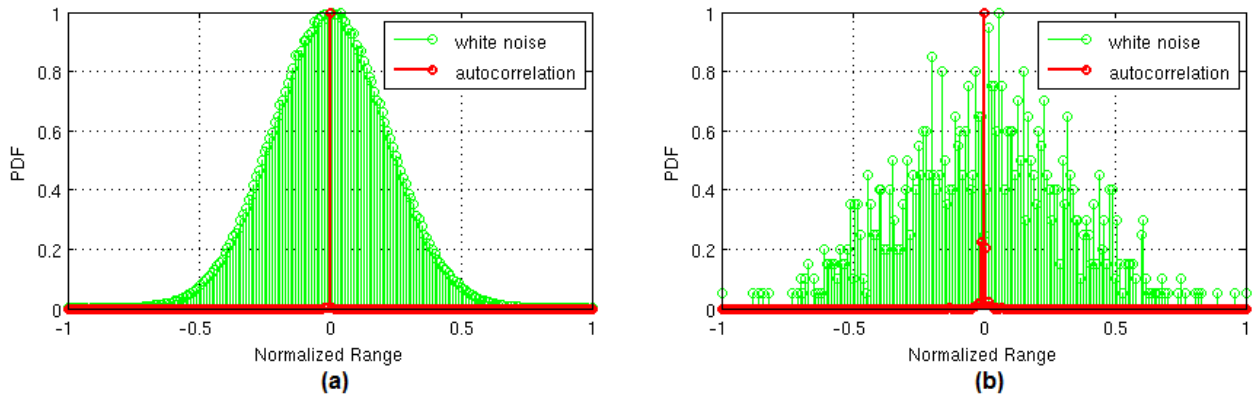


Figure 3.14 – Normalized PDF of white noise and its autocorrelation with (a) 10^6 samples, and (b) 10^3 samples.

When a signal is present, its additive thermal noise should also average out to zero when autocorrelated, while the signal itself retains its mean power as its periodicity is captured by and translates to the autocorrelated response. Thus, the PDF of the signal power should also approximate a delta function, but centered at the mean power of the signal. With two delta functions rather than two Gaussians, the P_{fa} and P_d are obviously drastically decreased and increased, respectively, for the same threshold and mean power separation between S_{in} and N_{therm} . This indicates that more robust sensing could be achieved with lower detection margins.

In Figure 3.10 (b), with a more limited number of samples, the PDF of the autocorrelation deviates from its ideal behavior and widens. However, since each autocorrelation point is already the average of many noise samples, this creates a layer of buffering against instantaneous variances in noise power. Therefore, the difference in the autocorrelation PDF is slight while the PDF of pure energy detection differs much more drastically from its ideal distribution, and the relative robustness of autocorrelation is maintained as the number of samples scale.

We simulated detection using autocorrelation through our model receiver, and the results are shown below. In Figure 3.15, the spectrum of the filtered pilot, where blue is signal and green is noise, is shown before and after autocorrelation. With pure energy detection, signal power is barely greater than noise and visually indistinguishable. With autocorrelation, an obvious improvement in detected SNR is demonstrated.

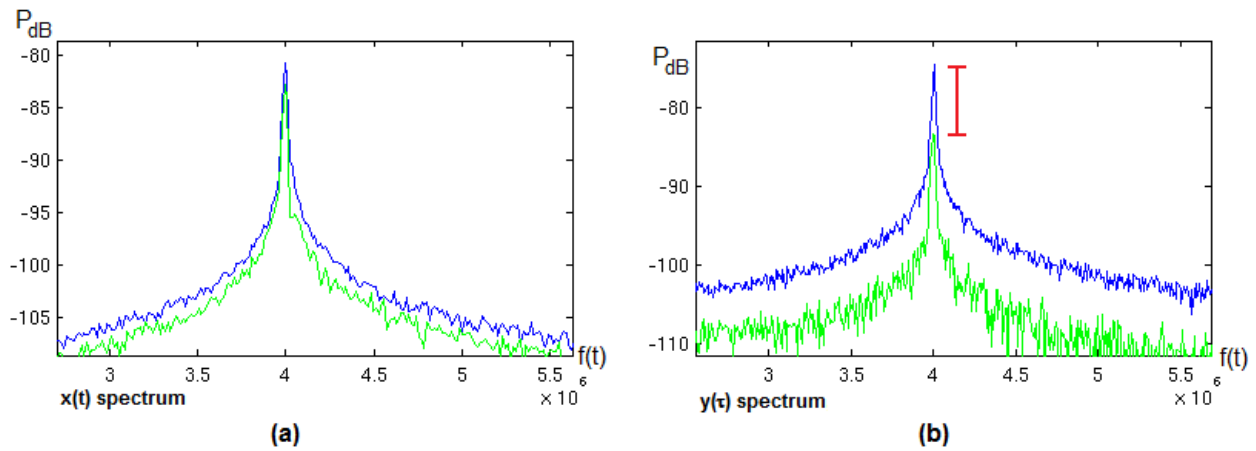


Figure 3.15 - Spectrum of detected pilot and noise for (a) energy detection, and (b) autocorrelation detection.

Next, we simulated detected SNR as a function of input power. Shown below in Figure 3.16, while the average detected SNR for pilot detection is acceptable for very weak signals, instantaneous variations can lower it to the undetectable region around 0 dB. Autocorrelation detection, on the other hand, maintains a robust margin.

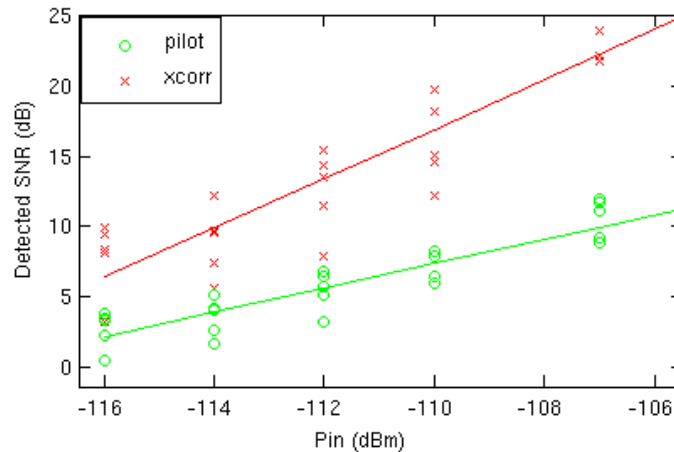


Figure 3.16 - Detected SNR vs. input power for pilot and autocorrelation detection.

3.6 – System Sensitivity

Simulation results for the sensitivities of the three detection modes are shown below in Figure 3.17. The receiver has been assumed to be ideal without additive electronic noise. However, noise-folding effects are included from simulation of subsampling behavior.

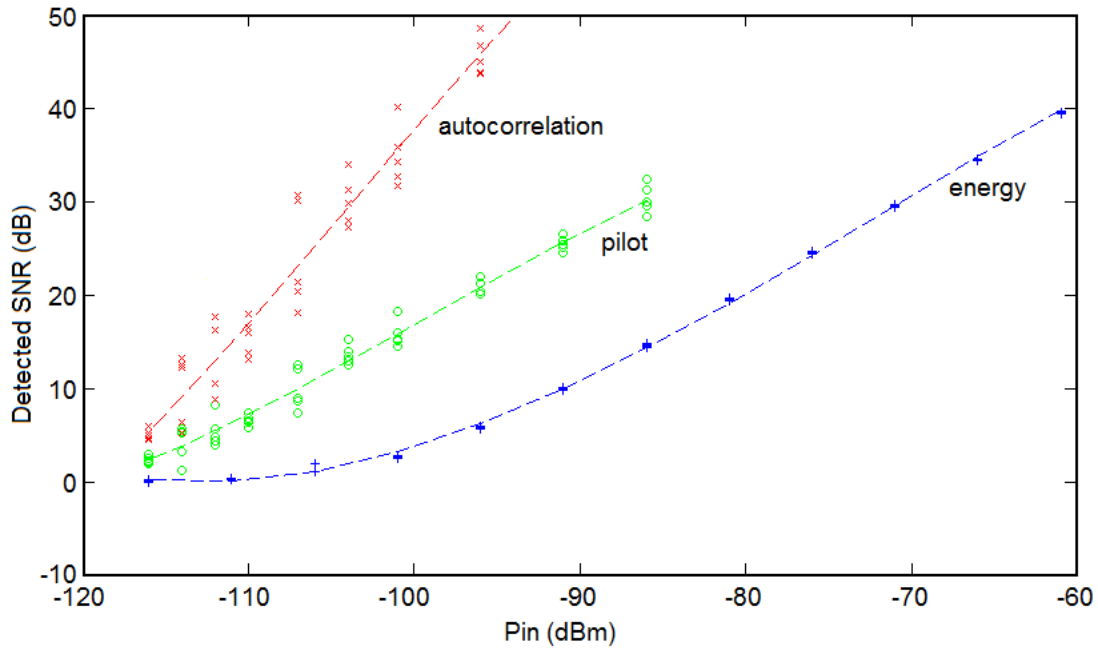


Figure 3.17 - Detected SNR vs. input power for all detection modes.

As expected, for channel energy detection, the detected SNR tracks input power until the noise floor is reached at around -105 dBm. Then, as signal power becomes much weaker than noise power, the detected SNR flattens to 0 dB since a signal buried in noise has approximately the same power as pure noise, and energy detection can no longer differentiate between the two scenarios. Pilot detection never reaches its own noise floor at -133 dBm, but at weak input signals, variations bring detected SNR very close to 0 dB. Finally, autocorrelation demonstrates improvements in detected SNR over the two other methods, and remain relatively robust at the weakest signal levels.

Chapter 4 – System Architecture

In this chapter, we look at specific implementation options for our proposed system. We specifically concentrate on the two main innovations of our system with respect to spectrum sensing – subsampling and analog autocorrelation. We then propose a detailed system architecture to realistically implement the theoretical detection modes described in Chapter 3.

4.1 – Subsampling

Downconversion using subsampling is one of the primary areas of novelty in our spectrum sensing system. Subsampling creates power savings from the use of lower sampling frequencies (and the associate clock generation and distribution networks), while maintaining high signal fidelity. Furthermore, for our purposes, subsampling allows ease of autocorrelation implementation in the analog and time domains. The concept of subsampling is shown below in Figure 4.1.

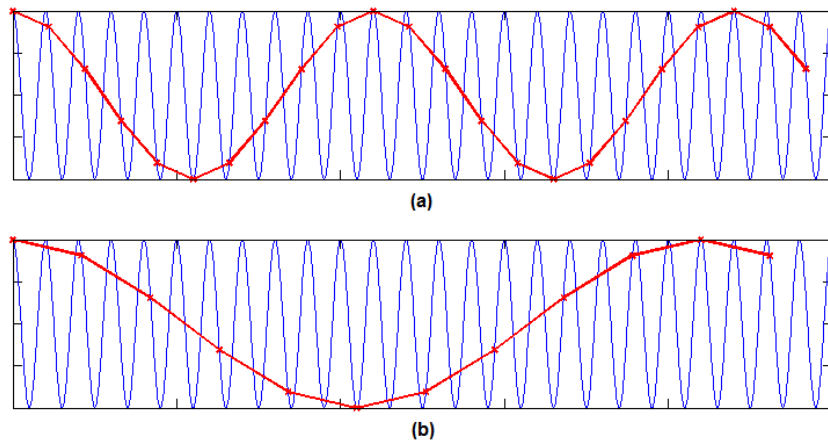


Figure 4.1 - Illustration of subsampling with (a) subsampling ratio of 1, and (b) subsampling ratio of 2.

In traditional Nyquist sampling, the number of points per cycle that are sampled scales linearly with the ratio of sampling frequency to signal frequency. Thus if we want to sample many points per cycle in order to not lose signal fidelity through the sampling process, we must oversample the signal with a much high sampling frequency. If the input frequencies are in the hundreds of MHz range as ours are, oversampling becomes very power-inefficient and unattractive.

Subsampling, on the other hand, uses a sampling frequency slightly below the signal frequency. As shown in Figure 4.1(a), the signal is consequently sampled only once per period, but each sample is of a different temporal location within the signal period. Given enough signal cycles, the entire characteristic of a periodic signal will be sampled, and we get a high-fidelity signal with numerous sampled points per cycle, albeit translated to a lower frequency. In Figure 4.1 (b), the same concept is demonstrated but with a subsampling ratio of 2, meaning that the signal is sampled once per every 2 cycles.

Based on the number of samples per cycle, the equivalent oversampling frequency of a subsampled signal is:

$$f_{seq} = \frac{fin * fs}{fin - N * fs}, \quad 0.5 * fin < N * fs < fin \quad (9)$$

where fs is the subsampling frequency, fin is the signal frequency, and N is the subsampling ratio. According to Equation (9), we can get a high f_{seq} with an extremely low fs if we choose a high N . More importantly, the linear relationship between samples per cycle and sampling frequency has been broken. To achieve a higher f_{seq} and therefore a higher number of samples per cycle, we would only need to place $N * fs$ closer to fin .

The translated lower frequency after subsampling is:

$$f_{sampled} = fin - N * fs, \quad 0.5 * fin < N * fs < fin \quad (10)$$

We can therefore easily choose N and fs so that $f_{sampled}$ will fall in the baseband, providing a convenient mixing step as well. This frequency aliasing effect, however, also provides the main problem of subsampling in the form of noise folding, illustrated below in Figure 4.2.

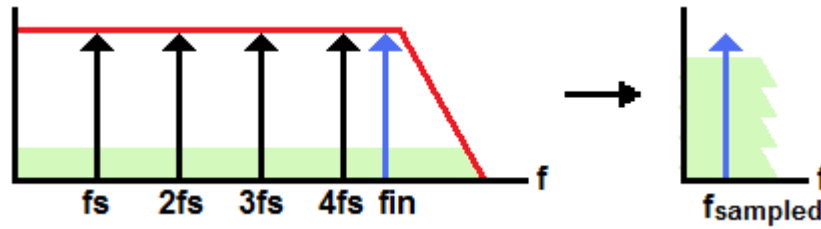


Figure 4.2 – Noise folding from subsampling.

The problem of noise folding arises from the subsampling ratio being a design decision that the system is blind to. For example, in Figure 4.2, we may intend to sample fin with $N = 4$, but everything around the frequencies $3 * fs$, $2 * fs$, and fs will be translated to baseband as well, because the system cannot know what we intended N to be. This causes the noise around each multiple of fs , shown in green, to be folded to the same $f_{sampled}$, severely degrading SNR.

Furthermore, in traditional Nyquist sampling, anti-alias filters can attenuate high frequency, out-of-band signals and noise so that they do not get translated by harmonics of the sampling clock. No such method can be used here as filter bandwidths must accommodate the input signal, which is at a higher frequency than the clock signal and its harmonics. While bandpass anti-alias filters can be used, the noise generated from the sampler itself will still have a low-pass characteristic and will not be filtered. This folding of sampler noise has resulted in subsampling mixers generally having prohibitively high noise figures, making them unattractive for low noise systems [18].

For the above reasons, we choose for our implementation the low subsampling ratio of 1, so that noise is folded only once. Furthermore, the folded noise is spread over the full spectrum up to the sampling frequency and is not concentrated purely in the baseband, which also improves noise figure. While the power benefits of subsampling are lessened with $f_s \approx f_{in}$, it is still a factor of 2 improvement over the minimum Nyquist sampling frequency of $2 * f_{in}$, and more critically, we still gain the benefit of a much higher f_{seq} .

Next, we consider the options in converting the signal to baseband. Standard receivers would set the mixing frequency to the carrier frequency, and use I and Q paths to reject aliased images. A Hilbert filter would then reconstruct the real baseband signal from the downconverted vestigial sideband in order to allow signal decoding [19]. However, since the pilot also falls on the carrier frequency, this provides several problems for our purposes. Firstly, the pilot power, as a DC component, becomes vulnerable to DC offset and low frequency flicker noise. Moreover, the pilot is in a 10 kHz bandwidth, indicating that a lengthy time will be required to collect enough samples for an accurate measurement.

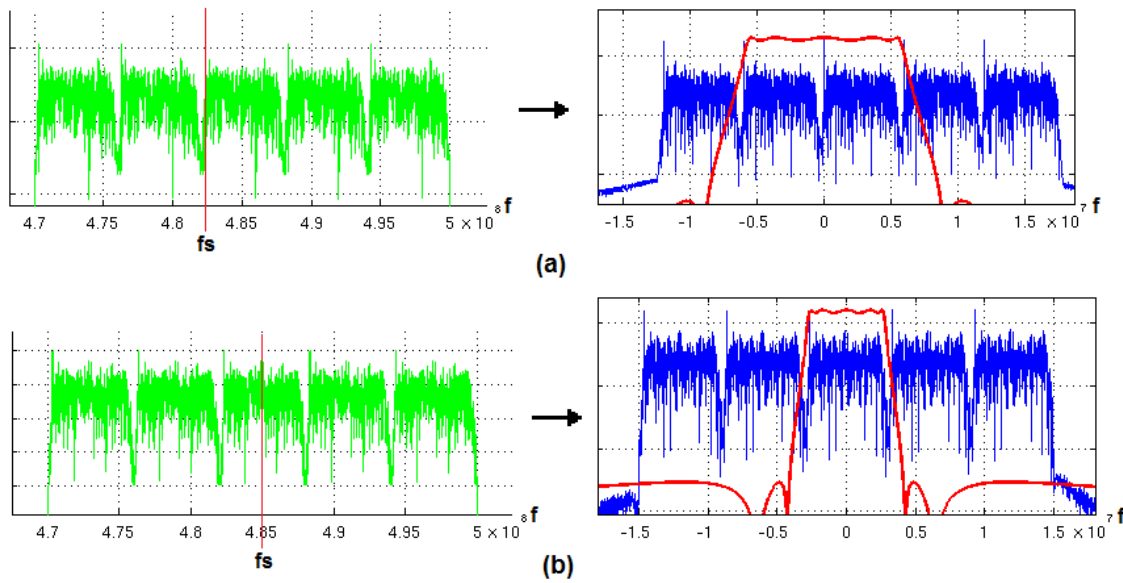


Figure 4.3 - Conversion to baseband with sampling frequency set to (a) carrier frequency, and (b) channel center frequency.

A better solution, shown in Figure 4.3, is to sample at the center frequency of each channel instead of at the carrier. This would give us true direct conversion where image rejection and I and Q paths are no longer necessary, as no out-of-channel signals would fall in-band. Without image rejection and Hilbert filters, the signal would fold onto itself to create an un-decodable 3 MHz baseband signal. However, since we only seek to detect the presence of a signal and not to receive it, a non-real signal is acceptable for our system.

While the channel undergoes true direct conversion, the pilot, on the other hand, falls at the low IF frequency of 2.69 MHz. This alleviates the flicker noise problem and eliminates the DC offset

problem altogether. Furthermore, pilot power can now be measured much more rapidly, and the pilot retains its periodicity to allow autocorrelation detection.

4.2 – Autocorrelation

Autocorrelation detection of the pilot in the analog and time domains is the other primary novelty of our system. Using the two sampler technique described in [17] and Section 3.2, we step the time shift between the two samplers and get the autocorrelation as a function of τ , shown in Figure 4.4.

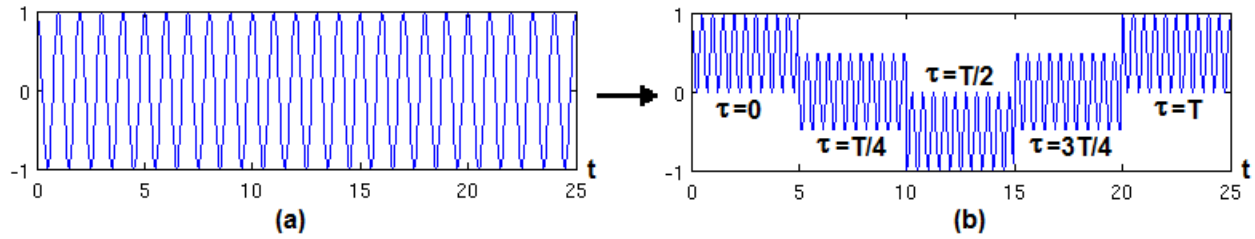


Figure 4.4 - Time-domain waveform of (a) a sinusoid, and (b) its autocorrelation.

Figure 4.4 (a) shows a sinusoid signal, and Figure 4.4 (b) shows 5 τ -steps encompassing 1 cycle of its autocorrelation in time domain. In post-processing, the mean at each τ -step would be calculated to recreate the sinusoidal behavior. This is shown in Figure 4.5 below, where we look at the effect of our autocorrelation technique on a moderately noisy sinusoid tone. While there is still some un-ideal behavior from residual noise in the autocorrelation in Figure 4.5 (b), we see that it is much cleaner than the original signal in Figure 4.5 (a), indicating an SNR improvement.

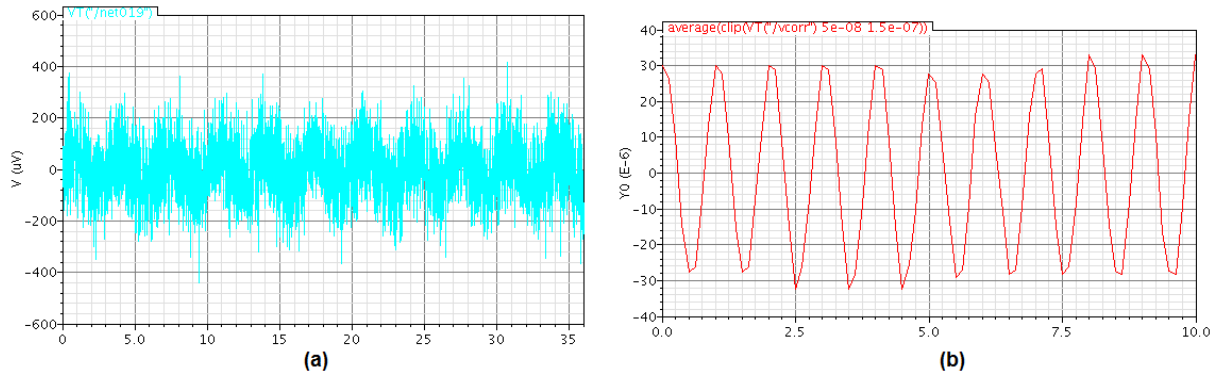


Figure 4.5 - Simulated time-domain waveforms of (a) a noisy pilot, and (b) its autocorrelation.

Stepping τ through multiple autocorrelation cycles as in Figure 4.5 (b) is useful if we eventually want to extract frequency information through a FFT; however, it is unnecessary for our purposes. Since the pilot is a pure sinusoid, one autocorrelation cycle is adequate to gain its entire characteristic and to make a decision. There should be no difference (statistically) between the autocorrelated results of, for example, $\tau = T/4$ and $\tau = 5T/4$, and thus, simulation has shown that performance could be better improved by simply averaging at each τ -step for a longer

period of time rather than incorporating more τ -steps and multiple autocorrelation cycles. This eliminates the need for long, variable delay lines.

To create the time shift between the two receiver paths, we delay the second sampling clock by τ , where all τ -steps are less than one sampling clock period. Due to the properties of subsampling, a phase shift in the sampling clock translates directly to an identical phase shift in the subsampled baseband signal. This is illustrated in Figure 4.6 below, where the blue RF signal is sampled on 4 different phases of the subsampling clock.

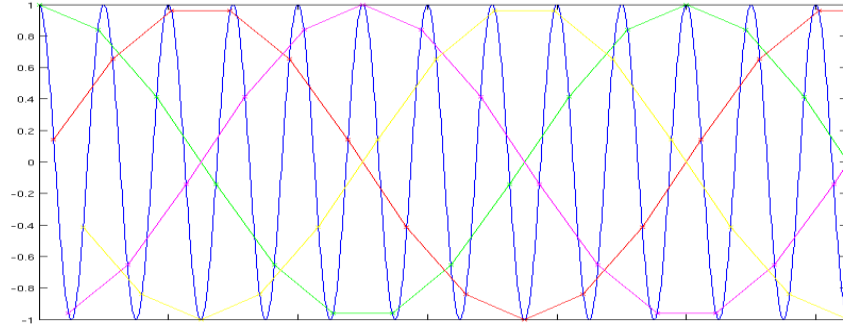


Figure 4.6 - Subsampled response to phase shifts in the sampling clock.

4.2.1 - Single-Bit Autocorrelation

To perform correlation, a multiplication circuit with high dynamic range, high linearity, and low noise would be required. Designing a traditional high-resolution analog multiplier to meet these specifications presents a challenging task. Alternately, a much easier method of performing multiplication is to first convert the high-resolution analog signals into low-resolution, single-bit digital signals using sigma-delta ($\Sigma\Delta$) modulators. Then, multiplication would simply be an XOR operation on the two bitstream outputs from the $\Sigma\Delta$ modulators.

The structure of a $\Sigma\Delta$ modulator is shown below in Figure 4.7.

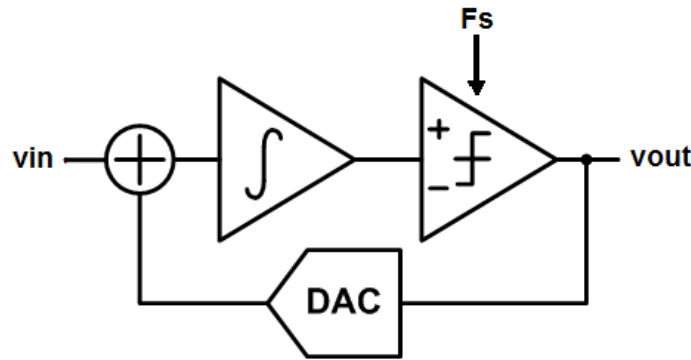


Figure 4.7 - Block diagram of a sigma-delta converter.

The $\Sigma\Delta$ modulator feeds back each digital bit and integrates the error in order to estimate the next bit output, with the comparator frequency F_s much greater than the highest frequency content of the input. This results in an oversampled bitstream that approximates the input signal, as shown

below in Figure 4.8. The blue lines representing output bits are all high when the signal (in red) is at its highest points, and the opposite occurs when the signal is at its lowest points. When the signal nears 0, the bits alternate between 0 and 1.

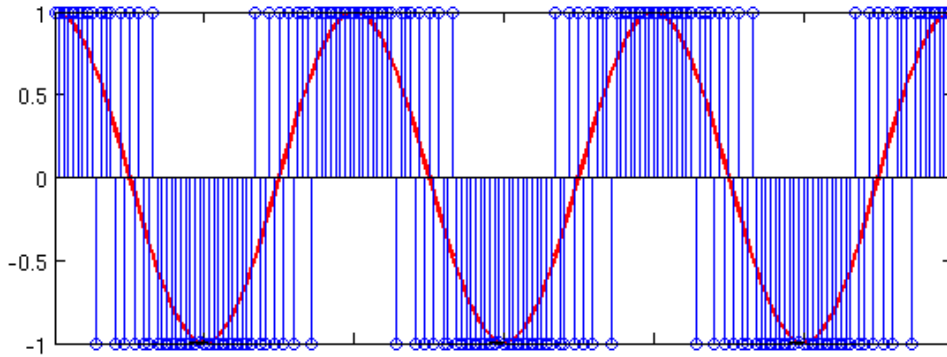


Figure 4.8 - Sigma-delta conversion of a sinusoid.

$\Sigma\Delta$ conversion also has a beneficial noise-shaping property, where the noise transfer function has a highpass characteristic, so that quantization noise is pushed to higher frequencies outside of baseband. This is illustrated below, where Figure 4.9 (a) shows the spectrum of the input signal and Figure 4.9 (b) shows the spectrum of the $\Sigma\Delta$ output. In a standard $\Sigma\Delta$ ADC, the noise at higher frequencies will eventually be filtered out through digital decimation, and thus noise-shaping leaves a high resolution baseband signal.

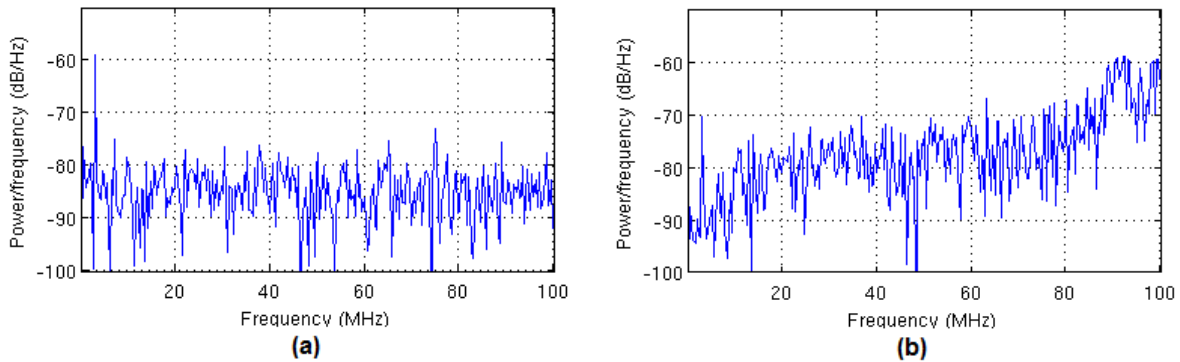


Figure 4.9 - Frequency spectrum of (a) sigma-delta input, and (b) sigma-delta output.

Applying $\Sigma\Delta$ conversion directly to correlation presents a fundamental problem due to this noise-shaping behavior. If we correlate the two bitstreams directly, the much higher quantization noise at high frequencies will mix with each other back to baseband, drastically raising baseband noise and degrading SNR. If we decimate the bitstreams to filter out high frequency noise, we are left with two high-resolution digital signals, the multiplication of which then require the complex digital processing we sought to avoid. We thus lose the advantage of bitwise multiplication that using $\Sigma\Delta$ modulators was originally intended to provide.

An alternative to decimation is to correlate one $\Sigma\Delta$ bitstream with the original analog signal, a technique proposed in [20]. Since the quantization noise in the $\Sigma\Delta$ output is completely

uncorrelated with the analog thermal noise, it should ideally not affect the final correlated result. While multiplying a bitstream with an analog signal is more complicated than a mere XOR operation, it is still trivial compared to full analog or full digital multiplication. The bitstream would simply act as a sign bit that selects between the inverted and non-inverted forms of the analog signal. The performance of the various correlation methods were simulated and shown below in Figure 4.10.

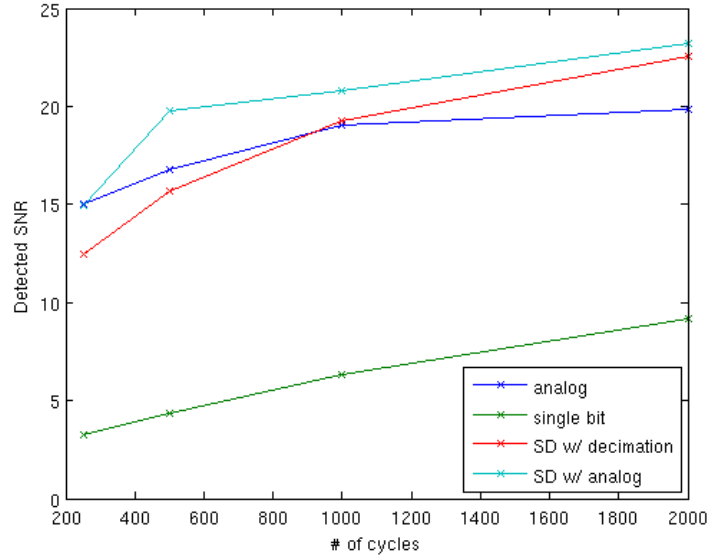


Figure 4.10 - Comparison of various autocorrelation implementations.

We compared the detected SNR of subsampled signals after correlation. Blue shows the baseline performance with an ideal analog multiplier. When both signals are converted to $\Sigma\Delta$ bitstreams and correlated, the detected SNR degrades severely, shown in green. The red shows the result of decimating before correlation, and we see that performance has been brought back up to around the analog baseline level. However, correlating a bitstream with analog signal, shown in cyan, can also achieve a performance comparable with and in fact a little bit better than the analog baseline, indicating that this hybrid correlation method would work well for our purposes. Finally, as expected, with more samples and more averaging, the detected SNR in all cases improve.

4.3 – System Architecture

The final system architecture is shown below in Figure 4.11.

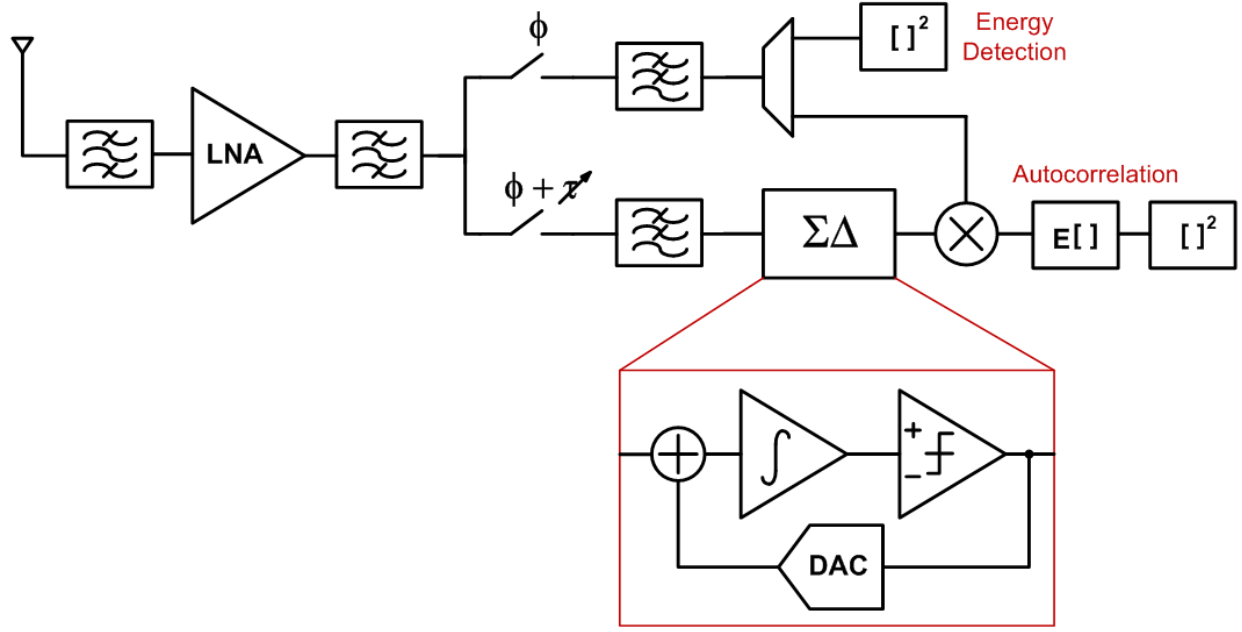


Figure 4.11 – Block diagram of proposed system.

After the antenna, we first need a fixed bandpass filter to attenuate potential out-of-band blockers, especially from the neighboring cellular band at around 800 MHz ~ 900 MHz. Next, the low noise amplifier (LNA) must provide wideband input impedance matching as well as high gain, which is critical for a low system noise figure. The total noise factor of a set of cascaded blocks is given by

$$F_{tot} = F_1 + \frac{F_2 - 1}{A_1} + \frac{F_3 - 1}{A_1 A_2} + \dots \quad (11)$$

where F_1, F_2, \dots and A_1, A_2, \dots are respectively the noise factors and gains of each block. Since the sampler especially will be a high noise block due to folding, a high LNA gain is required to lessen its impact on overall system noise figure. However, we also expect a wide dynamic range in input signal power levels, from -114 dBm up to -8 dBm. In order to not saturate the subsequent blocks or even the LNA itself, the front-end must also have variable gain.

Next, we consider that the sampler will fold not only noise, but also any interference around the harmonic frequencies of the sampling clock. Thus, another RF bandpass filter is needed before the sampler to provide additional interference attenuation. In fact, even if the first bandpass filter were ideal, there would still be harmonic distortion and intermodulation products generated by the LNA that would need to be attenuated prior to sampling. Since our target band of 400 MHz ~ 800 MHz is relatively wide, the second bandpass filter will be implemented as a narrow RF tracking filter in order to provide sufficient attenuation at harmonic frequencies.

After the second RF filter, the signal is split into two paths for our implementation of autocorrelation. The two-path method gives an additional benefit in that the additive noise of the

samplers as well as all the baseband blocks would be uncorrelated, which effectively decreases the system noise figure in autocorrelation mode. Of course, path-splitting from the antenna and having two completely separate RF front-ends, as in [14], would decrease the noise figure further. However, the power overhead was determined to be too high for our target specifications.

After the sampler, both paths are channel-filtered at baseband. Then, the first path goes through standard square-law energy and pilot detection. The second path goes through a $\Sigma\Delta$ modulator, whose raw bitstream output is correlated with the analog signal from the first path for autocorrelation detection.

Chapter 5 – System Implementation

Circuit implementation of the proposed system is currently ongoing. For the initial test-chip, we seek to implement the core of the system as a proof of concept for the subsampling autocorrelation technique. Many of the peripheral and support functions will therefore be generated off-chip. This includes clock generation, clock time-shift generation for autocorrelation, and back-end power measurements for energy detection. At the front-end, the first bandpass filter will also be implemented as an off-chip surface acoustic wave (SAW) filter. The SAW selects the UHF band and attenuates strong out-of-band blockers so they do not saturate the LNA.

In this chapter, we will describe the design of a few of the completed blocks, namely the LNA and the RF bandpass filter.

5.1 - LNA

For our LNA design, we adopted a common-gate (CG) architecture for wideband impedance match to the antenna. Furthermore, the LNA is made fully-differential for good IIP2 performance. The schematic of the LNA is shown below in Figure 5.1.

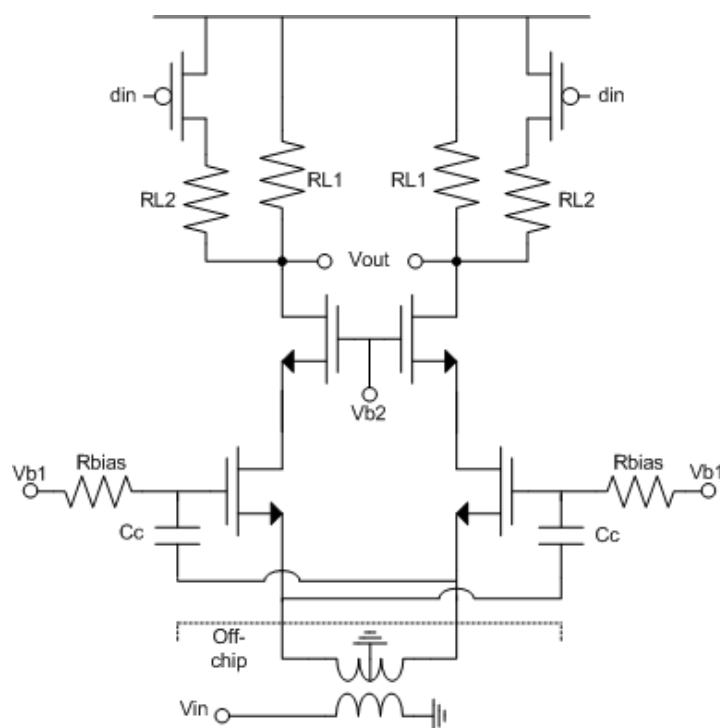


Figure 5.1 - Schematic of LNA.

An off-chip balun converts the incoming single-ended signal to differential. The cross-coupling capacitors act as gain-boosting amplifiers with a gain of 1, which effectively doubles the gm of

the input devices. Resistive loads are used for lower noise, and the cascode devices provide a high output impedance to not degrade the load resistance. The set of switched load resistors $RL2$ have lower resistance than $RL1$ and provide a low gain mode with improved linearity.

The theoretical gain and noise figure of the LNA are:

$$Av = 2 * gm * RL \quad (12)$$

$$F = 1 + \frac{\gamma}{2} + \frac{2}{gm * RL} \quad (13)$$

In practice, the gain and noise figure deviates from Equations (12) and (13) due to, of course, finite transistor output impedance, but also finite cross-coupling capacitance. The capacitors should ideally feedthrough the inverted signal un-changed, but in actuality, the feedthrough ratio is less than 1 due to the capacitive divider formed by Cc and the input device's gate capacitance. The resulting effective Gm , rather than being twice gm , is now:

$$Gm_{eff} = \left(1 + \frac{Cc}{Cc + C_{gate}}\right) * gm \quad (14)$$

Thus, Cc must be made sufficiently large to not significantly degrade performance.

The input impedance of a CG LNA is $1/gm$, or in our case, $1/Gm_{eff}$. The standard antenna impedance of 50Ω therefore limits gm to 20 mS. (The factor of 2 gm-boost is cancelled out by a factor of 2 impedance transformation through the balun). As seen in Equation (13), noise figure depends solely on $gm * RL$. Since gm is limited by the input impedance, and RL is limited by headroom, there is a fundamental lower limit to the noise figure achievable in a CG LNA architecture.

However, Equation (13) assumes a perfectly impedance-matched condition. If we intentionally mismatch the LNA input impedance, Rin , with antenna impedance, the noise figure becomes:

$$F = 1 + \alpha * \frac{\gamma}{2} + \frac{(\alpha + 1)^2}{2 * \alpha} * \frac{2}{gm * RL}, \quad \alpha = Rin/Rs \quad (15)$$

where Rs is the 50Ω antenna impedance. Assuming $gm * RL$ is sufficiently large so that $\gamma/2$ is the dominant term, a smaller Rin (corresponding to a larger gm) results in a smaller noise figure. On the other hand, an impedance mismatch causes more power to be reflected back to the antenna rather than delivered to the LNA. This reflection is measured with $S11$, which is:

$$S11 = \left| \frac{Zin - Rs}{Zin + Rs} \right| \quad (16)$$

A perfect match results in an $S11$ of 0, meaning all power is delivered and none reflected. However, an $S11$ of -10 dB has traditionally been considered acceptable, so some mismatch can be tolerated. Noise figure and $S11$ plotted as a function of Rin is shown below in Figure 5.2.

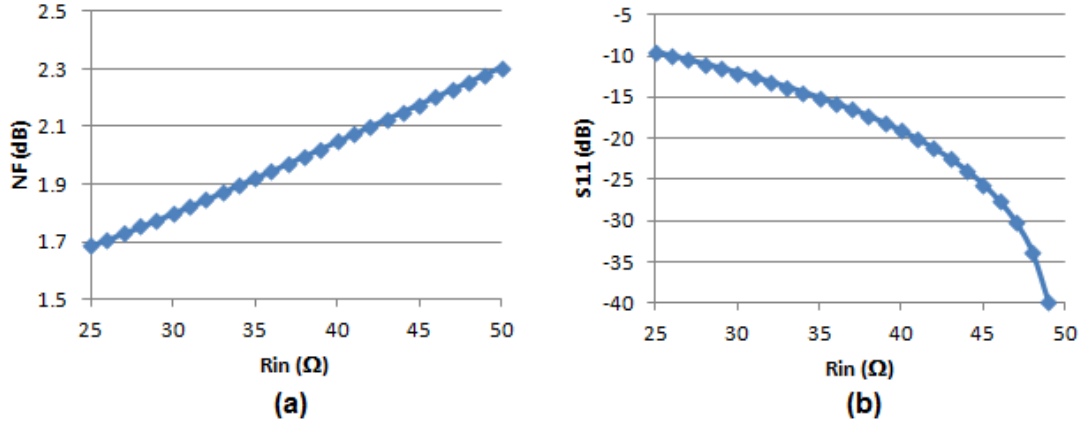


Figure 5.2 - (a) NF, and (b) S11 as a function of R_{in} .

We see that at the cost of spending power to make gm bigger, we can achieve a lower noise figure while still having sufficient matching. Our final design targets a gm of 28 mS and a $gm * RL$ of 9.5. The resulting frequency responses for gain and noise figure are shown below in Figure 5.3.

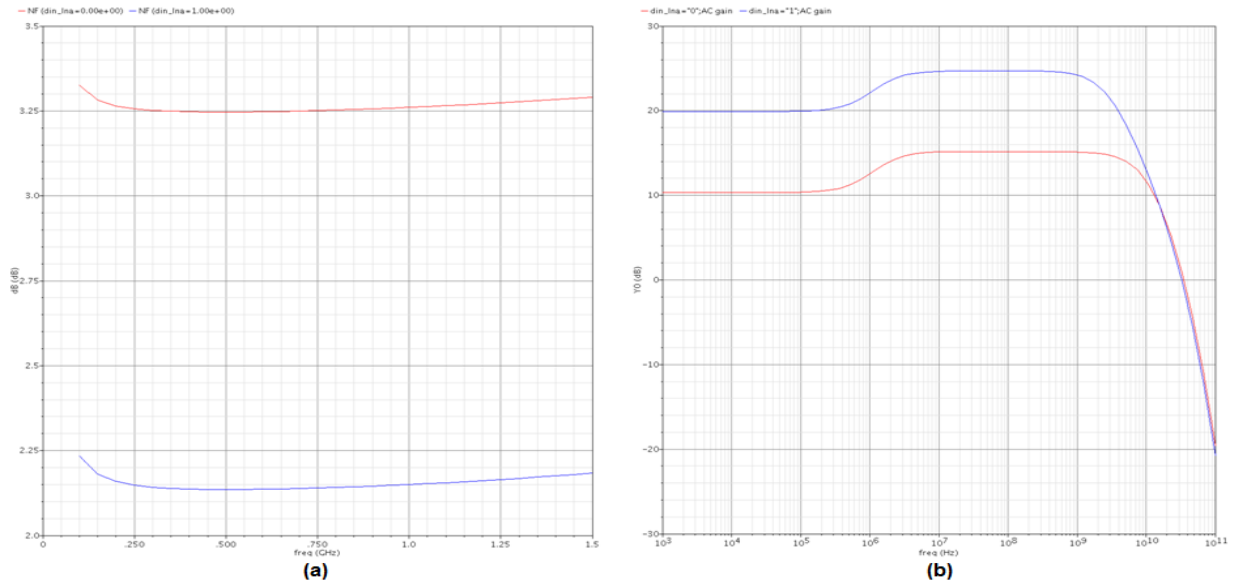


Figure 5.3 – Frequency response of LNA (a) NF, and (b) voltage gain for high- and low-gain modes.

The low frequency gain drop is a result of Cc becoming an open circuit, and the pole position is defined by Cc and R_{bias} . Simulated post-layout gain is 24.5 dB and 15 dB for high and low gain modes. Noise figure is 2.15 dB and 3.25 dB for high and low gain modes, including the ~ 0.35 dB of degradation from layout parasitic resistances. S_{11} is below -15 dB across the entire band for both gain modes. P_{1dB} is about -16 dBm and -6 dBm for high and low gain modes.

5.1.1 Bias

We use a replica bias scheme for the LNA, shown below in Figure 5.4.

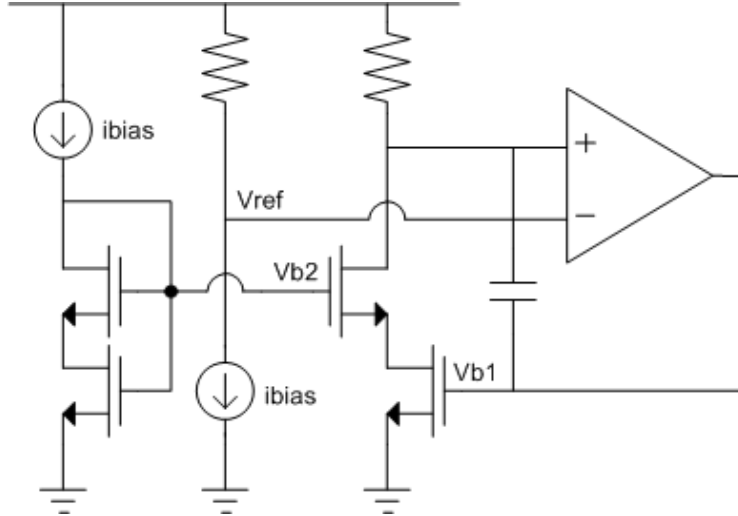


Figure 5.4 - Schematic of LNA bias.

The rightmost branch is a replica of one single-ended branch of the LNA at 1/6 its size. The middle branch generates a reference voltage for the output common mode, and the feedback loop generates bias voltage $Vb1$ to keep the output common mode constant. The leftmost branch generates the cascode bias.

The entire LNA block, including bias circuitry, consumes 6 mW.

5.2 – RF Tracking Filter

We adopt the Gm-C biquad architecture for our RF tracking filter. The biquad structure presents a high impedance to the previous stage LNA and does not have to be impedance matched, and the Gm-C implementation gives the lowest power. The structure of the biquad is shown below in Figure 5.5.

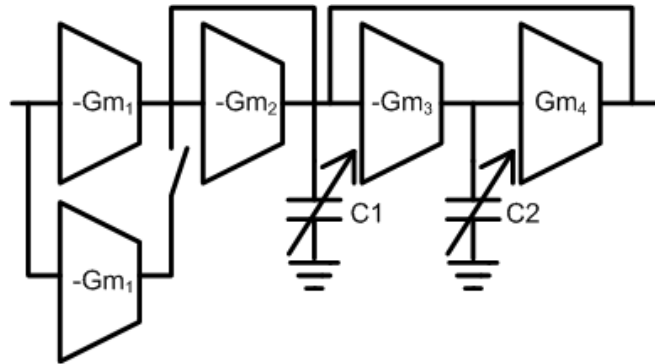


Figure 5.5 - High level schematic of RF tracking filter.

The transfer function of the biquad is given by:

$$H(s) = \frac{s * Gm1 * C2}{s^2 * C1 * C2 + s * Gm2 * C2 + Gm3 * Gm4} \quad (17)$$

From the transfer function, we can derive the following characteristics:

$$Av = Gm1/Gm2 \quad (18)$$

$$\omega_o = \sqrt{\frac{Gm3 * Gm4}{C1 * C2}} \quad (19)$$

$$Q = \sqrt{\frac{Gm3}{Gm2}} * \sqrt{\frac{Gm4}{Gm2}} * \sqrt{\frac{C1}{C2}} \quad (20)$$

$$F|_{\omega_o} \approx 1 + \frac{\gamma}{Av * Gm2 * Rs} \left(1 + \frac{1}{Av} + \frac{Gm4}{Av * Gm2} \left(1 + \frac{C1}{C2} \right) \right) \quad (21)$$

Varying $C1$ and $C2$ changes the center frequency of the bandpass filter, while varying $Gm1$ changes the gain. We implement high and low gain modes by switching an extra $Gm1$ section as shown in Figure 5.5. For lower noise, we choose $Gm4 = Gm2$, so that we are left with $Q = Gm3/Gm2 = C1/C2$.

The diode-connected $Gm2$ effectively acts as a resistance; thus, any finite output resistance on the output node will result in a smaller effective $1/Gm2$ resistance, which translates to a larger effective $Gm2$. This, in turn, degrades both Q and gain. This degradation can be significant since the output node sees the output impedances of 3 Gm-cells in parallel, with $Gm1$ being especially large if high gain is desired. Since there is not enough headroom for cascode devices, we boost the Gm-cell output impedances by using longer channel devices. We also add a Q-boosting negative-gm component to further compensate for the loss.

The feedback loop involving $Gm3$ and $Gm4$ has two poles at the same frequency, and thus needs to be carefully designed for stability. The phase margin of the loop, however, also corresponds with Q , where a higher Q translates to a loop closer to instability. To have a safety margin where there is at least 30° phase margin across all corners and process variations, we are limited to about 14 dB of rejection at the third harmonic frequency. Since more rejection is needed before the sampler, we cascade two Gm-C biquads to form a 4th order bandpass filter.

The schematic of a Gm-cell is shown below in Figure 5.6.

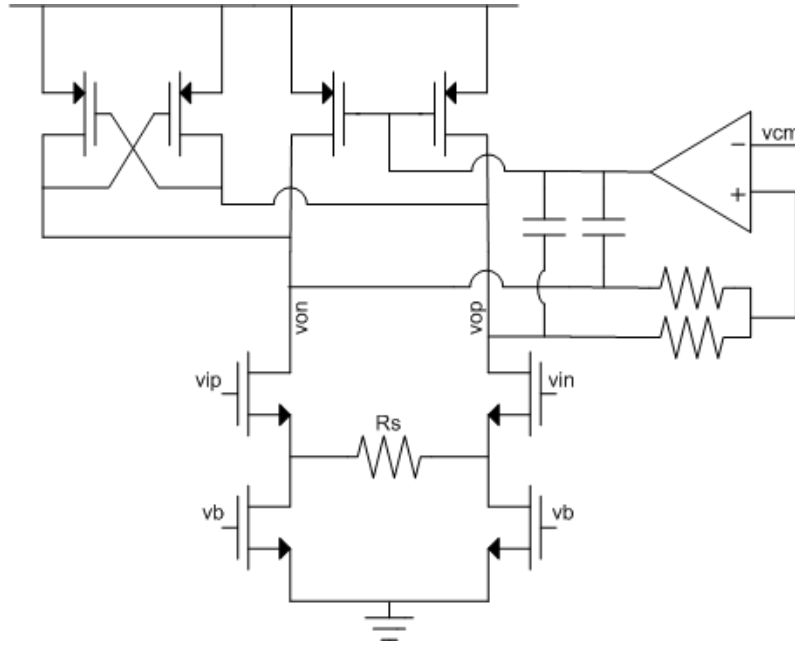


Figure 5.6 - Schematic of a Gm-cell.

The Gm-cells employ a resistively-degenerated common source architecture, with $gm * Rs$ set to 1 for balance between lower power and greater linearity. The output common mode is set by the common mode feedback (CMFB) loop controlling the load PMOS current sources. The additional cross-coupled PMOS pair forms a negative-gm cell that implements the Q-boost. While the CMFB and Q-boost are shown on an individual Gm-cell in Figure 5.6, they are implemented only once on each shared biquad node.

The frequency responses for gain and noise figure for one biquad are shown below in Figure 5.7.

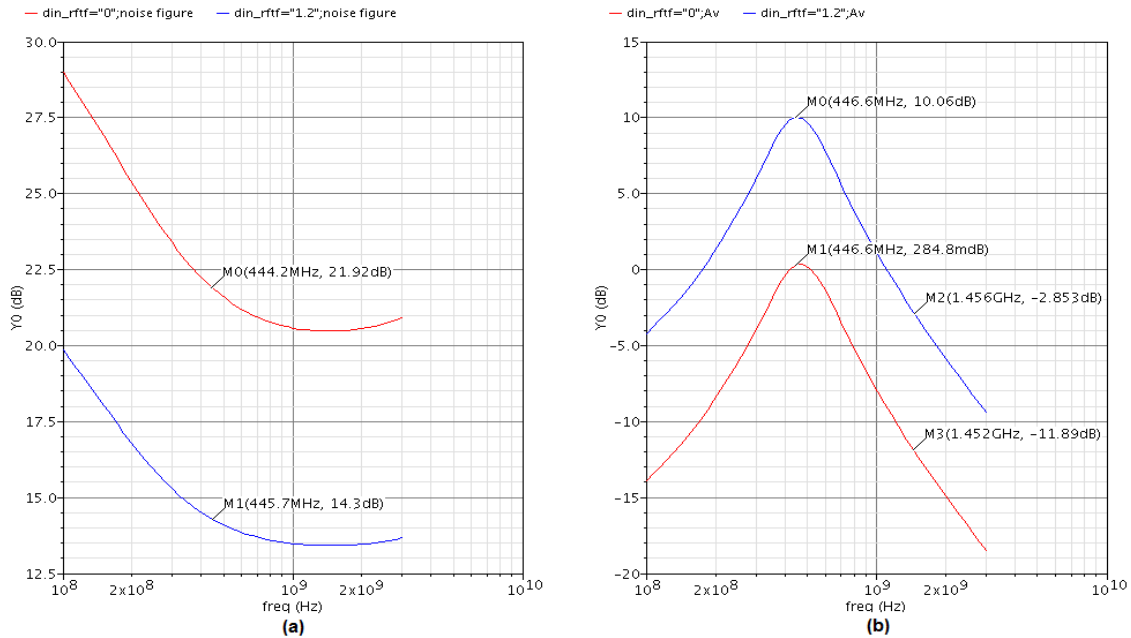


Figure 5.7 - Frequency responses of filter (a) NF, and (b) voltage gain.

The peak gain of each biquad is about 10 dB and 0 dB respectively for high and low gain modes. Rejection at the 3rd harmonic is 12 dB to 14 dB across all gain and frequency settings. The tunable capacitors of each biquad are implemented as binary-coded capacitor banks with 3-bit digital inputs and 8 frequency settings. A schematic of one capacitor bank is shown below in Figure 5.8.

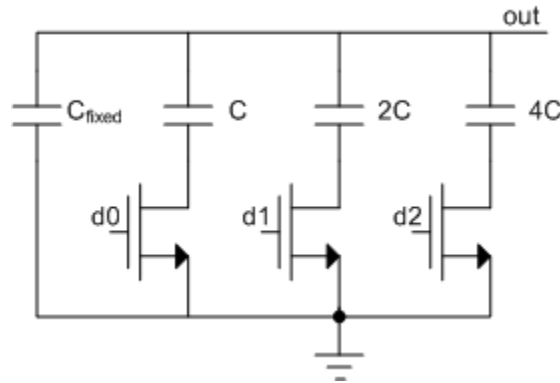


Figure 5.8 - Schematic of one filter capacitor bank.

The frequency response of one biquad at all of its frequency settings is shown below in Figure 5.9.

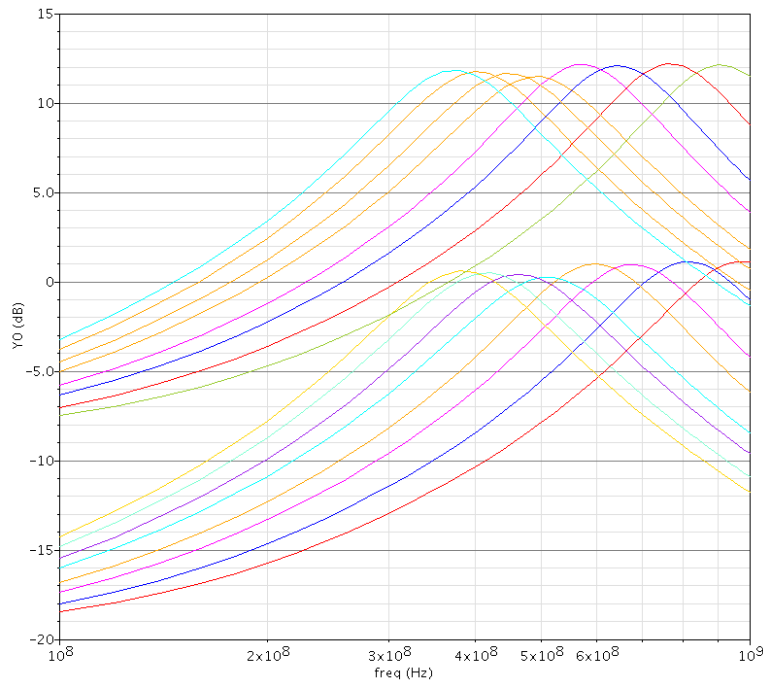


Figure 5.9 - Frequency-tuned response of RF tracking filter.

Since each biquad can contribute additional gain, we can tolerate a higher noise figure from the second biquad. The second biquad is thus designed to be half the size of the first one, consuming roughly half of the first's power at a cost of higher noise. The two cascaded biquads, including bias and peripheral circuitry, consume 9 mW in total.

Chapter 6 – Conclusion and Future Work

In this report, we reviewed the requirements and challenges of performing spectrum sensing in the TV band for mobile applications. We described the existing state of this area of research from both a technical and a policy perspective, and reviewed existing theoretical techniques for performing signal detection. We then proposed a novel physical implementation of a spectrum sensing system by using subsampling autocorrelation in the analog domain.

The most immediate goal of this project is to finish circuit implementation and tape out a test-chip. If the proof of concept is successful, there are many ways this system can be enhanced and expanded in future projects. For example, the off-chip peripheral and support components of this work could be moved on-chip, and doing so would generate a slew of challenges that would need to be addressed. Moving clock generation on-chip would make phase noise a concern, and the implementation of linear, high dynamic-range power detectors is also not trivial. Eliminating the front-end SAW filter would make the system vulnerable to strong out-of-band blockers, and we would need additional innovation in circuit and architecture design and/or detection algorithms to make the system resilient against out-of-band interference. In addition, this work is focused specifically on detection of DTV signals in the UHF band, and takes advantage of known characteristics of ATSC signals. When spectrum sensing eventually expands beyond the TV bands, signal detection in a signal agnostic regime will become necessary and a major area of possible future research.

Chapter 7 – References

- [1] "Connecting America: The National Broadband Plan," Federal Communications Commission, Apr. 01, 2009.
- [2] Julius Genachowski, "America's Mobile Broadband Future," in *International CTIA WIRELESS I.T. & Entertainment*, San Diego, 2009.
- [3] Joseph Mitola and Gerald Q. Maguire, "Cognitive Radio: Making Software Radios More Personal," *IEEE Personal Communications*, vol. 6, no. 4, pp. 13-18, August 1999.
- [4] Danijela Cabric, "Cognitive Radios: System Design Perspective," University of California, Berkeley, Ph.D. Thesis, 2007.
- [5] U.S. FCC, ET Docket 08-260, "Second Report and Order and Memorandum Opinion and Order," 2008.
- [6] U.S. FCC, ET Docket 10-174, "Second Memorandum Opinion and Order," 2010.
- [7] U.S. FCC, OET 08-TR-1005, "Evaluation of the Performance of Prototype TV-Band White Space Devices, Phase II," 2008.
- [8] Advanced Television Systems Committee, "ATSC Digital Television Standard, Part 2 - RF/Transmission System Characteristics (A/53, Part 2:2007)," 2007.
- [9] Stephen J. Shellhammer, "Spectrum Sensing in 802.22," in *Workshop on Cognitive Information Processing*, 2008.
- [10] Advanced Television Systems Committee, "ATSC Recommended Practice: Receiver Performance Guidelines (A/74:2010)," 2010.
- [11] Rahul Tandra, Shridhar Mubaraq Mishra, and Anant Sahai, "What is a Spectrum Hole and What Does it Take to Recognize One?," *Proceedings of the IEEE*, pp. 824-848, May 2009.
- [12] Jongmin Park et al., "A Fully Integrated UHF-Band CMOS Receiver With Multi-Resolution Spectrum Sensing (MRSS) Functionality for IEEE 802.22 Cognitive Radio Applications," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 1, pp. 258-268, January 2009.
- [13] Masaki Kitsunezuka et al., "A 30-MHz–2.4-GHz CMOS Receiver With Integrated RF Filter and Dynamic-Range-Scalable Energy Detector for Cognitive Radio Systems," *IEEE Journal*

- of Solid-State Circuits*, vol. 47, no. 5, pp. 1084-1093, May 2012.
- [14] Mark S. Oude Alink, Eric A.M. Klumperink, Michiel C.M. Soer, Andre B.J. Kokkeler, and Bram Nauta, "A 50MHz-to-1.5GHz Cross-Correlation CMOS Spectrum Analyzer for Cognitive Radio with 89dB SFDR in 1MHz RBW," in *Proc. IEEE Symp. New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, 2010, pp. 1-6.
- [15] Tsung-Han Yu, Chia-Hsiang Yang, Danijela Čabrić, and Dejan Marković, "A 7.4mW 200MS/s Wideband Spectrum Sensing Digital Baseband Processor for Cognitive Radios," in *Proc. IEEE Symposium on VLSI Circuits (VLSIC)*, 2011, pp. 254-255.
- [16] Danijela Cabric, Artem Tkachenko, and Robert W. Brodersen, "Experimental Study of Spectrum Sensing based on Energy Detection and Network Cooperation," in *Proc. First Int'l. Workshop on Technology and Policy for Accessing Spectrum (TAPAS)*, 2006.
- [17] Elad Alon, Valentin Abramzon, Bitá Neẓamfar, and Mark Horowitz, "On-Die Power Supply Noise Measurement Techniques," *IEEE Transactions on Advanced Packaging*, vol. 32, no. 2, pp. 248-259, May 2009.
- [18] Holly Pekau and James W. Haslett, "A 2.4 GHz CMOS Sub-Sampling Mixer with Integrated Filtering," *IEEE Journal of Solid-State Circuits*, vol. 40, no. 11, pp. 2159-2166, November 2005.
- [19] Supisa Lerstaveesin, Manoj Gupta, David Kang, and Bang-Sup Song, "A 48–860 MHz CMOS Low-IF Direct-Conversion DTV Tuner," *IEEE Journal of Solid-State Circuits*, vol. 43, no. 9, pp. 2013-2024, September 2008.
- [20] P.H. Dietz and L.R. Carley, "Analog/Digital Hybrid VLSI Signal Processing Using Single Bit Modulators," in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, 1993, pp. 136-139.