

Identification of Interfering Signals in Software-Defined Radio Applications Using Sparse Signal Reconstruction Techniques

Randy M. Yamada

Thesis submitted to the faculty of the Virginia Polytechnic Institute and State University in partial fulfillment of the requirements for the degree of

Master of Science
in
Electrical Engineering

Lamine Mili (Chair)
Jeffrey H. Reed
Amir I. Zaghloul

April 23, 2013
Arlington, VA

Keywords: software-defined radio, sparse signal reconstruction, interference, decimation, alias, spectrum recovery

Copyright 2013, Randy M. Yamada

Identification of Interfering Signals in Software-Defined Radio Applications Using Sparse Signal Reconstruction Techniques

Randy M. Yamada

ABSTRACT

Software-defined radios have the agility and flexibility to tune performance parameters, allowing them to adapt to environmental changes, adapt to desired modes of operation, and provide varied functionality as needed. Traditional software-defined radios use a combination of conditional processing and software-tuned hardware to enable these features and will critically sample the spectrum to ensure that only the required bandwidth is digitized. While flexible, these systems are still constrained to perform only a single function at a time and digitize a single frequency sub-band at time, possibly limiting the radio's effectiveness.

Radio systems commonly tune hardware manually or use software controls to digitize sub-bands as needed, critically sampling those sub-bands according to the Nyquist criterion. Recent technology advancements have enabled efficient and cost-effective over-sampling of the spectrum, allowing all bandwidths of interest to be captured for processing simultaneously, a process known as band-sampling. Simultaneous access to measurements from all of the frequency sub-bands enables both awareness of the spectrum and seamless operation between radio applications, which is critical to many applications. Further, more information may be obtained for the spectral content of each sub-band from measurements of other sub-bands that could improve performance in applications such as detecting the presence of interference in weak signal measurements.

This thesis presents a new method for confirming the source of detected energy in weak signal measurements by sampling them directly, then estimating their expected effects. First, we

assume that the detected signal is located within the frequency band as measured, and then we assume that the detected signal is, in fact, interference perceived as a result of signal aliasing. By comparing the expected effects to the entire measurement and assuming the power spectral density of the digitized bandwidth is sparse, we demonstrate the capability to identify the true source of the detected energy. We also demonstrate the ability of the method to identify interfering signals not by explicitly sampling them, but rather by measuring the signal aliases that they produce. Finally, we demonstrate that by leveraging techniques developed in the field of Compressed Sensing, the method can recover signal aliases by analyzing less than 25 percent of the total spectrum.

Acknowledgements

I would first like to thank my advisor, Dr. Lamine Mili, for support in developing this thesis through concept, critique, and advice, and for instruction in stochastic systems. I would also like to thank Dr. Jeffrey Reed and Dr. Amir Zaghoul for the instruction and insights provided during my time at Virginia Tech in software-defined radio and radar, as well as for review of this document.

I would like to extend great appreciation to Dr. Allan Steinhardt and David LeDoux for countless technical discussions, discovery, vision, critiques, ideas, and an unwavering approach to the hardest problems in MASINT and signal processing. I have grown substantially as an engineer as a direct result of your interactions and hope to continue to do so. I have no doubt that our exploits will have significant impact on this world.

I would also like to thank my friends and family for support and companionship, and for pushing me continue to better myself. To mom and dad, for the lifetime dedicated to giving me the best of everything. To Roger and Kevin, for the inspiration and awe. And finally, to my wonderful wife Esther, for review of this thesis and being the best thing that ever happened to me.

Contents

| | |
|---|-----|
| Acknowledgements..... | iv |
| List of Figures..... | vii |
| Chapter 1 Introduction | 1 |
| 1.1. Thesis Objective..... | 2 |
| 1.2. Original Contributions..... | 3 |
| 1.3. Thesis Organization..... | 3 |
| Chapter 2 Software-Defined & Cognitive Radio | 6 |
| 2.1. A Brief History of Radio..... | 6 |
| 2.2. Benefits of Software-Defined and Cognitive Radio | 8 |
| 2.3. Software-Defined Radio Architecture..... | 9 |
| 2.4. Data Converters..... | 10 |
| 2.5. Band-Sampling..... | 11 |
| Chapter 3 Simultaneously Sensing Disparate Frequency Bands | 13 |
| 3.1. Capturing the Spectrum..... | 15 |
| 3.1.1. Channelization of Sub-bands | 16 |
| 3.1.2. Decimation..... | 16 |
| 3.1.3. Cascaded Integrator-Comb (CIC) Filtering | 21 |

| | | |
|------------|--|----|
| 3.2. | Interference Mitigation..... | 23 |
| Chapter 4 | Sparse Signal Reconstruction | 25 |
| 4.1. | An Underdetermined Linear System..... | 25 |
| 4.2. | Similar Developments in the Field of Compressed Sensing | 26 |
| 4.3. | Methods Developed in the Field of Compressed Sensing..... | 27 |
| 4.4. | Open Source Solvers for the Spectrum Reconstruction | 30 |
| Chapter 5 | Analysis and Simulation | 31 |
| 5.1. | Band Sampling | 32 |
| 5.2. | Spectral Analysis..... | 35 |
| 5.3. | Sparse Signal Reconstruction..... | 37 |
| 5.3.1. | Energy from a Desired Signal Aliases Into Other Search Bands..... | 39 |
| 5.3.2. | Energy From an Undesired Interfering Signal Aliases Into Any Search Band | 46 |
| Chapter 6 | Conclusions and Future Work..... | 52 |
| References | | 55 |

List of Figures

| | |
|--|----|
| Figure 1: Block diagrams of radio frequency (RF) band-sampling (top) and analog RF signal selection and conditioning control (bottom). | 10 |
| Figure 2: An example of several frequency bands of interest that are neither adjacent or co-located, but are contained within an encompassing bandwidth. | 14 |
| Figure 3: Train of Dirac delta functions modeling the sampling of a signal. | 17 |
| Figure 4: Train of Dirac delta functions in the Fourier domain modeling the sampling of a signal. | 18 |
| Figure 5: Sampling modeled as convolution of the spectrum of a signal with an infinite train of Dirac delta functions. | 18 |
| Figure 6: Overlapping spectra causing aliasing. | 19 |
| Figure 7: The inverse relationship between sampling period and bandwidth in the time and frequency domains. | 20 |
| Figure 8: The spectral response of an example of a CIC decimator filter reducing the incoming data rate by a factor of 448 in two stages. | 23 |
| Figure 9: Example of a 107Hz narrowband signal digitized using a high sampling rate. | 32 |
| Figure 10: Single-sided spectra of the 106Hz narrowband signal after downconversion for analysis of the 102Hz–111Hz sub-band (solid blue), the 50Hz–60Hz sub-band (red), and the 520–530Hz sub-band (green) showing application of a notional low pass filter (dashed blue). | 33 |

Figure 11: Single-sided spectrum of the 102Hz–111Hz sub-band after decimation (frequency axis corrected to show relationship to the original spectrum) showing the true 107Hz signal. 34

Figure 12: Single-sided spectrum of the 50Hz–60Hz sub-band after decimation (frequency axis corrected to show relationship to the original spectrum) showing the aliased 107Hz signal. 34

Figure 13: Single-sided spectrum of the 520Hz–530Hz sub-band after decimation (frequency axis corrected to show relationship to the original spectrum). 35

Figure 14: A graphical representation of the underdetermined linear set of equations describing the true power spectral density (PSD) of R frequency bins with n measured responses resulting from the parallel decimation of R frequency bins. 36

Figure 15: Example of R frequency responses (in dB) from a set of 16 filters prior to downconversion and decimation, representative of the matrix A from Equation 3. 40

Figure 16: Coefficients of the true vector x (from Equation 3) representing the true power spectral densities at the frequencies corresponding to the filter coefficient locations. 40

Figure 17: Coefficients of the true vector x (from Equation 3) recovered using the Basis Pursuit Denoise algorithm implemented by SPGL1 [48] showing accurate confirmation that energy measured within a search band is due to a signal from that band despite an underdetermined system. 41

Figure 18: Results of a Monte Carlo simulation showing the percentage correct of coefficients of the true vector x (from Equation 3) recovered using the Basis Pursuit Denoise algorithm implemented by SPGL1 [48] based on the SNR of the signal in question and the ratio of monitored to total bandwidth when the signal in question is within a desired search band. 42

Figure 19: Results of a Monte Carlo simulation showing the percentage correct of coefficients of the true vector x (from Equation 3) recovered using the Stage-Wise

Orthogonal Matching Pursuit algorithm [45] implemented by SparseLab [49] based on the SNR of the signal in question and the ratio of monitored to total bandwidth when the signal in question is within a desired search band. 44

Figure 20: Results of a Monte Carlo simulation showing the percentage correct of coefficients of the true vector x (from Equation 3) recovered using the Matching Pursuit algorithm [38] implemented by SparseLab [49] based on the SNR of the signal in question and the ratio of monitored to total bandwidth when the signal in question is within a desired search band. 45

Figure 21: Example of R frequency responses (in dB) from a set of 16 filters prior to downconversion and decimation, representative of the matrix A from Equation 3. 46

Figure 22: Coefficients of the true vector x (from Equation 3) recovered using the Basis Pursuit Denoise algorithm [48] showing accurate recovery of an out-of-band interfering signal despite an underdetermined system. 47

Figure 23: Results of a Monte Carlo simulation showing the percentage correct of coefficients of the true vector x (from Equation 3) recovered using the Basis Pursuit Denoise algorithm implemented by SPGL1 [48] based on the SNR of the signal in question and the ratio of monitored to total bandwidth when the interfering signal is outside of any desired search bands. 48

Figure 24: Details of results of a Monte Carlo simulation using the Basis Pursuit Denoise algorithm implemented by SPGL1 [48] showing percent correct of recovery. 49

Figure 25: Results of a Monte Carlo simulation showing the percentage correct of coefficients of the true vector x (from Equation 3) recovered using the Stage-Wise Orthogonal Matching Pursuit algorithm [45] implemented by SparseLab [49] based on the SNR of the signal in question and the ratio of monitored to total bandwidth when the interfering signal is outside of any desired search bands. 50

Figure 26: Results of a Monte Carlo simulation showing the percentage correct of coefficients of the true vector x (from Equation 3) recovered using the Matching Pursuit

algorithm [38] implemented by SparseLab [49] based on the SNR of the signal in question and the ratio of monitored to total bandwidth when the interfering signal is outside of any desired search bands.

51

Chapter 1

Introduction

Radio frequency spectrum sensing technology has evolved from traditional analog means to complex digital analyses, and continues to receive attention as spectral awareness becomes increasingly important. From personal communications services to sensing (including weather radars, automotive radars, and imaging), the demand for spectrum bandwidth is ever-increasing [1]. The increased use of bandwidth requires that systems remain operable even in the noisiest of environments. A challenge in meeting demand under these conditions is to sense weak signals buried in noise, which allows us to increase the radio's physical operating range and save power in a number of applications. This challenge has necessitated the use of stringent analog filtering prior to digitization, ultimately limiting the overall bandwidth of the radio and therefore, jeopardizing its potential to be multi-purpose. A recent trend fueled by technology advancement in digitization for spectrum use permits significant over-sampling of the desired bandwidth and thereby, enables the once-analog filters to be realized digitally and allows complex processing including adaptive techniques. This over-sampling method does present drawbacks since higher sampling rates require more computing power to implement digital filtering, which can often

limit performance of radios. Less stringent filtering that uses less computing power can, for example, introduce significant aliasing when the incoming data stream is decimated. Aliasing of these ever-present strong signals can corrupt on-going measurements for weak signals in spite of digital filtering.

1.1. Thesis Objective

The main objective of this thesis is to develop an approach to simultaneously monitoring disparate frequency bands for weak signals using a band-sampling approach rather than the traditional critical sampling approach. The new approach that we propose applies sparse signal reconstruction techniques to band-sampled signals in the Fourier domain that have been decimated (as is often done with weak signals measurements). This work assesses the capability of determining if weak measured energy is in fact due to a signal outside of the frequency band of interest and is interfering with the measurement, or if the measured energy is truly the result of a signal within the frequency band of interest. System performance is based on the ratio of the observed bandwidth ratio to the decimation rate and the signal-to-noise ratio (SNR) of the interfering signal. The research presented in this thesis shows that a sensor being used to simultaneously monitor disparate frequency bands for weak signals can identify if detected energy is in fact due to a weak signal or strong interfering signal outside the observation band, effectively enhancing sensor performance and selectivity.

1.2. Original Contributions

Cognitive radios are generally designed to adapt (often in the analog or physical sense) to a desired center frequency, modulation structure, and communications protocol. This thesis expands on that characteristic to address multiple signals that may be of interest across disjoint frequency bands simultaneously [2]. Contrary to the critical sampling criteria imposed throughout the design of software radio [2], a band-sampling approach to spectrum sensing is proposed that digitizes far more bandwidth than required based on the Nyquist criterion. We consider a situation where monitored sub-bands may be both disparate (perhaps at opposite ends of the digitized bandwidth) and a small percentage of the overall bandwidth has been digitized (perhaps 25% or less), as described in [3]. Furthermore, we apply sparse signal reconstruction techniques based on approximating underdetermined systems of linear equations that have been developed for Compressed Sensing, to the identification of interfering, aliasing signals while monitoring multiple sub-bands. Finally, we apply and characterize these techniques in the Fourier domain to reconstruct information lost due to decimation. Although not developed in this thesis, the described techniques can also be extended and formulated to monitor many frequency sub-bands using fewer carefully chosen sub-bands by explicitly knowing the decimator response characteristics.

1.3. Thesis Organization

This chapter has introduced the motivation and concepts supporting the research described in this thesis. The increasing demand for spectrum has been briefly discussed with particular attention to weak signals existing among noise and large interfering signals. The main

objective of this thesis has also been described, outlining a construct within which interference will be identified. Finally, original contributions including novel approaches and algorithms to identify the interfering signals are described.

Chapter 2 presents a brief overview of software-defined & cognitive radios, with particular attention to the system hardware architecture. Chapter 2 also introduces data converters and identifies them as a major performance-defining component of the software-defined radio. Finally, Chapter 2 introduces the concept of band-sampling, which may become more prevalent as data converter technology advances.

Chapter 3 introduces the motivation behind simultaneously sensing disparate frequency bands. Chapter 3 also reviews approaches, benefits, and complexities associated with capturing the spectrum, specifically channelizing and decimating frequency sub-bands for analysis. Chapter 3 describes the design of a cascaded integrator comb (CIC) filter, which is commonly used in software-defined radio to mitigate interference, including performance limitations.

Chapter 4 introduces sparse signal reconstruction (i.e. Compressed Sensing), a research field that includes significant development towards estimating underdetermined linear systems of equations assuming sparse solutions. Chapter 4 also describes the parallels between the research described in this thesis and the developments in the field of Compressed Sensing. Finally, Chapter 4 identifies existing algorithms that could be applied to the construct formulated throughout this thesis, and reviews open source solvers and their instantiations of the described algorithms.

Chapter 5 describes detailed analysis and simulation of the concepts described throughout this thesis. Further, Chapter 5 applies the sparse signal reconstruction techniques to spectrum estimation for the purposes of interference identification.

Chapter 6 provides conclusions of the work performed and described throughout this thesis, and proposes direction for future research building on the ideas described herein.

Chapter 2

Software-Defined & Cognitive Radio

2.1. A Brief History of Radio

Conveying information wirelessly through the propagation of electromagnetic radiation, a significant feature of today's world, is based on numerous experiments and discoveries dating back to the early 19th century. James Clerk Maxwell first famously published 20 equations describing the behavior of electromagnetic radiation using fields and potentials in 1865 [4]. Subsequently, experimental physicist David Hughes demonstrated a working radio communication system in 1879 through experimentation [5], and physicist Heinrich Hertz conclusively proved Maxwell's theories through experiments performed in 1888 of propagating energy to and from spark gaps [6]. Building on these developments, numerous systems designed to embed information in propagating electromagnetic radiation have been proposed, notably by Guglielmo Marconi in 1896 [7] and Nikola Tesla in 1897 [8].

Since the proposals of Tesla and Marconi, research and development of electronic devices beginning with the vacuum tube radio frequency detector in 1906 [9] and the transistor in 1926 [10], enabled amplification of signals received in the form of electromagnetic radiation, thereby increasing radio operating ranges and information densities. Further, numerous methods to predictably affect electromagnetic radiation have been discovered and demonstrated to convey information, including combinations of varying the amplitude, frequency, and phase of the generated electromagnetic wave. Today, these combinations, or protocols, are meticulously defined, continuously refined, and maintained by organizations including the Institute of Electrical and Electronics Engineers (IEEE) and the European Telecommunications Standards Institute (ETSI).

Until the late 20th century, all radios were built for applications such as broadcast or two-way communication, and used hardware specifically designed and built to interface with very few protocols, if not one. Many of these radios can be mechanically tuned to alter parameters such as frequency, but were largely limited to a single function. The accelerating advances in computing during this time enabled the concept of digitizing signals and applying protocols in software to allow using different wireless interface protocols with the same hardware. Beginning in the late-1980s, the United States government funded the development of the SpeakEasy transceiver, one of the first software-defined radios, which was intended to allow interoperability between different communication protocols used by different branches of the armed forces [11]. These software-defined radios not only allow interoperability between existing radios, but enable potential forward compatibility with future protocols or protocol revisions. Further developing software-defined radios with parameters controlled by human input has led to the introduction of cognitive radio, which uses computer-controlled decisions based on measurements and

algorithmic learning to tune the software-defined radio for optimum performance [12]. Developments and recent research in cognitive radio are highlighted in [13].

2.2. Benefits of Software-Defined and Cognitive Radio

The communications industry and community has shown great interest in software-defined and cognitive radio [12] because of its potential for efficient spectrum use, flexible radio architectures, and forwards and backwards protocol compatibility [14]. Efficient spectrum use is often enabled via dynamic spectrum access, wherein uncoordinated point-to-point communications devices can make use of a wide variety of frequency bands that are traditionally unavailable (e.g. those bands that are licensed to other entities), but are currently unoccupied [15]. This requires that radio architectures have a keener ability to sense the surrounding environment, instant and autonomous control over configuration parameters, and powerful decision logistics for more robust use [15]. In turn, this flexibility necessitates advanced signal processing techniques, which are easily implemented in the digital domain and are inherently reconfigurable as fast as the code-bearing firmware can be overwritten. This code-bearing firmware can carry protocol updates and the latest in machine-learning techniques and algorithms to optimize performance. These key benefits to radio communications have been a driving force behind current research into software-defined and cognitive radio as described in [13].

2.3. Software-Defined Radio Architecture

The ideal software-defined radio has full digital access to the entire desired spectrum [2], typically requiring an antenna to transduce an electromagnetic wave into electrical potential difference, analog radio frequency (RF) selection and signal conditioning, an analog-to-digital converter (ADC) to quantize the analog electrical potential difference to digital values, and digital signal processing (DSP) to execute the desired computations. Access to the entire spectrum can be achieved either through RF band-sampling [16] or through complete control of flexible analog RF signal selection and conditioning [2] (as simplified in Figure 1). A key difference between these approaches is that RF band-sampling allows for simultaneous access to all of the desired frequency bands by digitizing the entire bandwidth of interest, whereas flexible analog RF signal selection and conditioning, and thus requires some time to tune to a desired band using digital control of analog components.

While a short time to tune from band to band may be acceptable in some software-defined radio applications, in other applications such as monitoring different communications bands for activity or sensing different types of emissions, simultaneous access to all of the radio bands—also referred to as instantaneous bandwidth—is necessary. Requirements for a wide instantaneous bandwidth have a greater effect on the specifications of the ADC (more generally, the data converter) necessary to achieve the instantaneous bandwidth than they have on any other component in the software-defined radio.

The objective of this thesis is to develop a method of identifying interfering signals in software-defined radios using the band-sampling method, where the interfering signals are the result of aliasing due to signal decimation. This is largely impacted by the specifications of the

data converter, rather than the analog preselect or the signal processing hardware. As such, we do not focus on the developments of the analog preselect or the signal processing hardware

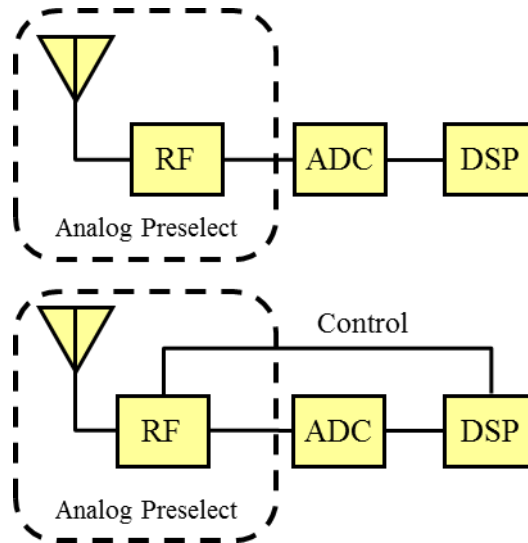


Figure 1: Block diagrams of radio frequency (RF) band-sampling (top) and analog RF signal selection and conditioning control (bottom).

2.4. Data Converters

At the heart of the cognitive radio, as defined by Mitola [13], is the data converter, which transitions energy between the analog and digital domains and enables complex digital signal processing, a requirement for realizing the benefits of a cognitive radio. The data converter plays a vital role in efficient spectrum use by defining limiting characteristics of spectrum sensing [2]. The latter involves measuring the electromagnetic conditions of the environment surrounding the user, which is essential prior to making informed decisions regarding transmitter parameters, including frequency, power, modulation, and bandwidth settings [17]. Spectrum sensing is described further in Chapter 3.

Data converters greatly impact the overall power consumption, dynamic range, bandwidth, and total cost of a radio [2]. Unlike the analog radio frequency (RF) front ends and digital signal processing back ends, data converters are by definition, a single transition point from the analog to the digital domain. While analog RF front ends and digital signal processing back ends can cascade multiple components with less cost, less complexity, and worse performance and realize overall acceptable performance, data converters cannot be efficiently cascaded to improve key performance parameters including resolution, sample rate, and dynamic range without complex and painstaking synchronization, if at all. These key performance parameters dictate the part of the spectrum that is selected and the fidelity to which it is digitized, which are critical to radio performance and the ability to sense the electromagnetic spectrum.

2.5. Band-Sampling

Recent advances in technology have brought the commercialization and cost-effective production of data converters that exceed 3GSps [16]. These data converters enable the software-defined radio to digitize a wide bandwidth and potentially decrease the use of analog signal conditioning and downconversion prior to digitization. This approach to over-sampling a signal and performing the frequency downconversion and filtering digitally is referred to as band-sampling. In contrast to critical sampling, where the sampling rate is matched to capture no more than the total information bandwidth of the expected signal (known as the Nyquist criterion), band-sampling digitizes (generally significantly) more bandwidth than the maximum information bandwidth of the signal.

Band-sampling reduces the analog complexity of the radio by relaxing performance requirements of filters, which prompts the need for a final downconversion to baseband and

tunable analog components, thereby likely reducing its overall size. Furthermore, if the radio makes use of more than one frequency band in relative proximity, band-sampling can enable the radio to capture the entire bandwidth and perform digital sub-band channelization, allowing the sub-bands to be used simultaneously. Despite the increased cost in size, weight, and power of these band-sampling data converters as compared to their critical-sampling counterparts, band-sampling data converters offer channel flexibility not attainable through analog means.

Chapter 3

Simultaneously Sensing

Disparate Frequency Bands

The rise in use of personal electronic devices based on wireless machine-to-machine communications (i.e., [18] and [19]) has increased the demand for electromagnetic spectrum allocation and use, specifically in the bands suited for radio communications. Personal electronic device manufacturers take advantage of open source communication protocol standards such as Wi-Fi (IEEE 802.11), Bluetooth (IEEE 802.15.1), and ZigBee (IEEE 802.15.4) and the industrial, scientific, and medical (ISM) radio bands, which are various frequency bands with different bandwidths across the radio spectrum for public use as designated by the International Telecommunication Union's Radiocommunication Sector (ITU-R) [20]. In addition to machine-to-machine communications, there is a rise in interest in the simultaneous measurement and estimation of point-to-point or broadcast communications signals and signals that are comparatively weak, such as emergency beacons. For example, a communications unit may communicate long range using the 900MHz ISM band, communicate locally using the

2.4GHz ISM band, and be required to monitor 1.6GHz for weak radio astronomy transmissions. These disparate frequency bands of interest exist inside of an encompassing bandwidth as shown in Figure 2, and have a comparatively small total monitored bandwidth as compared to the encompassing bandwidth. Furthermore, spectrum sensing and cognitive radio architectures can provide value to non-communications applications including radar and passive electromagnetic sensing, perhaps eliminating the divide between radio frequency equipment for communications and non-communications uses.




Figure 2: An example of several frequency bands of interest that are neither adjacent or co-located, but are contained within an encompassing bandwidth.

The agility to measure and estimate signals from relatively small, possibly dynamic, and generally dispersed frequency bands, requires simultaneous and accurate digitization of an encompassing, larger search bandwidth. Fortunately, as the encompassing bandwidth is generally static, a fixed analog RF front end to pre-select and downconvert the entire encompassing bandwidth is sufficient. It eliminates the need for tunable, un-scalable analog components in a software-defined radio. In general, fixed-frequency analog components associated with radio frequency downconversion have lower cost and power consumption than their digitally or mechanically tunable counterparts [21]. Thus, for the construct wherein

disparate frequency bands must be monitored simultaneously, we postulate that digitizing the encompassing bandwidth to capture all of the frequency bands of interest simultaneously is a viable approach.

3.1. Capturing the Spectrum

There are many software-defined radio architectures that rely on a tunable RF front end for spectrum channel pre-selection prior to digitization [22]. This is certainly a cost-effective approach to sensing disparate frequency bands, but this approach requires mechanical switching between frequency bands and, therefore, cannot monitor or make use of them simultaneously. Specifically, the mechanically tunable RF channel pre-selection approach reduces the instantaneous bandwidth of the radio and eliminates the capability of monitoring or making use of the frequency bands simultaneously.

If a given encompassing frequency bandwidth is digitized to simultaneously select disparate frequency bands, some digitized bandwidth that is not of interest will be selected, and thus the sample rate used is higher than necessary, as dictated by the Nyquist criterion [2]. Therefore, some fraction of the digitized samples can be discarded while still fully representing the information contained within each of the bandwidths of interest for band-limited signals [23]. While this method reduces the computational processing burden by decreasing the incoming data rate, it also creates spectral ambiguity and introduces image noise into the frequency band of interest in an effect known as aliasing [2]. We investigate the process of decimation and the relationship to aliasing in Chapter 5.

3

In many applications, making use of several different frequency bands requires that they be processed individually and with minimal noise from outside sources. Furthermore, if these bands are being used simultaneously, then the optimum architecture to be implemented is to process the frequency bands in parallel. As such, it is necessary that the incoming data is channelized, that is, the incoming data is prepared for further processing using specifications of each particular sub-band. For example, digital signal processing functions are often formulated to operate at baseband; therefore, the incoming data must be digitally downconverted to shift the sub-band of interest closer to baseband for analysis and processing [2]. This is often of concern in high data rate applications because every multiplication that must be performed at the full data rate significantly burdens the overall processing requirements.

3

The signal processing technique known as decimation is commonly applied in software radios to effectively reduce the incoming data rate without loss of fidelity of the information contained therein. Typically, the incoming signal is filtered prior to reducing the sample rate. Decimation avoids the negative impacts of noise and interfering signals that might reside at image frequencies when reducing the sampling rate—and therefore, information bandwidth—of an incoming signal [2]. Sampling theory can illustrate the concept of image frequencies and aliasing and their relationship to reducing the data sample rate. Discrete sampling of data can be modeled ideally as the multiplication of a signal with an infinite train of Dirac delta functions spaced equally apart by the sampling period (T_s), as illustrated in Figure 3 [24]. Formally, we have

$$x_{\delta}(t) = \sum_{n=-\infty}^{\infty} \delta(t - nT_s) . \quad (1)$$




Figure 3: Train of Dirac delta functions modeling the sampling of a signal.

The effect of sampling on the spectrum of the signal can be seen by observing the Fourier transform of the train of Dirac delta functions, which is simply a train of Dirac delta functions spaced equally apart by the inverse of the sampling period (f_s), as illustrated in Figure 4 [24].

Formally, we have

$$X_{\delta}(f) = \frac{1}{T_s} * \sum_{n=-\infty}^{\infty} \delta(t - nf_s) . \quad (2)$$

Since multiplication by the pulse train in the time domain amounts to applying convolution in the Fourier domain, sampling the signal in the frequency domain translates into having the baseband spectrum to be infinitely repeated with periodicity equal to the inverse of the sampling period as depicted in Figure 5.




Figure 4: Train of Dirac delta functions in the Fourier domain modeling the sampling of a signal.




Figure 5: Sampling modeled as convolution of the spectrum of a signal with an infinite train of Dirac delta functions.

The repeated spectrum shown in Figure 5 is not of serious concern because a filter can be constructed to retain a single copy of the spectrum while filtering out the other copies; however, as shown in Figure 6, if the bandwidth of the signal is too wide or the sample rate is too low,

then the repeated copies of the spectrum will overlap, causing aliasing of the energy and corrupting the measurement of the spectrum. Furthermore, because the aliasing is due to a lack of information needed to discern the origin of the energy contained in that frequency bin, filtering cannot separate the overlapping spectra.




Figure 6: Overlapping spectra causing aliasing.

Although the spectrum is infinitely replicated, it should be noted that the desired information is only contained within bandwidth equal to half of the sampling rate f_s . Furthermore, the inverse nature of the relationship between the sampling period and frequency indicates that, as the time between samples increases, the distance in frequency between repeated spectra decreases. As shown in Figure 7, reducing the sampling rate in the time domain (by, for example, discarding every other sample) has the negative effect of reducing the distance in frequency between replicated spectra.




Figure 7: The inverse relationship between sampling period and bandwidth in the time and frequency domains.

The desired course of action then is to reduce the distortions introduced by aliased energy on the measured bandwidth of interest through filtering; however, digital filtering can only attenuate energy from these image frequencies, not completely eliminate it. Moreover, the quality of the filtering is directly proportional to the computing time required to perform it, while the amount of time available to execute filtering is often a function of the incoming data rate. There are numerous designs resulting from engineering trades that facilitate an effective balance between decimation rates and filtering complexities. These designs are often viable if the first component specified in the design is the data converter while the latter is critically chosen to have the minimum sampling rate necessary for the application, thereby easing many other constraints [2]; however, the engineering trades are more challenging for the disparate frequency

situation described throughout this thesis since the data converter is operating at a rate higher than is required for each frequency sub-band, but is critically specified to capture the encompassing bandwidth as a whole. Furthermore, this challenge may be compounded by relatively large decimation rates, where the repeated spectra in the Fourier domain are shifted closer together by a large fraction. This shifting necessitates significant filtering prior to the decimation, which is often not possible at high data rates.

There has been significant research on classes of digital filters that are optimized for embedded computing technologies and capable of filtering in high data rate reduction situations, including polyphase filtering [25] and cascaded integrated-comb (CIC) filtering [26], described further in Section 3.1.3. In certain applications, such as sensing for weak signals, this suppression of noise is insufficient and can be further complicated with high incoming data rates. Furthermore, these designs do not sufficiently suppress “ghost” signals introduced through aliasing due to the power of the signal of interest, at least not at an acceptable cost of computation.

3 ~~Channelization~~

Digital channelization in concert with decimation is well explored. One method that has many desirable properties when implemented within software-defined radios is the Cascaded Integrator Comb (CIC) filter [26]. The CIC filter cascades the data decimation into stages allowing progressively more complex—and closer to ideal—finite impulse response (FIR) filtering while gradually reducing the data rate. This also allows superior suppression of interfering energy from out-of-band frequencies [26] that are possibly due to noise, intentional jamming signals, or unintentional communications traffic. Further, the CIC filter can be

implemented using only addition and subtraction, realizing computational savings as compared to band-pass FIR filtering, especially when implemented in embedded processing hardware.

Despite the performance and computational savings achieved by the CIC filter architecture, there are still situations wherein the computational burden is not advisable and the interference suppression performance is not adequate, such as with real-time sensing for weak signals. For example, the spectral response of a CIC filter is shown in Figure 8, which consists of four cascaded stages of integrators and combs and has an overall decimation factor of 448 (first decimated by a factor of eight then by a factor of 56). Theoretically, the expected maximum aliased energy is suppressed by $138dB$; however, this assumes adequate filtering at the full rate as well as after the initial decimation by eight. When dealing with very high initial sample rates to capture wide bandwidths and simultaneous parallel processing of several disparate frequency bands, it is likely that the computational resources for high-quality filtering prior to decimation and after the initial decimation stages are unavailable and energy present at image frequencies will corrupt measurements of frequency bands of interest through aliasing.

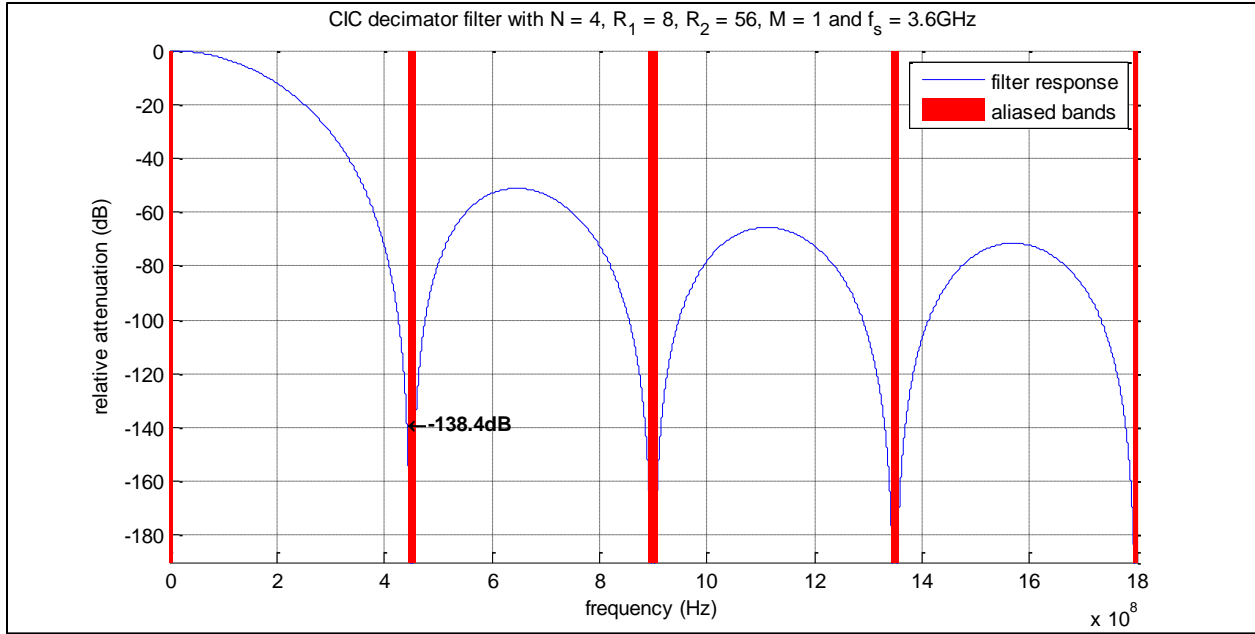


Figure 8: The spectral response of an example of a CIC decimator filter reducing the incoming data rate by a factor of 448 in two stages.

3.2. Interference Mitigation

Two types of receiver architectures are typically considered for application to cognitive radio: direct conversion and RF band-pass sampling receiver architectures [22]. Both receiver architectures require tunable hardware such as RF filters or local oscillators to maintain frequency agility while removing unwanted signals that may impact system performance. One approach to increasing radio flexibility is to choose the fastest possible data converter sample rate, thereby maximizing the digitized bandwidth of the receiver. For certain applications, such as personal communications services, the technology required for this approach is commercially available [16]. Digital filters are used to channelize the desired or allocated bandwidths with high-performance image-suppressing data decimation. Furthermore, this approach enables simultaneous analyzing and monitoring of disparate frequency bands with diversity allowing for

the implementation of new interference mitigation techniques for cognitive and software-defined radio.

Chapter 4

Sparse Signal Reconstruction

4.1. An Underdetermined Linear System

The decreasing costs and increasing performance of data converters has accelerated the generation of sensor data worldwide. In turn, the limiting factor of sensors has shifted from the data converters to the data communication, processing, and storage [27]. Decimation is commonly applied in software radios, as described in Section 3.1.2, to effectively reduce the incoming data rate without loss of fidelity of the information contained therein. Decimation is even more critical in software-defined radio systems designed to prosecute disjoint bandwidths across a large bandwidth. This process can be modeled as an underdetermined system of linear equations given by

$$y = Ax , \tag{3}$$

where A is a known $n * R$ sensing matrix representing a set of n parallel spectral filter responses (or Fourier coefficients) at R periodic frequency bins, x is a set of R true frequency power spectral densities, and y is the n set of measured power spectral densities.

Numerical techniques capable of solving the underdetermined linear equations subject to constraints have been explored in the field of Compressed Sensing and shown to be practical if x is sufficiently sparse [28], [29]. It is by these methods that the spectral leakage from a sparse (in the frequency domain sense) signal, namely a narrowband interfering signal, can be identified and therefore mitigated.

4.2. Similar Developments in the Field of Compressed Sensing

In this thesis, we make the following conjecture: if the desired bandwidth of electromagnetic spectrum is sufficiently sparse and sufficiently sensed, then measured energy can be determined to be (i.e. automatically classified as either) (i) desired, in-band signals or (ii) undesired, strong, out-of-band, interfering, aliased signals. Compressed Sensing has strong roots in data compression, particularly in striving to fully represent images and other data with less than the original data set (called sparse approximation), as is commonly done with the JPEG and MP3 compression standards [30]. These sparse approximation methods for compression of existing data—traditionally for communication or storage—have been applied to image capture systems in an attempt to extract the same meaningful data by sampling in a sparser basis but fully reconstructing the information afterwards [31]–[33]. We apply these techniques for reconstructing underdetermined systems under sparse signal constraints in order to identify interfering signals in software-defined radio applications. This is a new application of these

techniques to achieve interference mitigation. A different method, similar in name only, is being used in [34] to cancel interference in compressive sensing systems that is sparse in the time domain (inferring that the interference is not band-limited) rather than in the frequency domain. The application described throughout this thesis is a band-sampling software-defined radio system making use of compressive sensing techniques to identify interference in the frequency domain.

4.3. Methods Developed in the Field of Compressed Sensing

Invoking the application described here of identifying narrowband interfering signals in a spectrally sparse environment, it is assumed that the sparsest solution to the underdetermined linear system of equations is desired, that is, the sparsest solution to the underdetermined linear system of equations is the most accurate. The minimum l_1 -norm solution to the underdetermined system of linear equations $y = Ax$ is often the sparsest [35], as expressed by

$$\underset{x}{\text{minimize}} \quad \|x\|_1 \quad (4)$$

subject to

$$Ax = y.$$

A number of decomposition strategies and approaches to approximating an underdetermined system of equations have been proposed based on the least squares and basis pursuit approaches, including the Method of Frames [36], the Best Orthogonal Basis (signal expansion over an optimized wave-packet orthonormal basis) [37], Matching Pursuit [38], least absolute shrinkage and selection operator (Lasso) [39], and Basis Pursuit [40]. For example, the

Matching Pursuit [38] and Orthogonal Matching Pursuit [41] algorithms seek the solution with the minimum l_0 -norm, as given by

$$\underset{x}{\text{minimize}} \quad \|x\|_0 \quad (5)$$

subject to

$$Ax = y .$$

Linear programming algorithms developed to execute (5) search for a solution vector with the fewest non-zero components. Algorithms seeking the solution with the minimum l_0 -norm are not practical for the cases considered in this work since the components without signals consist of noise, rather than zeros, causing the algorithm to often fail, even with relatively high SNR. This failure occurs for any spectrum sensing application.

Approaches that are less computationally intensive than those seeking to minimize the l_0 -norm seek to minimize the l_1 -norm of the solution, which is desirable for seeking sparse solutions to underdetermined systems of equations [35]. Solutions minimizing the l_1 -norm tend to be more unique than those minimizing the l_2 -norm, that is, the Least Square regression, which tend to have several solutions approaching the optimum. Algorithms minimizing the l_1 -norm include the Iteratively Reweighted Least Squares [42] algorithm, the Least Angle Regression [43] algorithm, and the Lasso [39] algorithm, and basis pursuit methods. The latter comprise Basis Pursuit [40], Polytope Faces Pursuit [44], and Stage-Wise Orthogonal Matching Pursuit [45]. These algorithms solve the same optimization problem. Upon review of the different techniques in relation to the application considered in this work, the Lasso, Least Angle Regression, and Basis Pursuit Denoising optimizations are most promising and seek a solution to the following problem:

$$\underset{x}{\text{minimize}} \quad \frac{1}{2} \|y - Ax\|_2^2 + \lambda \|x\|_1, \quad (6)$$

where the coefficient λ allows for adjustment of the desired degree to which the solution is sparse.

In a real world spectrum sensing problem, sources of spectral leakage might be sparse (i.e. where strong, narrowband interfering signals are sparse), but each frequency bin that ambiguously contributes energy to the measured frequency bin of interest is only attenuated and will be non-zero. As such, the complex optimization shown in (6) can be shown as the closely related and more familiar problem expressed as

$$\underset{x}{\text{minimize}} \quad \|x\|_1 \quad (7)$$

subject to

$$\|Ax - y\|_2 \leq \sigma ,$$

allowing the solution to consist of sparse coefficients among noise, rather than sparse coefficients among zeros. It should be noted that the approach formulated in (7) is a minimization of the l_1 -norm of the solution vector that sufficiently reduces the l_2 -norm of the underdetermined system. The complex optimization in (7) allows for reconstruction while taking into account an estimate of the variance of the measured noise, σ , and the solution to fit an approximation of the linear system, rather than the exact system.

4.4. Open Source Solvers for the Spectrum Reconstruction

Because the application of the sparse reconstruction techniques—rather than the derivation of the solutions themselves—is the focus of this thesis, an efficient open-source solver is used to demonstrate the feasibility of the interference concepts. The reader is referred to [46], [47] for comprehensive reviews and comparisons of algorithms and solvers for l_1 regularization. We evaluated several open source solvers, including the University of British Columbia’s Spectral Projected Gradient for l_1 Minimization (SPGL1) [48], Stanford’s SparseLab [49], Technical University (TU) Darmstadt’s and TU Braunschweig’s Sparse Exact and Approximate Recovery (SPEAR) [50], and California Institute of Technology’s (with Stanford) l_1 -MAGIC [51], to develop the application of sparse reconstruction techniques to spectrum reconstruction that is described throughout this thesis. The results of this investigation are summarized in the next Chapter.

Chapter 5

Analysis and Simulation

The focus of the original content of this thesis is the application of known techniques of approximating underdetermined linear systems to identify interfering signals in software-defined radio applications. It is typically advised that the highest bandwidth signal expected for the software-defined radio dictate the sampling rate to minimize power consumption and that the sampling rate is to be chosen appropriately to prevent aliasing of adjacent channel interference [2]. The significant over-sampling of signal bandwidths to simultaneously monitor disparate frequency bands, which was discussed in Chapter 3, is not yet common practice, and hardware capable of sampling rates in excess of 1GSps are relatively new to the commercial marketplace. Furthermore, sparse signal reconstruction techniques have only seen active research and publication in the last 20 years. The system architectures described in Chapter 3 are simulated here to show proof of concept and initial performance estimates of the application of sparse signal reconstruction techniques to interference mitigation in software-defined radio.

5.1. Band Sampling

For ease of computation and display, these simulations are executed at reduced rates without loss of fidelity of the techniques. In this simulation, there are three desired bands of interest, each with 10Hz bandwidth. The first sub-band is centered at 106.5Hz, the second sub-band is centered at 55Hz, and the third sub-band is centered at 525Hz. The spectrum is over-sampled at a rate of 2.4kSps, well over the requisite 20Sps based on the Nyquist criterion. Figure 9 shows a relatively quiet, whitened (frequency calibrated) spectrum sampled at 1200Sps with a single narrowband signal of interest at approximately 107Hz and an SNR of approximately 40dB.

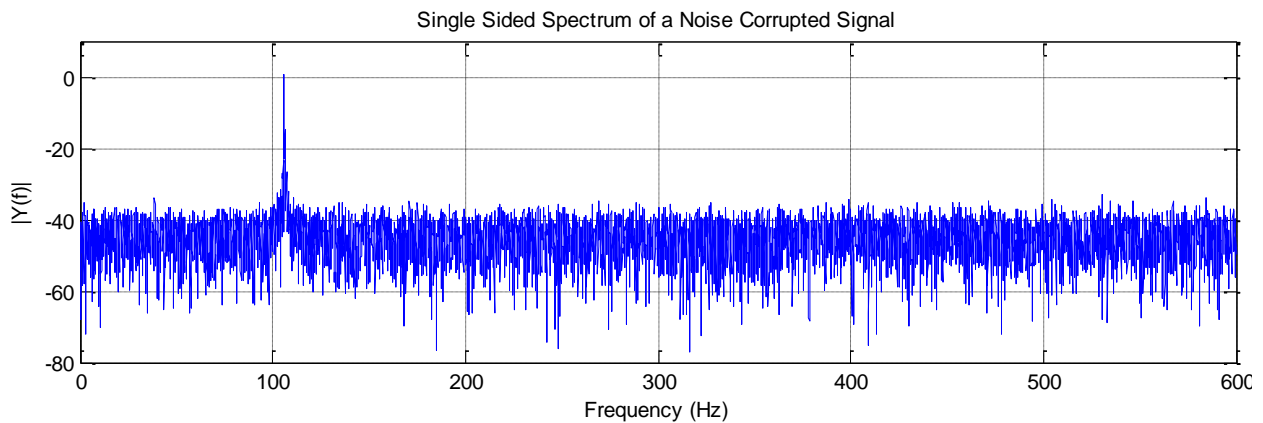


Figure 9: Example of a 107Hz narrowband signal digitized using a high sampling rate.

Once sampled, the digitized data is downconverted according the bands of interest (separately, and in parallel), filtered, and decimated in concert with the aforementioned Nyquist criterion. A comparison of the spectra after the independent, parallel digital downconversion is shown in Figure 10, along with the response of a notional low-pass filter that would be selected and tuned based on the application of the overall system. Traditionally the low-pass filter is expected to effectively eliminate the interfering signal in the 55Hz and 525Hz centered sub-

bands; however, the effectiveness of these filters is very dependent on both the computational resources available and the knowledge of the location of the interfering signal. It is also presumed that in a real system, a detailed summary view of the spectrum would not be available for monitoring in real-time.




Figure 10: Single-sided spectra of the 106Hz narrowband signal after downconversion for analysis of the 102Hz–111Hz sub-band (solid blue), the 50Hz–60Hz sub-band (red), and the 520–530Hz sub-band (green) showing application of a notional low pass filter (dashed blue).

After the digital downconversion, a decimation rate of 64 is applied to each of the sub-bands so that the bandwidth of each sub-band is approximately 1/38 of the original sampling rate. Firstly, the observed spectrum of the sub-band centered around 106.5Hz (Figure 11) correctly shows a 107Hz signal with sufficient SNR to be deemed a detection of a signal of interest. Secondly, due to aliased energy from the 107Hz signal, the observed spectrum of the sub-band centered around 55Hz (Figure 12) incorrectly shows a signal at 51Hz with sufficient SNR to be deemed a detection of a signal of interest. Finally, the observed spectrum of the sub-band centered on 525Hz (Figure 13) correctly shows no discernible signals with sufficient SNR to be deemed a detection of a signal of interest. This is a relevant example, particularly to weak

signals detection situations where the interference is not obvious (for example, showing significant SNR at all sub-bands), not known a priori, and not adequately attenuated by filtering.




Figure 11: Single-sided spectrum of the 102Hz–111Hz sub-band after decimation (frequency axis corrected to show relationship to the original spectrum) showing the true 107Hz signal.

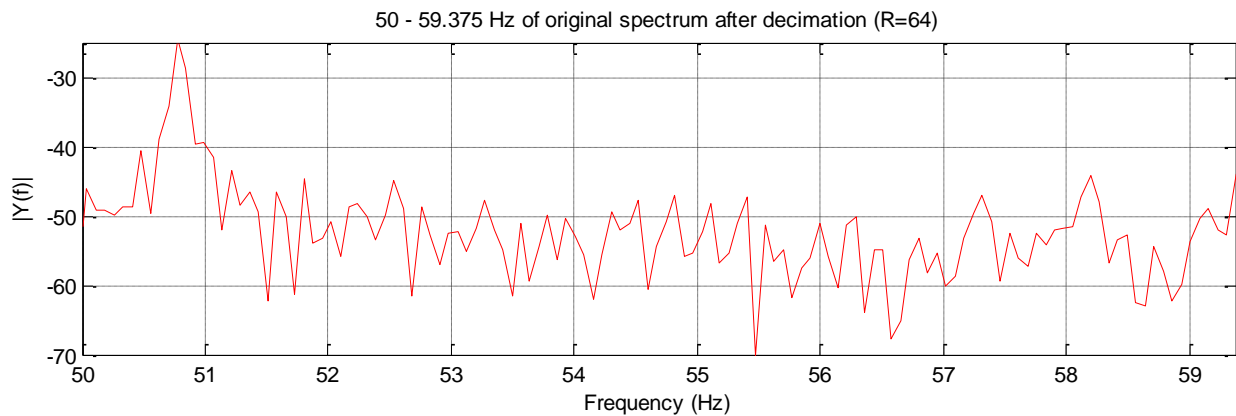


Figure 12: Single-sided spectrum of the 50Hz–60Hz sub-band after decimation (frequency axis corrected to show relationship to the original spectrum) showing the aliased 107Hz signal.




Figure 13: Single-sided spectrum of the 520Hz–530Hz sub-band after decimation (frequency axis corrected to show relationship to the original spectrum).

5.2. Spectral Analysis

Unfortunately, once the decimation takes place, the original information accurately describing the spectrum (as in Figure 9) is lost, hence the term "alias." Even in the absence of signals, each frequency bin of the decimated spectra are the sums of the energy in that frequency bin and the energy in each of its aliased bins, all conditioned by the response of the filter at those respective bins (as shown in Figure 14). The spectrum capture situation simulated above is a simplified example taken directly from a weak signals sensing application. This differs slightly from traditional communications applications in that interfering signals are not high-powered, out-of-band signals that manifest by exceeding the dynamic range of the analog to digital converter. Rather, interfering signals are signals that, despite having been filtered, have enough power to exceed the SNR threshold of the weak signals detector. As shown above, a signal that is of interest in one band can still act as an interfering signal in other bands. In theory, the addition of a taper would increase the out-of-band rejection; however, this comes at a loss of signal gain in the band of interest, which is unacceptable for weak signal reconstruction and detection. This restriction requires focus on adaptive solutions.

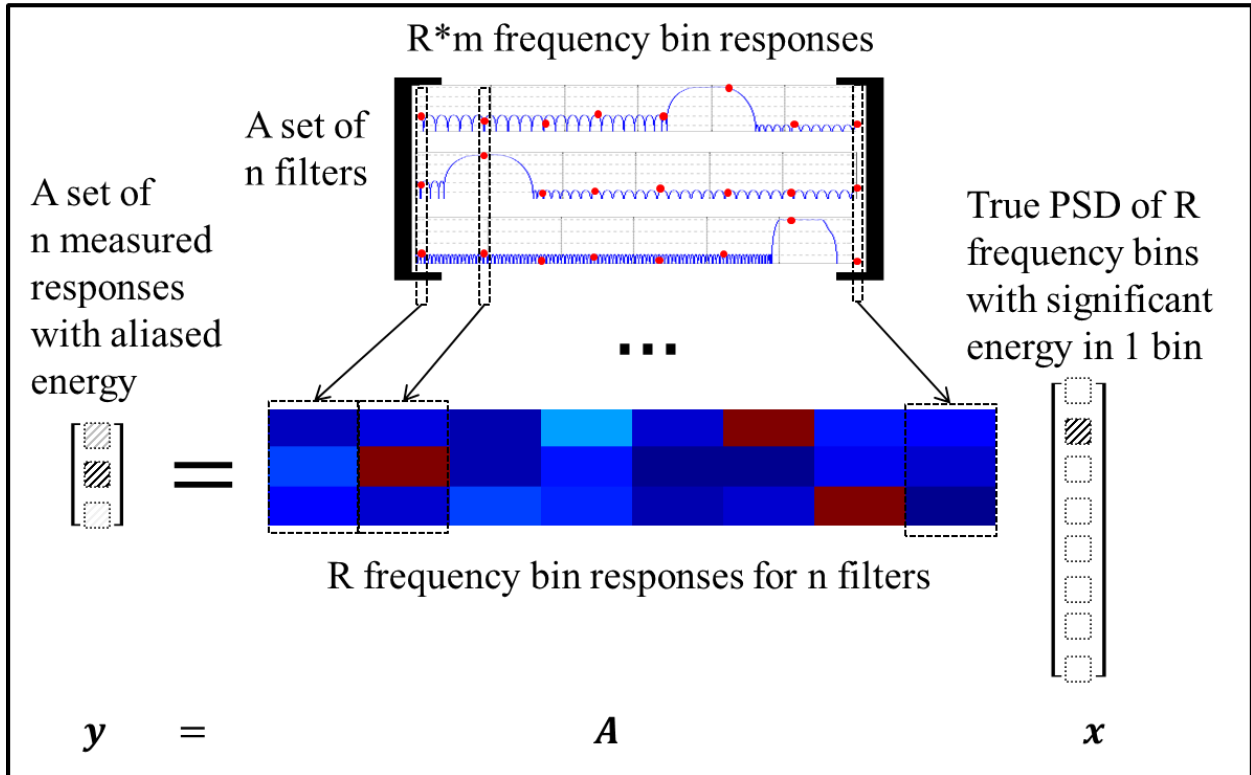


Figure 14: A graphical representation of the underdetermined linear set of equations describing the true power spectral density (PSD) of R frequency bins with n measured responses resulting from the parallel decimation of R frequency bins.

Upon analyzing the spectra shown in Figure 10 through Figure 12, a simple threshold detector easily detects two signals, one at approximately 107Hz (which is a true signal) and one at approximately 51Hz (which is not a true signal but the result of aliasing). One application of the original content of this thesis is the development and presentation of a technique to investigate the detections (from the n measured responses) by reconstructing the true power spectral density of the R desired and aliased frequency bins.

In order to simplify the construct, it is assumed that the filter response at each bin is known and accurate, and that the rate of decimation is equal across the filter sets. The approach described below is still valid for differing combinations of decimation rates; however,

performance is expected to vary since the density of information in the system will be lower (there will be zeros contained in matrix A). As described in Chapter 4, the aliasing is modeled as an underdetermined set of linear equations, $y = A x$, where y represents the bin under examination of the measured spectrum, A represents a matrix of the R responses of n filters (where m is also the decimation rate), and x represents the actual spectral responses at the appropriate bins across the originally sampled spectrum.

5.3. Sparse Signal Reconstruction

The spectrum sensing scenario described above is modeled as a system that is monitoring some set of frequency sub-bands within the total digitized spectrum. For simplicity, it is assumed that each sub-band is of equal bandwidth such that the decimation rates are equal. It would be trivial to include sub-bands with unequal bandwidths because doing this would simply add non-zero and possibly large terms to the filter matrix; however, the simulation results shown throughout this section indicate that adding additional columns to the matrix such that the ratio of matrix rows to matrix columns decreases would reduce the overall performance of all of the algorithms. Each of the monitored sub-bands in the Fourier domain provides a series of measurements, where each measurement represents the estimated power spectral density within that frequency bin. Unfortunately, each estimate includes a linear combination of one frequency bin from the sub-band and $R - 1$ other frequency bins from other sub-bands (where R is the decimation rate) where the coefficient for the desired sub-band is hopefully larger than the others.

In this section, two particular situations are modeled and analyzed. In the first situation, the signal of interest is contained within a desired search band, and the sparse signal reconstruction techniques are used to confirm that the combinations of energy measured are indeed from the signal in question and not aliases. In the second situation, the signal of interest is not contained within a desired search band, and the sparse signal reconstruction techniques are used to determine, from the combinations of energy measured, whether the signal in question is in fact an alias. For each situation, hundreds of Monte Carlo simulations are executed using a range of monitored bandwidths and a variety of solution techniques. Techniques assessed include Basis Pursuit Denoising [48] (an l_1 minimization approach with noise allowance), LASSO [39] (an l_1 minimization approach with noise allowance), Polytope Faces Pursuit [44] (an l_1 minimization approach), Iteratively Reweighted Least Squares [42] (an l_1 minimization approach), and Matching Pursuit [38] (an l_0 minimization approach). Table 1 shows a summary of the algorithms for which results are included in this thesis.

Table 1: A summary of the algorithms used within this thesis.

| Algorithm | Formulation | Primary Open-Source Package Used |
|--|---|----------------------------------|
| Basis Pursuit Denoising | $\begin{aligned} & \text{minimize} \quad \ x\ _1 \\ & x \\ & \text{subject to} \\ & \ Ax - y\ _2 \leq \sigma \end{aligned}$ | SPGL1 [48] |
| Orthogonal Stage-Wise Matching Pursuit | $\begin{aligned} & \text{minimize} \quad \ x\ _1 \\ & x \\ & \text{subject to} \\ & Ax = y \end{aligned}$ | SparseLab [49] |
| Matching Pursuit | $\begin{aligned} & \text{minimize} \quad \ x\ _0 \\ & x \\ & \text{subject to} \\ & Ax = y \end{aligned}$ | SparseLab [49] |

3 EXPERIMENTAL

In this section, a desired signal within a sub-band of interest is modeled. Energy from this signal may alias into the other measurements as dictated by the filter responses of the other sub-bands. Further, the author is interested in identifying if the energy measured is in fact the result of the desired signal in that particular sub-band, rather than signals in other sub-bands, desired or not. First, random frequency responses from a set of N filters are generated where the pass-bands are not overlapping. Next, R evenly spaced bins are selected from each filter and are formed into a matrix, as shown in Figure 15. Throughout these simulations a decimation rate (R) of 128 is used, which is low enough to ease computing, large enough to show important transitions in bandwidth, and realistic for weak signals sensing applications.




Figure 15: Example of R frequency responses (in dB) from a set of 16 filters prior to downconversion and decimation, representative of the matrix A from Equation 3.

The randomly generated filter matrices are applied to a vector representing the true power spectral densities at the frequencies corresponding to the filter coefficient locations. An example of true power spectral densities is shown in Figure 16. This computation represents the digital sub-band channelization, filtering, and decimation. The resulting vector, or measured vector, consists of the measured power spectral density at each of the filter center frequency locations.




Figure 16: Coefficients of the true vector x (from Equation 3) representing the true power spectral densities at the frequencies corresponding to the filter coefficient locations.

The measured vector can provide measurements at the sub-band of interest, but as shown in Section 5.1, the measured vector contains noise from aliasing bands despite the strength of the

image filtering. The sparse signal reconstruction techniques described in Section 5.3 are applied to these measured vectors to attempt to recreate the original or true vectors representing the true power spectral densities. Figure 17 shows an example of a spectrum that is decimated by a factor of 128 and where 16 sub-bands are monitored (as denoted by red dashed boxes). The image shows the original true power spectral densities as red asterisks and the estimated power spectral densities using the Basis Pursuit Denoise algorithm implemented by SPGL1 [48] as blue circles. In this example, there is a true signal with sufficient SNR within a monitored sub-band. The algorithm requires only the measured vector and the full filter matrix response to construct an estimate. This simulation was repeated for a number of different sub-band-to-total-bandwidth ratios, interfering signal SNRs, and reconstruction techniques.-

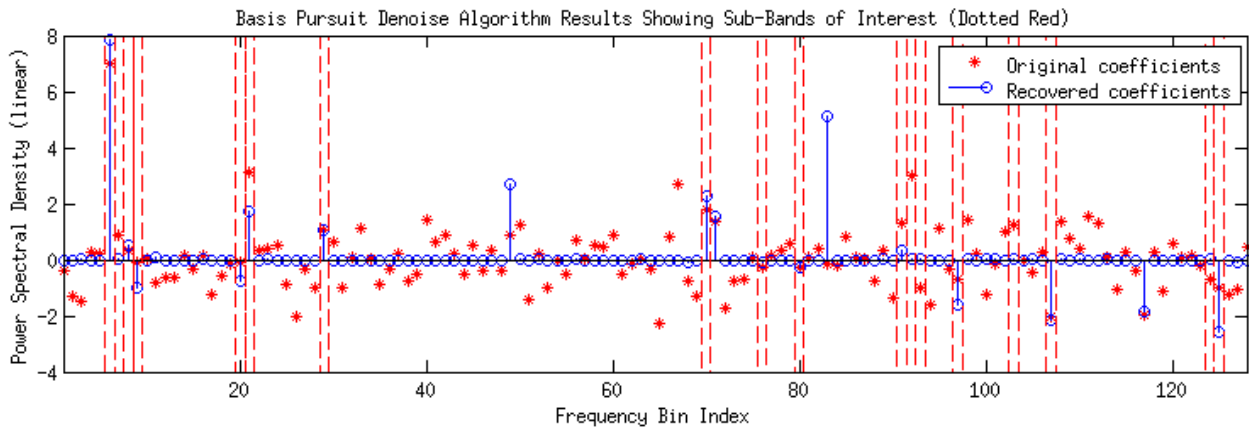


Figure 17: Coefficients of the true vector x (from Equation 3) recovered using the Basis Pursuit Denoise algorithm implemented by SPGL1 [48] showing accurate confirmation that energy measured within a search band is due to a signal from that band despite an underdetermined system.

First, the Basis Pursuit Denoise Algorithm as implemented by SPGL1 [48] is tested to estimate the original sensed spectrum. A summary of the results is shown in Figure 18 for different SNRs and monitored bandwidths. The special case where only a single sub-band is

monitored (bottom row of image in Figure 18) is trivial since increased aliasing would only add in magnitude to the correct result. Further, since the signal of interest is contained within a monitored sub-band, lower SNR cases show good performance at low ratios of monitored bandwidth, which is largely due to the size of the filter matrix in that the problem is severely underdetermined. Finally, it is shown that even with weak SNRs (less than 10dB) the Basis Pursuit Denoise algorithm can correctly rule out the presence of aliasing in a spectral estimate.




Figure 18: Results of a Monte Carlo simulation showing the percentage correct of coefficients of the true vector x (from Equation 3) recovered using the Basis Pursuit Denoise algorithm implemented by SPGL1 [48] based on the SNR of the signal in question and the ratio of monitored to total bandwidth when the signal in question is within a desired search band.

Second, the Stage-Wise Orthogonal Matching Pursuit algorithm [45] implemented by SparseLab [49] is tested to estimate the original sensed spectrum using the same conditions as above for comparison against an algorithm that does not take noise into account. A summary of the results is shown in Figure 19 for different SNRs and monitored bandwidths. One limitation of this class of algorithm is shown in the case where only a single sub-band is being monitored and contains energy (bottom row of image in Figure 19). The Stage-Wise Orthogonal Matching Pursuit algorithm is unable to find a solution with the acceptable tolerance since the model does not account for noise and the measured energy is the result of measured noise and aliased noise and signal. Further, the effect of assuming a noiseless system is apparent in that greater SNRs are required for correct results and the variance in the results is more pronounced as compared to the Basis Pursuit Denoise algorithm. Finally, it should be noted that the Stage-Wise Orthogonal Matching Pursuit algorithm [45] as implemented by SparseLab [49] is less computationally intensive than the Basis Pursuit Denoise Algorithm as implemented by SPGL1 [48], and though less effective, may be sufficient for some applications.




Figure 19: Results of a Monte Carlo simulation showing the percentage correct of coefficients of the true vector x (from Equation 3) recovered using the Stage-Wise Orthogonal Matching Pursuit algorithm [45] implemented by SparseLab [49] based on the SNR of the signal in question and the ratio of monitored to total bandwidth when the signal in question is within a desired search band.

Finally, the Matching Pursuit algorithm [30] implemented by SparseLab [41] is tested to estimate the original sensed spectrum using the same conditions as above for completeness, even though minimizing the l_0 -norm is considered an NP-hard problem. As such, computation of the Matching Pursuit algorithm is very intense and requires careful selection of iteration stopping criteria. Even though the stopping criteria may be set properly, the Matching Pursuit algorithm is not robust to noise as shown in a summary of the results (Figure 21) for different SNRs and

monitored bandwidths. In the case where only a single sub-band is monitored, the solution is trivial since the Matching Pursuit algorithm selects the strongest signal detected as true. As expected, once the signal-to-noise ratio of the signal exceeds the variance of the noise, the correct answer is selected without fail. As additional sub-bands are included in the measurement, the susceptibility of this algorithm to noise is shown, seldom exceeding a forty percent success rate.




Figure 20: Results of a Monte Carlo simulation showing the percentage correct of coefficients of the true vector x (from Equation 3) recovered using the Matching Pursuit algorithm [38] implemented by SparseLab [49] based on the SNR of the signal in question and the ratio of monitored to total bandwidth when the signal in question is within a desired search band.

2 FIGURE

In this section, an undesired interfering signal outside of all sub-bands of interest is modeled. Energy from this interfering signal will alias to some degree into the measurements of the monitored sub-bands as dictated by the collective filter responses. The methods described throughout this thesis propose the ability to recover the presence of an interfering signal even though the corresponding sub-band is not monitored. As in Section 5.3.1, random frequency responses from a set of N filters are generated where the pass-bands are not overlapping. Again, R evenly spaced bins are selected from each filter and are formed into a matrix, as shown in Figure 21. All parameters are kept constant between the two sets of simulations except for the location of the signals in question (in either a monitored or unmonitored sub-band).




Figure 21: Example of R frequency responses (in dB) from a set of 16 filters prior to downconversion and decimation, representative of the matrix A from Equation 3.

As in Section 5.3.1, the measured vector, containing noise from aliasing bands despite the strength of the image filtering, is combined with the knowledge of the filter response matrix. The sparse signal reconstruction techniques described in Section 5.3 are again applied to these measured vectors to attempt to recreate the true vectors representing the true power spectral densities. Figure 22 shows an example of a spectrum that is decimated by a factor of 128 and

where 16 sub-bands are monitored (as denoted by red dashed boxes). The image shows the original true power spectral densities as red asterisks and the estimated power spectral densities using the Basis Pursuit Denoise algorithm implemented by SPGL1 [48] as blue circles. This example shows an interfering signal successfully identified within an unmonitored sub-band.

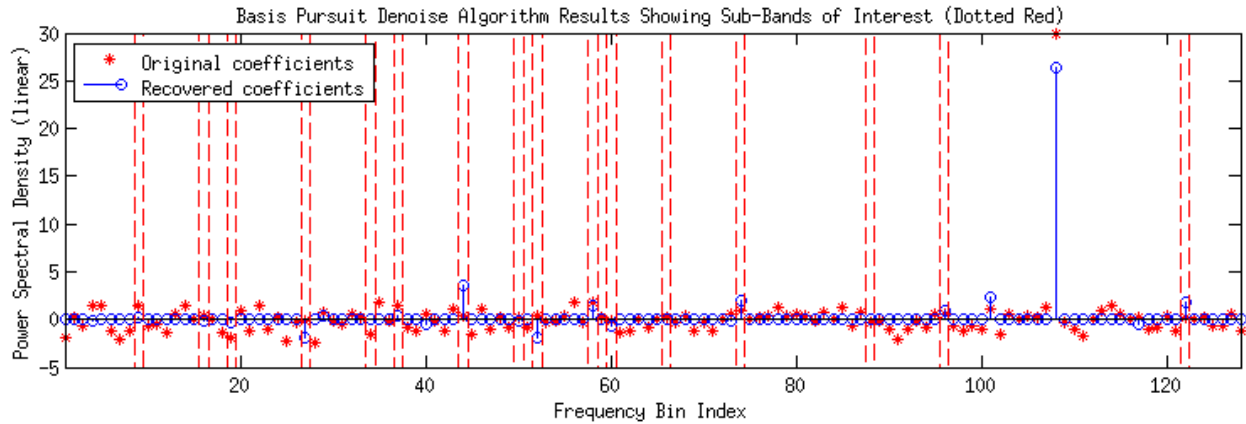


Figure 22: Coefficients of the true vector x (from Equation 3) recovered using the Basis Pursuit Denoise algorithm [48] showing accurate recovery of an out-of-band interfering signal despite an underdetermined system.

First, the Basis Pursuit Denoise Algorithm as implemented by SPGL1 [48] is tested to estimate the original sensed spectrum. A summary of the results is shown in Figure 23, revealing a crisp performance bound for this particular filter set. As expected, recovery of the interfering signal is more likely as either the SNR is increased or the number of monitored sub-bands is increased. Despite significantly increasing the SNR of the interfering signal, Figure 23 indicates that identifying the true sub-band of the interfering signal is unlikely without monitoring at least 10% of the total bandwidth (for each simulation, monitored sub-bands were uniformly randomly distributed across the total bandwidth). Further, as the number of monitored sub-bands approaches the total bandwidth, performance approaches traditional detection theory, as

expected. A detailed view of the Basis Pursuit Denoise results is shown in Figure 24, highlighting the crisp, rapid transition between near-certain recovery and near-certain failure.




Figure 23: Results of a Monte Carlo simulation showing the percentage correct of coefficients of the true vector x (from Equation 3) recovered using the Basis Pursuit Denoise algorithm implemented by SPGL1 [48] based on the SNR of the signal in question and the ratio of monitored to total bandwidth when the interfering signal is outside of any desired search bands.




Figure 24: Details of results of a Monte Carlo simulation using the Basis Pursuit Denoise algorithm implemented by SPGL1 [48] showing percent correct of recovery.

Proceeding as in Section 5.3.1, the Stage-Wise Orthogonal Matching Pursuit algorithm [45] implemented by SparseLab [49] is tested to estimate the original sensed spectrum including an interfering signal using a less robust algorithm. A summary of the results is shown in Figure 25 for comparison to Figure 24. As before, it is apparent that the Stage-Wise Orthogonal Matching Pursuit algorithm and other minimum l_1 -norm algorithms without noise compensation are far less effective than those with noise compensation, including the Basis Pursuit Denoise algorithm. The shape of the success region in both cases follows an exponential curve; however, the transition region from success to failure for the Stage-Wise Orthogonal Matching Pursuit algorithm is far larger. The Stage-Wise Orthogonal Matching Pursuit algorithm [45] has

increasingly better performance with significantly larger SNRs since the noise values effectively approach zero, which the algorithm assumes. This algorithm as implemented by SparseLab [49], though less effective than the Basis Pursuit Denoise Algorithm as implemented by SPGL1 [48], may be sufficient and desirable for some applications.




Figure 25: Results of a Monte Carlo simulation showing the percentage correct of coefficients of the true vector x (from Equation 3) recovered using the Stage-Wise Orthogonal Matching Pursuit algorithm [45] implemented by SparseLab [49] based on the SNR of the signal in question and the ratio of monitored to total bandwidth when the interfering signal is outside of any desired search bands.

Finally, for completeness, both algorithms are compared to the Matching Pursuit algorithm [30] implemented by SparseLab [41]. As expected, the results displayed in Figure 26

reveal a significant degradation of performance for the latter algorithm in both probability of success and transition distance between failure and success. Although computationally complex, the conceptually simple minimization of the l_0 -norm is quickly defeated as the contribution of noise power is assumed to be negligible. The ineffectiveness and computational complexity of minimizing the l_0 -norm of the solution subject to the constraints described above indicate that a higher-order norm minimization is preferred for the problem described throughout this thesis, particularly because of its robustness to noise in the measurements.




Figure 26: Results of a Monte Carlo simulation showing the percentage correct of coefficients of the true vector x (from Equation 3) recovered using the Matching Pursuit algorithm [38] implemented by SparseLab [49] based on the SNR of the signal in question and the ratio of monitored to total bandwidth when the interfering signal is outside of any desired search bands.

Chapter 6

Conclusions and Future Work

The techniques introduced in Chapter 4 are applied in the Fourier domain via a novel approach to interference identification within a band-sampling software-defined radio. They are suitable for applications requiring the detection of weak signals, as described in Chapter 3 and demonstrated in Chapter 5 by means of several open-source algorithm implementations developed primarily for the Compressed Sensing field. Specifically, we show that these algorithms can successfully reconstruct interfering signals from measured aliasing effects predicted only by the known filter responses prior to incoming signal decimation. Analysis and simulation results reveal that the Basis Pursuit Denoising, LASSO, and other algorithms seeking equivalent mathematical optimization, are able to identify interfering signals that are sparse in the frequency domain, given sufficient SNR and instantaneous bandwidth.

Furthermore, analysis and simulation results show that, while not ideal, algorithms that simply minimize the l_1 -norm subject to some constraints could be used in some cases. Finally,

as spectrum availability becomes scarcer and personal electronic devices proliferate, additional applications requiring pristine and flexible access to disparate frequency bands will surface, requiring the development of new interference identification and mitigation techniques.

Applying sparse signal reconstruction techniques for interference identification leaves several interesting research topics open for future work. These include:

- **Interference mitigation:** Once an interfering signal is identified within an unmonitored sub-band, adaptive techniques may be developed to temporarily suppress the effects of the interfering signal. They may include frequency suppression through dynamic re-computation of filters, additive filters (that is, filters applied to incoming data prior to digitization), or multi-channel beam-forming techniques.
- **Spectrum sub-space monitoring:** As discussed in Chapter 2, instantaneous bandwidth can limit the performance of a software-defined radio. As such, monitoring additional bandwidth beyond what is traditionally dictated by the Nyquist sampling criterion is of interest. The ideas presented here regarding estimation of unmonitored frequency sub-bands could be applied in reverse, that is, a set of filters and sub-bands could be designed such that monitoring some subset of the total bandwidth could enable detection of sparse interfering signals across the entire bandwidth. Therefore, a set of sub-bands could be designated and set aside purely for interference detection across the entire instantaneous bandwidth, regardless of the sub-bands currently in use.
- **Adaptive interference identification:** Implementations of the algorithms studied throughout this thesis were manually tuned to the known conditions of the simulation for performance. These algorithms could be reviewed and implemented in an adaptive sense, automatically assessing incoming data and estimating input parameters for maximum

performance. This would be important for cognitive radio applications, where knowledge of the surrounding environment is important input for the cognitive reasoning engine.

- Degrees of sparseness: The effect of the degree of sparseness on the effectiveness of the algorithms described throughout this thesis should be reviewed. In the Fourier domain, the degree of sparseness may include the signal-to-noise ratio of the interfering signal and the statistical distribution of the noise.
- Filter exploration: Other commonly used digital filters in software-defined radio should be explored and applied to this construct. Filter side lobe responses dictate the majority of the elements in the filter matrix. There may be a filter type or design principle that exhibits side lobe patterns that would be beneficial to the sparse signal reconstruction algorithm. For example, a Chebyshev-type filter that has amplitude-matched side-lobes might at first glance seem undesirable because the filter matrix would be low rank and contain insufficient information to reconstruct the underdetermined linear system of equations. In contrast, a Chebyshev-type filter designed to have unique response with respect to the particular frequency bin at each sub-band (i.e., where the null to null pass-band of the filter is mismatched to the sub-band bandwidth) could perform quite nicely.

References

- [1] “Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2011–2016,” Cisco Systems, Inc., White Paper, Feb. 2012.
- [2] J. H. Reed, *Software Radio : A Modern Approach to Radio Engineering*. Upper Saddle River, NJ: Prentice Hall, 2002.
- [3] J. Mitola, “The software radio architecture,” *Communications Magazine, IEEE*, vol. 33, no. 5, pp. 26–38, 1995.
- [4] J. C. Maxwell, “A Dynamical Theory of the Electromagnetic Field,” *Philosophical Transactions of the Royal Society of London*, vol. 155, no. 0, pp. 459–512, Jan. 1865.
- [5] “Prof. D. E. Hughes’ Research in Wireless Telegraphy,” *The Electrician*, vol. 40, pp. 40–41, 1899.
- [6] J. J. Fahie, *A History of Wireless Telegraphy*, 3rd ed. Dodd, Mead, and Company, 1902.
- [7] G. Marconi, “Transmitting Electrical Signals,” U.S. Patent 0,586,19313-Jul-1897.
- [8] N. Tesla, “System of Transmission of Electrical Energy,” U.S. Patent 0,645,57620-Mar-1900.
- [9] Lee De Forest, “The Audion; A New Receiver for Wireless Telegraphy,” *Transactions of the American Institute of Electrical Engineers*, vol. XXV, pp. 735–763, Jan. 1906.
- [10] J. E. Lilienfeld, “Method and Apparatus for Controlling Electric Currents,” U.S. Patent 1,745,17528-Jan-1930.
- [11] *Handbook of RF and wireless technologies*. Amsterdam ; Boston: Newnes, 2004.
- [12] J. Mitola and G. Q. Maguire, “Cognitive radio: making software radios more personal,” *IEEE Personal Communications*, vol. 6, no. 4, pp. 13–18, Aug. 1999.
- [13] A. B. MacKenzie, J. H. Reed, P. Athanas, C. W. Bostian, R. M. Buehrer, L. A. DaSilva, S. W. Ellingson, Y. T. Hou, M. Hsiao, Jung-Min Park, C. Patterson, S. Raman, and C. da Silva, “Cognitive Radio and Networking Research at Virginia Tech,” *Proceedings of the IEEE*, vol. 97, no. 4, pp. 660–688, Apr. 2009.
- [14] M. D. Gallagher, F. R. Wentland, E. Drocella, and M. Settle, *Facilitating Opportunities for Flexible, Efficient, And Reliable Spectrum Use Employing Cognitive Radio Technologies*. 2005.
- [15] Q. Zhao and B. M. Sadler, “A survey of dynamic spectrum access,” *Signal Processing Magazine, IEEE*, vol. 24, no. 3, pp. 79–89, 2007.
- [16] Texas Instruments, “RF-Sampling and GSPS ADCs.” 2012.
- [17] M. J. Marcus, “Unlicensed cognitive sharing of TV spectrum: The controversy at the federal communications commission,” *Communications Magazine, IEEE*, vol. 43, no. 5, pp. 24–25, 2005.
- [18] G. Lawton, “Machine-to-machine technology gears up for growth,” *Computer*, vol. 37, no. 9, pp. 12–15, 2004.
- [19] K. Pahlavan, T. H. Probert, and M. E. Chase, “Trends in Local Wireless Networks,” *IEEE Communications Magazine*, vol. 33, no. 3, pp. 88–95, Mar-1995.
- [20] ITU, “Radio Regulations Articles.” 2012.
- [21] D. Schmidt and A. Gupta, “Evolving wireless technology & multi mode requirements for client platforms,” in *Radio and Wireless Symposium, 2006 IEEE*, 2006, pp. 1–2.

- [22] A. Bourdoux, J. Craninckx, A. Dejonghe, and L. Van der Perre, "Receiver architectures for software-defined radios in mobile terminals: The path to cognitive radios," in *Radio and Wireless Symposium, 2007 IEEE*, 2007, pp. 535–538.
- [23] C. E. Shannon, "Communication in the Presence of Noise," *Proceedings of the IRE*, vol. 37, no. 1, pp. 10–21, Jan. 1949.
- [24] B. Sklar, *Digital communications : fundamentals and applications*. Upper Saddle River, N.J.: Prentice-Hall PTR, 2001.
- [25] S. Weiss and R. W. Stewart, "On the optimality of subband adaptive filters," pp. 59–62.
- [26] E. B. Hogenauer, "An Economical Class of Digital Filters for Decimation and Interpolation," *Transactions of Acoustics, Speech, and Signal Processing, IEEE*, vol. ASSP-29, no. 2, pp. 155–162, Apr. 1981.
- [27] R. G. Baraniuk, "More Is Less: Signal Processing and the Data Deluge," *Science*, vol. 331, no. 6018, pp. 717–719, Feb. 2011.
- [28] D. L. Donoho, "Sparse nonnegative solution of underdetermined linear equations by linear programming," *Proceedings of the National Academy of Sciences*, vol. 102, no. 27, pp. 9446–9451, Jun. 2005.
- [29] E. J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [30] *Compressed sensing: theory and applications*. Cambridge; New York: Cambridge University Press, 2012.
- [31] R. G. Baraniuk, "Compressive sensing [lecture notes]," *Signal Processing Magazine, IEEE*, vol. 24, no. 4, pp. 118–121, 2007.
- [32] E. J. Candès, "Compressive sampling," in *Proceedings of the International Congress of Mathematicians: Madrid, August 22-30, 2006: invited lectures*, 2006, pp. 1433–1452.
- [33] D. L. Donoho, "Compressed sensing," *Information Theory, IEEE Transactions on*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [34] M. A. Davenport, P. T. Boufounos, and R. G. Baraniuk, "Compressive domain interference cancellation," DTIC Document, 2009.
- [35] D. L. Donoho and Y. Tsaig, *Fast solution of l_1 -norm minimization problems when the solution may be sparse*. Department of Statistics, Stanford University, 2006.
- [36] I. Daubechies, "Time-frequency localization operators: a geometric phase space approach," *IEEE Transactions on Information Theory*, vol. 34, no. 4, pp. 605–612, Jul. 1988.
- [37] R. R. Coifman and M. V. Wickerhauser, "Entropy-based algorithms for best basis selection," *IEEE Transactions on Information Theory*, vol. 38, no. 2, pp. 713–718, Mar. 1992.
- [38] S. G. Mallat and Zhifeng Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.
- [39] R. Tibshirani, "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 58, no. 1, pp. 267–288.
- [40] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic Decomposition by Basis Pursuit," *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 33–61, Jan. 1998.
- [41] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in

- Signals, Systems and Computers, 1993. 1993 Conference Record of The Twenty-Seventh Asilomar Conference on*, 1993, pp. 40–44.
- [42] I. Daubechies, R. DeVore, M. Fornasier, and C. Sinan Gunturk, “Iteratively re-weighted least squares minimization for sparse recovery,” *eprint arXiv:0807.0575*, Jul. 2008.
 - [43] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani, “Least angle regression,” *The Annals of statistics*, vol. 32, no. 2, pp. 407–499, 2004.
 - [44] M. Plumbley, “Recovery of sparse representations by polytope faces pursuit,” *Independent Component Analysis and Blind Signal Separation*, pp. 206–213, 2006.
 - [45] D. L. Donoho, I. Drori, Y. Tsaig, and J.-L. Starck, *Sparse solution of underdetermined linear equations by stagewise orthogonal matching pursuit*. Department of Statistics, Stanford University, 2006.
 - [46] M. Schmidt, G. Fung, and R. Rosales, “Optimization methods for l_1 -regularization,” *University of British Columbia, Technical Report TR-2009-19*, 2009.
 - [47] G.-X. Yuan, K.-W. Chang, C.-J. Hsieh, and C.-J. Lin, “A comparison of optimization methods and software for large-scale l_1 -regularized linear classification,” *The Journal of Machine Learning Research*, vol. 9999, pp. 3183–3234, 2010.
 - [48] E. van den Berg and M. P. Friedlander, *SPGL1: A solver for large-scale sparse reconstruction*. 2007.
 - [49] D. Donoho, I. Drori, V. Stodden, Y. Tsaig, and M. Shahram, *SparseLab*. 2007.
 - [50] D. A. Lorenz, M. E. Pfetsch, A. M. Tillmann, and C. Kruschel, *SPEAR – Sparse Exact and Approximate Recovery*. 2013.
 - [51] E. J. Candès and J. Romberg, *l1-MAGIC*. 2005.