

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Maximizing available spectrum for cognitive radios

Permalink

<https://escholarship.org/uc/item/722103kg>

Author

Mishra, Shridhar Mubaraq

Publication Date

2009

Peer reviewed|Thesis/dissertation

Maximizing available spectrum for cognitive radios

by

Shridhar Mubaraq Mishra

A dissertation submitted in partial satisfaction

of the requirements for the degree of

Doctor of Philosophy

in

Engineering—Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Robert W. Brodersen, Chair

Professor Anant Sahai

Professor Howard Shelanski

Fall 2009

The dissertation of Shridhar Mubaraq Mishra, titled Maximizing available spectrum for cognitive radios, is approved:

Chair _____ Date _____

_____ Date _____

_____ Date _____

University of California, Berkeley

Maximizing available spectrum for cognitive radios

Copyright © 2009

by

Shridhar Mubaraq Mishra

Abstract

Maximizing available spectrum for cognitive radios

by

Shridhar Mubaraq Mishra

Doctor of Philosophy in Engineering—Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Robert W. Brodersen, Chair

Cognitive radios have been proposed to address the dual problem of spectrum under-utilization and the need for vast swathes of new spectrum in the 0-3GHz band for wireless data services. Previous research in the area of cognitive radios has concentrated on signal processing (SP) innovations for improved detection with a limited number of samples (a.k.a detection sensitivity). However the SP perspective alone is unable to recover much unused spectrum. In this thesis we focus on two fronts to maximize the recovery of unused spectrum. First, we take a spatial perspective on spectrum usage. Second, we will look at advanced algorithms which use other radios and frequencies to aid the detection process.

The spatial perspective is necessitated by the FCC’s decision to open up TV bands for ‘white space’/opportunistic spectrum use. Examination of the FCC’s choice of parameters reveals the political and engineering tradeoffs made by the FCC. The FCC’s rules applied to the database of TV transmitters and the population density of the United States as per the census of 2000 reveal that the geo-location rules enable an average of 9 white space channels per person. This number does not change significantly even if the secondaries take a purely selfish approach and operate only where the pollution from TV transmitters is low. The corresponding political tradeoff sacrifices 1 people-channel for broadcast use to gain 8 people-channels for white space use. Spectrum sensing is another mechanism to recover the available opportunity. Here, the FCC’s sensing rule of -114dBm is very conservative and yields only a average of 1 white space channel per person.

This spatial analysis points to the inability of traditional detection metrics in predicting the spatial performance of sensing. To overcome this problem, we propose the twin metrics of Fear of Harmful Interference (F_{HI}) and Weighted Probability of Area Recovered ($WPAR$) to quantify the performance of sensing relative to geo-location. *Fear of Harmful interference* F_{HI} , captures the safety to the primary users and is largely the fading-aware probability of missed detection with modifications to allow easier incorporation of system-level uncertainty. *Weighted Probability of Area Recovered* $WPAR$, captures the performance of spectrum sensing by appropriately weighting the probability of false alarm (P_{FA}) across different spatial locations. These metrics give a unifying framework in which to compare different spectrum-

sensing algorithms. Reasonable parameters for this metric (most crucially, something which can be interpreted as a spatial discount factor) are obtained from the TV database and US population data.

Cooperative Sensing, in which secondaries use sensing results from multiple nearby radios to decide on whether the band is free to use, has been proposed as a mechanism to improve the performance over a single radio. However, the motivation for using cooperative sensing is different for recovering a temporal as opposed to recovering a spatial hole. For a temporal hole, cooperative sensing is used to gather as much energy in the primary signal as possible, thus maximizing the gap between the means of the signal present and signal absent hypotheses. For a spatial hole there is a mapping between the nominal signal strength at a given distance from the transmitter and the distance itself. By cooperating the hope is to converge to this nominal signal level. This difference has major implications on the choice of cooperation rule for different scenarios. A median rule is proposed as a robust cooperation rule for spatial holes. This rule can be implemented as a hard combining rule and is robust to uncertainties in the fading models and untrusted radio behavior.

Unfortunately all cooperation suffers when there is a loss in spatial diversity. To combat shadowing correlation across space, assisted/calibrated detection in the form of multiband sensing is proposed. In multiband sensing, detection results from nearby frequencies is used to aid detection in the frequency band of interest. Cooperation across multiband radios is robust against channel correlation and provides gains by weeding out radios afflicted by adverse propagation environments. A mobile, wideband testbed was designed and used to capture TV signals in the 500-700MHz band at various locations in Berkeley, CA. Analysis of these measurements reveals high shadowing and multipath spread correlation across frequencies. These new insights into the spectral environment are used to design detectors that perform better than a single band detector.

The ability to coexist with primaries of different scales (high power TV transmitters and low power Wireless Microphones) is an important requirement for a cognitive radio system. When sensing a large scale primary, a small scale secondary user can make its own decision about transmission based on the sensing results from its neighborhood. This assumption fails when the scale of the primary is comparable to the scale of the secondary user. In this scenario, we need to decouple sensing from admission control – a sensor network is required to perform the sensing and localize the primary. For small primaries, the environment over which the sensing results are valid is small which imposes certain minimum density requirements for sensor nodes. Location information of the primary and secondary users is key for such an admission control algorithm to operate successfully.

Acknowledgements

I guess we spend a lifetime learning about ourselves. Life is a lesson in self discovery. In this quest, graduate school is probably the best time for introspection. We pitch ourselves against the best and come out better, wiser and more appreciative of the special qualities that make us who we really are. This learning process is greatly enhanced by our companions in this journey – our advisors, fellow graduate students, research collaborators and even the coffee vendor whose coffee helped us remain awake at ungodly hours.

During my six years in Berkeley a few of people have had a tremendous impact on me in general and on my research in particular (its the 80-20 rule!). Primary among these have been my advisors: Bob and Anant. Bob has always been a role model weaving plentiful biking vacations into a hectic schedule. From him, I've learned as much about research as I have learned about how to live life. He gave me a free hand to explore different problems but always nudged me into the right path when the time was right. Anant taught me that defining the 'right' problem was as important as finding the 'right' solution. Anant's extraordinary ability in seeing connections between diverse concepts provided critical insights into the cognitive radio problem. I would spend days working on a problem and walk into his office self assured that I all bases covered, but he would always managed to ask questions for which I had no answer!

Nothing beats working on interesting problems with fellow students. Rahul Tandra and I spend many hours working on problems together. It was a great experience and led to some very exciting research. I also had very meaningful discussions about my research with Pukit Grover, Kristen Woyach, Danijela Cabric, Hari Palaiyanur, David Qiu and Kate Harrison. Many new ideas emerged from these discussions that have found their way into various papers. Discussions with Prof. Adam Wolisz (TU Berlin/UC Berkeley), John Notor (Cadence), Steve ten Brink (Wionics), Steve Shellhammer (Qualcomm), Gerald Chouinard

(CRC Canada) and Jeff Schiffer (Intel) were greatly valued and help set the course of my research at various times.

The best part of graduate school is that learning and discussions are not limited to the classroom or one's research. Engrossing conversations with other students provided great learning (mostly) and entertainment (sometimes). For these experiences I have to acknowledge Jing Yang, Simone Gambini, Rahul Shah, Aarsh Parsa, Vinayak Nagpal, David Chen, Zhengya Zhang, Jana Van Greunen, Louis Alacron, Ehsan Adabi, Mark Wagner, Henry Chen, Animesh Kumar, Cindy Liu, Tanya Roosta, Chintan Thakkar, Seng Oon Toh, Renaldi Winoto, Omar Bakr, Jonathan Tsao, Lakshmi Subramanian, DeLynn Bettencourt, Kris Eswaran, Bobak Nazar and Ashkan Borna.

Everyone at the BWRC played a vital part in providing the infrastructure to make this research happen. Sue Mellers and David Qiu deserve special thanks for helping in developing the testbed. Brian Richards was the go to person for any Simulink/Matlab/Xilinx issues and Tom Boot was universally helpful with all administration related issues. A special thanks to Kristen, Hari, Pulkit and Avni for proof reading this thesis.

Finally, I would like to thank my family and friends for supporting me through this period. My parents developed in me a love for learning. They were ecstatic to see me go to graduate school though I don't think that six years in graduate school was what they had envisioned. My wife, Avni, has been a source of great support and encouragement through these years (where the encouragement was sometimes provided via cleverly disguised variants of the same question, "When are you getting done?"). It always reminded me of PhD comics. They say married graduate students finish earlier than single ones. Now I understand why!

To
my family

Contents

1	Introduction	1
1.1	Motivation: The apparent scarcity problem	1
1.2	Mechanisms to solve the scarcity problem	2
1.3	A historical perspective on Opportunistic Spectrum Access	3
1.4	Defining a spectrum hole	4
1.5	Organization of this thesis	6
2	Spatial spectrum hole/White space: Definition and evaluation	9
2.1	Prior work in estimating available white space	11
2.2	White space using the pollution viewpoint	12
2.2.1	Adjacent channel considerations in the pollution viewpoint	13
2.2.2	White space calculations using actual population density	13
2.2.3	Limitations of this evaluation	14
2.3	White space using the protection viewpoint	14
2.3.1	Defining white space for the protection viewpoint	16
2.3.2	Calculating the available white space from the protection viewpoint	19
2.3.3	Impact of Secondary power on white space availability	22
2.3.4	Available white space considering adjacent channels	22
2.3.5	Available white space considering multipath fading	24

2.4	A principled way to choose the protection margin	24
2.4.1	Impact of adjacent channels	26
2.4.2	Opportunistic spectrum use versus reallocation	26
2.4.3	Impact of the pollution viewpoint	31
2.5	Impact of FCC rules on white space availability	31
2.5.1	Spectrum sensing: SNR as a proxy for distance	34
2.6	Discussion	43
2.6.1	Impact of wireless Microphones	43
2.6.2	Towards a framework to generate tradeoff plots	46
2.7	Summary	46
3	Metrics	48
3.1	Traditional Metrics	50
3.2	Impact of fading	51
3.3	Towards New Metrics	54
3.4	Spatial Metrics	57
3.4.1	Safety: Controlling the Fear of Harmful Interference	58
3.4.2	Performance	58
3.4.3	Quantile Models	59
3.4.4	Single radio performance	60
3.4.5	Comparing metrics to empirical performance	60
3.5	Summary	63
4	Cooperation	64
4.1	Related work	65
4.2	Cooperation for temporal holes	65
4.2.1	Cooperation strategies	65
4.2.2	Performance of temporal cooperation strategies	70

4.2.3	Temporal strategies in the real world	71
4.3	Cooperation for spatial holes	76
4.3.1	Maximum-likelihood detector: soft-decision combining	78
4.3.2	OR-rule detector: hard decision combining	80
4.3.3	The Median Rule	82
4.3.4	Evaluating the relative performance of hard decision rules	85
4.4	Cooperation with Uncertainty	88
4.4.1	Performance of cooperation under quantile models	92
4.4.2	Performance of cooperation under loss of independence	99
4.4.3	Performance of cooperation with unreliable radios	101
4.5	Summary	104
5	Assisted Detection	106
5.1	Related work	107
5.2	Single Tower Multiband Sensing	108
5.2.1	Problem Setup	108
5.2.2	Propagation Model	109
5.2.3	Single Band Detection	109
5.2.4	Multiband detection	111
5.2.5	Cooperative gains from Multiband sensing	116
5.2.6	Moving from anchor to multiple primaries	123
5.3	Experimental results	125
5.3.1	Shadowing Correlation	125
5.3.2	Spread of received signal strengths	127
5.3.3	Multipath variation and its correlation across channels	128
5.3.4	Multiband sensing in the real world	134
5.4	Summary	140

6	Coexistence with small scale primaries	141
6.1	Related Work	142
6.2	Can secondary users protect similar scale primary users?	143
6.3	Admission Control	153
6.3.1	Secondary Pollution Area	153
6.3.2	Location uncertainty of the primary	156
6.3.3	Location uncertainty of the secondary	156
6.4	Discussion	159
6.5	Summary	159
7	Conclusions and Future Work	161
7.1	Advances	161
7.2	Future work	162
7.2.1	Near Term	162
7.2.2	Long term	163
A	Example for Evaluating Whitespace	174
A.1	Calculating the pollution radius	174
A.2	Calculating the protection radius r_p	175
A.3	Calculating the no-talk radius (r_n)	176
A.4	Calculating the no-talk radius (r_n) for adjacent channels	176
A.5	Calculating r_n using the FCC method	177
A.6	Calculating the radius for single-threshold sensing rules	177
A.7	Incorporating multipath into the path loss model	177
B	Testbed	179
B.1	Overall Architecture	179
B.2	Antenna	180

B.3	Filtering and Gain	180
B.4	Locations and Measurement methodology	182

Chapter 1

Introduction

1.1 Motivation: The apparent scarcity problem

Wireless systems have reached a point where the next big gains in capacity will come from improving the utilization of allocated spectrum. This fact is best understood by contrasting the allocation view of spectrum to the usage view. Figure 1.1(a) shows the spectrum allocation chart of the Federal Communications Commission (FCC). All frequencies have been allocated – spectrum seems to be fully utilized and additional spectrum is hard to come by. The scarcity view of spectrum is further reinforced by the 19 billion dollar auction in the 700MHz [1]. However, Figure 1.1(b) shows the spectrum usage measured at the Berkeley Wireless Research Center over a period of ten minutes. The usage is concentrated in the cellular, ISM and TV bands. Similar measurements made at different locations in the world ([2–5]) shows extremely low levels of utilization (a range of 1% utilization to 17% utilization across different locations). The gap between allocation and usage is real and must be addressed.

So what is the reason for the large gap between allocation and use? Spectrum allocation occurs on a scale of years and continents. In order to ensure acceptable interference to wireless systems, regulators have to keep large frequency (guard bands) and spatial buffers between systems. Furthermore, since regulators do not have a mechanism to control fine grained temporal access to spectrum, spectrum is allocated for multiple years to a single entity (company/government). These buffers are overheads that have to be sacrificed to ensure that individual systems can be designed without concern for the presence of other systems – interference is controlled by design. These temporal, frequency and spatial buffers appear as unutilized spectrum. The essential problem is that the buffer overhead consumes a large portion of the available spectrum¹.

¹This is akin to a Medium Access Layer (MAC) where the signalling overhead is 95%!




Figure 1.1: FCC spectrum allocation and the measured usage at the Berkeley Wireless Research Center. The measurements were performed in the 0-2.5GHz frequencies over a period of 10 minutes. Source: Dr. Ada Poon

1.2 Mechanisms to solve the scarcity problem

Clearly the current command and control regulatory regime over-provisions spectrum and needs to be changed. There are a number of mechanisms to bridge the gap between allocation and use, all of which aim to enable spectrum sharing between the primary owners of the spectrum² and secondary users who want to use the spectrum.

Market Mechanism: Spectrum is auctioned by the primary users. Secondaries bid for spectrum and are guaranteed access to the spectrum as per the terms of the auction. For this mechanism to work the primaries have to be willing to auction their spectrum and the auction mechanism has to be fairly automated to enable fine-grained reuse.

Joint design of primary and secondary waveforms: The primary and secondary jointly design their waveforms to ensure that the primary receives appropriate quality of service while the secondary can gain some data rate. This mechanism is good for new primary users – legacy primaries may not be willing to redesign their systems.

Opportunistic spectrum use by secondaries: The secondary user can reuse the spectrum when the primary is not using the spectrum (temporal reuse) or when the secondary user is far enough from the primary (spatial reuse) to ensure that it will not interfere with the primary. The primary does not explicitly know of the presence of the secondary unless

²A primary user of the spectrum has been allocated the spectrum by the concerned regulatory authority (the Federal Communications Commission (FCC) in the United States).

it suffers interference from the secondary.

For any of these mechanisms to be implementable in practice, flexible, frequency-agile hardware is required. Such hardware must be able to operate in multiple frequency bands and employ flexible waveform generation [6].

In this thesis we will focus on opportunistic spectrum reuse. If all technical challenges are solved, this mechanism will enable us to reuse spectrum on a fine-grained level. Furthermore, this is the only mechanism that offers a purely technical solution to the ‘apparent spectrum scarcity problem’ without the need to redesign existing primary systems.

1.3 A historical perspective on Opportunistic Spectrum Access

In 2000, in his PhD thesis, Mitola envisioned cognitive radios as being capable of reasoning about their environment and user needs [7]. Cognitive radios should be able to “detect user communications needs as a function of use context” and “provide radio resources and wireless services most appropriate to those needs”. Thus a cognitive radio is able to select the best/cheapest service for a radio transmission and is even able to delay or bring forward certain transmissions depending on the currently available resources.

In 2002, the FCC spectrum Policy Task Force Report reported low spectral utilization and identified cognitive radios as a candidate to sense usage in primary user bands — the task force’s definition of cognitive radios was limited to spectral awareness [8]. As a consequence of the task force report the FCC proposed the opening up of the TV bands for opportunistic spectrum use (following the DTV transition in February, 2009) and issued a Notice of Proposed Rule making to that effect [9].

At around the same time DARPA embarked on the XG project which aimed to address the problem of having radios with regulable kernels, which can be controlled via policy rules. This allows radios to transcend regulatory borders with simple policy changes. XG radios are capable of sensing the environment to determine unused spectrum (opportunity awareness), use policy constraints to allocate these channels (opportunity allocation) among a group of radios that together form an XG-domain and to determine mechanisms for using those channels (opportunity use) [10, 11].

Finally, on November 4th, 2009, the FCC issued a notice and order to allow opportunistic spectrum usage in the TV bands [12].

1.4 Defining a spectrum hole

For opportunistic spectrum reuse to work, a secondary user must be able to reliably identify a spectrum hole. A spectrum hole is defined as a {time, location, frequency band} tuple that a secondary user can use while keeping interference to all primary systems that use the specified frequency band within an acceptable level (we will make the notion of acceptable interference precise in the subsequent chapters). Figure 1.2 figure shows the operation of two primary users (on the same frequency band) over time and space (the depiction of space is restricted to the line joining the two primary users). Both grey and yellow are time-space regions that the primaries can use when secondary users are not present. To accommodate secondary transmission the primaries need to sacrifice some time and space usage. This region is shown in yellow. The green and light blue regions together are a potential spectrum hole of which the light blue region is the one that can be recovered by the secondaries. The region in green is lost due to usage characteristics of the secondary and imperfection/budgets in the secondary's ability to reliably identify the presence of the primary³.

The emphasis of this thesis is two fold: firstly, we examine ways to specify the extent of the yellow/grey region in Figure 1.2 and secondly we will examine advanced sensing techniques to recover the region in green and light blue. A major focus of this thesis is on recovering spatial domain holes even though we shall present multiple techniques to recover temporal domain holes as well. A purely spatial domain hole occurs when the primary is always transmitting. There are a number of reasons to focus extensively on spatial domain holes:

- The first bands to be opened up by the FCC for opportunistic reuse are the Digital TV bands [12]. Digital TV are always 'on' transmitters hence only spatial spectrum holes are possible. This is further discussed in Chapter 2 and Chapter 3.
- Most research has concentrated on temporal holes since they lend themselves well to classical detection theory. Spatial holes have been left unexamined.
- Even when dealing with time domain primaries (primaries which are transmitting a fraction of the time) we need to determine the appropriate detection threshold which allows us to recover a majority of the time-space hole that is available. This viewpoint is discussed in Chapter 4 and Chapter 5.
- When the transmitting primary (which we aim to identify) and the object of a secondary's interference is the same device (on the same frequency band), we can utilize

³One can imagine a hypothetical database mechanism which accurately predicts a primary's temporal usage at every point of time and hence reduces the temporal green region to zero. Note that the spatial green region can never be reduced to zero since the secondary must always maintain a buffer between its transmitter and the primary receiver.




Figure 1.2: This figure shows the operation of two primary users (on the same frequency band) over time and space (the depiction of space is restricted to the line joining the two primary users). Both grey and yellow are time-space regions that the primaries can use when secondary users are not present. To accommodate secondary transmission the primaries need to sacrifice some time and space usage. This region is shown in yellow. The green and light blue regions together are a potential spectrum hole of which the region in light blue is the one that can be recovered by the secondaries. The region in green is lost due to usage characteristics of the secondary and imperfection/budgets in the secondary's ability to reliably identify the presence of the primary.

propagation reciprocity and deal with a simplified problem. However, in most real world problems (reusing Digital TV bands is a classic case) we need to sense the transmitter but are concerned with not interfering with primary receivers that are not co-located with the primary transmitters and may be passive listeners. Furthermore these interference and sensing regions depend on the relative scales of the secondary and primary systems, respectively. In such cases the spatial perspective emphasized in this thesis (especially Chapter 6) becomes important.

1.5 Organization of this thesis

The remainder of this thesis is organized as follows:

In Chapter 2 we will formally define a spatial spectrum hole and evaluate the ‘amount’ of spatial spectral holes available in the Digital TV bands using actual DTV transmitter location information. We first start by considering the more classical communication theory viewpoint which we call the ‘pollution’ viewpoint. The pollution viewpoint reflects the cognitive radios’ perspective: where is interference from the primary tolerable? The primary user’s viewpoint (which is the viewpoint of concern in this thesis) is called the protection viewpoint: where can a secondary operate without generating excessive ‘harmful interference’? The final available spatial spectrum hole is the intersection of these viewpoints. We will also evaluate the size of the spatial spectrum hole (or ‘white space’) resulting from the FCC’s rulings [12] (‘white space’ is the term used by the FCC for spatial spectrum holes). The interplay of these two perspectives is discussed in this chapter. Finally, this chapter also quantifies the political tradeoff between person-channels gained for potential secondary use versus person-channels lost for primary users as we vary the protection advanced to the primary users. The calculations used to arrive at the results of this chapter are presented in Appendix A. This chapter is based on papers [13, 14] and [15]. The paper [16] extends this work to calculate the resulting white space capacity available per person. The Matlab code to produce the figures in [13] is available in [17].

Chapter 2 makes us realize that a single radio spectrum sensing is unable to recover large swathes of the available spatial spectrum hole as compared to a geo-location detector. We need a mechanism to capture this loss without having to perform extensive simulations. In Chapter 3 we start by discussing the temporal spectrum hole metrics of Probability of False Alarm (P_{FA}), Probability of Missed Detection (P_{MD}) and sample complexity. We see how these metrics fail to predict a single radio’s inability to recover spatial spectrum holes. Furthermore these metrics miss some important tradeoffs between detectors. To overcome these drawbacks we propose two new metrics to capture the spectrum hole recovery potential of a detection algorithm. The first metric, namely the *Fear of Harmful Interference* F_{HI} captures the safety to the primary users. This is largely the fading-aware probability of missed detection, with some modifications to allow easier incorporation of system-level

uncertainty. The second metric, namely the *Weighted Probability of Area Recovered WPAR*, captures the performance of spectrum sensing by appropriately weighting the P_{FA} across different spatial locations. These metrics give a unifying framework to compare different spectrum sensing algorithms. We show how to obtain reasonable parameters for specifying these metrics (most crucially, something which can be interpreted as a spatial discount factor) from real-world data. These metrics are used to capture the performance of an energy detector with and without noise uncertainty. Furthermore, the performance of the FCC's spectrum sensing rules relative to geo-location is predicted by these metrics to a reasonable degree. This chapter is based on the paper [18] and overlaps partially with [19].

To overcome the issues with single radio sensing we can use cooperative sensing. In cooperative sensing, results from multiple radios are combined to make a common decision about whether to transmit or not. In Chapter 4 we will first discuss cooperative sensing using the traditional metrics. Multiple combining rules are presented – these rules differ in whether they exploit diversity or gather energy. We examine experimental data from measurements we performed in Berkeley, CA in the Digital TV band to evaluate how these rules work in practice. For recovering spatial spectrum holes, the secondary user of the spectrum must ensure that it meets the target Fear of Harmful Interference (F_{HI}) even with uncertainties in deployment, fading, and radio behavior. We examine our experimental data to determine the impact of shadowing correlation uncertainty on cooperative gains. We also examine the impact of incomplete knowledge about the fading distribution – we may only know the fading distribution to a few quantiles. The final uncertainty of concern is unpredictable radio behavior (used to model malicious/unreliable/unmodeled radios). The uncertainty aims to encompass all radio/fading behavior for which we don't have an explicit model. This chapter is based on the papers [18, 20, 21] and [22].

Chapter 2 teaches us the pitfalls of a static sensing threshold. Ideally, we would like to change the sensing threshold dynamically based on secondary's fading environment. Chapter 5 shows us how. We start with a simple example of assisted detection where two transmitters (on different frequencies) are mounted on the same tower – one is an anchor and is always transmitting while the other may or may not be transmitting (this is the frequency that the secondary wishes to reuse). In this example we show the gains from sensing on two bands (called multiband sensing) to overcome the SNR wall⁴. These gains seen in this toy example rely on the fact that shadowing is highly correlated across frequencies. Simple OR rule cooperation among radios performing multiband sensing can help weed out 'unreliable' radios and improve the secondary system's performance. These gains are achievable even if we remove the assumption of an anchor node – we replace the anchor node by multiple primaries some of which may be 'on' or 'off'. Next, we examine experimental data to determine exciting attributes of the Digital TV spectral environment - that shadowing and multipath

⁴In [48] the author has shown that the presence of noise uncertainty limits the SNR of the signal that can be detected reliably. This SNR is called the SNR wall.

spread is highly correlated across frequencies. We use these attributes to design detectors that perform much better than singleband detectors. Finally, we analyze the FCC database to determine the statistics of tower and transmitter placement. These statistics are used to develop a two-dimensional Poisson/Matèrn Hard Core model of tower and transmitter placements. This model is used to show gains from utilizing the fact that proximity to one tower can be used to infer distance from other towers. This chapter is based on the papers [23] and [24].

In Chapter 6 we discuss the problem of co-existing with small scale primaries (Wireless microphones are a good example though many low power, low height repeaters also exist). In the presence of noise uncertainty, the detection footprint of these transmitters is limited while the interference footprint of the secondaries can be large. Taking into account asymmetric fading and propagation uncertainties requires the secondary to be orders of magnitude lower in power than the primary user. When sensing a large scale primary, a small scale secondary user can make its own decision about transmission based on the sensing results from its neighborhood. This assumption fails when the scale of the primary is comparable to the scale of the secondary user. In this scenario, we need to decouple sensing from admission control - a sensor network is required to perform sensing. For small primaries, the environment over which the sensing results are valid is small which imposes certain minimum density requirements for sensor nodes. Collective sensing is used to localize the primary while a separate admission control algorithm decides on which secondaries can safely transmit. Location information of the primary and secondary users is key for such an admission control algorithm to operate successfully. In the case of a large primary, location uncertainty did not impact results significantly since decisions are made about a primary that is very far as compared to the inter sensor distances. This is no longer valid for a small primary and hence exact location information is paramount. With location uncertainty of primary and secondary users, the effective primary user footprint can increase significantly. This chapter is based on the paper [25].

Finally, in Chapter 7 we recapitulate the main themes of the thesis and discuss future steps to enhance upon the insights developed in this thesis. Appendix A uses a simple example of the KCNS tower in San Francisco to illustrate the calculations used in Chapter 2. Appendix B discusses the architecture and design of the mobile, wideband testbed that was developed to measure DTV signals in the 500-700MHz bands.

Chapter 2

Spatial spectrum hole/White space: Definition and evaluation

In this chapter we will make precise the notion of a spatial spectrum hole. We will start by defining spatial spectrum holes using the classical communication theory viewpoint (which we refer to as the ‘pollution’ viewpoint). Using this viewpoint we will evaluate the size of a spatial spectrum hole in the Digital Television (DTV) bands. The DTV bands are of particular interest since the Federal Communication Commission (FCC), in its notice and order issued on November 14, 2008, has decided to open up these bands for the operation of secondary devices [12]).

Following that, we shall define a spatial spectrum hole using the ‘protection’ viewpoint. The ‘protection’ viewpoint stipulates that a secondary radio can only operate in locations where it cannot generate harmful interference to the primary system). As with the pollution viewpoint we shall use the protection viewpoint to evaluate the size of the spatial spectrum hole in the DTV bands. Both viewpoints exclude regions around a transmitter and the actual size of the spatial spectrum hole or ‘white space’¹ is the intersection of the spatial holes from these two viewpoints.

To quantify the white space available under the ‘pollution’ viewpoint, we can set limits on the amount of acceptable interference (above noise level) that a secondary can tolerate. For the protection viewpoint, we need a similar (and consistent) notion of ‘harmful interference’ and the protection that we guarantee to the primary system. The notion advanced in this chapter is the idea of an ‘eroded fading margin’ (or protection margin) that parallels the concept of interference temperature [26]. At every distance from the primary transmitter there exists a fading margin which *protects* the receivers from bad fading events. Once a secondary starts transmission, it raises the noise floor and hence erodes some of that fading margin.

¹White space is the term used for a spatial spectrum hole in FCC parlance. In the remainder of this chapter the terms white space and spatial spectrum hole shall be used interchangeably.

Next we propose a principled way for regulators to choose the eroded fading margin (protection margin) for primaries that can be eroded by secondary operation. This approach quantifies the political tradeoff between person-channels gained for potential white-space usage versus person-channels lost for broadcasters in the DTV bands as we vary the protection margin. A person-channel is defined as one person being able to use one TV channel for broadcast or for whitespace operations. If person i can see N_i broadcast channels then the total number of person-channels available for broadcast use in the United States is $\sum_{p=1}^P N_p$, where P is the number of people in the United States.

We will also evaluate the resulting white space from fixed threshold sensing rules (these will be explained in the chapter). The conservatism of fixed threshold rules for white-space detection will be revealed. This chapter is based on papers [13, 14] and [15]. The paper [16] extends this work to calculate the resulting white space capacity available per person. The Matlab code to produce the figures in [13] is available in [17].

The main concepts and results presented in this chapter are three-fold:

- The actual white space available in the continental United States is estimated based on the ‘pollution’ viewpoint in Section 2.2 and the ‘protection’ viewpoint in Section 2.3. In defining the notion of a spatial spectrum hole under the ‘protection’ viewpoint we will introduce the concept of *eroded fading margin*. We also evaluate the available white space using a variety of detection approaches and with respect to the actual population density². Furthermore, we also look at white space availability as the scale (tower-height/power) of the primary is varied. In doing so we negatively answer the question: “Is there enough white space for another 1MW DTV transmitter?”
- Section 2.4 proposes a principled way for regulators to make the essentially political choice of the protection margin. This method is based on examining the tradeoff between the number of person-channels gained for potential white space use and those lost for potential broadcast use as the protection margin is varied. Based on estimates of broadcast use and projections for white-space market penetration, regulators can choose the appropriate margin. We will also discuss a framework to generate tradeoff plots using different models of towers/population density, propagation and detector performance.
- Section 2.5 evaluates the FCC’s sensing specification of -114dBm for ATSC signals and show how a ‘one size fits all’ rule of this nature has to be conservative and cannot recover most white space. This further validates the spatial perspective discussed in this thesis.

²All data files, Matlab code and automated scripts to generate the tables and graphs in this chapter can be found at [17].

2.1 Prior work in estimating available white space

There has been prior work in estimating the amount of white space available. The general problem is that both the language and methodology used by previous work does not properly distinguish between the pollution and protection viewpoints and this is the source of much confusion. In [27], the author estimates the average amount of white space available per person to be 214MHz. This is based on the fact that the FCC estimates that the average American can receive 13.3 channels (~ 80 MHz). Hence the authors contend that remaining spectrum is white space (~ 214 MHz). This line of reasoning tends to overestimate white space availability. For example, the FCC website [28] reveals that Berkeley, CA (Zip Code 94704) can receive 23 DTV channels. This would imply that the remaining 24 channels are available for white space usage. However, if we calculate the available white space in Berkeley, CA, only 5 channels are available for white space use when the FCC's white space rules are applied. This is because of multiple reasons: the FCC rules extend protection to adjacent channels and requires a no-talk radius which is larger than the Grade-B protected contour [12]. Furthermore, [28] only consider the high power DTV stations; low power TV stations and TV booster stations are ignored, but these must also be protected.

In [29], the authors claim to use the actual population data to quantify the amount of white space available under different scenarios. The authors extract transmitter details from the FCC database (they did not use the High Power DTV transmitter list available from the FCC) as was done to prepare the plots in [30]. However, the authors assume that all locations beyond a station's Grade B contour can be used as white space. The authors have estimated the amount of white space under different scenarios. For scenario X (all DTV, Class A stations and TV translators; co-channel rules only) [29] estimates the median bandwidth to be ~ 180 MHz while our estimate of the median bandwidth per person is 126MHz. Similarly for scenario Z (all DTV, Class A stations and TV translators; co-channel and adjacent channel rules) [29] estimates the median bandwidth to be ~ 78 MHz while our estimate of the median bandwidth is ~ 36 MHz. This major discrepancy can be best understood by a deeper inspection of the author's use of population density data. The authors estimate the size of each census block to be around 16 square miles. Each census block contains an average of 1300 people. For each transmitter the number of people in its Grade B region can be estimated by taking the number of census blocks that fit into its Grade B contour times 1300. While this may sound reasonable, the effective consequence is that the author's analysis assumes a uniform population density across the United States! For a uniform population density, for scenario X (all DTV, Class A stations and TV translators; co-channel rules only) our estimate of the median bandwidth per person is ~ 186 MHz while that for scenario Z (all DTV, Class A stations and TV translators; co-channel and adjacent channel rules) it is ~ 102 MHz. These numbers closely resemble the results in [29] (the slight discrepancy in this case is due to the fact that [29] uses the FCC database which yields a much higher number of towers – 12339 versus 8071).

New America Foundation has also estimated the amount of white space available in major cities [31]. This study also, overestimates the amount of white space available (for example, they assume that 19 white space channels will be available in San Francisco). Since the methodology for computing available white space has not been described, we cannot explain the discrepancy between our work and [31].

2.2 White space using the pollution viewpoint

The pollution-oriented view of white space takes the perspective of a rational self-interested cognitive radio — a band is attractive when it has a low noise-floor. Implicit in this view is the idea that primary user transmissions are considered as raising the noise floor from the point of view of the secondary user. Each television tower can thus be viewed as having a pollution radius around it in which the band is unattractive for secondary use. It must be pointed out that there exist theoretical approaches where a secondary transmitter would prefer to be closer to the primary transmitter so as to decode the primary signal and use dirty-paper-coding techniques (DPC) and simultaneously boost the primary signal in the direction of interference [32, 33]. However, it has also been shown that this approach is not robust since simple phase uncertainty can significantly lower the performance of such schemes [34]. Knowledge of the primary user's codebook is useful if the secondary receiver can actually decode the primary signal and use multiuser detection. Otherwise, it has been shown that the secondary system is forced to treat the primary transmission as noise [35].

To calculate the available white space using the pollution viewpoint, we used the FCC's transmitter databases which list the latitude, longitude, effective radiated power (ERP) and transmitter height above sea level for all licensed high power and low power transmitters [36, 37]. These databases were combined with the ITU-R recommendations on propagation (ITU-R P-1536-3) to calculate the values of the pollution radius for all towers.

Appendix A uses the example of the KCNS transmitter on TV channel 39 to illustrate the procedure used to calculate the value for the pollution radius. A location is considered available for white space operation in a given TV channel if it does not fall within the pollution radius of any transmitter of that channel. Table 2.1 quantifies the available white space as the average number of white space channels available per location for various values of the acceptable pollution level. The 15dB pollution level was inferred from current IEEE 802.11g systems; a survey of IEEE 802.11g systems shows that raising the interference level to 15dB above noise only deteriorates the data rate from 54Mbps to 12Mbps [38]. The 5dB pollution level was calculated by analyzing the business case for Wireless ISPs for rural areas³.

³Estimated annual costs for each wireless ISP tower is \$22000. For a 50% profit margin per tower and a monthly subscription rate of \$25 per month we need ~ 147 installations $((\$44000/12)/\$25)$. The average rural population density in the US is around 6.6 people per km^2 [39]. Assuming 4 people per family and

Table 2.1 breaks down the TV channels into four frequency bands: Low VHF (LVHF⁴ - channels 2, 5 and 6), High VHF (HVHF - channels 7 to 13), Low UHF (LUHF - channels 14 to 51) and High UHF (HUHF - channels 52 to 69). As a fraction of the total number of channels in a band, there is significant white space available in the low VHF bands which runs counter to our intuition that propagation at low frequencies is rather good. The main reason for this is the remarkably low number of DTV stations in the low VHF bands – after the DTV transition, there will be only 31 high power DTV stations in channels 2, 5 and 6. Similarly, there is significant white space in the high UHF bands since there are no high power TV stations in these bands. However, these bands cannot be used for white space devices since they have been reallocated to uses other than television. From a white space recovery standpoint, the main band of interest is the low UHF band – for these channels and a tolerable interference level of 5dB, the average number of white space channels per location is ~ 16 which is substantial.

2.2.1 Adjacent channel considerations in the pollution viewpoint

Another consideration for secondary operation is interference received from TV transmitters on adjacent channels. For example, if a secondary device operates on channel 39, it would not want to be too close to TV transmitters in channels 38 and 40. The amount of interference that can be tolerated on adjacent channels depends on the characteristics of the secondary user's receiver (receiver filter etc). As an illustration, suppose that the channel-select filter can attenuate adjacent channels by 20dB or 40dB. The impact can be seen in Table 2.1. For the lower UHF channels, with a tolerable adjacent channel pollution level of 45dB above noise and a tolerable co-channel pollution of 5dB, the average number of white space channels per location drops from 16.2 (assuming perfect channel-select filters) to 15.6.

2.2.2 White space calculations using actual population density

Another way to evaluate available white space is to consider the number of people that are potential users of white space devices. This number is more valuable for businesses to determine revenue potential and also for regulators interested in the public good. To calculate available white space in terms of population, we used the US Census data of 2000 which lists the population density per zip code [40]. Furthermore, the region occupied by a zip code is specified as a polygon [41]. Using this data we can determine the average number of white space channels available per person. For low UHF channels and a tolerable pollution level of 5dB, the average number of white space channels is ~ 9 which is roughly half of

a 50% penetration, we need to service a area of $178km^2$ i.e. service radius 7.5 km. At this distance, the secondary received power is -94dBm. For a minimum of 1Mbps service the operational SNR for 802.11g is 7dB, hence we can tolerate -101dBm of interference which is 5dB above the noise floor.

⁴3 and 4 are excluded because people use them to connect their VCRs to their televisions. [12]

the number assuming a uniform population density across the country. This is because TV transmitters are constructed to serve areas of high population density and hence excluding areas around TV towers significantly reduces the population that can be served by white space devices. This observation, however, is not true for the low VHF channels. For the low VHF channels the amount of white space increases as we move from a uniform population density assumption to the actual population density. This is because the high power DTV stations (after the June '09 transition) in channels 2, 5 and 6 are located in areas of low population density – the actual population density around these regions is almost half the uniform population density.

The impact for the UHF channels is even more severe when interference from adjacent channels is considered – for a main channel tolerable pollution level of 5dB and an adjacent channel pollution level of 45dB, the average number of white space channels per person is 5.75. This is a third (15.6 to 5.75) of the numbers obtained assuming a uniform population density.

2.2.3 Limitations of this evaluation

There are a few limitations of this study that should be pointed out. Firstly, to calculate the available white space we assumed that the all the licensed transmitters in the FCC high power DTV transmitters database [36] and the master low power transmitter database are all transmitting [37] and neglected some of the clauses from the FCC ruling ⁵. If some of the transmitters are off, then sensing rules can help reclaim the additional white space. Secondly, we assume that the ITU propagation models predict the reality on the ground to a fair degree [42]. Thirdly, we are overestimating the number of people served by broadcasters today by assuming that everyone in the noise-limited contour can and does receive a TV signal.

2.3 White space using the protection viewpoint

In the previous section we have seen a spatial spectrum hole defined using the pollution viewpoint. In this section we will define a spatial spectrum hole using the protection viewpoint.

⁵We neglected the more stringent emission requirements for the 602-620MHz bands (Section 15.709 [12]). We also neglected the locations of cable headends, fixed broadcast auxiliary service (BAS) links, and PLMRS/CMRS devices (Section 15.712 [12]). We believe that the clauses we are neglecting will result in only a minor change in the results.

<i>Interference level above noise</i>	<i>By Area</i>				<i>By Population</i>			
	LVHF	HVHF	LUHF	HUHF	LVHF	HVHF	LUHF	HUHF
5dB Main channel only	1.69	1.74	16.3	15.9	1.83	1	8.71	15.5
5dB Main channel and 25dB adjacent channels	1.54	1.02	11.9	14.9	1.55	0.355	3.88	13.5
5dB Main channel and 45dB adjacent channels	1.64	1.49	14.9	15.7	1.69	0.61	5.75	14.4
10dB Main channel only	1.94	2.59	20	16.4	2.08	1.68	11.8	16
10dB Main channel and 30dB adjacent channels	1.8	1.69	15.5	15.7	1.83	0.67	5.87	14.2
10dB Main channel and 50dB adjacent channels	1.9	2.3	18.6	16.3	1.96	1.09	8.24	15.1
15dB Main channel only	2.16	3.43	23.3	16.8	2.29	2.42	15	16.3
15dB Main channel and 35dB adjacent channels	2.04	2.5	19.1	16.3	2.07	1.07	8.18	14.7
15dB Main channel and 55dB adjacent channels	2.13	3.17	22.2	16.7	2.17	1.65	11.2	15.6

Table 2.1: Comparison of available white space based on tolerable pollution levels for various frequency bands (Low VHF (LVHF - channels 2, 5 and 6), High VHF (HVHF - channels 7 to 13), Low UHF (LUHF - channels 14 to 51) and High VHF (HVHF - channels 52 to 69)). White space is measured as the average number of white space channels available per location/person in the continental United States.

2.3.1 Defining white space for the protection viewpoint

Co-located primary transmitter/receiver, TDD mode

First, let us start with a very simple case: the secondary radio wishes to coexist with a primary system consisting of a co-located transmitter-receiver pair operating on the same frequency (for example, a cellular system in Time Division Duplex (TDD) mode) and a second transmitter-receiver pair located at a certain distance away. Without fading the noise limited radius r_{nl} is the maximum distance between the communicating pair at which the target (δ) SINR is maintained at the receivers i.e. for all distances $r \leq r_{nl}$,

$$P_t - l_{pp}(r) - N_0 \geq \delta$$

where P_t is the transmit power in dBm, $l_{pp}(r)$ is the nominal path loss between the primary transmitter and the primary receiver separated by a distance r (in dB), N_0 is the noise level in dBm and δ is the target SINR in dB.

At every distance between the communicating pair, there exists a fading margin ($\psi(r)$) which ‘protects’ the receivers from bad fading events.

$$\psi(r) = P_t - l_{pp}(r) - N_0 - \delta \quad (2.1)$$

Once a secondary user starts transmission it raises the noise floor. Regulators/system designers must decide the amount of the protection margin ψ_t that they are willing to erode to enable secondary operation. However, corresponding to the selection of ψ_t is a protected distance r_p where the original fading margin was equal to ψ_t . With cognitive radio transmission, this margin is reduced to zero but if the communicating pair keep within r_p they can still communicate.

The cognitive radio can transmit if it does not degrade the SINR below δ . If P_s is the transmit power of the secondary, then the secondary can transmit at a distance r from the primary receiver if:

$$P_t - l_{pp}(r_p) - 10 \log_{10} \left(10^{\frac{P_s - l_{sp}(r)}{10}} + 10^{\frac{N_0}{10}} \right) \geq \delta \quad (2.2)$$

$$\psi_t + N_0 - 10 \log_{10} \left(10^{\frac{P_s - l_{sp}(r)}{10}} + 10^{\frac{N_0}{10}} \right) \stackrel{(a)}{\geq} 0 \quad (2.3)$$

where P_s is the secondary transmit power, $l_{sp}(r)$ is the path loss from the secondary transmitter to the primary receiver. Inequality (a) results from the application of (2.1) and by replacing $\psi(r_p)$ by ψ_t . How does a cognitive radio determine the value of $l_{sp}(r)$? In this particular case channel reciprocity comes to the rescue. The path loss between the primary transmitter and the secondary receiver is the same as the path loss between the secondary transmitter and the primary receiver i.e. $l_{sp}(r) = l_{ps}(r)$. If the secondary knows the power

of the primary transmitter (P_t) and can reliably estimate the primary's received power P_r then it can calculate l_{ps} as $P_t - P_r$. The final spatial spectrum hole or white space consists of all locations where equation 2.3 is true for both receivers in the communicating pair.

Three interesting aspects emerge in the above example:

- A spatial spectrum hole is defined completely on the basis of signal powers – the secondary does not need to know the actual distance from the transmitters.
- Sensing/detection of the primary signal is naturally required in this definition of a spectrum hole.
- If signal reciprocity holds, interference at the receivers can be accurately calculated by estimating the received primary power at the cognitive radio.

Co-located primary transmitter/receiver, FDD mode

If the system is Frequency Division Duplexed (FDD) (the two transmitter-receiver links operate on different frequencies) then $l_{sp}(r)$ cannot be estimated completely from $l_{ps}(r)$ and we must budget for differences in fading between the two frequencies. In such a setup sensing/detection can still be performed but interference can only be guaranteed probabilistically.

Non co-located primary transmitter/receiver

Finally we arrive at the case where receivers are not co-located with the transmitters and their locations are unknown. This is the setup in Digital TVs. Figure 2.1 illustrates this scenario. The noise limited (r_{nl}) coverage area is the broadcast coverage area with no interference from other transmissions. Within this area, the nominal TV SNR is greater than the target SINR (δ), *i.e.* for all distances $r \leq r_{nl}$ ($P_t - l_{pp}(r) - N_0 \geq \delta$). As in the previous scenario, at every distance from the TV transmitter there exists a fading margin ($\psi(r)$) which ‘protects’ the TV receiver from bad fading events ($\psi(r) = P_t - l_n(r) - N_0 - \delta$). As described earlier, corresponding to the selection of ψ_t is a protected radius r_p where the original fading margin was equal to ψ_t . With cognitive radio transmission, this margin is reduced to zero but all TV receivers within r_p are still nominally protected – they can still receive the TV signal by positioning their antennas appropriately.

$$r_p = l_{pp}^{-1}(P_t - \psi_t - N_0 - \delta)$$

It is important to note that there is very little correlation between l_{ps} and l_{sp} in this case. Since we don't know the position of the receivers we have to assume that a receiver can exist at any location within r_p (in the worse case all receivers are at r_p). The protected radius together with the transmit power/height of the secondary transmitter can then be used to calculate the no-talk radius for cognitive radio operation (r_n). The cognitive radio can only

transmit if it is outside the no-talk area of the primary transmitter in question. To calculate this we must first calculate the value of $r_n - r_p$ *i.e.* the distance beyond the protected radius where the secondary can safely transmit. In other words we need to determine the distance r_n such that a transmission from r_n results in a SINR of δ at r_p . We started by calculating the allowable interference level at r_p (I_{r_p}) as:

$$I_{r_p} = 10 \log_{10} \left(10^{\frac{P_t - l_{pp}(r_p) - \delta}{10}} - 10^{\frac{N_0}{10}} \right) \quad (2.4)$$

Next we calculate $r_n - r_p$ as:

$$r_n - r_p = l_{sp}^{-1}(P_s - I_{r_p})$$

where P_s is the transmit power (in dBm) of the secondary user, I_{r_p} is the maximum tolerated interference power (in dBm) from (2.4) and $l_{sp}(r)$ is the optimistic path loss from the secondary to the primary.

For each 6MHz TV channel, ‘white space’ can be defined as all area that is not within any relevant tower’s no-talk radius. Note that the lack of location information of the receivers forces us to adopt a distance based definition of a spatial spectrum hole where sensing/detection of the transmitter power does not play any role (except in acting as a proxy for distance as discussed later).




Figure 2.1: Definition of spectrum ‘white space’ from a primary protection perspective. The ‘extra no-talk area’ is the extra space where there are no protected primary users and yet secondary users are not permitted to transmit.

It is worth noting that the idea of the eroded fading margin is in line with the idea of interference temperature proposed in [8, 26, 43]. Any such approach requires a propagation model from the secondary transmitter to the primary receiver. The difference here is that

the original interference temperature idea sought to limit the amount of extra interference at all primary receivers while the eroded fading margin here is also used to determine *which* receivers are deemed protected. The key difference is that for receivers that are not protected, we are permitted to raise the potential interference as high as we would like. Notice however that moving from the eroded margin to the protected receivers requires us to know more about the primary users' characteristics: we must know the propagation model between the primary transmitters and receivers (i.e. deployment model), the height/power of the primary transmitter, *and* the minimal required SINR δ (i.e. modulation characteristics) for the primary system. It is this additional knowledge about the primary users that makes this whitespace approach to cognitive spectrum use possible.

2.3.2 Calculating the available white space from the protection viewpoint

As in Section 2.2, we used the FCC's transmitter databases and the ITU propagation curves to calculate the white space [42,44]. Appendix A.2 uses the example of the KCNS transmitter on channel 39 to illustrate the procedure used to calculate the values for r_p . We assumed that the primary transmitter to primary receiver path loss ($l_{pp}(r)$) is the ITU's F(50, 90) path loss *i.e.* the protection radius (r_p) is the distance where the SNR is δ for 50% of the locations, 90% of the time. To determine the maximum possible white space available we assumed a hypothetical 'zero Watt' transmitter which only causes harmful interference if it is right next to a TV receiver *i.e.* for this hypothetical transmitter $r_n - r_p$ is defined to be zero. Table 2.2 shows the white space available for such a transmitter for various values of the eroded fading margin ϕ_t . About 19 channels per person are available on average for white space use in the lower UHF bands even with a erosion of ~ 1 dB of protection margin. (A 1dB erosion in protection margin translates into an effective 20% reduction in the transmit power from the link-margin perspective of a primary user.)

Next, Table 2.3 calculates the white space for a 4W⁶ with a height above average terrain (HAAT) of 30m: both are the maximum allowed by the FCC for fixed white space devices [12]. For this secondary user, the value of $r_n - r_p$ is calculated using the procedure outlined in Appendix A.3. We assume that the optimistic path loss between the secondary transmitter and the primary receiver ($l_{sp}(r)$) is the ITU's F(50, 10) path loss *i.e.* the secondary has to be at a distance ($r_n - r_p$) beyond the protection radius to ensure that the signal would be higher than I_{r_p} for 50% of the locations only 10% of the time. This ensures that secondary users beyond r_n only interfere for at most 10% of the time. With this scale of a secondary user, the average number of white space channels per person is reduced by 4 to 14.9.

⁶ [12] actually specify a 1W maximum transmit power with a 6dBi antenna which translates (in a worst case scenario) into 4W transmit power towards the primary.

<i>Margin</i>	<i>Detection Scheme</i>	<i>By Area</i>				<i>By Population</i>			
		LVHF 2,5,6	HVHF 7-13	LUHF 14-51	HUHF 52-69	LVHF 2,5,6	HVHF 7-13	LUHF 14-51	HUHF 52-69
<i>0.1dB</i>	Geolocation	2.5	4.67	26.7	16.9	2.55	3.64	18.3	16.4
	Geolocation with adjacent channels	2.24	2.58	16.6	15.1	2.22	1.16	6.39	13.6
<i>1dB</i>	Geolocation	2.52	4.75	27	17	2.56	3.71	18.7	16.4
	Geolocation with adjacent channels	2.26	2.69	17.1	15.3	2.24	1.22	6.74	13.7
<i>2.73dB</i>	Geolocation	2.55	4.9	27.8	17.1	2.58	3.84	19.4	16.5
	Geolocation with adjacent channels	2.31	2.89	18.2	15.5	2.28	1.35	7.43	13.9

Table 2.2: Maximum possible white space using a hypothetical “zero Watt” transmitter for various values of the protection margin that is eroded. The 1dB (2.73dB) protection margin is the average (maximum) protection margin eroded across all towers using the FCC’s definition of harmful interference. For adjacent channels, it is assumed that the acceptable interference level is 27dB more than what is acceptable co-channel.

<i>Environment</i>	<i>Detection Scheme</i>	<i>By Area</i>				<i>By Population</i>			
		LVHF 2,5,6	HVHF 7-13	LUHF 14-51	HUHF 52-69	LVHF 2,5,6	HVHF 7-13	LUHF 14-51	HUHF 52-69
<i>No Multipath</i>	Geolocation	1.52	2.85	22.4	16.2	1.71	2.11	14.9	15.8
	Geolocation with adjacent channels	1.39	1.55	14.1	14.5	1.46	0.64	5.37	13.2
<i>Multipath</i>	Geolocation	1.56	2.73	21.9	16.1	1.75	2	14.4	15.8
	Geolocation with adjacent channels	1.42	1.47	13.8	14.4	1.48	0.609	5.2	13.2

Table 2.3: Comparison of the average number of white space channels available per user in various environments and frequency bands for a 4W secondary transmitter with a HAAT of 30m and a 1dB erosion of the fading margin of primary receivers.

2.3.3 Impact of Secondary power on white space availability

As seen earlier, the available white space is determined by the power control rule selected by the secondary – we lose 4 white space channels as we go from a hypothetical ‘zero watt’ transmitter to a 30m high 4W transmitter. How does white space availability scale with the transmit power of the secondary? Figure 2.2 helps us answer this question. If we assume that the secondaries are 1MW transmitters (90dBm), then there is not much white space available to accommodate them. On the other hand, portable (100mW) and fixed white space devices (4W) have plenty of white space available. All calculations assumed a transmitter HAAT of 30m. The white space availability is significantly less at higher transmit powers if we assume that transmitter height scales with the transmit power (a 1MW transmitter would sit on a higher tower).

2.3.4 Available white space considering adjacent channels

The FCC is concerned that the operation of cognitive radios in a given band will significantly impact TV receivers in adjacent bands. Hence they require that a cognitive radio should only transmit in a given band if TV receivers in the channel and in the adjacent channels can be protected. Consider a TV receiver on channel 39 attempting to receive the KCNS transmission from Sutro tower. Assume channel 40 is empty and being used by the secondary user. The first source of concern is that the secondary signal’s ‘skirts’ falling in channel 39 itself will harm the TV receiver. This can be dealt with by engineering the secondary transmission ‘skirts’ to an acceptable level.

The second source of concern is that the secondary’s transmission in channel 40 itself will impact the TV receiver. Each TV receiver has a requirement for the interference level in adjacent bands which depends on the characteristics (receive filter, etc.) of the TV receiver. This cannot be altered and hence must be budgeted for. To determine the impact of this rule on available white space, we first determined the no-talk radius for adjacent bands (this procedure is detailed in Appendix A.4). We considered a location/person to be available for white space operation in a given TV channel only if it is out of the no-talk area of all transmitters in its channel as well as out of the adjusted no-talk area of all transmitters in adjacent channels. This rule significantly reduces the available white space as can be seen in Tables 2.2, 2.3 and 2.4. For example, for the low UHF bands the maximum number of average white space channels per person drops from around 19 to around 7 assuming a 1dB tolerable erosion, dropping further to only about 5 for a 4W transmitter at 30m HAAT.

A way to avoid the problem of adjacent channels is to position the secondary transmission away from the channel boundary. This can be achieved by finding contiguous bands of open channels and then positioning the transmission in the center of this contiguous frequency band. Figure 2.3 shows the fraction of the population that can harness contiguous channels of a given length. It is interesting to note that the fraction of people that can find two




Figure 2.2: Impact of the power control rule for the secondary user. If we assume secondary users are 1MW transmitters (90dBm), then there is not much white space available to accommodate them. On the other hand small scale primaries (4W) have plenty of white space available. (all calculations assumed a transmitter HAAT of 30m). The maximum number of channels is 47 (channels 2 to 51 excluding 3,4 and 37).

free contiguous channels (without concerns for adjacent channels) is higher than the fraction that finds a single free channel while considering adjacent channels (See Figure 2.3 – the fraction of the population that can find 2 contiguous empty channels with only co-channel considerations is 0.97 while only 94% of the population can use one channel with adjacent channel considerations.). This suggests that it may be worthwhile to search for 2 contiguous channels and then create a waveform in the center of the 12MHz band.




Figure 2.3: Investigation of the number of contiguous channels available in different frequency bands (Lower VHF, Higher VHF, and Lower UHF) using different protection rules (a) Co-channel interference only (b) Considering adjacent channels as well.

2.3.5 Available white space considering multipath fading

The ITU-R propagation model does not incorporate multipath fading. To study the impact of multipath we created new path loss tables for different multipath scenarios (See Appendix A.7). To remain conservative we neglected multipath for calculating r_p . For calculating $r_n - r_p$ we assumed Rayleigh fading. The impact of multipath is to increase $r_n - r_p$ which reduces the available white space.

2.4 A principled way to choose the protection margin

As we add the potential for using white-space devices for some people, we potentially sacrifice the utility for other people who have been watching over-the-air broadcast television. This

tradeoff can be changed by modifying the protection margin offered to TV users (ψ), but it is essentially a political tradeoff since it involves weighing and balancing the interests across different parties. As we adjust the margin we change the protection radius r_p and the no-talk radius r_n . Figure 2.4(a) shows the average number of channels available for white space usage per user as the erosion margin is varied. With zero eroded margin, there is no white space. However, a zero margin also does not mean that all users can view all possible broadcast television channels – the average number of reliable broadcast channels per person is limited to 23 channels. As we increase the eroded margin, the number of broadcast channels diminishes slightly but there is a huge gain in the number of white space channels per user.

This tradeoff is better studied by examining the cumulative and marginal (instantaneous) gain-loss tradeoff curves. Let P_m^W be the average number of person-channels⁷ available for white space usage at an eroded margin of m dB. Similarly let P_m^B be the number of person-channels available for broadcast television usage at a margin of m dB. Then the red and blue curves in Figure 2.4(a) are the values of P_m^W and P_m^B respectively for different values of m . For a given margin, $\frac{P_m^W - P_0^W}{P_m^B - P_0^B}$ represents the cumulative gain-loss tradeoff for each margin setting. The values assuming a uniform population distribution and the actual population density are shown in Figure 2.4(b) as the solid and dotted black curves respectively. Further, $\frac{dP_m^W/dm}{dP_m^B/dm}$ is the marginal (instantaneous) gain-loss tradeoff at an eroded margin of m dB. The values assuming a uniform population distribution and the actual population density are shown in Figure 2.4(b) as the solid and dotted red curves respectively. These instantaneous tradeoff curves tell us the number of person-channels gained for potential white space use versus those lost for broadcast television viewing by varying the erosion margin settings around m .

At small erosion margin ($\sim .1dB$) the tradeoff is skewed towards white space usage – we gain a lot of white space person-channels for a few broadcast person-channels lost. For instance, if we increase the erosion margin infinitesimally beyond 0.1dB, we gain 20 white space person-channels for every broadcast person-channel lost. This unequivocally provides not just the motivation for allowing white space usage, but also provides regulators with a way to choose the tradeoff point by using estimates of social utility and market penetration.

We can repeat this exercise by considering the margin of protection as a fraction of the original noise-limited broadcast area that we are giving up. Let r_{nl} represent the noise limited radius of operation and r_p be the protected radius of operation, then we can measure the area margin eroded as the fraction $1 - (\frac{r_p}{r_{nl}})^2$. With this definition of the margin we can retrace the trade-off curves as shown in Figures 2.4(c-d).⁸

⁷If the total population is P and N_p is the number of channels that can be received by person p , then $\frac{1}{P} \sum_{p=1}^P N_p$ is the average number of person-channels available per person.

⁸The marginal (instantaneous) person-channels gained versus those lost shown in Figure 2.4(d) is a polynomial fitted version of the actual. This fitting procedure is robust for $r_p \geq 0.02$.

Due to variations in tower power and height, a single erosion margin translates into a range of protected radius values. The range spread and medians (across towers) of the area lost can be seen in Figure 2.5. The histogram of r_p values shows a bimodal distribution whose median increases with the margin. The spread in the area lost also increases with the margin as can be seen from the divergence of the 90th and 10th percentiles.

2.4.1 Impact of adjacent channels

Figure 2.6(a), shows the impact of adjacent channels on the average number of white space channels per person. Broadcast use does not change but the number of white space channels available is much lower when adjacent channels are considered. However as shown in Figure 2.6(b), the increase in the number of white space channels per margin increase is much larger when adjacent channels are considered. This is because increasing the margin reduces the protected radius of co-and-adjacent channels which greatly increases the number of people able to utilize a particular channel for white space use.

2.4.2 Opportunistic spectrum use versus reallocation

At this point we can ask the question, it is better to facilitate white space use using opportunistic spectrum use (where the secondary senses for the presence of the primary) or to reallocate spectrum between broadcast and white space use. Figure 2.7 compares the engineering and political tradeoffs of opportunistic spectrum access to those of a spectrum reallocation scheme. In the reallocation scheme, we assign channels to white space use in increasing order of their current broadcast use (i.e. channels with the lowest broadcast use are reassigned first). In Figure 2.7(a) the average number of broadcast channels per person is plotted against the average number of white space channels per person. The total number of channels per person (broadcast+white space) is higher when using opportunistic access than reallocation up to the point of 12 (72MHz) white space channels (this cross over point can be extended by using a hybrid scheme where we reallocate 2-3 channels and make the rest available for opportunistic use). A similar trend is visible in Figure 2.7(b) where we plot the the number of people-channels gained for white space use per every person-channel lost for broadcast use against the average number of white space channels available per person. At the FCC's operating point of 9 white space channels per person, opportunistic access gives a gain-loss ratio of 7:1 while reallocation gives a ratio of 4:1 (the discrepancy between the gain-loss predicted by using the eroded protection margin in Figure 2.6(b) and here (11 versus 7) is due to the fact that we are using the median protection margin as a representative number which is dominated by the UHF channels. For VHF channels the actual margin is closer to 2.7dB which would reduce the tradeoff considerably).




Figure 2.4: (a) Average number of channels available for white space/broadcast usage per person as the protection margin (in dB) is varied. (b) The cumulative/marginal (instantaneous) gain-loss tradeoff as the protection margin is changed. (c) Average number of channels available for white space/broadcast usage per person as the margin in terms of relative area lost is varied (d) The cumulative/marginal (instantaneous) gain-loss tradeoff for white space usage as the relative area lost is varied. The light blue and dark blue lines at the top of (b) and (d) represents the range of margins achievable under a combined pollution/protection viewpoint.






Figure 2.5: Variation in the broadcast area lost as the protection margin is varied.



(a)



(b)

Figure 2.6: (a) Average channels per person (using the actual population density) as the margin is varied. The figure contrasts the average number for the case when adjacent channels are neglected to the one in which adjacent channels are considered. Taking adjacent channels into account reduces the number of available channels for white space usage (b) The cumulative/marginal (instantaneous) gain-loss tradeoff as the protection margin is changed. Again, the figure contrasts cumulative/marginal tradeoff when the adjacent channels are neglected to the one in which adjacent channels are considered. The marginal tradeoff with adjacent channels starts worse off but quickly overtakes the marginal tradeoff without adjacent channels.




Figure 2.7: The engineering and political tradeoff of opportunistic access is compared to a spectrum reallocation scheme. In the reallocation scheme, we assign channels to white space in increasing order of their current broadcast use (i.e. channels with the lowest broadcast use are reassigned first). In (a) the average number of broadcast channels per person (using the actual population density) is plotted against the average number of white space channels per person while in (b) the number of people-channels gained for white space use per every person-channel lost for broadcast use is plotted against the average number of white space channels per person. To provide less than 12 white space channels per person (72MHz of spectrum) opportunistic communications is preferred.

2.4.3 Impact of the pollution viewpoint

The combined impact of the protection and pollution perspective is studied in Figure 2.4. Depending on the chosen value for the erodible margin and the tolerable pollution level, the perspective that dominates white space availability can change. In Figure 2.4(a) the average number of channels per person for white space usage increases as the erosion margin is increased, but saturates after a certain point where the pollution viewpoint starts to dominate. Physically, this means that the additional white-space opened up is not considered to be worth using. The same goes when we consider the margin in terms of relative broadcast area lost (Figure 2.4(c)). This limit on the achievable tradeoff can also be viewed in Figures 2.4(b) and (d) where the tradeoff points achievable by the pollution viewpoint are marked out. Increasing the erosion/area margin beyond these points is not advisable since the additional white space recovered is too polluted to be worthwhile.

Since the pollution viewpoint tends to limit the number of usable white space channels, there is an optimal margin beyond which increasing the margin leads to diminishing returns. Figure 2.8 shows the number of useful channels (white space + broadcast) minus the number initial broadcast channels (broadcast channels per person at a margin of 0) for various values of the margin. This quantity attains a maximum at a particular margin setting which is a function of the amount of pollution that the white space device can tolerate.

Another way to determine the impact of the pollution viewpoint is to consider the overhead (in terms of channels unutilized) due to pollution and protection as the margin is varied. Figure 2.9(a) and 2.9(b) show this overhead for a tolerable interference of 15dB above noise and 5dB above noise, respectively. The co-channel protection overhead is the value of $r_n - r_p$ which shrinks rapidly as the margin is increased. The co-channel pollution overhead increases with the margin since the area outside r_n cannot be used due to pollution requirements. The overhead due to adjacent channels works in the other way. In this case the protection overhead has a larger impact than the pollution overhead.

2.5 Impact of FCC rules on white space availability

The FCC has developed a slightly different approach to determine the no-talk radius around a tower. They assume that the protected radius is the same as the Grade B contour and the interference level at the protected radius should be 23dB lower than the signal level. Furthermore, the FCC specifies the value of $r_n - r_p$ as a distance that should be used irrespective of the channel frequency (See Section 15.712 in [12]). The details of using this methodology to calculate available white space is given in Appendix A.5. Table 2.4 shows the white space available using the FCC style definition of the no talk radius. These calculations were performed for a 4W, 30m (HAAT) secondary transmitter. Comparing Table 2.4 with Table 2.3 we can see that the FCC style rules predict very similar white space values for the lower UHF bands but not for the VHF bands. For the values of r_n and r_p calculated






Figure 2.8: The total number of useful channels (white space + broadcast) minus the number of initial broadcast channels. This quantity attains a maximum at a particular margin settings which is a function of the amount of pollution that the white space device can tolerate.



(a)



(a)

Figure 2.9: Number of useful channels per person (broadcast and white space) and the amount of overhead due to protection and pollution as the protection margin is varied. (a) For a tolerable interference of 15dB above noise (b) For a tolerable interference of 5dB above noise.

using the FCC method, the average margin eroded in the low VHF bands is $\sim 2.73\text{dB}$ while for the high VHF bands it is $\sim 2.3\text{dB}$ ⁹. Since white space increases with increasing margin (See Table 2.2) the FCC method shows more white space in these bands as compared to the margin method. For the low UHF bands the margins used in both methods are comparable and hence the white space numbers are similar.

Figure 2.10 shows the number of white space channels at each location (as per the FCC's rules) in the continental United States using a color coded system. The color black is used for locations where no white space channels exist. Interestingly these are locations on hill tops where multiple transmitters and relays are mounted. The color red is used to denote locations where a single white space channel exists while yellow denotes areas with two white space channels. Most metros see between 3-8 white space channels (these areas are marked in shades of blue).

2.5.1 Spectrum sensing: SNR as a proxy for distance

For the discussion in Section 2.3, we see that one of main tasks of the secondary system is to determine its position with respect to the primary transmitter and to start transmission only if it is reasonably sure that it is outside the no-talk radius (r_n). One of the methods that the secondary can use for determining its distance from the primary is to estimate the primary's received signal strength and use it as a proxy for distance from the primary transmitter. So the problem becomes: at what level (the detection sensitivity) must the secondary detect the primary system to be reasonably sure that it is outside the no-talk radius? If P_t is the transmit power of the primary and α is the attenuation exponent, then the secondary can transmit if the received power from the primary at the secondary is less than $P_t r_n^{-\alpha}$ i.e.

$$P_r \underset{\text{use}}{\overset{\text{don't use}}{\geq}} P_t r_n^{-\alpha}, \quad (2.5)$$

where P_r is the received primary power at the secondary radio.

The above analysis assumed that if a system can detect the received power with enough certainty, then it can deterministically determine its position and recover all the area beyond the no-talk radius. In reality there are uncertainties that make this task difficult:

- **Multipath:** Multipath occurs because of the constructive/destructive addition of different paths of the transmitted wave. A radio within the no-talk may experience destructive multipath which may cause it to see a weak signal and wrongly infer that

⁹It should be noted that the FCC did not set the rules by first setting the margin to be eroded. The rules were set in good faith to ensure that TV receivers were protected as before (i.e. at the grade B contour). Furthermore, the FCC intended the additional interference there to be below the noise floor. 3dB would have been at the thermal noise floor.

<i>Environment</i>	<i>Detection Scheme</i>	<i>By Area</i>				<i>By Population</i>			
		LVHF 2,5,6	HVHF 7-13	LUHF 14-51	HUHF 52-69	LVHF 2,5,6	HVHF 7-13	LUHF 14-51	HUHF 52-69
<i>No Multipath</i>	Geolocation ([12] Section 15.712)	2.4	4.13	23.8	16.4	2.5	3.24	16.1	15.9
	Geolocation with adjacent channels ([12] Section 15.712)	2.17	2.35	14.9	14.7	2.21	1.1	5.82	13.3
	Geolocation	1.87	3.35	22.4	16.1	2.08	2.54	14.8	15.7
	Geolocation with adjacent channels	1.71	1.89	14.1	14.4	1.81	0.82	5.36	13.1
	-114dBm rule	0.985	0.409	7.7	13.8	1.13	0.167	2.57	13.6
	-114dBm rule with adjacent channels	0.631	0.0505	2.63	9.83	0.639	0.004	0.284	8.87
<i>Multipath</i>	Geolocation	1.86	3.2	21.9	16	2.07	2.42	14.4	15.6
	Geolocation with adjacent channels	1.7	1.8	13.8	14.3	1.8	0.775	5.18	13
	-114dBm rule	1.09	0.524	8.85	14.3	1.25	0.225	3.29	14
	-114dBm rule with adjacent channels	0.73	0.0705	3.19	10.4	0.74	0.008	0.38	9.46

Table 2.4: Comparison of the average number of white space channels available per user in various environments, frequency bands and the methodology specified in the FCC report for a 4W 30m tall transmitter. LVHF denotes the lower VHF TV bands (channels 2-6), HVHF denotes the higher VHF TV bands (channels 7-13), LUHF denotes the lower UHF TV bands (channels 14-51) and HVHF denotes the higher VHF TV bands (channels 52-69) using the FCC’s definition of white space. The top two rows use the definition of $r_n - r_p$ from Section 15.712 in [12] while the next two rows calculate $r_n - r_p$ as per Appendix A. The -114dBm sensing rule for co-channels translates into a -110dBm rule for adjacent channels.




Figure 2.10: Visual representation of the available white space are all locations in the continental United States (as per the FCC's rules) The color black is used for locations where no white space channels exist. Interestingly these are locations on hill tops where multiple transmitters and relays are mounted. The color red is used to denote locations where are single white space channel exists while yellow denotes areas with two white space channels. Most metros see between 3-8 white space channels (these areas are marked in shades of blue).

it is outside the no-talk. Multipath is well modeled and hence we have some hope to incorporate its effects into the analysis.

- **Shadowing:** Shadowing occurs because of attenuation from obstacles in the path between the transmitter and the receiver. Shadowing could be very large (imagine a radio that is in a basement and hence sees a very weak signal). Shadowing is not well modeled and hence we need to budget for it.

One way to deal with shadowing and multipath is to budget for them explicitly – incorporate them into our decision rule. Let Δ (in dB) be our budget for shadowing and multipath, then rule (2.5) becomes:

$$P_r \underset{\text{use}}{\overset{\text{don't use}}{\geq}} P_t r_n^{-\alpha} 10^{\frac{-\Delta}{10}} \quad (2.6)$$

In Equation (2.6), the value of Δ impacts the secondaries ability to guarantee non-interference to the primary as well as to recover area for its own operation. If Δ is large then the secondary acts conservatively and only declares a location usable when the signal is very weak. In normal circumstances such a weak signal will occur very far from the TV transmitter and the secondary is unable to recover much area (See Figure 2.11) but it is able to ensure non-interference to the primary. If Δ is small then there is a chance that the secondary will not see the weak signal and will declare the location usable. In this case the secondary will be interfering with the primary more often even though it gets to use a major portion of the area (See Figure 2.11).

In its ruling the FCC has also proposed a detection sensitivity of -114dBm for ATSC signals. Given that the nominal signal level at the no-talk radius (r_n) is around -96dBm, this detection level means that the FCC has chosen a Δ of 18dB. To evaluate the amount of white space recovered by the -114dBm rule we determined the distance at which the signal dropped to the -114dBm level for 50% of the locations, 50% of the time (F(50, 50) rule). Ideally a cognitive radio user can be expected to adjust his/her sensing antenna so as to catch the worst fades (remember: if the signal decays below -114dBm the user *can* use the channel). But such manoeuvring cannot be performed for all transmitters and channels simultaneously. Hence considering the F(50, 50) propagation model is reasonable. From Table 2.4 it is clear that the -114dBm rule reduces the available white space by a factor of three. This effect is even more dramatic when white space by population is considered. In this case the available white space is reduced by a factor of five.

The impact of the -114dBm rule is visible in Figures 2.12(a) and (b) which visually depict the available white spaces for various rules for channel 7 and channel 39. For both channels, the -114dBm style rules with adjacent channel considerations knocks out a considerable amount of area. For co-channel considerations, the pollution viewpoint is more restrictive while for adjacent channel considerations the protection viewpoint is more restrictive.







Figure 2.11: If the budget for multipath and shadowing is small (Δ small), then the secondary can recover a majority of the area outside the no talk area. If the budget for multipath and shadowing is large (Δ large), then the secondary cannot use a lot of the area outside the no talk.



(a)



(b)



(c)

Figure 2.12: (a) Visual representation of the available white space for channel 5 using various usage guidelines. (b) Visual representation of the available white space for channel 40 using various usage guidelines. (c) Legend for the visual representation. Notice that pollution dominates protection within the channel while it is the other way around for adjacent-channel considerations.

The inability of the -114dBm rule to recover any area in the west and east coasts of the continental USA is shown in Figure 2.13 by plotting the number of white space channels that are available (according to FCC rules) on a hypothetical trip from Berkeley, CA to Washington, D.C. along Interstate 80.




Figure 2.13: Number of channels available on a trip from Berkeley, CA to Washington, D.C. The availability of white spaces is greatly reduced at the two coasts where most of the population lives. The analysis uses the FCC’s rules and neglects the pollution viewpoint.

This contrast is also visible in Figure 2.14 which contrasts the average channels per location per the FCC’s sensing rules (as seen in Figure 2.10) with the sensing rules. Sensing leads to may areas with high population density seeing no white space channels

Figure 2.15 shows the probability of getting at least a given number of white space channels using various detection and inference rules. This plot enables one to determine the quality of service that can be offered to white space users. For example, if we need 5 channels available, then we can only ensure this for 26% of the population using the -114dBm rule.

The tradeoff analysis using person-channels gained versus those lost can be extended to fixed threshold (eg. -114dBm) rules. As seen earlier, setting the threshold rule in a conservative manner results in a significant loss of area. For every setting of the threshold rule, we gain white space for people outside the region where the signal falls below the threshold for 50% of the locations, 50% of the time (the F(50, 50) point). A secondary at such a location will cause the primary receiver to see degradation in its SINR. Areas around the primary where the SINR is still above the target SINR will be able to receive a DTV signal but regions close to the noise limited radius will not. By varying the detection threshold we can examine the tradeoff between person-channels gained for white space versus those lost for broadcast. Figure 2.16(a) shows the average number of channels available for broadcast television usage as the threshold is varied.






Figure 2.14: Average number of white space channels per location as per the FCC geo-location (above) and sensing (below) rules. The contrast between the two images visually drives home the point that the FCC's sensing rules are too conservative.



(a)



(a)

Figure 2.15: Probability of getting at least a given number of channels using FCC style geolocation and fixed threshold (-114dBm) rules for uniform and actual population densities. (a) Adjacent channels are ignored. (b) Adjacent channels are considered

The reasonable interpretation of the -114dBm rule is that it represents the fear of the FCC – that the -114dBm rule is really a -94dBm rule (for instance) with a 10dB sensing margin for deep fading events. In other words, for mobile secondary users, the FCC cannot trust the placement and hence assumes that there will exist a secondary radio which nominally receives a signal of -94dBm but ends up seeing a 10dB fading event when trying to sense it. This is rational because the assumption is that there are many such mobile devices deployed and any one of them could transmit and harm TV receivers. On the other hand, any particular secondary user cannot assume deep fading events and hence under the -114dBm rule is unable to access vast amounts of white space.

The average number of broadcast channels resulting from this viewpoint can be seen in Figure 2.16(a) for sensing fading margins of 0dB, 10dB and 20dB. For a 0dB fading margin, the average number of broadcast users remains relatively flat for a range of detection thresholds (-120dBm to -100dBm). This is because the secondary interference from these distant locations is still much lower than the DTV noise floor. Since the -114dBm rule corresponds to large distances, the tradeoff point is much higher (500-1000 person-channels gained to 1 lost - see Figure 2.16(b)) for a fading margin of 0dB. For sensing fading margins of 10dB and 20dB, the tradeoff curves plummet considerably before they rise again. This is because low population densities in rural areas result in few white space users being added. However the broadcasters' fear of deep fading events means that they expect a large number of suburban broadcast users to lose service.

Figures 2.17(a) and (b) shows how 'one size fits all' style rules like the -114dBm style rule do not yield consistent area loss across towers. Figures 2.17(a) and (b) show the histogram and median/percentiles of the relative area lost as we change the threshold rule. The wide spread in the 10th and 90th rules show the problem with fixed threshold rules – they effect different towers in a vastly different manner.

2.6 Discussion

2.6.1 Impact of wireless Microphones

Wireless microphones also need protection from secondary interference [12]. One way to protect wireless microphones would be to identify all concert halls, indoor arenas and opera houses in the United States and to exclude a region around them that is equivalent to a -114dBm sensing rule. According to Wikipedia ([45]), there are 177 major concert halls, 1030 indoor arenas and 46 Opera houses. Excluding these venues would exclude a total of 1.4% of the total area of the United States (this assumes that all wireless microphones occupy the same 6MHz channel).




Figure 2.16: (a) Average number of channels available for white space/broadcast usage per person as the detection sensitivity (in dBm) and the assumed sensing fading margin (in dB) is changed. (b) The cumulative/incremental gain-loss tradeoff as the detection sensitivity (in dBm) and sensing fading margin (in dB) is varied. These plots are for the actual population density.




Figure 2.17: Relative area lost for a range of sensing threshold rules

2.6.2 Towards a framework to generate tradeoff plots

The tradeoff analysis in Sections 2.4 and 2.5 was performed with the actual locations and heights of the transmitters, the actual population density (accurate to the granularity of a zip code), the ITU propagation models ([42]), and a geo-location (Section 2.4) or sensing (2.5) detector. Models for the three constituents of this analysis (model of towers/population density, propagation and detector performance) can be changed (simplified/complexified) to get different tradeoff curve. For example we could change the detector type to that of an energy detector with a finite number of samples. Instead of using the location of all towers and population density we could use a representative tower and have an exponentially decaying model for the population density. Similarly, we could use a simple polynomial signal decay model for path loss. Such a methodology is described further in Chapter 3.

2.7 Summary

In this chapter we discussed the notion of a spatial spectrum hole using the ‘pollution’ (stay away from the primary to avoid its interference) and ‘protection’ (stay away from the primary to avoid interference to it). We saw how the lack of knowledge of the receiver’s locations leads us to a definition of a spatial spectrum hole based solely on distance. For the case of Digital TVs as primaries, we quantify the amount of white space (name given to a spatial spectrum hole by the FCC) that is actually available for secondary use. The FCC’s high and low power databases gave details of all licensed TV (digital and analog) stations and their latitude, longitude, effective radiated power (ERP) and transmitter heights above sea level. The ITU-R recommendations on propagation (ITU-R P-1536-3) and the US Census data of 2000 which lists the population density per zip code were used to get the white space availability per location/person. We compared the amount of white space resulting from the pollution and protection viewpoints. Engineering intuition suggests that the erosion margin should be set such that white space resulting from the two viewpoints is matched.

Quantifying white space as the average number of white space channels per location, we have shown that the main channels of relevance are the lower UHF channels (channels 14-51) in which we could get ~ 15 channels per person. However this number drops significantly (to ~ 5) when adjacent channels also have to be protected. Furthermore, the amount of white space is a function of the scale of the secondary. While we could get almost twenty-two channels per person (even without the upper UHF bands) for a hypothetical ‘zero Watts’ secondary, the number of white space channels for a 1MW secondary is almost zero.

Fixed threshold rules (for example the -114dBm rule proposed by the FCC for ATSC signals) are very conservative and result in almost no channels per person (especially when adjacent channels are considered). Under these observations, while geo-location may seem to be the way to go, sensing can be used to complement geo-location to determine channels which transmitters are not actually transmitting.

Lastly, we set forth a principled way to set the erosion margin. This principled approach considers the person-channels gained for white space use versus those lost from broadcast use. Based on the relative weighing of white space person-channels and broadcast person-channels, regulators can set forth the appropriate protection margin to be sacrificed for secondary use. Using the FCC’s report, we have “reverse engineered” the tolerable erosion margin as around 1dB. This will allow the community to propose new sensing strategies to follow the spirit of this margin while not being as conservative as the -114dBm rule.

Chapter 3

Metrics

In Chapter 2 we witnessed the poor performance of sensing with respect to geo-location in recovering white space. This was the result of conservatism in setting the value of the sensing threshold (or detection sensitivity). The traditional view of detection is focussed on obtaining good detection sensitivity with a given number of samples; hence the large literature on novel sensing algorithms. The problem with detection sensitivity is that it does not take into account fading and is unable to predict the performance of algorithms in recovering spatial spectrum holes.

In this chapter we shall introduce new metrics which capture the impact of fading and can predict the performance of sensing with respect to geo-location. These metrics are required to expedite research into algorithms that perform well with spatial spectrum holes – with these metrics we have a mechanism to compare algorithms without having to perform extensive Monte Carlo simulations as we did in Chapter 2. Chronologically though, the metrics grew out of the need to showcase the performance gains of cooperative sensing to an IEEE 802.22 audience deeply set in the ‘detection sensitivity’ paradigm (I worked with Gerald Chouinard on performance evaluation of cooperative sensing [46]). The only way to market cooperative sensing was to examine its role in reducing false positives (or false alarms). What was missed in that analysis was that for spatial spectrum holes (unlike temporal spectrum holes) false positives do not result from noise triggering the detection threshold – false positive also involve the primary signal at a lower level.

This chapter is organized as follows: First, in Section 3.1 the traditional binary hypothesis testing story is recapitulated with the conceptually central role being played by the traditional detection metric of *sensitivity*¹. It is well known that the sensitivity of detectors can be improved by increasing the sensing time and so the sample complexity gives us a natural way to compare different spectrum sensors.

Next in Section 3.2 we look at how fading affects different detectors. This reveals that a head-to-head comparison of the sensitivity of two detectors can be misleading. Instead, the

¹The sensitivity of a detector is the lowest SNR for which a given target probability of error can be met.

possibility of fading has to be incorporated into the “signal present” hypothesis itself. The logical tradeoff is then between the effective probability of missed detection (it matters not whether the miss is due to an unfortunate fade or a quirk in the noise realization) and the sample complexity. Here, we show a surprising example. Whereas traditionally, the coherent detection of a pilot tone is considered to have better asymptotic sample complexity than an energy detector, this need not be true when fading is considered. The coherence bandwidth matters. If the primary signal is sufficiently wideband, then the simple energy detector can have *better* asymptotic sample-complexity behavior than a coherent pilot-tone detector!

The bigger conceptual challenge comes in trying to understand “false alarms.” The traditional hypothesis-testing formulation would say that a false alarm consists of when the detector says that we should not use the channel when the primary is truly absent. But this is not the problem actually facing a cognitive radio. It wants to avoid saying that we are close to the primary when we are indeed far enough away. The “signal absent” hypothesis needs to be modified in some reasonable way.

At this point, a spatial perspective is essential and while the resulting formulation can fit into the traditional binary hypothesis testing framework, it is useful to reconceptualize the problem in terms of two new metrics, first introduced in [18]. These new metrics are presented in Sections 3.3 and 3.4. The first metric, namely the *Fear of Harmful interference* F_{HI} , captures the safety to the primary users. This is largely the fading-aware probability of missed detection introduced earlier, with some modifications to allow easier incorporation of system-level uncertainty. The second metric, namely the *Weighted Probability of Area Recovered* ($WPAR$), captures the performance of spectrum sensing by appropriately weighting the probability of false alarm (P_{FA}) across different spatial locations. These metrics give a unifying framework in which to compare different spectrum-sensing algorithms. We show how to obtain reasonable metric parameters (most crucially, something which can be interpreted as a spatial discount factor) from real-world data. The tradeoff between $WPAR$ and F_{HI} is thus the correct *ROC* curve for spectrum sensing. However, the probabilistic uncertainty underlying the hypotheses is non-ergodic and so the tradeoffs are interesting even if we allow an infinite number of samples. The new metrics show that fading uncertainty forces the $WPAR$ performance of single-radio sensing algorithms to be very low for desirably small values of F_{HI} , even with infinite samples. The physical reason for such a poor performance is that a single radio cannot distinguish whether it is close to the primary user and severely shadowed, or if it is far away and not shadowed. The new metrics are able to predict the FCC’s operating point (as determined in Chapter 2) fairly accurately. This chapter is based on the paper [18] and overlaps partially with [19].

3.1 Traditional Metrics

We first consider the traditional formulation of the spectrum sensing problem as a binary hypothesis test.

Let $x[n]$ denote the sampled version of a band-limited signal with power P in dBm and let the additive noise be $w(n) \sim \mathcal{N}(0, 10^{\frac{N_0}{10}})$ where N_0 is the Noise power in dBm. The two hypotheses are:

$$\begin{aligned} \text{Signal absent } \mathcal{H}_0 : y[n] &= w[n] & n = 0, \dots, N-1 \\ \text{Signal present } \mathcal{H}_1 : y[n] &= x[n] + w[n] & n = 0, \dots, N-1 \end{aligned} \quad (3.1)$$

Based on the assumptions we make about the signal, we get different test statistics ($T(\mathbf{y})$). The test statistic is compared with a threshold λ to make a decision:

$$T(\mathbf{y}) \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \lambda,$$

The secondary would like to limit the probability that it will mis-detect the primary and hence potentially cause harmful interference. The probability of causing harmful interference is bounded from above by the probability of misdetecting the primary which is calculated as:

$$P_{MD} = \mathcal{P}(T(\mathbf{y}) \leq \lambda | \mathcal{H}_1) \quad (3.2)$$

where the probability is calculated over the distribution of noise and the signal.

Similarly the secondary would like to maximize the probability of recovering a spectrum opportunity when the primary is not on. This probability can be computed as $(1 - P_{FA})$ where P_{FA} is the probability that the detector false alarms with only noise present:

$$P_{FA} = \mathcal{P}(T(\mathbf{y}) > \lambda | \mathcal{H}_0) \quad (3.3)$$

The number of samples needed to achieve a target (P_{MD}, P_{FA}) is a function of the SNR $(P - N_0)$ i.e.

$$N = \mathbf{f}(P_{MD}, P_{FA}, SNR)$$

Stated differently, for a given N, P_{MD}, P_{FA} the minimum SNR that can be detected is called the *Detection Sensitivity*.

The actual function $\mathbf{f}(\cdot)$ (and test statistic) depends of the assumptions made about the signal $x[n]$. Two cases are of particular interest:

- The signal has zero mean. In this case the test statistic ($T(\mathbf{y}) = \frac{1}{N} \sum_{n=1}^{N-1} |y[n]|^2$) is the energy detector and the required number of samples to achieve the target P_{MD}, P_{FA} scales as $\mathcal{O}(10^{\frac{-SNR}{5}})$ [47, 48].
- The signal has a known mean (pilot/preamble sequence - $s[0], \dots, s[N-1]$): The test statistic is $T(\mathbf{y}) = \frac{1}{N} \sum_{n=0}^{N-1} s[n] y[n]$. and the required number of samples to achieve the target P_{MD}, P_{FA} scales as $\mathcal{O}(10^{\frac{-SNR}{10}})$ [47, 48].

The performance of the detector for the two signal types is shown in Figure 3.1. The figures suggests that for low SNR operations, the coherent detector is always preferred i.e. a coherent detector is more sensitive than a energy detector.

3.2 Impact of fading

At this point we have to ask the question – *At what level should we set the detector sensitivity?* We have seen that is easier to detect signals at high SNR (See Figure 3.1). As we move away from the transmitter, the signal becomes weaker. As discussed in Chapter 2, while we are within the no-talk radius of the tower we must say “do not use this channel”. If there was a one-to-one mapping between distance and received signal strength we could easily determine the required sensitivity. But there is shadowing and multipath — we need to protect the primary users even under adverse fading and shadowing conditions. It is important to maintain an appropriately low probability of mis-declaring that we are outside whenever we are actually inside. We may hope to average over the multipath fading since it changes every coherence time. However, the coherence time is itself uncertain since it depends on physical motion. Furthermore, behavior of different detectors may be effected differently by the details of the fading distribution.

In Figure 3.2 we incorporate the role of fading in the plot of Figure 3.1. We introduce the fading distribution at r_n into $\mathcal{H}_1$ (signal present). The threshold of each detector is set so that the P_{FA} is met. Notice that so far, fading has no effect on the “signal absent” hypothesis since there is nothing to fade! The average P_{MD} is calculated as: $\int_{-\infty}^{\infty} P_{MD}(p) dF_{r_n}(p)$ where $F_{r_n}(p)$ is the cumulative distribution function (CDF) of the received signal strength at the no-talk radius. The performance of the radiometer depends on the amount of multipath averaging it can count on (for example Digital TV occupies a band of 6MHz and the coherence bandwidth is significantly smaller ([49]) hence it can count on a frequency diversity gain [50]) (the diversity order specifies the number of taps in the channel filter). With a diversity order




Figure 3.1: The performance of the energy and coherent detectors are compared. For low SNR operation the coherent detector is preferred.




Figure 3.2: Number of samples versus the target value of the average P_{MD} while holding $P_{FA} = 0.1$. We assume a DTV transmitter with a 1MW transmit power and a r_n of 156km. The average received power is -96dBm with a noise power of -106 dBm. With a diversity order of just 2 the energy detector performs better than the coherent detector. The model assumes log-normal multipath and no shadowing.

of two² or more the radiometer performs better than the coherent detector at all desired P_{MD} . This example illustrates a major point – taking the fading distribution into account is important since it impacts the choice of the detector to be used.

3.3 Towards New Metrics

P_{MD} averaged over fades better captures safety for the primary and reveals issues that the sensitivity metric alone does not. We now turn our attention to rethinking $\mathcal{H}_0$. Traditionally, this has been viewed as the “signal absent” hypothesis and modeled as receiving noise alone. However, that does not reflect what we actually care about for cognitive radio systems. We only want to verify that the local primary user is absent: it is perfectly fine for there to be some distant primary transmission if we are beyond that tower’s no-talk radius.

We’ve seen in Chapter 2 that our threshold impacts how much area we can recover for cognitive radio operation. To capture this very fundamental tradeoff in our mathematical model we have to rely on asymmetry. The true position of the cognitive radio is unclear. For $\mathcal{H}_1$, it was natural to take the worst-case position of being just within r_n and then evaluate P_{MD} averaging over fading. That is the most challenging point in terms of sensitivity. Suppose we took the same approach to $\mathcal{H}_0$. We would then evaluate P_{FA} immediately outside r_n . After all, if we can recover this location then we can recover all the area greater than r_n . This approach is fatally flawed since the distribution of the signal strength just within r_n and outside of r_n are essentially the same. No interesting tradeoff is possible. We miss the fundamental fact that we must give up some area immediately outside of r_n to avoid having a cognitive radio use the channel within r_n .

Simply averaging over r (distance from the primary transmitter) also poses a challenge. The interval (r_n, ∞) is infinite and hence there is no uniform distribution over it. This mathematical challenge corresponds to the physical fact that if we take a single primary-user model, the area outside r_n that can be potentially recovered is infinite. With an infinite opportunity, it does not matter how much of it that we give up! We need to come up with probability distribution over r or in other words, we need to weight/discount area outside r_n appropriately. Weighting area by “utility” is a possibility, but as discussed in [18], this would tightly couple the evaluation of sensing with details of the business model and system architecture. It is useful to find an approximate utility function that decouples the evaluation of the sensing approach from all of these other concerns.

Two discounting approaches can be considered:

- We are using an overtly single-primary user model to capture the reality of having many primary users reusing a particular frequency. As we move away from any specific

²Since there is also shadowing (common to all taps) in real life, the net diversity order can easily be a fractional value and does not have to be larger than two.




Figure 3.3: Voronoi regions of the transmission towers for TV channel 39.


tower, there is a chance that we may enter the no-talk zone for another primary tower transmitting on the same frequency. As discussed in [18], this can be viewed as a spatial analogy to “drug-dealer’s discounting” in which money in the future is worth less than money in the present because it is uncertain whether the drug dealer will survive into the future because of the arrival of the police or a rival gang [51].

Figure 3.3 shows the Voronoi regions³ for the transmitters on channel 39. Figure 3.4(b) illustrates the result. The X-axis is the distance from the tower while the Y-axis is the natural logarithm of the percentage of the circle (of that radius) that is included in the transmitter’s Voronoi region. Beyond 400km, the mean is dominated by rare large values. The natural logarithm of the included fraction decays linearly with distance *i.e.* the included fraction decays exponentially with distance. The decay rate is roughly 110×10^{-4} per km.


- Presumably, we want cognitive radios to be useable by people. Since TV towers are often located to serve areas of high population density, areas around the no-talk region are more valuable than areas far away. This can be viewed as a spatial analogy to “banker’s discounting” in which money in the future is worth less in present units.

Population densities can be modeled as decaying exponentially as one moves away from the central business district [52]. To validate this, we calculated the distribution of the

³Ideally we would like to construct the received-signal-strength Voronoi region – a received signal strength Voronoi around a transmitter would be all points where the $F(50, 50)$ signal strength from that transmitter is higher than from any other transmitter. Such a Voronoi region is hard to compute.



(a)



(b)

Figure 3.4: (a) shows the decay of population density with distance from a Digital TV transmitter. The average, 90th and 10th percentile are plotted. The decay rate is around 8×10^{-4} people/km² per km. (b) shows the fraction of a circle (of a given radius) that is included in a tower's Voronoi region. The average, 90th percentile and median are plotted (over towers). The decay rate is around 110×10^{-4} per km.

population density at a given distance for all TV towers (taking into account both the high power and low power DTV databases as discussed in Chapter 2). If Q_i^r is the number of people in a radius r around tower i , the average population density at a given distance r is given by:

$$\frac{\sum_{i=1}^{N_T} Q_i^{r+\delta} - Q_i^r}{2\pi r \delta N_T} \quad (3.4)$$

where N_T is the total number of transmitters.

The top graph in Figure 3.4(a) shows the average population at a given distance from the tower. The decay rate is roughly 8×10^{-4} people/km² per km. We see that this is much smaller than the discounting induced by the frequency reuse.

Now we have a way of discounting the area recovered and presenting the cognitive radio's ability to recover area as a single number. In this thesis we will use the Voronoi approach to provide the area discounting factor. This means that we can compare the performance of geo-location and sensing in their ability to recover area but not people.

3.4 Spatial Metrics

The discussion so far has resulted in a new hypothesis-testing problem. In both of the hypotheses, the received signal $y[n] = \sqrt{P(r)}Hx[n] + w[n]$ but the two hypotheses potentially differ in $P(\cdot)$ (the path-loss and transmit power function), r (the distance from the primary user to the cognitive radio), H (fading), and w (the noise process).

Both hypotheses could agree on common models for $P(\cdot)$, H , and w , but there is guaranteed to be a difference in r .

$$\begin{aligned} \text{Safe to use } \mathcal{H}_0 &: r \sim w(r)r \\ \text{Unsafe } \mathcal{H}_1 &: r \in [0, r_n] \end{aligned} \quad (3.5)$$

where $w(r)$ must satisfy $\int_{r_n}^{\infty} w(r) r dr = 1$ and $w(r) = 0$ if $r < r_n$. Following the discussion in the previous section, the numerical results in this paper have been computed using an exponential weighting function, $w(r) = \kappa \exp(-\kappa r)$.

From the above, the asymmetry between the two hypotheses is clear. $\mathcal{H}_0$ is a well-defined probability and so P_{FA} can be calculated for a detector. Meanwhile, $\mathcal{H}_1$ has r as a non-probabilistic uncertainty and so we would have to require something like $\sup_{r \in [0, r_n]} \mathcal{P}_{Fr}(T(\mathbf{Y}) < \gamma | \mathcal{H}_1) \leq P_{MD}$. The resulting mixed ROC curve for a detector reveals the fundamental tradeoffs.

However, such a formulation mixing worst-case and Bayesian uncertainties in different ways across the two hypotheses is novel. Using the traditional names P_{FA} and P_{MD} in this context is likely to lead to confusion. Therefore, in [18] we gave them new more descriptive names that better reflected their roles.

3.4.1 Safety: Controlling the Fear of Harmful Interference

The idea behind the safety metric to measure the worst-case safety that the cognitive radio can assure the primary (the worst case is calculated over the fading distribution negotiated between the cognitive radio and the primary – for example in Figure 3.2 the primary and cognitive radio agree on a single distribution of fading).

To capture the strength of the safety guarantee under uncertainty we define a new metric called the *Fear of Harmful Interference*. This is the same as the average P_{MD} but takes into account all uncertainties in location and fading.

Definition 1 Fear of Harmful Interference (F_{HI}) metric is

$$F_{HI} = \sup_{0 \leq r \leq r_n} \sup_{F_r \in \mathbb{F}_r} \mathcal{P}_{F_r}(D = 0 | R = r). \quad (3.6)$$

where $D = 0$ is detector's decision declaring that the cognitive radio is outside the no-talk radius and $\mathbb{F}_r$ is the set of possible distributions for $P(r), H, W$ at a distance of r from the primary transmitter. The outer supremum is needed to issue a uniform guarantee to all protected primary receivers and also reflects the uncertainty in cognitive radio deployments. The inner supremum reflects the non-probabilistic uncertainty in the distributions of the path-loss, fading, noise, or anything else.

3.4.2 Performance

Next we consider a metric to deal with the cognitive radio's performance — its ability to identify spectrum opportunities. From a traditional perspective, this is basically a weighted P_{FA} . Every point at a radial distance $r > r_n$ is a spectrum opportunity. For any detection algorithm, there is a probability associated with identifying an opportunity there, called the probability of finding the hole $P_{FH}(r)$:

$$P_{FH}(r) = \mathcal{P}_{F_r}(D = 0 | R = r), \quad r > r_n. \quad (3.7)$$

where F_r represents the propagation, fading, and noise models believed by the cognitive radio designer. A worst-case perspective could also be used here if needed, but it is reasonable to believe that the designer's uncertainty about this could be placed into a Bayesian form. They have no reason to lie to themselves.

Parameter	Value
Protected radius r_p	135km
No-talk radius r_n	150.3km
Weighting Function $w(r)$	$\kappa \exp(-\kappa r)$ $\kappa = 110 \times 10^{-4}/\text{km}$
Transmit Power	1000kW
HAAT	500m

Table 3.1: Single tower parameters.

Definition 2 *The Weighted Probability of Area Recovered (WPAR) metric is*

$$WPAR = \int_{r_n}^{\infty} P_{FH}(r) w(r) r dr, \quad (3.8)$$

where $w(r)$ is a weighting function that satisfies $\int_{r_n}^{\infty} w(r) r dr = 1$.

The name $WPAR$ reminds users of the weighting of performance over spatial locations that is fundamental to the cognitive radio context. $1 - WPAR$ is the appropriate analog of P_{FA} and measures the sensing overhead [30].

The final parameters for the single tower model (to be used for all subsequent calculations) are shown in Figure 3.1

3.4.3 Quantile Models

To compute F_{HI} , we cannot always use the nominal model for shadowing and multipath as it is important to model the fact that the primary user does not trust this model completely. Instead, it is possible that the primary user trusts only a quantized version (or a coarse histogram) of the fading distribution. Mathematically, we model this as a class of distributions ($\mathbb{F}_r$) that satisfy given quantile constraints.

Definition 3 *A single quantile model $\mathbb{F}_r$ is a set of distributions for the received signal power defined by a single number $0 \leq \beta \leq 1$ and a function of r denoted $\gamma(r, \beta)$. A distribution $F_r \in \mathbb{F}_r$ iff*

$$\mathcal{P}_{F_r}(P < \gamma(r, \beta)) = \beta. \quad (3.9)$$

A k -quantile model is a set of distributions $\mathbb{F}_r$ for the received signal power defined by a list of numbers $(\beta_1 < \beta_2 < \dots < \beta_k)$ and a corresponding list of functions $(\gamma_1(r, \beta_1), \dots, \gamma_k(r, \beta_k))$. A distribution $F_r \in \mathbb{F}_r$ iff $\forall i \leq k$

$$\mathcal{P}_{F_r}(P < \gamma_i(r, \beta_i)) = \beta_i. \quad (3.10)$$

For consistency, the quantiles are chosen so that the nominal Gaussian $\mathcal{N}(\mu(r), \sigma^2)$ is always one of the possible distributions for P .

$$\gamma(r, \beta) = \mathcal{Q}^{-1}(1 - \beta)\sigma + \mu(r), \quad (3.11)$$

where $\mathcal{Q}^{-1}(\cdot)$ is the inverse of the standard Gaussian tail probability function.

Figure 3.5 shows a picture of the distributions allowed under the quantile model (5 learned quantiles) defined in this section. The set of allowed Cumulative Distribution Functions (CDF's) for P under our quantile model is precisely the set of all possible non-decreasing curves sandwiched between the upper and lower bounds shown in Figure 3.5. The dashed (black) curve in the figure shows that nominal Gaussian CDF for P , and the quantile constraints can be thought of as samples of the nominal CDF (the triangle points (in red) in the figure).

3.4.4 Single radio performance

Consider a single radio running a radiometer to detect whether the frequency band is used/unused. As the uncertainty in the fading can be non-ergodic, the F_{HI} vs $WPAR$ tradeoff for a single-radio detector is interesting even when the number of samples is infinity. In the rest of the paper we assume that the detectors have infinite samples. The test-statistic for a radiometer with infinite samples is $T(\mathbf{Y}) := \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=0}^{N-1} |y[n]|^2 = 10^{\frac{P(r)H^2}{10}} + \sigma_w^2$, where $P(r)H^2$ (in dBm) is the average received signal power at distance r . Therefore, the perfect radiometer decides whether the band is used/unused according to the following rule

$$\begin{aligned} T(\mathbf{Y}) &= 10^{\frac{P(r)H^2}{10}} + \sigma_w^2 \begin{matrix} \stackrel{D=1}{\geq} \\ \stackrel{D=0}{\leq} \end{matrix} \lambda \\ &\Leftrightarrow P(r)H^2 \begin{matrix} \stackrel{D=1}{\geq} \\ \stackrel{D=0}{\leq} \end{matrix} 10 \log_{10}(\lambda - \sigma_w^2) =: \tilde{\lambda}. \end{aligned} \quad (3.12)$$

The curve in Figure 3.6 shows the F_{HI} vs $WPAR$ tradeoff for a single user running a radiometer. The $WPAR$ performance at low F_{HI} is bad even for the perfect radiometer. This captures the physical intuition that guaranteeing strong protection to the primary user forces the detector to budget for deep fading events. Unlike in traditional communication problems where there is no harm if the fading is not bad, here there is substantial harm since a spectrum opportunity is left unexploited.

3.4.5 Comparing metrics to empirical performance

To determine the quality of the newly defined metrics we can compare the performance of sensing relative to geo-location as predicted by the metrics to calculated numbers from




Figure 3.5: The quantile model for the received signal power (P) distribution. The dashed (black) curve is the nominal Gaussian CDF for P , and the triangle points (in red) show the quantile constraints on the CDF. The dashed-dotted (magenta) curve is the upper bound and the solid (blue) curve is the lower bound on the allowable CDF for P . The actual CDF can lie anywhere in between, and must pass through the 5 triangle points (quantile constraints).




Figure 3.6: F_{HI} versus $WPAR$ tradeoff of a single user running an energy detector with infinite samples. The predicted performance of sensing for the FCC's F_{HI} value of 1.69×10^{-4} is .47 (the green dot in the plot). The same number computed based on actual data is 0.36 (the red dot on in plot). This discrepancy can be accounted for by not taking into account power and height variations in our model.

Chapter 2. We will use channel 39 as an example. For channel 39, geo-location recovers 64% of the area while sensing recovers 23%. Hence sensing recovers 0.35 fraction of the area as compared to geo-location. Figure 3.6 shows the performance of energy detector with infinite samples. The empirical performance of sensing with respect to geo-location⁴ is calculated as 0.47 i.e. the metrics are optimistic in predicting sensing performance. This is because we are neglecting variations in transmit power and HAAT among the Digital TV transmitters.

3.5 Summary

In this chapter, we showed that the triad — Probability of Missed Detection (P_{MD}), Probability of False Alarm (P_{FA}) and number of samples, do not provide the right insights into the spectrum hole recovery process. The traditional metrics suggest that the coherent detector is always preferred over the energy detector at low SNRs. However, simply taking into account fading changes the qualitative choice of the rule to be used. For spatial hole recovery, the new metrics of Fear of Harmful Interference (F_{HI}) and Weighted Probability of Area Recovered ($WPAR$) are proposed. F_{HI} is the fading aware Probability of Missed Detection with modifications to take into fading/radio behavior uncertainties. $WPAR$ captures the performance of spectrum sensing relative to geo-location. The parameters of these metrics are derived from the FCC's Digital TV databases.

⁴We can calculate the F_{HI} specified by the FCC by determining the gap between the median signal strength at r_n and the sensing threshold (-114dBm).

Chapter 4

Cooperation

In Chapter 2 we saw how very sensitive detectors were unable to recover spectrum real estate when compared to geo-location. To capture this gap we introduced the metrics of Fear of Harmful Interference (F_{HI}) and Weighted Probability of Area Recovered ($WPAR$). In Chapter 3, we also saw that frequency diversity helped a energy detector outperform a coherent detector. In this section we are counting on diversity again (spatial diversity in this case) to outperform a single radio detector — we look at cooperative sensing as a mechanism to recover area while ensuring F_{HI} compliance.

Based on whether we are trying to recover a temporal hole or spatial hole, the motivation for cooperating might be different. For a temporal hole we are using cooperative sensing to gather as much energy in the signal as possible (either through combining or by looking for the secondary with the maximum signal strength) and thus creating a large gap between means for the signal present and the signal absent hypotheses. For recovering spatial holes the problem is different. For a spatial hole there is a mapping between the nominal signal strength at a given distance from the transmitter and the distance itself. By using multiple radios we are hoping to converge to this nominal signal level. This difference between the usage of cooperative sensing for temporal and spatial holes has major implications on the choice of the optimal cooperative rule and is examined in depth in this chapter.

The organization of this chapter is as follows: In Section 4.1 we will review existing cooperative sensing literature by identifying the major themes of cooperative sensing research. In Section 4.2 we present cooperation techniques for recovering temporal spectrum holes. This section reveals that the performance of cooperation strategies suffers due to the presence of radios that are excessively faded/shadowed. Furthermore, analyzing cooperation with our Digital TV measurements shows that robustness of a strategy is vital for good real-world performance. Echoing Section 4.2, Section 4.3 presents cooperation techniques for recovering spatial spectrum holes. The median rules performs as well as the averaging rule and is shown to be robust to misbehaving/misconfigured users and uncertain fading models (in the form of quantile models). Alas, all cooperation strategies are non-robust to loss in spatial

diversity. This chapter is based on the papers [18, 20, 21] and [22].

4.1 Related work

Cooperative sensing has been a very active area of research in cognitive radios as well as for sensor networks [53, 54]. Several groups have proposed cooperation among cognitive radios as a tool to improve performance [20, 55–61]. Given the numerous papers in this field we list the major research themes in Table 4.1 and representative references. Gains from cooperation can either be viewed as diversity gains where multiple radios reduce the collective probability of getting a bad fade [20, 60] or as a mechanism to reduce sensing overhead [55, 56, 61]. Dealing with uncertainty (in the form of correlated data measurements and/or malfunctioning/malicious radios) forms a major component of this research. In addition, design of optimal cooperative sensing schemes under various constraints (communication/synchronization constraints for example) is also an active area of research.

4.2 Cooperation for temporal holes

4.2.1 Cooperation strategies

We consider a group of M cognitive radios at a distance r from the primary transmitter. The received signal at the M radios is given by $P_1, \dots, P_M$. The radios run a detection algorithm that outputs a binary decision D about the state of the primary band: 1-used/0-unused. Let the hypothesis $\mathcal{H}_0$ denote the situation that the primary is absent and hypothesis $\mathcal{H}_1$ denote the situation that the primary is present. The *probability of harmful interference* $\mathcal{P}_{F_r}(D = 0|\mathcal{H}_1)$ at a distance r is the probability that a secondary user is unable to detect the primary¹ at a distance r . Similarly, the probability of recovering a potential white space location (*probability of finding a hole*) is $\mathcal{P}_{F_r}(D = 1|\mathcal{H}_0)$. Here F_r is the probability distribution of the combined multipath and shadowing-induced fading at a distance r from the primary transmitter.

In this section we will consider distributions where F_r takes on a single tuple only i.e. $F_r(P_1, \dots, P_M) = (p_1, \dots, p_M)$ with probability 1. This enables us to focus on certain cases of interest that highlight the differences between the different rules.

We assume that each of the M radios collects N samples each. The goal is to distinguish between the following hypothesis:

¹Here we are ignoring the case that r is so large that we could be outside the no-talk radius of the transmitter.

<i>Research Theme</i>	<i>Main idea/goals</i>	<i>References</i>
Cooperation as diversity	Cooperation can be seen as providing diversity gains by reducing sensitivity requirements for individual radios.	[20, 60]
Cooperation as gains in degrees of freedom	Cooperation reduces sensing time or lowers false alarms for the same level of detection.	[55, 56, 61]
Impact of/Dealing with correlation	Determine the impact of channel correlation on cooperation gains as well as mechanisms of dealing with correlation uncertainty.	[20, 59, 62]
Impact of/Dealing with malicious/lying users	Determine the impact of incorrect sensing responses and mechanisms for weeding out misbehaving users.	[20, 63–65]
Cooperation and Communication	Determine the impact of communications/synchronization constraints on cooperation performance.	[55, 56, 66–68]
Fusion rules	Investigate various soft/hard combining rules.	[21, 58, 61, 69–71]
Utilizing sparsity/multiband information	Utilize multiple frequency bands for cooperative gains.	[23, 72]

Table 4.1: Description of various research thrusts in the area of cooperative spectrum sensing.

$$\begin{aligned}\mathcal{H}_0 : y_m[n] &= w_m[n] & n = 0, \dots, N-1, \quad m = 1, \dots, M \\ \mathcal{H}_1 : y_m[n] &= x_m[n] + w_m[n] & n = 0, \dots, N-1, \quad m = 1, \dots, M\end{aligned}$$

We assume that all $w_m[n]$ are independent and identically distributed as $\mathcal{N}(0, \sigma_{w_m}^2)$ and $N_{0_m} = 10 \log_{10}(\sigma_{w_m}^2)$. Similarly we assume that all $x_m[n]$ are independent and distributed as $\mathcal{N}(0, \sigma_{s_m}^2)$ and $P_m = 10 \log_{10}(\sigma_{s_m}^2)$. Furthermore we assume that two radios are not synchronized and hence $x_m[n]$ is independent of $x_{m'}[n']$ for all $0 \leq n, n' \leq N-1$ and $1 \leq m, m' \leq M$ i.e. the time at which the two radios start sensing is not closely co-ordinated.

LLRT optimal rule – Averaging with discounting

First we derive the optimal rule using the Neaman-Pearson hypothesis testing model. The likelihood ratio is given by:

$$\frac{\prod_{m=1}^{M-1} \prod_{n=0}^{N-1} \frac{1}{\sqrt{2\pi(\sigma_{s_m}^2 + \sigma_{w_m}^2)}} \exp\left(\frac{-y_m[n]^2}{2(\sigma_{s_m}^2 + \sigma_{w_m}^2)}\right)}{\prod_{m=1}^M \prod_{n=0}^{N-1} \frac{1}{\sqrt{2\pi\sigma_{w_m}^2}} \exp\left(\frac{-y_m[n]^2}{2\sigma_{w_m}^2}\right)} \underset{H_0}{\overset{H_1}{\geq}} \lambda$$

Taking log on both sides and renaming constants gives:

$$\sum_{m=1}^M \frac{\sigma_{s_m}^2}{\sigma_{w_m}^2 (\sigma_{s_m}^2 + \sigma_{w_m}^2)} \sum_{n=0}^{N-1} \frac{1}{N} y_m[n]^2 \underset{H_0}{\overset{H_1}{\geq}} \lambda'$$

which gives the final test statistics:

$$T(\mathbf{y}) = \sum_{m=1}^M \frac{\sigma_{s_m}^2}{\sigma_{w_m}^2 (\sigma_{s_m}^2 + \sigma_{w_m}^2)} \frac{1}{N} \sum_{n=0}^{N-1} y_m[n]^2. \quad (4.1)$$

Define, $SNR_m \triangleq \frac{\sigma_{s_m}^2}{\sigma_{w_m}^2}$. Then under $\mathcal{H}_1$, $\sum_{n=0}^{N-1} \frac{1}{N} y_m[n]^2$ is distributed as $\mathcal{N}(\sigma_{s_m}^2 + \sigma_w^2, \frac{2}{N}(\sigma_{s_m}^2 + \sigma_w^2)^2)$ for large N [47, 48]. Hence,

$$T(\mathbf{y}|\mathcal{H}_1) \sim \mathcal{N}\left(\sum_{m=1}^M SNR_m, \frac{2}{N} \sum_{m=1}^M SNR_m^2\right)$$

Similarly under $\mathcal{H}_0$, $\sum_{n=0}^{N-1} \frac{1}{N} y_m[n]^2$ is distributed as $\mathcal{N}(\sigma_w^2, \frac{2}{N} \sigma_w^4)$ for large N . Hence,

$$T(\mathbf{y}|\mathcal{H}_0) \sim \mathcal{N}\left(\sum_{m=1}^M \frac{SNR_m}{SNR_m + 1}, \frac{2}{N} \sum_{m=1}^M \left(\frac{SNR_m}{1 + SNR_m}\right)^2\right)$$

If we define, $\mu_0 = \sum_{m=1}^M \frac{SNR_m}{SNR_m + 1}$, $\sigma_0 = \sqrt{\frac{2}{N} \sum_{m=1}^M \left(\frac{SNR_m}{SNR_m + 1}\right)^2}$, $\mu_1 = \sum_{m=1}^M SNR_m$, $\sigma_1 = \sqrt{\frac{2}{N} \sum_{m=1}^M SNR_m^2}$, we can calculate the probability of missed detection (and hence causing harmful interference) $P_{HI} = P_{MD}$ and the probability of finding a temporal spectrum hole $P_{FH} = 1 - P_{FA}$,

$$P_{FH} \approx 1 - Q\left(\frac{\lambda' - \mu_0}{\sigma_0}\right) \text{ and } P_{HI} \approx 1 - Q\left(\frac{\lambda' - \mu_1}{\sigma_1}\right) \quad (4.2)$$

If we take the case of where the noise ($\sigma_{w_m}^2$), is the same for all m , (4.1) becomes:

$$T(\mathbf{y}) = \sum_{m=1}^M \frac{SNR_m}{1 + SNR_m} \sum_{n=0}^{N-1} \frac{1}{N} y_m[n]^2. \quad (4.3)$$

i.e the energy from each radio is combined using a discounting factor. The discounting factor is a function of the SNR of radio m . Energy of secondaries that see a high SNR is used as is (since $\frac{SNR_m}{1+SNR_m} \sim 1$ for large SNR_m) while the energy of low SNR secondaries is discounted by the SNR. (since $\frac{SNR_m}{1+SNR_m} \approx SNR_m$ for $SNR_m \ll 1$). Similar discounting is proposed in [62]. This *averaging with discounting rule* is the sensing equivalent of the Maximum Ratio Combining (MRC) rule used for MIMO channels (in MRC the channel estimate h^* is assumed to be known and used as the weighting factor [50]) while in *averaging with discounting* rule the SNRs are assumed to be known and a function of the SNR is used for weighting.).

Optimal Rule with SNR estimates – Simple averaging rule

The problem in implementing the optimal rule is in knowing the actual SNRs at each radio. The solution to this problem is to insert the ML estimate of the SNR into (4.3) – the GLRT test statistic. To do this we proceed as follows; we first calculate the empirical received

power as: $\hat{Q}_m = \frac{1}{N} \sum_{n=0}^{N-1} y_m[n]^2$. Inserting this into (4.3) gives:

$$T(\mathbf{y}) = \sum_{m=1}^M \frac{(\hat{Q}_m - \sigma_{w_m}^2)}{\sigma_{w_m}^2}$$

The resulting rule is a simple averaging rule. It combines the energy across radios irrespective of their operational SNRs.

Under $\mathcal{H}_1$, $\hat{Q}$ is distributed as $\mathcal{N}(\sigma_{s_m}^2 + \sigma_w^2, \frac{2}{N}(\sigma_{s_m}^2 + \sigma_w^2)^2)$ for large N . Using the distribution of $\hat{Q}_m$, $1 \leq m \leq M$, the distribution of $T(\mathbf{y})$ can be derived under the two hypothesis:

$$\begin{aligned} T(\mathbf{y}|\mathcal{H}_1) &\sim \mathcal{N}\left(\sum_{m=1}^M SNR_m, \frac{2}{N} \sum_{m=1}^M (1 + SNR_m)^2\right) \\ T(\mathbf{y}|\mathcal{H}_0) &\sim \mathcal{N}\left(0, \frac{4}{N}\right) \end{aligned}$$

As with the LLRT optimal rule, the resulting P_{HI} and P_{FH} can be calculated.

OR detection rule

In the OR rule, we declare the primary to be present if any of the radios declares the primary to be present. If $P_{HI_m} = P_{MD_m}$, $1 \leq m \leq M$ and $P_{FH_m} = 1 - P_{FA_m}$, $1 \leq m \leq M$ are the probability of harmful interference and probability finding a spectrum hole for radio m respectively, then the system probability of detection and false alarm (assuming independence across the radios) are given by:

$$\begin{aligned} P_{HI} &= \prod_{m=1}^M P_{HI_m} \\ P_{FH} &= \prod_{m=1}^M P_{FH_m} \end{aligned}$$

It is easy to show that the best performance of the OR rule is obtained when all the radios have the same probability of false alarm. This is the same phenomenon seen in [73] where the best performance is obtained when the radios estimate their own noise and set their own thresholds.

Max detection rule

Another form of the OR rule is the max rule. Each radio computes its test statistic $T(\mathbf{y}_m)$ and the overall test statistic is the max of the individual test statistic i.e.

$$T(\mathbf{y}) = \max_{1 \leq m \leq M} T(\mathbf{y}_m) - \sigma_{w_m}^2$$

Subtracting the noise power is useful when the two radios has different noise figures.

Median detection rule

Taking the median of the test statistics computed by each radio is another way of combining measurements from various radios. Each radio computes its own test statistic $T(\mathbf{y}_m)$. Let $T(\mathbf{y}_{(1)}), \dots, T(\mathbf{y}_{(M)})$ represents the order statistics of $T(\mathbf{y}_1), \dots, T(\mathbf{y}_M)$ i.e. $T(\mathbf{y}_{(k-1)})$ represents the k^{th} smallest member of $T(\mathbf{y}_1), \dots, T(\mathbf{y}_M)$. The median rule can be implemented as:

$$T(\mathbf{y}) = \begin{cases} T(\mathbf{y}_{(\frac{M+1}{2})}) & M \text{ odd} \\ \frac{T(\mathbf{y}_{(\frac{M}{2})}) + T(\mathbf{y}_{(\frac{M}{2}+1)})}{2} & M \text{ even} \end{cases}$$

4.2.2 Performance of temporal cooperation strategies

In this section we will highlight the difference between the various rules for 2 radios under the two cases:

1. Equal SNRs — the two radios have equal SNRs, i.e. $[P_1 - N_0, P_2 - N_0] = [-10dB, -10dB]$.
2. Unequal SNRs — the two radios have unequal SNRs, i.e. $[P_1 - N_0, P_2 - N_0] = [-10dB, -15dB]$.

Figure 4.1 shows the performance of the different rules for the two cases. The LLRT Optimal/ *Averaging with discounting* rule performs the best in both scenarios since it weighs the energy estimate with the actual signal strength. When the SNRs are equal, the rule serves to increase the number of samples by a factor of M . This can be seen by eliminating λ' from (4.2). In doing so we get,

$$P_{HI} = 1 - Q \left(\frac{Q^{-1}(1 - P_{FH}) - \sqrt{\frac{NM}{2}} SNR}{SNR + 1} \right)$$

Comparing to the equation for a single radio [48] we see that the number effective of samples has increased by a factor of M .

The simple averaging rule performs the best when both radios see the same SNR. Averaging collects as much energy as the LLRT optimal rule and hence has the same performance. When the SNRs are unequal, the simple averaging rule does not distinguish between the high and low SNR user and ends up collecting noise from the low SNR radio. In fact, the rule's performance is worse than that of a single radio at the higher SNR. The performance of the median rule is the same as the simple averaging rule.

As stated earlier, the best performing OR rule has the same performance as the max rule. When the SNRs at the two radios is very different, both rules perform worse than a single radio at the higher SNR. The OR/Max rule thrive on diversity and count on the performance of the best user. Hence their performance is comparable to the performance of the simple averaging rule when unequal SNRs are involved.

One point is important to understand: *Cooperation is adversely affected by the presence of radios in deeply faded/shadowed environments.* Eliminating the results of these radios from the cooperation process would boost performance. This issue is revisited in Chapter 5 where we use cooperation among multiband radios to weed out badly faded/shadowed radios.


4.2.3 Temporal strategies in the real world

We build a mobile, wideband testbed to capture DTV signals in the 500-700MHz bands. The methodology to collect the data and details of the experimental apparatus can be found in Appendix B. These experiments were performed at multiple locations around Berkeley, CA. Table 4.5 lists these locations and their shadowing characteristics. These locations provide a rich set of spectral environments - indoors, indoors near windows, outdoor with line-of-sight and outdoors with shadowing obstacles of different sizes.


A total of 12 channels were reliably detected by the setup. Table 4.3 lists these channels and their transmission attributes. Note that these channels have different transmit powers and are located at two different towers ((10 from Sutro tower and 2 from Mt. San Bruno). Figure 4.2 shows a snap shot of the channels taken on Dr. Sahai's deck. The 12 channels are clearly visible. Also visible are the pilot signals located on the left edge of each 6MHz band. The dynamic range (40dB) of the testbed does not allow us to measure extremely low power signals. This dynamic range issue is apparent in Table 4.4 which compares the number of signals captured by the testbed with the predicted number of signals that can be detected/excluded.

If we assume that these 12 are the only primaries that we need to protect then we can treat all other signals as noise/interference and aim to reuse those bands. This sets up nice hypothesis testing problem in the real world in which we can evaluate the performance of cooperative sensing.

We assume that each measurement location represents a cognitive radio making mea-



a



b

Figure 4.1: Cooperation between two radios using various rules with (a) the two radios at the same SNR (-10dB) and (b) the radios are at different SNRs (-10dB, -15dB). The *averaging with discounting rule* performs the best in both cases. The simple averaging rule works best when the two radios are at the same SNRs but does distinguish between the high and low SNR radio and ends up collecting noise from the low SNR radio. Performance of the OR/Max rule is comparable to the performance of the simple averaging rule when unequal SNRs are involved.

Location Name	Shadowing environment	Typical gain	Time of Measurement	Number of measurement points
BWRC Lab	Room with no windows	40dB	Afternoon (10am-2pm) and Night (2am-6pm)	296
BWRC Classroom	Room on 2nd floor with windows	20dB	6pm-10pm	70
BWRC Basement	First floor room with no windows	40dB	7pm-11pm	72
Dr. Sahai's Deck	Line-of-sight to Sutro in the West	0dB	3pm-8pm	67
Milvia and California streets	1 or 2 story houses on either side	20dB	9pm-1am	188
Oxford Street	Large buildings blocking Sutro	20dB	2am-4am	81

Table 4.2: Experimental locations and their attributes

Station Name	Channel	Transmitter Location	Transmit Power (kW)	Predicted Receive Power (dBm)	HAAT (m)
KOFY	19	Sutro	383	-27dBm	418
KTSF	27	San Bruno	403	-33dBm	500
KPIX	29	Sutro	1000	-23dBm	506
KQED	30	Sutro	709	-25dBm	509
KMTP	33	Sutro	50	-27dBm	491
KFSF	34	Sutro	150	-33dBm	419
KRON	38	Sutro	712	-24dBm	446
KCNS	39	Sutro	1000	-25dBm	428
KKPX	41	San Bruno	418	-27dBm	1000
KCSM	43	Sutro	576	-29dBm	428
KTVU	44	Sutro	811	-29dBm	433
KBCW	45	Sutro	400	-30dBm	446

Table 4.3: Experimental locations and their attributes




Figure 4.2: Power Spectral Density of a measurement performed at Dr. Sahai’s house. The 12 channels from Mt Sutro and Mt San Bruno are clearly visible. Also visible are the pilot signals located on the left edge of each 6MHz band.

	Number of Channels
Detected by the testbed	12
Detectable channels	32
Channels to be excluded by the FCC’s geo-location co-channel rules	34
Channels to be excluded by the FCC’s geo-location co-channel and adjacent channel rules	42
Channels that will be excluded by the -114dBm co-channel sensing rules	46
Channels that will be excluded by the -114dBm co-channel and adjacent channel sensing rules	47

Table 4.4: Comparison of the testbed results with the results using the ITU propagation curves and the FCC DTV database.

Location Name	Shadowing environment
BWRC Lab	Room with no windows
BWRC Classroom	Room on 2nd floor with windows
BWRC Basement	First floor room with no windows
Dr. Sahai's Deck	View of Sutro in the West
Milvia and California streets	1 to 2 story houses on either side
Oxford Street	Large buildings blocking Sutro

Table 4.5: Experimental locations and their attributes

measurements in all 33 bands. These radios form a ‘Berkeley’ network of radios which will perform cooperative sensing to make a decision for each of the 33 bands. Each radio chooses a random set of $M - 1$ neighbors to exchange their measurements with. The measurements are combined using the rules mentioned earlier. Since we do not know the true SNRs at each radio, we cannot use the *Averaging with discounting* rule. However we will apply the Averaging, max and Median rules.

Figure 4.3 shows the performance of different rules for the case of $M = 5$ cooperating radios. P_{HI} at a given threshold, is defined as the fraction of ‘on’ channels times measurement locations that are mis-detected. Similarly, P_{FH} at a given threshold, is defined as the fraction of ‘off’ channels times measurement locations that are detected as being ‘off’. While the performance advantage of the averaging rule over the max rule is expected from the previous section, the outstanding performance of the median rule is a surprise. From statistics, we know that the median is a robust statistic since it is not effected by outliers. It is this property of the median that is helping in this case. For the ‘off’ channels we see noise as well as interference (signals in the lab, spurious emission, low strength TV signals) in the band. The interference signals acts as outliers with respect to the noise distribution and hence are well handled by the median rule.

4.3 Cooperation for spatial holes

For spatial holes, cooperation is looked upon as a tool to increase the Weighted Probability of Area Recovered (*WPAR*). As in Section 4.2 we assume that a group of M cognitive radios are listening to the primary signal on a given frequency band. As for temporal holes, the radios runs a detection algorithm that outputs a binary decision D about the state of the primary band: 1-used/0-unused. For simplicity, we assume that each radio gets a perfect estimate² of the received primary power P_i (in dB) $i = 1, \dots, M$. We make this assumption

²Shadowing and fading results in a certain received primary power at the cognitive radio. We can converge to this power estimate by running a radiometer for very long sensing times, i.e., $N \rightarrow \infty$. This assume a




Figure 4.3: Cooperation between hypothetical radios operating at each measurement point in Berkeley, CA. Each radio is assumed to choose a set of 4 neighbors among the Berkeley network and cooperate with them. The performance of the Median rule is better than the performance of the averaging/max rule due to its ability to deal with outliers.

to isolate the gains due to cooperation from those due to a longer effective sampling time.

Each of the received signal strengths is written as $P_m = p_t - (10 \log_{10} r_m^\alpha + S_m + M_m)$, where p_t is the transmit power of the primary signal, r_m is the distance from the m^{th} radio to the TV tower, and S_m and M_m are respectively the losses due to shadow and multipath fading at the m^{th} radio. All cooperating radios are assumed to be located at approximately the same distance from the TV tower, i.e., $r_m = r$ for all $m = 1, 2, \dots, M$. This models the case when the scale of cooperation is much smaller than the scale of the primary transmissions.³ This assumption also guarantees that all the cooperating radios are trying to identify the same spectrum hole in space.

To start with, shadowing and multipath are modeled to be independent across the different radios. It is safe to assume that the $\{M_m\}$ are independent⁴ of each other since multipath is independent at distances on the order of a few wavelengths [50]. By contrast, shadowing is independent only on a much larger spatial scale [74]. Even though independence might not be an accurate modeling assumption, we first analyze cooperative gains under this best-case assumption. Section 4.4.2 computes the loss in performance if the shadowing is not independent.

4.3.1 Maximum-likelihood detector: soft-decision combining

Our goal is to find the optimal estimate of the distance r , given the vector of received power observations $(P_1, P_2, \dots, P_M)$. When the model is completely known, the optimal detector is the ML detector. We assume a nominal Gaussian model for both the shadowing and multipath distribution, i.e., $P_m \sim \mathcal{N}(\mu(r), \sigma^2)$, where $\mu(r)$ is $\mathbb{E}[p_t - (10 \log_{10} r_m^\alpha + S_m + M_m)]$ (in our plots the parameter α is chosen to make the propagation behave quantitatively similar to the ITU's F(50, 90) curves for F_{HI} and the F(50, 50) curves for P_{FH} [42]). Under this model, the mean of the received power is dependent on its distance from the TV tower and the standard deviation is independent of the distance from the tower. Using this model, it is easy to see that the ML detector is equivalent to

$$\frac{1}{M} \sum_{m=1}^M P_m \underset{D=0}{\overset{D=1}{\gtrless}} \lambda. \quad (4.4)$$

This detector computes the average received signal power on a dB scale (This is an example of a soft-decision combining rule since the radios have to send their received power

slow fading situation.

³In reality, the radial footprint of the cooperating radios has to be dealt with as a minor increase in the no-talk radius r_n . However, we assume that the footprint of cooperation is much smaller than the margin $(r_n - r_p)$ and thus ignore this small effect.

⁴This is not true strictly speaking. In general, M_m 's are conditionally independent given the shadowing environment since the shadowing environment can determine if there is or is not a strong line-of-sight path.

values to a central combiner rather than just send 1-bit decisions) and compares it to a threshold λ . The frequency band is declared free if the mean signal power is less than λ .

Assuming that $P_m \sim \mathcal{N}(\mu(r), \sigma^2)$, we have $\frac{1}{M} \sum_{m=1}^M P_m \sim \mathcal{N}(\mu(r), \frac{\sigma^2}{M})$. Therefore if we assume that the primary user also trusts the nominal model,

$$\begin{aligned}
 F_{HI} &= \max_{0 \leq r \leq r_n} \mathcal{P} \left(\frac{1}{M} \sum_{m=1}^M P_m \leq \lambda | r_{actual} = r \right) \\
 &\stackrel{(a)}{=} \mathcal{P} \left(\frac{1}{M} \sum_{m=1}^M P_m \leq \lambda | r_{actual} = r_n \right) \\
 &= 1 - \mathcal{Q} \left(\frac{\lambda - \mu(r_n)}{\frac{\sigma}{\sqrt{M}}} \right).
 \end{aligned} \tag{4.5}$$

where (a) is because the expression is maximized when $r_{actual} = r_n$.

The detector threshold λ must be chosen such that $F_{HI} \leq F_{HI}^{target}$. Hence (4.5) gives

$$\lambda = \frac{\sigma}{\sqrt{M}} \mathcal{Q}^{-1}(1 - F_{HI}^{target}) + \mu(r_n). \tag{4.6}$$

For this choice of λ , the probability of finding a hole is

$$\begin{aligned}
 P_{FH}(r) &= \mathcal{P} \left(\frac{1}{M} \sum_{m=1}^M P_i \leq \lambda | r_{actual} = r \right), \quad r > r_n \\
 &= 1 - \mathcal{Q} \left(\frac{\lambda - \mu(r)}{\frac{\sigma}{\sqrt{M}}} \right).
 \end{aligned} \tag{4.7}$$

The WPAR can be computed by substituting (4.7) into (3.8) i.e.

$$\begin{aligned}
 WPAR &= \int_{r_n}^{\infty} P_{FH} w(r)(r) r dr \\
 &= \int_{r_n}^{\infty} \left(1 - \mathcal{Q} \left(\frac{\lambda - \mu(r)}{\frac{\sigma}{\sqrt{M}}} \right) \right) w(r) r dr
 \end{aligned} \tag{4.8}$$

where λ is given in (4.6).

Figure 4.4 shows the performance of the maximum likelihood detector for several values of

the number of cooperating radios M . It is clear that the performance significantly improves even with a few cooperating radios ($M=5$). If $M \rightarrow \infty$, all the area is eventually recovered.

It is interesting to compare the spatial ML rule to the LLRT optimal temporal rule – the *Averaging with discounting* rule as discussed in Section 4.2.1. In the spatial domain, there is a mapping from the mean of the signal ($\mu(r)$) and the actual distance r . The spatial ML rule aims to converge to that mean to predict its distance from the tower (r) with accuracy. In the temporal domain, the LLRT optimal rule assigns a low weighting to radios with low SNRs thereby ensuring that the gap between the mean of the test statistic for the case of primary present and primary absent is as large as possible.

4.3.2 OR-rule detector: hard decision combining

The spatial OR rule is the same as the temporal OR rule (as discussed in Section 4.2.1). Each radio compares its received power to a threshold. It tentatively declares the band free to use if its received power is below the threshold. Then, each radio sends its tentative 1-bit sensing decision to the central combiner (there are other ways to fuse decision based on the radio topology [75]). The global decision is to use the band only if all the sensors declare the band to be free.

The safety/performance of the OR rule is easy to compute. Assume that each radio uses a detector threshold $\lambda_{radio,M}$ (detection threshold for a single radio assuming a total of M cooperating radios). Then, the fear of harmful interference for each radio is given by

$$\begin{aligned} F_{HI,radio} &= \mathcal{P}(P_m \leq \lambda_{radio,M} | r_{actual} = r_n) \\ &= 1 - \mathcal{Q}\left(\frac{\lambda_{radio,M} - \mu(r_n)}{\sigma}\right). \end{aligned}$$

The fear of harmful interference for the single radio with infinite samples can be computed by knowing a single quantile of the distribution (namely $F_{r_n}(\lambda_{radio,M})$, where F_{r_n} is the distribution of the fading at r_n). The system of cognitive radios causes harmful interference only if every radio individually fails to detect the primary user, and so by independence:

$$F_{HI,system} = (F_{HI,radio})^M.$$

In order to meet the target F_{HI} , each radio must choose a $\lambda_{radio,M}$ satisfying

$$\begin{aligned} \mathcal{Q}\left(\frac{\lambda_{radio,M} - \mu(r_n)}{\sigma}\right) &= 1 - [F_{HI}^{target}]^{\frac{1}{M}} \\ \Rightarrow \lambda_{radio,M} &= \sigma \mathcal{Q}^{-1}\left(1 - [F_{HI}^{target}]^{\frac{1}{M}}\right) + \mu(r_n). \end{aligned} \tag{4.9}$$




Figure 4.4: Performance of the ML detector in (4.4) with complete knowledge of the fading/shadowing distribution.

As $M \rightarrow \infty$, it is immediately clear that for any given target F_{HI} , the term $[F_{HI}^{target}]^{\frac{1}{M}} \rightarrow 1$ and so the threshold (and thus target quantile) approaches the case of extremely favorable fading.

For such a choice of $\lambda_{radio,M}$, the probability of finding a hole at a radial distance of r is given by

$$\begin{aligned} P_{FH,radio}(r) &= \mathcal{P}(P_m \leq \lambda_{radio,M} | r_{actual} = r) \\ &= 1 - \mathcal{Q}\left(\frac{\lambda_{radio,M} - \mu(r)}{\sigma}\right). \end{aligned}$$

The system finds a hole only if all the radios find a hole.

$$\begin{aligned} P_{FH,system}(r) &= (P_{FH,radio})^M \\ &= \left[1 - \mathcal{Q}\left(\frac{\lambda_{radio,M} - \mu(r)}{\sigma}\right)\right]^M. \end{aligned} \quad (4.10)$$

Substituting (4.10) in (3.8), we get the WPAR for the OR rule.

4.3.3 The Median Rule

We now present the spatial counterpart of the temporal Median Rule. Let $P_{(1)}, \dots, P_{(M)}$ represent the order statistics of the received powers $P_1, \dots, P_M$ i.e $P_{(k)}$ is the k^{th} smallest member of $P_1, \dots, P_M$.

$$P_{med}(r) = \begin{cases} P_{(\frac{M+1}{2})} & M \text{ odd} \\ \frac{P_{(\frac{M}{2})} + P_{(\frac{M}{2}+1)}}{2} & M \text{ even} \end{cases}$$

In this section, however, we will take a slightly different approach to the median rule and implement it as a hard decision rule. Being able to implement a rule as a hard decision rule has two advantages:

- Its allows the radio to perform a simple threshold comparison and ship a single decision bit to the controller.
- Curtailing each radio output to a single bit restricts the error causing potential of a radio and allows us to bound the performance of the system if a fraction of the radio are misbehaving (see Section 4.4.3).

In general, lets consider the $(\lceil \zeta M \rceil)^{th}$ percentile rule (remember, that the median is implemented with $\zeta = 0.5$). Since ζM may not be an integer, we use the $\lceil \cdot \rceil$ function. For

even M , this does not give the correct value of the median but the difference diminishes as $M \rightarrow \infty$.

We compare this quantity to a threshold and compute F_{HI} and P_{FH} except in this case will write the threshold as $F_{r_n}^{-1}(p_{r_n})$ where $F_{r_n}(\cdot)$ is the CDF of the distribution of the received signal power at r_n . The Fear of Harmful Interference ($F_{HI,\zeta}$) and Probability of Finding a Hole at distance r ($P_{FH,\zeta}(r)$) are given by:

$$\begin{aligned} F_{HI,\zeta} &= \mathcal{P}_{F_r} (P_{(\lceil \zeta M \rceil)} \leq F_{r_n}^{-1}(p_{r_n}) | r = r_n) \\ &\stackrel{(a)}{=} \sum_{i=\lceil \zeta M \rceil}^M \binom{M}{i} (p_{r_n})^i (1 - p_{r_n})^{M-i} \end{aligned} \quad (4.11)$$

where (a) is because, for $P_{\lceil \zeta M \rceil}$ to be less than $F_{r_n}^{-1}(p_{r_n})$ at least $\lceil \zeta M \rceil$ of the $P_1, \dots, P_M$ have to be less than $F_{r_n}^{-1}(1 - p_{r_n})$.

and,

$$\begin{aligned} P_{FH,\zeta}(r) &= \mathcal{P}_{F_r} (P_{(\lceil \zeta M \rceil)} \leq F_{r_n}^{-1}(p_{r_n}) | r = r_n) \\ &\stackrel{(b)}{=} \sum_{m=\lceil \zeta M \rceil}^M \binom{M}{m} (p_r)^m (1 - p_r)^{M-m} \end{aligned} \quad (4.12)$$

where (b) is due to the following facts:

- Our model of shadowing+fading is such that the distribution of fading+shadowing at different distances from the tower form a location-scale family i.e. the distribution of fading+shadowing at two distances d_1 and d_2 ($F_{d_1}(\cdot)$ and $F_{d_2}(\cdot)$) are related as: $F_{d_1} \left(\frac{x - \mu(d_1)}{\sigma(d_1)} \right) = F_{d_2} \left(\frac{x - \mu(d_2)}{\sigma(d_2)} \right)$ where $\mu(\cdot)$ and $\sigma(\cdot)$ are the mean and standard deviation of fading+shadowing and serve as the location and scale parameters, respectively. Furthermore, we have assumed that $\sigma(d_1) = \sigma(d_2)$ i.e. the variance in fading+shadowing does not change with distance⁵. Hence we get the following equality:

$$F_{d_1}(x - \mu(d_1)) = F_{d_2}(x - \mu(d_2)) \quad (4.13)$$

- Equation (4.13) allows us to calculate p_r as:

⁵In reality, shadowing should increase with distance since there are larger chances of the signal going through obstacles. However, there is no good model for this increase in shadowing with distance.

$$p_r = F_r \left(F_{r_n}^{-1}(p_{r_n}) + \mu(r_n) - \mu(r) \right) \quad (4.14)$$

We can cast the above problem into a hard detection problem in the following manner: Define,

$$T(\mathbf{P}) = \frac{1}{M} \sum_{m=1}^M \mathbf{1}_{\{P_m > F_{r_n}^{-1}(p_{r_n})\}} \quad (4.15)$$

i.e. we compute the fraction of radios that trigger the detection threshold. Then the corresponding $F_{HI}, (F_{HI,hd})$ is given by:

$$\begin{aligned} F_{HI,hd} &= \mathcal{P}_{F_{r_n}}(T(\mathbf{P}) \leq 1 - \zeta | r = r_n) \\ &\stackrel{(a)}{=} \sum_{m=0}^{\lceil (1-\zeta)M \rceil} \binom{M}{m} (1 - p_r)^m p_r^{M-m} \\ &\stackrel{(b)}{=} \sum_{m=0}^{\lceil (1-\zeta)M \rceil} \binom{M}{M-m} (1 - p_r)^m p_r^{M-m} \\ &\stackrel{(c)}{=} \sum_{m'=\lceil \zeta M \rceil}^M \binom{M}{m'} (1 - p_r)^{M-m'} p_r^{m'} \end{aligned} \quad (4.16)$$

where (a) is because $\sum_{m=1}^M \mathbf{1}_{\{P_m > F^{-1}(p_{r_n})\}}$ forms a Binomial distribution with parameters $M, 1 - p_{r_n}$ and hence the probability that it is less than $\lceil 1 - \zeta M \rceil$ is given by the standard Binomial expansion formula. (b) is because $\binom{M}{m} = \binom{M}{M-m}$ and (c) is formed by making the substitution $m' = M - m$.

In reality (4.16) is only true for ζM being an integer. However, only when ζM is an integer are changes in probability values expected hence the detection algorithm is only expected to select such values of ζ .

Similarly the probability of finding a hole is calculated as:

$$\begin{aligned} P_{FH,hd}(r) &= \mathcal{P}_{F_r}(T(\mathbf{P}) \leq 1 - \zeta | r) \\ &= \sum_{m=\lceil \zeta M \rceil}^M \binom{M}{m} (1 - p_r)^{M-m} p_r^m \end{aligned} \quad (4.17)$$

Comparing, (4.11) with (4.16) and, (4.12) with (4.17) we see that:

$$\begin{aligned} P_{HI,hd} &= P_{HI,\zeta} \\ P_{FH,hd} &= P_{FH,\zeta} \end{aligned}$$

i.e. there is a one-to-one mapping between the percentile test statistic and corresponding counting rule. We will utilize this correspondence for the rest of this chapter.

Figure 4.5 compares the performance of the median rule (implemented as using (4.15) with $\zeta = 0.5$) to the performance of the OR and averaging rule. The median rule has the best area recovery performance amongst all hard-decision rules.

4.3.4 Evaluating the relative performance of hard decision rules

In Figure 4.5, we saw that the median rule outperforms the OR rule. Is this always the case? In this section we seek the answer to this question. We look at the two user case for ease of understanding. With two users, only two hard decision rules are possible – the OR rule (the system declares the band unusable if any of the radios declares the band unusable) and the AND rule (the system declares the band usable if any of the radios declares the band usable). We setup these rules to meet the target Fear of Harmful Interference (F_{HI}). We will concentrate on the case of infinite samples and ignore all uncertainties in calculating F_{HI} . Furthermore, we assume that the distribution of fading+shadowing is a location-scale family as expressed in (4.13).

Let $\lambda_{OR,radio}$ and $\lambda_{AND,radio}$ be the thresholds for individual radios for OR and AND combining rules respectively. These thresholds have to be set to meet the target F_{HI} (F_{HI}^{target}).

For the OR rule,

$$\begin{aligned} \mathcal{P}_{F_r}(P_1 \leq \lambda_{OR,radio}, P_2 \leq \lambda_{OR,radio} | r = r_n) &\leq F_{HI}^{target} \\ F_{r_n}(\lambda_{OR,radio})^2 &\leq F_{HI}^{target} \\ \lambda_{OR,radio} &\leq F_{r_n}^{-1} \left(\sqrt{F_{HI}^{target}} \right). \end{aligned} \quad (4.18)$$

Similarly for the AND rule,

$$\begin{aligned} \mathcal{P}_{F_r}(P_1 \geq \lambda_{AND,radio}, P_2 \geq \lambda_{AND,radio} | r = r_n) &\geq 1 - F_{HI}^{target} \\ (1 - F_{r_n}(\lambda_{AND,radio}))^2 &\geq 1 - F_{HI}^{target} \\ \lambda_{AND,radio} &\leq F_{r_n}^{-1} \left(1 - \sqrt{1 - F_{HI}^{target}} \right). \end{aligned} \quad (4.19)$$




Figure 4.5: Performance of infinite-sample cooperation using different fusion rules. The ML Rule performs the best while the OR rule performs the worst. The Median rule (majority-vote) has the best performance among the hard-decision rules.

Next we will calculate the probability that a radio finds a hole using the $\lambda_{OR,radio}$ threshold ($P_{FH,OR,radio}(r)$):

$$P_{FH,OR,radio}(r) = \mathcal{P}_{F_r}(P_m \leq \lambda_{OR,radio}) \quad (4.20)$$

$$\stackrel{(a)}{=} F_{r_n}(\lambda_{OR,radio} - \mu_r + \mu_{r_n}) \quad (4.21)$$

where (a) follows from applying (4.13) to (4.20).

Similarly for the AND rule,

$$\begin{aligned} P_{FH,AND,radio}(r) &= \mathcal{P}_{F_r}(P_m \leq \lambda_{AND,radio}) \\ &= F_{r_n}(\lambda_{AND,radio} - \mu_r + \mu_{r_n}) \end{aligned} \quad (4.22)$$

And finally, the system level probabilities of finding a hole for the two rules are given by:

$$\begin{aligned} P_{FH,OR}(r) &= P_{FH,OR,radio}(r)^2 \\ P_{FH,AND}(r) &= 1 - (1 - P_{FH,AND,radio}(r))^2 \end{aligned}$$

The relationship between the above quantities is shown in Figure 4.6. For the OR rule to perform better than the AND rule (i.e. $P_{FH,OR}(r) > P_{FH,AND}(r)$), the following relationship must hold between $P_{FH,OR,radio}(r)$ and $P_{FH,AND,radio}(r)$.

$$P_{FH,AND,radio}(r) \leq 1 - \sqrt{1 - P_{FH,OR,radio}(r)^2} \quad (4.23)$$

$$(4.24)$$

For example, if $P_{FH,OR,radio}(r)$ is 0.7 then $P_{FH,AND,radio}(r)$ must be less than 0.286 (a difference of 0.414) for the OR rule to outperform the AND rule. If however, $F_{HI}^{target} = 0.01$, the difference in values between $\sqrt{F_{HI}^{target}}$ and $1 - \sqrt{1 - F_{HI}^{target}}$ is just 0.005. Hence if the OR rule has to outperform the AND rule for $F_{HI,t} = 0.01$, the slope of the function $F_{r_n}(\cdot)$ has to increase significantly when moving from $\lambda_{OR,radio}$ to $\lambda_{OR,radio} - \mu(r) + \mu(r_n)$. This case is illustrated in Figure 4.6 where the normal distribution is used for F_{r_n} . An F_{HI}^{target} of 0.01, for the OR rule, translates into a radio level F_{HI} of 0.1. The same for the AND rule translates into a radio level F_{HI} of 0.005. These values are marked on the y-axis in Figure 4.6. These radio level F_{HI} requirements translates into threshold requirements $\lambda_{OR,radio}$ and $\lambda_{AND,radio}$ as per (4.18) and (4.19), respectively. At some distance $r > r_n$

we wish to determine $P_{FH,OR,radio}(r)$ and $P_{FH,AND,radio}(r)$. Since the fading+shadowing distribution forms a location-scale family, we can evaluate the two quantities by using the distribution at r_n as expressed in (4.21) and (4.22). This is shown in the figure by moving the threshold by $\mu(r_n) - \mu(r)$ on the x-axis. This shift translates into a large difference between $P_{FH,OR,radio}$ and $P_{FH,AND,radio}$ due to the large slope of F_{r_n} at that point which enables the OR rule to outperform the AND rule. This suggests that the slope of the function F_{r_n} has a lot to do with the relative performance of various rules.

Figure 4.7 plots $P_{FH,OR}(r) - P_{FH,AND}(r)$ against F_{HI}^{target} and r for two cases of F_{r_n} . In the first case F_{r_n} is $\mathcal{N}(\mu(r_n), \sigma^2)$ and in the second case, F_{r_n} is $\mathcal{U}[\mu(r_n) - \sqrt{3}\sigma, \mu(r_n) + \sqrt{3}\sigma]$. The mean and variance of the distribution in two cases is the same. The large slope of the normal distribution is key for the OR achieving a better performance for the normal distribution. The opposite is true for the uniform distribution where the single slope at all values helps the AND rule.

Since the slope of F_{r_n} is the density function at r_n , is there a requirement on the density function to predict which rule will perform best? For this we need to move to the percentile rules has discussed in Section 4.3.3 i.e. we compute the $\lceil \zeta M \rceil^{th}$ percentile and use it for detection. Let $\tilde{P}_{\lceil \zeta M \rceil}$ represent the estimate of the $\lceil \zeta M \rceil^{th}$ percentile. According to [76], the mean and variance of this estimate is given by:

$$\mathbb{E}[\tilde{P}_{\lceil \zeta M \rceil}] = F_{r_n}^{-1}(\zeta) - \frac{\zeta(1-\zeta) F_{r_n}''(F_{r_n}^{-1}(\zeta))}{2(M+2) (F_{r_n}'(F_{r_n}^{-1}(\zeta)))^3} + \mathcal{O}(1/M^2) \quad (4.25)$$


$$\mathbb{V}ar[\tilde{P}_{\lceil \zeta M \rceil}] = \frac{\zeta(1-\zeta)}{(M+2) (F_{r_n}'(F_{r_n}^{-1}(\zeta)))^2} + \mathcal{O}(1/M^2) \quad (4.26)$$

Examining (4.26) we see that the variance of the estimate is proportional to $(\zeta(1-\zeta))$ which is maximized at $\zeta = .5$ which suggests that the extreme percentiles should exhibit the fastest convergence. However, the term $F_{r_n}'(F_{r_n}^{-1}(\zeta))$ in the denominator is a function of the distribution. Figure 4.8 plots the value of the $\frac{\zeta(1-\zeta)}{(F_{r_n}'(F_{r_n}^{-1}(\zeta)))^2}$ for various distributions. For the exponential distribution the 0th percentile converges fastest, for the uniform it is the extremal percentiles, for the normal its the median and for the Rayleigh (in dB units) its the 80th percentile.


4.4 Cooperation with Uncertainty

For a single radio uncertainty in noise+interference is of concern since it imposes a lower limit on the SNR of the primary signal that can be reliably detected [19]. Similarly, for multiple radios uncertainties can limit the achievable WPAR since the target F_{HI} needs to be met under all uncertainties. Three types of uncertainties are discussed in this section: incomplete





(a)



(b)

Figure 4.7: Difference in the performance of the OR and AND rules ($M = 2$) for (a) Normal distribution (b) Uniform distribution. The large slope around the median is key to the OR achieving a better performance for the Normal distribution. The opposite is true for the Uniform distribution where the single slope at all values helps the AND rule.




Figure 4.8: The variance of the quantile estimators for various quantiles for different distributions – normal, uniform and exponential. For the exponential distribution the 0th percentile converges fastest, for the uniform it is the extremal quantiles, for the normal its the mean and for the Rayleigh (in dB) its the 80th percentile.

fading model, loss of spatial diversity and misbehaving/unreliable/malicious radios.

4.4.1 Performance of cooperation under quantile models

The improvements with cooperation illustrated in Figure 4.4 assume complete consensus regarding the fading distribution. In reality it is likely that the primary user of the channel does not trust the nominal Gaussian models for shadowing and fading distributions. The cost of addressing this distrust of primary users is a reduced performance for the same value of safety.

Under the independent fading assumption, it is illustrative to use the quantile models discussed in Section 3.4.3 for each received power P_m . We start with a single quantile that can be optimized. Let the class of marginal distributions satisfying the β th quantile constraint be denoted by $\mathbb{F}_r$. For simplicity, consider the two cooperating radios case, i.e., $M = 2$.

Performance of the ML detector under quantile models

The maximum-likelihood estimate detector under uncertain fading distributions (even for a single-quantile uncertainty model) does not even make sense. Hence, we do not attempt to solve for the best possible detector under modeling uncertainties. Instead, we continue to work with the averaging detector given in (4.4). As discussed in Section 4.3.1, this detector is the ML detector under perfectly modeled Gaussian fading.

For this detector, we can show that for a given choice of quantile β , the best choice of λ that minimizes F_{HI} is $\lambda = \gamma(r_n, \beta)$, where $\gamma(r_n, \beta)$ is the β th quantile threshold in (3.11). For this choice of λ , the fear of harmful interference is given by

$$\begin{aligned} F_{HI} &= \sup_{0 \leq r \leq r_n} \sup_{F_r \in \mathbb{F}_r} \mathcal{P}_{F_r} \left(\frac{P_1 + P_2}{2} \leq \lambda \mid r_{actual} = r \right) \\ &= 1 - (1 - \beta)^2. \end{aligned} \quad (4.27)$$

The expression for F_{HI} in (4.27) can easily be derived graphically from Figure 4.9. In this figure, the P_1, P_2 plane is divided into four quadrants as marked by the dashed-dotted lines (red). The single quantile constraint on the marginal distributions can be written as probability mass constraints within each quadrant. The averaging detector in (4.4) for a fixed λ can be drawn as a straight line dividing the P_1, P_2 plane into two half planes (the solid (blue) line in Figure 4.9). If the received power (P_1, P_2) falls in the shaded region, the band is declared ‘free to use’, otherwise the band is declared ‘used’. Hence, the probability of harmful interference is the supremum of the probability mass in the shaded region, where the supremum is taken over all distributions $F \in \mathbb{F}_\beta$. Similarly, the probability of finding a hole is the probability mass in the unshaded region, under the nominal distribution.

If $\lambda < p_{th}$, where $p_{th} = \gamma(r_n, \beta)$ (the detector line is on the left of the black dot in the




Figure 4.9: The averaging detector for two user cooperation under the single-quantile fading model. The solid (black) box in the perimeter represents the P_1, P_2 plane, the solid (blue) line represents the 2-user ML detector in (4.4), and the dashed (red) lines represent the quantiles describing the distribution of P_1 and P_2 . The shaded area represents the region of the received power pairs (P_1, P_2) for which the detector declares the band unused and the unshaded area represents the region where the detector declares the band as used. The threshold p_{th} in the figure is used to denote $\gamma(r_n, \beta)$.

figure), then F_{HI} is the sum of probabilities in quadrants *II*, *III*, and *IV*. This is because one can always choose a distribution that satisfies the quantile constraints and puts all the probability mass in quadrants *II*, *III*, and *IV* within the shaded region. Thus, in this case $F_{HI} = 1 - (1 - \beta)^2$. On the other hand if $\lambda \geq p_{th}$, then $F_{HI} = 1$. Therefore, the optimal choice of λ for a given quantile β that minimizes F_{HI} and maximizes $WPAR$ is $\lambda = \gamma(r_n, \beta)$.

Assuming that the β th quantile for the marginal distribution is the same as that of the nominal Gaussian $\mathcal{N}(\mu(r_n), \sigma^2)$, we have $\gamma(r, \beta) = \mu(r) + \sigma \mathcal{Q}^{-1}(1 - \beta)$. To evaluate the $WPAR$, the nominal fading distribution can be assumed so we have:

$$\begin{aligned}
 P_{FH}(r) &= \mathcal{P}\left(\frac{P_1 + P_2}{2} \leq \lambda | r_{actual} = r\right) \\
 &= 1 - \mathcal{Q}\left(\frac{\lambda - \mu(r)}{\frac{\sigma}{\sqrt{2}}}\right) \\
 &= 1 - \mathcal{Q}\left(\frac{\mu(r_n) + \sigma \mathcal{Q}^{-1}(1 - \beta) - \mu(r)}{\frac{\sigma}{\sqrt{2}}}\right). \tag{4.28}
 \end{aligned}$$

The $WPAR$ can be computed by substituting $P_{FH}(r)$ from (4.28) into (3.8).

Figure 4.10 plots the performance of the averaging detector under the single-quantile model for the fading distribution. The dashed curve (blue) is the performance of the averaging detector when the fading distribution is completely known. The solid curve (black) is the performance of the averaging detector under minimal knowledge of the fading distribution, i.e., with knowledge of a single quantile. From the figure it is clear that the 2-user averaging detector is highly non-robust to uncertainties in the fading distribution. This shows that blindly using the form of the ML detector (averaging) assuming complete knowledge can be disastrous under modeling uncertainties.

The performance of the averaging detector improves if we assume multiple quantile knowledge for shadowing and fading distributions. The performance is shown in Figure 4.11 and it is clear that the performance of the averaging detector improves as we learn more quantiles about the fading distribution. However, the first few quantiles learned give more performance improvement than the later ones — with performance approaching that of a fully trusted nominal model as the number of trusted quantiles increase.

Figure 4.12 shows the performance of the averaging detector under the single-quantile model for $M \geq 2$, with $F_{HI} = 10^{-2}$. The black curve is the averaging detector with single-quantile knowledge, the light blue curve is the ‘OR-rule’ and the purple curve is the median rule and the dark blue curve is the averaging detector with complete trust in the nominal distributional (in this case the averaging detector is the ML detector). Note that whereas the performances of the OR rule, median and the averaging detector under complete distributional knowledge improve with increasing M , the averaging detector with single-quantile knowledge does worse as the number of cooperating radios M increases! This is because the number of quantiles contributing towards F_{HI} increases exponentially with the number of




Figure 4.10: Averaging detector for two user cooperation: performance under complete knowledge of the fading/shadowing distribution versus performance under the single-quantile uncertainty model for fading/shadowing distribution.




Figure 4.11: Performance of the averaging detector under varying quantile knowledge for the shadowing and fading distributions. The quantiles were chosen to minimize F_{HI} for a given WPAR value.

users.⁶ This shows the non-robustness of blindly using the form of the ML detector even under uncertainties.

Performance of the OR rule under quantile models

As compared to the averaging rule, the OR rule is much more robust to uncertainties in the fading distribution. By substituting (4.9) into (4.10) we see that under the single-quantile model of uncertainty and nominal Gaussian fading, the WPAR tends to 1 for the OR rule as the number of cooperating users increases. Figure 4.10 compares the performance of the OR-rule detector with the averaging detector with complete knowledge, and the averaging detector with single-quantile knowledge for the case of two cooperating radios ($M = 2$) while Figure 4.12 compares the same for the case of $M > 2$.

Gains by using the OR rule are accomplished by taking the single quantile to correspond to ever more favorable fading realizations. This is problematic since it involves achieving a consensus regarding the rare best fading events — this is as implausible as achieving a consensus regarding the rare worst fading events. In addition, there is a very natural deployment scenario — outdoors on a rooftop — in which the best fading events cannot be too good. This is a little counterintuitive, but remember that multipath fading can result in both destructive and constructive interference. Indoors or in an urban canyon, the best-case fading corresponds to lucky constructive interference. Outdoors, with a dominant line-of-sight path, such constructive interference cannot occur.

Strangely enough, when cooperation is involved, it is this possibility of a clean line-of-sight path that requires the uncertainty model $\mathbb{F}_r$ to impose a bound on how lucky the fading can be. This effectively caps λ_{radio} to the fade that corresponds to a single line-of-sight path. Once the number of cooperating users has reached a point that they can support the desired F_{HI} using that particular quantile, there is no further benefit to increasing the number of users if the OR rule is used. In fact, the performance will drop if cooperating users are blindly added as there is an increased chance of a single user (who happens to be in a rich multipath environment) getting a very lucky constructive fade and thereby deciding that they are within the no-talk radius. The kinked-light-blue curve in Figure 4.12 illustrates what happens if the uncertain fading model includes the possibility for a line-of-sight path at the 10%-best quantile.

⁶To understand why the averaging detector is so vulnerable to uncertainties of this form, remember that the empirical average is very sensitive to outliers. A single very negative number can dominate the entire average. Quantile models can be thought of as histograms. As such, they do not impose any restriction on how negative the rare bad fading can be since the outermost bin of a histogram includes everything from $-\infty$ on up. Consequently, the averaging detector cannot afford even a single user experiencing a fade from that lowermost bin.




Figure 4.12: Performance of detectors as a function of the number of cooperating users M for a fixed target $F_{HI} = 10^{-2}$.

Performance of the Median rule under quantile models

Like the OR rule, the median rule is robust to uncertainties in the fading distribution. From (4.15) it is obvious that knowledge of a single quantile p_{r_n} is needed for all quantile based rules. Figure 4.12 compares the performance of the Median detector with the averaging and OR rules for $M > 2$.

For a fixed, F_{HI} , as $M \rightarrow \infty$ the required quantile p_{r_n} needed for the Median rule tends to 0.5. This means that median rule is unaffected by the loss of constructive multipath that plagues the OR rule at large M .

4.4.2 Performance of cooperation under loss of independence

We have shown that safety/performance can improve significantly if radios cooperatively sense for the primary user as compared to sensing individually. This assumed that the channels from the primary transmitter to the individual secondary radios are independent. However, the primary user might not trust this assumption since all the cognitive radios may be behind the same obstacle and hence see correlated shadowing. This situation is seen in our Digital TV measurements where the variation among measurements made in the BWRC lab is small as compared to the measurements on Oxford street (See Figure 4.13). All the hypothetical radios in the BWRC lab are within the same walls. This can provide diversity against multipath but not shadowing.

Loss of diversity implies that the detector needs to set its thresholds conservatively, leading to a loss in WPAR performance. The issue of correlated-shadowing was also discussed in [62], where the authors examine the performance of their proposed linear-quadratic detector with correlation uncertainty. The proposed detector is shown to have better probability of detection than a simple counting rule for correlation values greater than 0.4.

As before, let $(P_1, P_2, \dots, P_M)$ be the received powers at the M secondary users. To isolate the effect of dependent shadowing, we assume that the marginal distributions for P_i are completely known, but there is some uncertainty in the correlation across users. For ease of analysis, we assume that $(P_1, P_2, \dots, P_M)$ is a jointly Gaussian random vector with marginals given by $P_i \sim \mathcal{N}(\mu(r), \sigma^2)$, where r is the common radial distance from the primary transmitter. Further, the $M \times M$ covariance matrix C has entries $C(i, j)$ given by

$$C(i, j) = \begin{cases} \rho\sigma^2 & \text{if } i \neq j \\ \sigma^2 & \text{if } i = j \end{cases}$$

where the correlation coefficient ρ is uncertain within known bounds, i.e., $\rho \in [0, \rho_{max}]$, with $0 \leq \rho_{max} \leq 1$.

Under this uncertain correlation model it is easy to show that the averaging detector in (4.4) is the ML detector no matter what the value of ρ is. Further, it is straightforward to show that to meet a low F_{HI} constraint, the averaging detector must design its λ for the




Figure 4.13: Variation among measurements of channel 19 at the BWRC lab is small (5dB) while the variation for the same channel among Oxford street measurements is large (20dB).

worst case correlation, $\rho = \rho_{max}$. For this choice of ρ we have $\frac{1}{M} \sum_{i=1}^M P_i \sim \mathcal{N}(\mu(r), \frac{1}{M}[1 + (M-1)\rho_{max}]\sigma^2)$. Therefore,

$$F_{HI} = 1 - \mathcal{Q} \left(\frac{\lambda - \mu(r_n)}{\sqrt{\frac{1}{M}[1 + (M-1)\rho_{max}]\sigma^2}} \right)$$

From the above equation we can choose a λ such that the target F_{HI} requirement is met. Given this λ we compute the *WPAR* performance assuming the nominal model, which corresponds to complete independence, $\rho = 0$, i.e., $\frac{1}{M} \sum_{i=1}^M P_i \sim \mathcal{N}(\mu(r), \frac{1}{M}\sigma^2)$.

Figure 4.14 shows the performance of the averaging detector designed for different values of ρ_{max} . It is clear that as the amount of uncertainty in the correlation increases, the performance of the averaging detector decreases. The OR and median rule see similar loss in performance. Even a small amount of correlation results in a significant drop in performance. As the number of users increases, this particular model of correlation is even more harmful. This can be seen by giving a simple interpretation to this correlation – fading for any user is the sum of a common random fading and a fade local to this user. It is clear that no amount of cooperation can overcome the non-spatially-ergodic common fade. Without a way to combat the fear of such non-spatially-ergodic shadowing uncertainty, there is no way to safely recover the full spectrum hole.

4.4.3 Performance of cooperation with unreliable radios

For cooperative sensing, trust issues arise naturally given the usage model:

- Sensing a frequency band, consumes energy and time which may have been better used for transmission. Hence users have an incentive to either not sense at all or to sense for a shorter duration than stipulated.
- For an individual user, there may be a valid reason to report detection results in a certain way. They can be pro-primary (i.e. they always assert the presence of the primary) and hence deny others the opportunity to take advantage of the available bandwidth or be pro-secondary (i.e. they always assert the absence of the primary) in which case they always aim to use the channel at any cost.
- Radios may fail in unpredictable ways or be simply malicious. For such users, we need to budget for worse case performance.
- Our fading models are not accurate. While quantile models provide a way to budget for this, there are still uncertainties that we do not have a model for. Uncertain user behavior provides a mechanism to budget for those issues.




Figure 4.14: Performance of the averaging detector with varying amount of correlation uncertainty, ρ_{max} . These plots correspond to the case of ten cooperating users, i.e., $M = 10$.

We assume that of the M radios cooperating, a fraction ($\gamma \in [0, 1]$) may be unreliable. The performance of any rule will depend on whether these γM radios are pro-primary radios (and hence declare the band ‘used’) or pro-secondary (and hence declare the band ‘unused’). For soft combining rules there is no mechanism to bound the performance of the system with such radios (though we may rely on outlier detection algorithms to remove results from unreliable radios, the performance of such algorithms cannot be modeled explicitly).

For hard decision rules we can explicitly model the performance of the rules under different behaviors of the unreliable radios. In this section, performance of any $k = \lceil \zeta M \rceil$ out of M rule will be calculated.

To calculate Fear of Harmful Interference we have to assume worse case uncertainty. In the case of unreliable radios, the worse case uncertainty is when the unreliable radios behave as pro-secondary radios i.e. they always declare the band free. For the system to erroneously declare the band free (a harmful interference event), fewer than k of the remaining $\lfloor M(1 - \gamma) \rfloor$ radios must declare the band used. Hence the corresponding F_{HI} is given by:

$$F_{HI} = \sum_{i=0}^{k-1} \binom{\lfloor M(1 - \gamma) \rfloor}{i} \mathcal{Q} \left(\frac{\lambda - \mu(r_n)}{\sigma} \right)^i \left(1 - \mathcal{Q} \left(\frac{\lambda - \mu(r_n)}{\sigma} \right) \right)^{\lfloor M(1 - \gamma) \rfloor - i}.$$

where λ is the detection threshold used by the individual radio.

For calculating Probability of Finding a Hole, we consider three cases depending on the behavior of the unreliable radios: Reliable behavior, pro-secondary behavior, pro-primary behavior.

If the unreliable radios behave reliably, we have M reliable radios and hence the Probability of Finding a Hole at distance r , $P_{FH}(r)$ is given by:

$$P_{FH}^r(r) = \sum_{i=0}^{k-1} \binom{M}{i} \mathcal{Q} \left(\frac{\lambda - \mu(r)}{\sigma} \right)^i \left(1 - \mathcal{Q} \left(\frac{\lambda - \mu(r)}{\sigma} \right) \right)^{M-i}$$

If the unreliable radios behave as pro-secondary radios then they declare the band as unused. For the system overall to declare the band unused, the number of reliable users declaring the band used must be less than $k - \lceil \zeta M \rceil$ i.e.

$$P_{FH}^{ps}(r) = \sum_{i=0}^{k-1} \binom{\lfloor M(1 - \gamma) \rfloor}{i} \mathcal{Q} \left(\frac{\lambda - \mu(r)}{\sigma} \right)^i \left(1 - \mathcal{Q} \left(\frac{\lambda - \mu(r)}{\sigma} \right) \right)^{\lfloor M(1 - \gamma) \rfloor - i}$$

If the unreliable radios behave as pro-primary radios then they declare the band as used. If $k \leq \lceil M\gamma \rceil$, then the system will always declare the band used and the corresponding

WPAR is 0. If $k \geq \lceil M\gamma \rceil + 1$, then the number of reliable users declaring the band used must be less than $k - \lceil M\gamma \rceil$ for the system to recover a hole. Hence the Probability of Finding a Hole for pro-primary unreliable radios is given by:

$$P_{FH}^{pp}(r) = \begin{cases} \sum_{i=0}^{k-\lceil M\gamma \rceil-1} \binom{\lfloor M(1-\gamma) \rfloor}{i} \mathcal{Q}\left(\frac{\lambda-\mu(r)}{\sigma}\right)^i \left(1 - \mathcal{Q}\left(\frac{\lambda-\mu(r)}{\sigma}\right)\right)^{\lfloor M(1-\gamma) \rfloor-i} & \text{if } k \geq \lceil M\gamma \rceil + 1 \\ 0 & \text{otherwise} \end{cases}$$

Figure 4.15 shows the performance of various rules k -out-of- M rules with 10% of radios being unreliable. The median ($k = \lfloor \frac{M}{2} \rfloor$), ($k = 2\gamma$) and the OR rule are displayed. The three curves for each rule are the pro-primary $WPAR$ ($WPAR^{pp}$), the pro-secondary $WPAR$ ($WPAR^{ps}$) and the $WPAR$ with the users behaving reliably. The median rule has the least performance gap between the different curves while the OR rule has the maximum gap. In fact the performance of the OR rule is reduced to zero when the radios behave pro-primary.

4.5 Summary

In this chapter, we presented cooperative sensing for recovering time and space domain spectrum holes. For time domain holes, the *Averaging with discounting* detector performs the best but requires knowledge of the true SNRs at all radios. All other cooperation rules suffer from the presence of radios in adverse conditions. The median rule performs well with Digital TV data captured in Berkeley, CA since it is robust to outliers. The median rule is also the rule of choice for recovering spatial holes. It works well for fading fading distributions where the convergence of the median is fastest among other percentile rules. Furthermore it is robust to uncertainties in the fading distribution and to the presence of unreliable/malicious radios.




Figure 4.15: Performance of various k -out-of- N rules with 10% of radios being unreliable. The median ($k = \lfloor \frac{M}{2} \rfloor$), ($k = 2\gamma M$) and the OR rule are displayed. The three curves for each rule are the pro-primary $WPAR$ ($WPAR^{pp}$), the pro-secondary $WPAR$ ($WPAR^{ps}$) and the $WPAR$ with the users behaving reliably. The median rule has the least performance gap between the different curves while the OR rule has the maximum.

Chapter 5

Assisted Detection

As seen in Chapter 2, the metric of detection-sensitivity, while being useful for algorithm design is not indicative of a system's ability to recover spectrum holes in space. In-fact a sensing based approach recovers nine times less people than an approach based on geo-location. The main reason for this loss is the conservative setting of the sensing threshold to an absolute value of -114dBm – which is 8dB below the thermal noise. Hence there is a need to move to sensing strategies that are aware of their actual environments and can adopt a threshold dynamically. In this chapter we propose *multiband sensing* as a sensing strategy which can calibrate to the spectral environment. This strategy relies on a small number of relatively well establish environmental factors that are well modeled and can be relied upon:

- Shadowing is highly correlated across frequencies. Shadowing results from the absorption of radio waves through material. The absorption properties of most materials does not change significantly across close-by frequencies.
- A single tower may host transmitters for multiple channels. For example, Sutro tower in San Francisco broadcasts 11 Digital TV stations [77]
- Multipath spread is correlated to the shadowing environment.

This chapter is based on the papers [23] and [24]. The main contributions that emerge in this chapter are as follows:

- We show gains from multiband sensing using a very simple model of a primary co-located with an anchor transmitter (the anchor transmitter is always 'on'). Gains can be had when the primary is 'off' since we could adjust the detection threshold for the primary based on the received power from the anchor. Next we show that the gains persist when the anchor is replaced by multiple primaries, not all of which are off. Multiband sensing provides noise calibration gains, cooperative gains and channel

calibration gains in different forms and environments. These connections are explored and expounded upon. Gains from multiband sensing were first shown theoretically in [23] for a simple model.

- We backup theoretical analysis using data from the FCC database and experimental data captured in Berkeley, CA. This data is used to provide evidence to support the environmental facts on which multiband sensing is based. The measurements were performed using a wideband receiver under various spectral environments – indoors, outdoor non-line-of-sight (NLOS) and outdoor line-of-sight (LOS).

The chapter is organized as follows: In Section 5.1 we will review related work in the areas of multiband sensing and assisted detection. Following this, in Section 5.2 we start by presenting a simple version of multiband sensing for a single tower with two transmitters one of which is always on (called the anchor node). For this simple model we show the gains realizable from utilizing shadowing correlation. These gains are recoverable with multiple transmitters on the same tower even though all of them may not be on. We show how multiband sensing provides noise calibration and cooperative gains for recovering temporal holes. Finally in Section 5.3 we analyze the FCC database of transmitters and actual measurement data captured in Berkeley, CA to expose interesting facets of the DTV spectral environment. The hardware setup and procedure used to collect data is discussed in Appendix B. Real world data is used to support the conjecture that shadowing across frequencies is highly correlated and multipath spread is a function of the shadowing environment and can be used to provide further gains in detection. Furthermore, multipath spread measured in a small frequency band is a very good predictor of the multipath spread in the overall band.

5.1 Related work

Simultaneously detecting primary users over multiple bands has been explored by many researchers for a variety of reasons. A wideband detector is capable of capturing multiple bands simultaneously and hence reduces sensing time dramatically. However, such a wideband detector also needs a very large dynamic range to deal with a range of primary powers. In [78] the authors discuss techniques to actively cancel high power primary/secondary transmissions in order to decode weak primary signals. In [79], the authors propose to use multiple bands to perform joint detection of the noise and the signal.

Wideband operation also enables the secondary to optimize throughput through power allocation [80]. Since wideband operation by a secondary requires dealing with multiple primaries, joint detection of multiple primaries is performed. Joint detection of multiple primaries is also used in [81] to perform fast sensing. The assumption is made that most bands are unused so a False Discovery Rate based technique can be used to separate the ‘on’ and ‘off’ primaries using few samples.

Later in this chapter we propose a detector that uses the fact that signals undergo multipath distortion while noise does not. [82] utilize a similar insight and propose the Covariance based detector.

The general idea of calibrated/assisted detection is used in Assisted GPS operation where a assistance server helps a GPS lock on to satellites.

5.2 Single Tower Multiband Sensing

In this section we set up the multiband sensing problem for a single tower with multiple identical transmitters.

5.2.1 Problem Setup

The problem setup is similar to the setup described in Section 4.2.1. A cognitive radio is at a distance r from a primary transmitter. The cognitive radio is unaware of its location and the tower's location. The cognitive radio wishes to reuse a frequency band whose transmitter is hosted on the tower. The secondary user runs a detection algorithm that outputs a binary decision D about the state of the primary band: 1-used/0-unused. Similarly let the hypothesis $\mathcal{H}_0$ denote the situation that the primary is absent and hypothesis $\mathcal{H}_1$ denote the situation that the primary is present. The *probability of harmful interference* $\mathcal{P}_{F_r}(D = 0|\mathcal{H}_1)$ at a distance r is the probability that a secondary user is unable to detect the primary at a distance r ¹. Similarly, the probability of recovering a potential white space location (*probability of finding a hole*) is $\mathcal{P}_{F_r}(D = 0|\mathcal{H}_0)$. Here F_r is the probability distribution of the combined multipath and shadowing-induced fading at a distance r from the primary transmitter.

Formally, we define the *Probability of Harmful Interference* ($P_{HI}(r)$) at a distance r as:

Definition 4

$$P_{HI}(r) = \mathcal{P}_{F_r}(D = 0|\mathcal{H}_1). \quad (5.1)$$

Similarly, we define the *Probability of Finding A Hole* ($P_{FH}(r)$) as:

Definition 5

$$P_{FH}(r) = \mathcal{P}_{F_r}(D = 0|\mathcal{H}_0). \quad (5.2)$$

¹Here we are ignoring the case that r is so large that we could be outside the no-talk radius of the transmitter. This aspect is discussed in Chapter 3

Since P_{FH} is used to measure average utility of the secondary system we may have faith in a single model for multipath and shadowing. Hence the absence of any worst case evaluations.

In the remainder of this paper we will simplify the notation to $P_{HI} = \mathcal{P}(D = 0|\mathcal{H}_1)$ and $P_{FH} = \mathcal{P}(D = 1|\mathcal{H}_0)$ by dropping all references to r unless required.

5.2.2 Propagation Model

The received primary signal strength P (in dBm) at a distance r can be modeled as

$$P = p_t - (l(r) + S + M), \quad (5.3)$$

where p_t is the power of the transmitted signal (in dBm), $l(r)$ is the loss in power due to attenuation at a distance r from the primary transmitter, S is the loss due to shadowing and M is the loss due to multipath fading. Unless specifically mentioned, we assume that all powers are measured in dB scale. We assume that $l(r) = 10 \log_{10}(r^\alpha)$, and α is the true attenuation exponent.

For convenience, S and M are assumed to be independent of r and to follow a nominal model: that is Gaussian ($S \sim \mathcal{N}(\mu_S, \sigma_S^2)$) on a dB scale and ($M \sim \mathcal{N}(\mu_M, \sigma_M^2)$).

This implies that $P \sim \mathcal{N}(\mu(r), \sigma_S^2 + \sigma_M^2)$, where $\mu(r) = p_t - (l(r) + \mu_S + \mu_M)$. For the plots in this paper, $\mu_S = \mu_M = 0$ dB and the standard deviation $\sigma_S^2 + \sigma_M^2 = 5.5$ dB were chosen to match standard assumptions in the IEEE 802.22 literature [83].

5.2.3 Single Band Detection

For single band detection, the detection problem was described in 3.1. We revisit the problem here to illustrate the impact of noise uncertainty:

$$\begin{aligned} \mathcal{H}_0 : y[n] &= w[n] & n = 1, \dots, N \\ \mathcal{H}_1 : y[n] &= x[n] + w[n] & n = 1, \dots, N \end{aligned}$$

where $w[n] \sim \mathcal{N}(0, 10^{\frac{N_w}{10}})$ and N_w is the noise power in dBm. Similarly we assume that the received signal $x[n] \sim \mathcal{N}(0, 10^{\frac{P}{10}})$

To incorporate uncertainties in noise and interference we consider N_w to be in a set $[N_0 - \Delta, N_0 + \Delta]$. This is a simplified version of the model proposed in [48].

The optimal detector for the problem specified in ?? is an energy detector i.e.

$$T(\mathbf{y}) = \frac{1}{N} \sum_{n=1}^N |y[n]|^2 \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \lambda, \quad (5.4)$$

For this detector, we can calculate P_{HI} and P_{FH} as follows:

$$P_{HI} = \max_{N_w \in [N_0 - \Delta, N_0 + \Delta]} \mathcal{P}[T(\mathbf{y}) < \lambda | \mathcal{H}_1] \quad (5.5)$$

Similarly, we can calculate P_{FH} as:

$$P_{FH} = \min_{N_w \in [N_0 - \Delta, N_0 + \Delta]} \mathcal{P}[T(\mathbf{y}) < \lambda | \mathcal{H}_0] \quad (5.6)$$

To see the impact of noise uncertainty, let us first consider the case of infinite samples. In this situation the test statistic converges to the mean:

$$\begin{aligned} \text{Under } \mathcal{H}_1 & : T(\mathbf{y}) = 10^{\frac{P}{10}} + 10^{\frac{N_w}{10}} \\ \text{Under } \mathcal{H}_0 & : T(\mathbf{y}) = 10^{\frac{N_w}{10}} \end{aligned}$$

The corresponding P_{HI} and P_{FH} are given by:

$$\begin{aligned} P_{HI} &= \max_{N_w \in [N_0 - \Delta, N_0 + \Delta]} \mathcal{P}[10^{\frac{P}{10}} + 10^{\frac{N_w}{10}} < \lambda] \\ &\stackrel{(a)}{=} \mathcal{P}[10^{\frac{P}{10}} + 10^{\frac{N_0 - \Delta}{10}} < \lambda] \\ &= \mathcal{P}[P < 10 \log_{10}(\lambda - 10^{\frac{N_0 - \Delta}{10}})] \\ &= \mathcal{P}[P < \lambda_a] \end{aligned}$$

where (a) is because $N_0 - \Delta$ maximizes P_{HI} and $\lambda_a = 10 \log_{10}(\lambda - 10^{\frac{N_0 - \Delta}{10}})$. Since $P \sim \mathcal{N}(\mu(r), \sigma_S^2 + \sigma_M^2)$ we can calculate $\lambda_a = \mu(r) + \mathcal{Q}^{-1}(1 - P_{HI}) \sqrt{\sigma_S^2 + \sigma_M^2}$.

and,

$$\begin{aligned} P_{FH} &= \max_{N_w \in [N_0 - \Delta, N_0 + \Delta]} \mathcal{P}[10^{\frac{N_w}{10}} < \lambda] \\ &= \mathcal{P}[10^{\frac{N_0 + \Delta}{10}} < \lambda] \\ &= \mathcal{P}[10^{\frac{N_0 + \Delta}{10}} < 10^{\frac{\mu(r) + \mathcal{Q}^{-1}(1 - P_{HI}) \sqrt{\sigma_S^2 + \sigma_M^2}}{10}} + 10^{\frac{N_0 - \Delta}{10}}] \end{aligned}$$

For $P_{FH} > 0$ we need

$$\begin{aligned}
 10^{\frac{N_0+\Delta}{10}} - 10^{\frac{N_0-\Delta}{10}} &< 10^{\frac{\mathcal{Q}^{-1}(1-P_{HI})\sqrt{\sigma_S^2+\sigma_M^2}+\mu(r)}{10}} \\
 \frac{10 \log_{10}(10^{\frac{N_0+\Delta}{10}} - 10^{\frac{N_0-\Delta}{10}}) - \mu(r)}{\sqrt{\sigma_S^2 + \sigma_M^2}} &< \mathcal{Q}^{-1}(1 - P_{HI}) \\
 1 - \mathcal{Q}\left(\frac{10 \log_{10}(10^{\frac{N_0+\Delta}{10}} - 10^{\frac{N_0-\Delta}{10}}) - \mu(r)}{\sqrt{\sigma_S^2 + \sigma_M^2}}\right) &< P_{HI}
 \end{aligned}$$

i.e. for $P_{FH} > 0$, P_{HI} must be greater than the probability that the signal falls below the detection sensitivity defined as: $10 \log_{10} \left(10^{\frac{N_0+\Delta}{10}} - 10^{\frac{N_0-\Delta}{10}} \right)$. This P_{HI} versus P_{FH} curve is plotted for $\Delta = 1\text{dB}$ in Figure 5.1 as the blue line.

Similarly we can calculate, P_{HI} and P_{FH} for the case of finite samples as by computing the probabilities of missed detection and false alarms at a given received power (See [?] for details) and then taking the expectation over the distribution of shadowing+fading at a distance r :

$$\begin{aligned}
 P_{HI} &= \mathbb{E}_{F_r} \left[1 - \mathcal{Q} \left(\frac{\lambda - (10^{\frac{P}{10}} + 10^{\frac{N_0-\Delta}{10}})}{\sqrt{\frac{2}{N}(10^{\frac{P}{10}} + 10^{\frac{N_0-\Delta}{10}})}} \right) \right] \\
 P_{FH} &= \mathbb{E}_{F_r} \left[1 - \mathcal{Q} \left(\frac{\lambda - (10^{\frac{P}{10}} + 10^{\frac{N_0+\Delta}{10}})}{\sqrt{\frac{2}{N}(10^{\frac{P}{10}} + 10^{\frac{N_0+\Delta}{10}})}} \right) \right]
 \end{aligned}$$

The tradeoff between P_{HI} and P_{FH} for 100 samples is plotted in Figure 5.1 as the green curve. We can get a better tradeoff by time sharing between the points $(P_{HI}, P_{FH}) = (0, 0)$ and $(P_{HI}, P_{FH}) = (0.1736, 0.8)$ as shown by the purple curve.

5.2.4 Multiband detection

In this section we will define the working of the Multiband detector and evaluate its performance under noise uncertainty.

Defining the multiband detector

We assume that we are trying to detect a primary that is either on or off. Co-located with the primary are other primary transmitters that are always on. These primary users are called ‘anchors’ – they are always on and their positions with respect to the primary user is also fixed.




Figure 5.1: Performance of singleband sensing for infinite and 200 samples. For infinite samples, singleband sensing is unable to recover any spectrum hole for $P_{HI} < 0.118$. This figure also shows the performance of multiband sensing for finite samples.

Furthermore, assume that the received signal from the primary and the anchor nodes can be expressed in the form equivalent to the Equation (5.3):

$$\begin{aligned} P_P &= p_t - (L + S_P + M_P) \\ P_A &= p_t - (L + S_A + M_A) \end{aligned}$$

where p_t (the transmit power) is the same for the anchor and the primary, L is the path loss which is also the same for both (assuming they are co-located). S_P and S_A is the shadowing seen by the primary and the anchor signals respectively. We assume that the shadowing is completely correlated across the bands (i.e. $S_A = S_P$) and the multipath is completely independent. As before S_P, S_A, M_P, M_A are assumed to be independent of r and to follow a nominal model: that is Gaussian ($S_x \sim \mathcal{N}(\mu_S, \sigma_S^2)$) on a dB scale and ($M_x \sim \mathcal{N}(\mu_M, \sigma_M^2)$). ρ is the resulting correlation between the random variables ($S_P + M_P$) and ($S_A + M_A$).

To keep things concrete we work with energy detection, i.e., the cognitive radio can measure the energy in both the primary and the anchor band. The question is whether having access to an anchor improves the performance of a single user energy detector. The Multiband detector works as follows:

- Run an energy detector for both the primary and anchor bands.
- Let the energy detection thresholds for the primary bands be λ_P and let the corresponding threshold for the anchor bands be λ_A .
- Compute the empirical estimate of the energy in each band. Let $T(\mathbf{y}_A)$ denote this test-statistic for the anchor, and $T(\mathbf{y}_P)$ be the corresponding test-statistic for the primary band. Now compare these test-statistics to the corresponding thresholds.
- Let $D_P = 1$ denote the decision when the energy estimate in the primary band exceeds λ_P , and $D_P = 0$ otherwise. Similarly, declare D_A for the anchor band.
- Given the individual decisions in each band, make a global decision of whether the primary is present or absent. This decision is made as shown in Table 5.1. Here $D = 1$ denotes the global decision that the primary band is used and $D = 0$ denotes the global decision that the primary band is empty.

The case when neither the primary nor the anchor is detected $D_A = D_P = 0$ is particularly interesting. In this case, we are unsure of our shadowing environment and so declare the primary present just to be safe.

Anchor band decision	$D_A = 1$	$D_A = 1$	$D_A = 0$	$D_A = 0$
Primary band decision	$D_P = 1$	$D_P = 0$	$D_P = 1$	$D_P = 0$
Global Decision	$D = 1$	$D = 0$	$D = 1$	$D = 1$

Table 5.1: Multiband energy detection algorithm

For the multiband detector the performance are given by

$$\begin{aligned} P_{HI} &= \mathcal{P}(T(\mathbf{y}_A) > \lambda_A, T(\mathbf{y}_P) \leq \lambda_P | \mathcal{H}_1) \\ P_{FH} &= \mathcal{P}(T(\mathbf{y}_A) > \lambda_A, T(\mathbf{y}_P) \leq \lambda_P | \mathcal{H}_0) \end{aligned}$$

For any two achievable points (P_{HI}^1, P_{FH}^1) and (P_{HI}^2, P_{FH}^2) we can achieve all points on the line $(\theta P_{HI}^1 + (1 - \theta)P_{HI}^2, \theta P_{FH}^1 + (1 - \theta)P_{FH}^2)$ joining these two points by randomization according to $0 \leq \theta \leq 1$. The performance of multiband energy detection for the single anchor case can be characterized by the set of all achievable (P_{HI}, P_{FH}) pairs. Let $\mathbb{R}_{MB}(\rho)$ denote this region for a given frequency correlation coefficient ρ . Formally, we can write:

$$\begin{aligned} \mathbb{R}_{MB}(\rho) = \text{Convexhull} \quad \{ & (P_{HI}(\lambda_A, \lambda_P, \rho), P_{FH}(\lambda_A, \lambda_P, \rho)) \\ & : 0 \leq \lambda_A, \lambda_P \leq \infty \} \end{aligned}$$

Although every point in the region $\mathbb{R}_{MB}(\rho)$ is achievable, the interesting performance points are the Pareto optimal points. A point is defined as being Pareto optimal if there is no other point that is ‘better’ than it in both directions. In multiband sensing the point $(0, 1)$ is the ‘best point’. We call this set of Pareto optimal points the multiband performance frontier and denote it by $\mathbb{F}_{MB}(\rho)$. Note that by definition $\mathbb{F}_{MB}(\rho)$ is concave and $(0, 0)$ and $(1, 1)$ are the end points of this curve.

Performance under independent noise uncertainty

We will first evaluate the role of Multiband sensing in providing robustness. In Section 5.2.3 we saw that the Single Band detector was unable to recover any spectrum hole for small values of P_{HI} . Can the Multiband detector provide any gains against noise uncertainty? If so, under what circumstances are these gains realizable?

First we first evaluate the performance of multiband sensing under noise uncertainty with infinite samples, where both bands see independent noise uncertainty. For this case, we can evaluate P_{HI} as:

$$\begin{aligned}
 P_{HI} &= \max_{\substack{N_{wP} \in [N_0 - \Delta, N_0 + \Delta] \\ N_{wA} \in [N_0 - \Delta, N_0 + \Delta]}} \mathcal{P}[10^{\frac{P_A}{10}} + 10^{\frac{N_{wA}}{10}} > \lambda_A, 10^{\frac{P_P}{10}} + 10^{\frac{N_{wP}}{10}} \leq \lambda_P] \\
 &= \mathcal{P}[10^{\frac{P_A}{10}} + 10^{\frac{N_0 + \Delta}{10}} > \lambda_A, 10^{\frac{P_P}{10}} + 10^{\frac{N_0 - \Delta}{10}} \leq \lambda_P]
 \end{aligned}$$

Similarly,

$$\begin{aligned}
 P_{FH} &= \min_{\substack{N_{wP} \in [N_0 - \Delta, N_0 + \Delta] \\ N_{wA} \in [N_0 - \Delta, N_0 + \Delta]}} \mathcal{P}[10^{\frac{P_A}{10}} + 10^{\frac{N_{wA}}{10}} > \lambda_A, 10^{\frac{N_{wP}}{10}} \leq \lambda_P] \\
 &= \mathcal{P}[10^{\frac{P_A}{10}} + 10^{\frac{N_0 - \Delta}{10}} > \lambda_A, 10^{\frac{N_0 + \Delta}{10}} \leq \lambda_P]
 \end{aligned}$$

Clearly to get $P_{FH} > 0$, we need to set $\lambda_A \leq 10^{\frac{N_0 - \Delta}{10}}$ and $\lambda_P \geq 10^{\frac{N_0 + \Delta}{10}}$. For this setting, the lowest P_{HI} obtainable is $P_{HI} = \mathcal{P}(P_P \leq 10^{\frac{N_0 + \Delta}{10}} - 10^{\frac{N_0 - \Delta}{10}})$ which is the same performance one would get from single band sensing.

Performance under the dependent noise uncertainty

Next we look at the case where the noise uncertainty is the same in the anchor and primary bands.

Clearly, P_{HI} and P_{FH} for this case are:

$$\begin{aligned}
 P_{HI} &= \max_{N_w \in [N_0 - \Delta, N_0 + \Delta]} \mathcal{P}[10^{\frac{P_A}{10}} + 10^{\frac{N_w}{10}} > \lambda_A, 10^{\frac{P_P}{10}} + 10^{\frac{N_w}{10}} \leq \lambda_P] \\
 P_{FH} &= \min_{N_w \in [N_0 - \Delta, N_0 + \Delta]} \mathcal{P}[10^{\frac{P_A}{10}} + 10^{\frac{N_w}{10}} > \lambda_A, 10^{\frac{N_w}{10}} \leq \lambda_P]
 \end{aligned}$$

We can show noise calibration gains from Multiband sensing in the following manner. Set $\lambda_A = 10^{\frac{P_A}{10}} + 10^{\frac{N_w}{10}} - \epsilon$, $\epsilon > 0$ i.e. a shade below the measured detector power in the anchor band and set $\lambda_P = \lambda_A - \Gamma$ (where Γ is a parameter to change the tradeoff between P_{HI} and P_{FH}).

For $\Gamma = 0$ we get $P_{FH} = 1$ and $P_{HI} = \mathcal{P}(10^{\frac{P_A}{10}} - 10^{\frac{P_P}{10}} > -\epsilon)$. If the anchor and the primary are complete correlated, then we can set Γ to 0 to obtain $P_{HI} = 0$.

Figure 5.2 shows the performance of the Multiband detector with infinite samples and a Noise uncertainty of 1dB. The curve is traced by varying the value of Γ . Except for the degenerate case of complete correlation between the anchor and the primary, the actual correlation (ρ) between the anchor and the primary does not impact the performance of the Multiband detector in a big way (compare the green, black and red curves in Figure 5.2). The real gains are for very low values of P_{HI} where we can calibrate away the noise and get non-zero values of P_{FH} ².

In the remainder of this chapter we will ignore the gains from noise calibration. This topic has been discussed in [19]. For the case of independent noise in both bands we can write the performance of the Multiband detector explicitly as shown in (5.7) and (5.8).

Figure 5.1 compares the performance of the Multiband detector with performance of the single Band detector for 100 samples³. For singleband detection, the only way to obtain low P_{HI} is to lower the detection threshold λ_P below $10^{\frac{N_0+\Delta}{10}}$. However, setting a threshold in the noise uncertainty region adversely affects the system performance in the form of decreased P_{FH} (see the single radio curve in Fig. 5.1). For multiband detection, low P_{HI} can be obtained by setting $\lambda_P > 10^{\frac{N_0+\Delta}{10}}$ and making λ_A sufficiently large. This avoids a catastrophic degradation in P_{FH} at low P_{HI} .

5.2.5 Cooperative gains from Multiband sensing

Cooperation has been proposed to reduce sensitivity requirements and/or provide gains in the number of samples required to achieve a target P_{HI}, P_{FH} . In this section we will examine how multiband detection can be used to provide gains while doing cooperative sensing. For purposes of illustration, we shall use the OR rule in which the system decides that the primary band is vacant if and only if each of the individual radios declare that the band is vacant. We assume that the system consists of M radios each of which runs the multiband energy detector described. The radios cooperate to decide whether the primary band is used or empty. Assume also that each of the radios in the system are homogeneous, i.e., the detection thresholds λ_A and λ_P are same for all the radios in the system.

As a baseline, consider the strict OR rule in which the system decides that the primary band is vacant if each of the individual radios declare that the band is vacant. Since each individual radio is running a multiband energy detector, it declares the primary band empty if they detect the anchor and find the primary band empty. Under this rule the probability of harmful interference of the system is given by:

²We can get a better performance by timesharing but we ignore time sharing in the curves in Figure 5.2)

³We get a hundred samples for estimating the energy of the anchor and another hundred samples for estimating the energy of the primary




Figure 5.2: Performance of multiband sensing is compared to the performance of singleband sensing for infinite samples. For infinite samples, Multiband sensing helps us get gains in the form of noise calibration which lets us recover spectrum holes for a very low values of P_{HI} .

$$P_{HI} = \mathbb{E}_{F_r(P_A, P_P)} \left[\mathcal{Q} \left(\frac{\lambda_A - (10^{\frac{P_A}{10}} + 10^{\frac{N_0 - \Delta}{10}})}{\sqrt{\frac{2}{N}(10^{\frac{P_A}{10}} + 10^{\frac{N_0 - \Delta}{10}})}} \right) \left(1 - \mathcal{Q} \left(\frac{\lambda_P - (10^{\frac{P_P}{10}} + 10^{\frac{N_0 - \Delta}{10}})}{\sqrt{\frac{2}{N}(10^{\frac{P_P}{10}} + 10^{\frac{N_0 - \Delta}{10}})}} \right) \right) \right] \quad (5.7)$$

$$P_{FH} = \left(1 - \mathcal{Q} \left(\frac{\lambda_P - 10^{\frac{N_0 + \Delta}{10}}}{\sqrt{\frac{2}{N} 10^{\frac{N_0 - \Delta}{10}}}} \right) \right) \mathbb{E}_{F_r(P_A)} \left[\mathcal{Q} \left(\frac{\lambda_A - (10^{\frac{P}{10}} + 10^{\frac{N_0 + \Delta}{10}})}{\sqrt{\frac{2}{N}(10^{\frac{P}{10}} + 10^{\frac{N_0 + \Delta}{10}})}} \right) \right] \quad (5.8)$$

$$\begin{aligned} P_{HI}^{OR} &= \mathcal{P}(T(\mathbf{y}_{\mathbf{A}_1}) > \lambda_A, \dots, T(\mathbf{y}_{\mathbf{A}_M}) > \lambda_A, T(\mathbf{y}_{\mathbf{P}_1}) \leq \lambda_P, \dots, T(\mathbf{y}_{\mathbf{P}_M}) \leq \lambda_P) \\ &\stackrel{(a)}{=} \left[\prod_{i=1}^M \mathcal{P}(T(\mathbf{y}_{\mathbf{A}_i}) > \lambda_A, T(\mathbf{y}_{\mathbf{P}_i}) < \lambda_P | \mathcal{H}_1) \right] \end{aligned}$$

where (a) occurs when the channel to each radio is independent of the other radios. Here the probability is taken over the joint distribution of the random variables:

$(P_{A_1}, P_{P_1}, \dots, P_{A_N}, P_{P_N})$, where P_{A_i} is the received signal strength of the anchor signal at the i^{th} radio, and P_{P_i} is the received signal strength of the primary signal at the i^{th} radio.

Similarly, the probability of finding a hole of the system is given by:

$$P_{FH}^{OR} = \left[\prod_{i=1}^M \mathcal{P}(T(\mathbf{y}_{\mathbf{A}_i}) > \lambda_A, T(\mathbf{y}_{\mathbf{P}_i}) < \lambda_P | \mathcal{H}_0) \right]$$

An analysis of the above equations reveal that the gains from cooperation for this strict OR rule display the following trends:

- P_{HI}^{OR} decreases as M increases.
- P_{FH}^{OR} decreases as M increases.

There is no difference between the behavior of singleband and sultiband sensing radios using the strict OR rule as the number of radios in scaled.

Multiband cooperation with abstention

Recall that the multiband energy detection algorithm in Section 5.2.4 declared that the primary is present when the radio fails to detect the anchor. The reasoning behind this was the following: if the radio does not see the anchor which is always “ON” then it is most likely deeply shadowed to the primary too. Since a deeply shadowed radio’s detection results are unreliable, being conservative requires declaring that the primary might be “ON”.

The above discussion suggests a way to weed out deeply shadowed radios. We introduce a ternary decision scheme, i.e., the radio declares one of the following three decisions: $D = 1$ – the primary is present, $D = 0$ – the primary is absent and $D = \text{Abstain}$ – the primary abstains from making any decision. The modified detection algorithm is summarized in Table 5.2.

Anchor band decision	$D_A = 1$	$D_A = 1$	$D_A = 0$	$D_A = 0$
Primary band decision	$D_P = 1$	$D_P = 0$	$D_P = 1$	$D_P = 0$
Global Decision	$D = 1$	$D = 0$	$D = \text{Abstain}$	$D = \text{Abstain}$

Table 5.2: Modified multiband energy detection algorithm

Given that each radio makes a ternary decision, the OR rule for cooperation is modified as follows: the system declares that the band is safe to use if each of the qualified radio (radios that do not abstain) declares that the primary is absent. If all the radios abstain or any radio votes that the primary is present, then the system declares that the primary band is unsafe.

We now derive the performance of this new OR rule with abstentions for cooperation. For a given i , define S_i to be the set of all subsets of $\{1, 2, \dots, N\}$ of cardinality i . For each $\mathbf{u} \in S_i$, define $\hat{\mathbf{u}} = \{1, 2, \dots, N\} \setminus \mathbf{u}$.

The harmful interference event can be written as a disjoint sum of events parametrized by the number of abstaining users, $i = 0, 1, \dots, N - 1$. By the above definitions S_i denotes the set of possible combinations of i abstaining radios. Also, if $\mathbf{u}$ denotes the set of radios that abstain, then $\hat{\mathbf{u}}$ denotes the radios that don’t abstain.

For a given $\mathbf{u} \in S_i$, let $\mathcal{P}_{ABS}(\mathbf{u})$ denote the probability that the radios in the set $\mathbf{u}$ abstain and let $\mathcal{P}_{MD}(\hat{\mathbf{u}})$ denote the probability that the radios in the set $\hat{\mathbf{u}}$ mis-detect the primary. By definition we have:

$$\mathcal{P}_{HI}^{ORa} = \sum_{i=0}^{M-1} \binom{M}{i} \mathcal{P} \left(T(\mathbf{y}_{\mathbf{A}_{\mathbf{u}(\mathbf{k})}}) \leq \lambda_A \right)^i \mathcal{P} \left(T(\mathbf{y}_{\mathbf{A}_{\mathbf{u}(\mathbf{k})}}) > \lambda_A, T(\mathbf{y}_{\mathbf{P}_{\mathbf{u}(\mathbf{k})}}) \leq \lambda_P | \mathcal{H}_1 \right)^{M-i} \quad (5.9)$$

$$\mathcal{P}_{FH}^{ORa} = \sum_{i=0}^{M-1} \binom{M}{i} \mathcal{P} \left(T(\mathbf{y}_{\mathbf{A}_{\mathbf{u}(\mathbf{k})}}) \leq \lambda_A \right)^i \mathcal{P} \left(T(\mathbf{y}_{\mathbf{A}_{\mathbf{u}(\mathbf{k})}}) > \lambda_A, T(\mathbf{y}_{\mathbf{P}_{\mathbf{u}(\mathbf{k})}}) \leq \lambda_P | \mathcal{H}_0 \right)^{M-i} \quad (5.10)$$

$$\mathcal{P}_{ABS}(\mathbf{u}) = \prod_{k=1}^i \mathcal{P}(T(\mathbf{y}_{\mathbf{A}_{\mathbf{u}(\mathbf{k})}}) < \lambda_A)$$

$$\mathcal{P}_{MD}(\mathbf{u}) = \prod_{k=1}^i \mathcal{P}(T(\mathbf{y}_{\mathbf{A}_{\mathbf{u}(\mathbf{k})}}) > \lambda_A, T(\mathbf{y}_{\mathbf{P}_{\mathbf{u}(\mathbf{k})}}) < \lambda_P | \mathcal{H}_1)$$

Using the above notation, the probability of harmful interference can be written as:

$$\mathcal{P}_{HI}^{ORa} = \left[\sum_{i=0}^{M-1} \sum_{\mathbf{u} \in S_i} \mathcal{P}_{ABS}(\mathbf{u}) \cdot \mathbb{P}_{MD}(\hat{\mathbf{u}}) \right]$$

For simplicity, consider the situation where the joint distribution is symmetric. In this case $\mathcal{P}_{ABS}(\mathbf{u}) \cdot \mathbb{P}_{MD}(\hat{\mathbf{u}})$ is identical for all $\mathbf{u} \in S_i$. Hence, in the symmetrical case P_{HI} and P_{FA} can be written as shown in Equations (5.9) and (5.10).

Figure 5.3 compares the performance of the OR rule with the performance of the OR rule with abstentions for the case of spatially uncorrelated radios. We start with a single radio performance point P_{HI}, P_{FH} and plot the performance of the system as the number of cooperating radio increases. With the strict OR rule, both P_{HI} and P_{FH} reduces with M . With the *OR* rule with abstentions on the other hand, P_{HI} reduces with M while P_{FH} increases with M up to a critical point i.e. we are getting gains in *both* P_{HI} and P_{FH} . After the critical point, P_{HI} starts increasing and P_{FH} starts decreasing. Intuitively, the early gains from cooperation are due to reducing the probability of having everyone abstain. With many users, this is no longer the dominant source of missing a spectrum hole and the false alarms become significant.




Figure 5.3: Performance of the strict OR and OR rule with abstentions with increasing number of independent users. We start with a single radio performance point P_{HI}, P_{FH} and plot the performance of the system as the number of cooperating radio increases. With the strict OR rule, both P_{HI} and P_{FH} reduces with M . With the OR with abstentions rule on the other hand, P_{HI} reduces with M while P_{FH} increases with M up to a critical point. For the strict OR rule, loss of independence results in P_{FH} that is much worse than a single radio and a P_{HI} that is marginally better than a single user. On the other hand, for the OR rule with abstentions, the resulting (P_{HI}, P_{FH}) operational point is very similar to the single user operational point.

We now characterize the critical threshold for the number of cooperating users for which we get the best performance, i.e., maximum P_{FH} . We know that:

$$\begin{aligned} P_{FH}^{ORa} &= \sum_{i=0}^{M-1} \binom{M}{i} (1-\beta)^i \cdot [\beta(1-\eta)]^{M-i} \\ &= [(1-\beta) + \beta(1-\eta)]^M - (1-\beta)^M \\ &= [(1-\beta\eta)]^M - (1-\beta)^M \end{aligned}$$

where $\beta = \mathcal{P}(T(\mathbf{y}_A) > \lambda_A)$ is the probability of detecting the anchor and $\eta = \mathcal{P}(T(\mathbf{y}_P) > \lambda_P | \mathcal{H}_0)$ is the probability of false alarm. For notational simplicity define $P_{FH}^{ORa} =: f(M)$. Minimizing $f(M)$ with M gives us the critical number of cooperating users.

Let

$$\bar{M} = \left\lceil \frac{\log \left[\frac{\log(1-\beta)}{\log[(1-\beta\eta)]} \right]}{\log \left[\frac{(1-\beta\eta)}{(1-\beta)} \right]} \right\rceil$$

Then, the optimal number is given by

$$M^* = \arg \min \{f(\bar{M}), f(\bar{M} + 1)\}$$

Furthermore, it is clear that the optimal number of cooperating users M^* increases as η decreases – we can cooperate with more radios with the individual false alarms of the radios in small. In practice, the number of cooperating users can be kept at this optimal number by having the most qualified users (those who see the anchors the strongest) vote.

In plotting the curves in Figure 5.3 we assumed that shadowing is spatially uncorrelated, i.e., the received signal strength is independent across space. The figure also shows the loss in performance if the spatial independence assumption is false, i.e., shadowing is completely correlated. In such cases, increasing the number of users is only like increasing the number of samples taken.

Figure 5.3 shows the impact of losing independence across radios. For the strict OR rule, loss of independence results in P_{FH} that is much worse than a single radio and a P_{HI} that is marginally better than a single user. If secondaries relied on cooperation to help them achieve the target P_{HI} they could not convince the broadcasts of meeting this P_{HI} value under all spectral environments. On the other hand, for the OR with abstentions rule, the resulting (P_{HI}, P_{FH}) operational point is very similar to the single user operational point. This implies that secondaries could be engineered of achieve the target P_{HI} using a single user and then use cooperation in the field to get P_{FH} gains. If the secondaries ended up in an adverse environment the only loss would be in P_{FH} of the system, P_{HI} would not be

affected.

5.2.6 Moving from anchor to multiple primaries

So far we have investigated the impact of a single anchor on the detection of a given primary. There are two issues that remain within the example of co-located primary users: the advantage of using multiple anchors and what to do if there are no *a priori* known anchors. In the case of a single anchor, we had to move the anchor detection threshold to a very large value in order to obtain very small values for probability of harmful interference (P_{HI}). Indirectly, this meant that we had to rely on a model for the joint distribution of the anchor and primary SNRs for those extreme values. While this is possible, we do not trust the model for such extreme values. Multiple anchors allow the use of frequency diversity to avoid having to look at such extreme cases. If we assume that shadowing is constant across frequency and multipath varies, then having a large number of anchors locks down the shadowing value. This leaves only the multipath variation that needs to be accounted for in detecting the primary. This is a part of the model that can be trusted much more.

If we relax our model to include ordinary primaries rather than anchor nodes, then we have to deal with the issue that these transmitters may or may not be on. In this case we would need some knowledge of the rough number of primaries that are transmitting in a given area (for example, we could know that at least 5 of the transmitters on the Sutro tower in San Francisco are ‘ON’ at any given time). With this knowledge, the ‘qualified’ voters would be ones that see a certain number (say α) of primaries (5 in the case of Sutro tower). Since there is chance that less than α primaries may be active we can either absorb that event into our probability of missing a spectrum opportunity ($1 - P_{FH}$) or set the threshold for a qualified voter to be less than α . In the scenario of interest, we are interested in making a decision for a single frequency band while measuring the energy in all neighboring bands.

First let us define the following quantities with the implicit assumption of symmetry of distributions of various primaries/anchors:

- N_A - number of anchors.
- t - number of anchors that need to trigger for the Multiband detector to be qualified.
- P_{on} - probability that the anchor is on.
- $P_{off} = 1 - P_{on}$ - probability that the anchor is off.
- $P_{on}(i) = \binom{N_A}{i} P_{on}^i P_{off}^{N_A-i}$ - probability that i anchors are on and the remaining $N_A - i$ are off.
- P_{FA}^A - Probability that the radio false alarms for a given anchor ($\mathcal{P}(T(\mathbf{y}_A) > \lambda_A | \mathcal{H}_0)$)

- P_{FA}^P - Probability that the radio false alarms for the target primary ($\mathcal{P}(T(\mathbf{y}_P) > \lambda_P | \mathcal{H}_0)$)
- $P_A(m) = \binom{N_A-i}{m} (P_{FA}^A)^m (1 - P_{FA}^A)^{N_A-i-m}$ - probability that of the $N_A - i$ anchors that are not on, m cause a false alarm and hence trigger the detector.

The probability of finding a hole is given by:

$$P_{FH} = \sum_{k=t}^{N_A} \sum_{i=0}^{N_A} \sum_{j=\max(0, k+i-N_A)}^{\min(i, k)} \binom{i}{j} P_{on}(i) P_A(k-j) \mathcal{P} \left(T(\mathbf{y}_{\mathbf{A}_i^j}) > \lambda_A, T(\mathbf{y}_{\mathbf{A}_{j+1}^i}) \leq \lambda_A \right) (1 - P_{FA}^P). \quad (5.11)$$

Note that the notation $\mathcal{P} \left(T(\mathbf{y}_{\mathbf{A}_i^j}) > \lambda_A, T(\mathbf{y}_{\mathbf{A}_{j+1}^i}) \leq \lambda_A \right)$ is a short form for writing:

$$\mathcal{P} \left(T(\mathbf{y}_{\mathbf{A}_1}) > \lambda_A, \dots, T(\mathbf{y}_{\mathbf{A}_j}) > \lambda_A, T(\mathbf{y}_{\mathbf{A}_{j+1}}) \leq \lambda_A, \dots, T(\mathbf{y}_{\mathbf{A}_i}) \leq \lambda_A | \mathcal{H}_1 \right)$$

In Equation (5.11) a radio is considered ‘qualified’ when it detects $k, k \geq t$ anchors. i anchors are on and of these j exceed the detector threshold (λ_A). Of the remaining $N_A - j$ anchors that are off, $k - j$ cause a false alarm. If the ‘qualified’ radio does not false alarm for the off primary, a spectrum hole has been found.

Similarly, the probability of harmful interference (P_{HI}) is given by

$$P_{HI} = \sum_{k=t}^{N_A} \sum_{i=0}^{N_A} \sum_{j=\max(0, k+i-N_A)}^{\min(i, k)} \binom{i}{j} P_{on}(i) P_A(k-j) \mathcal{P} \left(T(\mathbf{y}_{\mathbf{A}_i^j}) > \lambda_A, T(\mathbf{y}_{\mathbf{A}_{j+1}^i}) \leq \lambda_A, T(Y_P) \leq \lambda_P \right). \quad (5.12)$$

Figure 5.4 shows the impact of moving from a single anchor that is always on to multiple anchors that may or may not be on. The low P_{HI} , low P_{FH} corner is achieved by placing the anchor threshold high. When the anchor is ‘off’ the threshold is never triggered which leads to a lower P_{HI} , P_{FH} point – the anchor being ‘off’ reduces the P_{HI} and P_{FH} simultaneously. Similarly, the high P_{HI} , high P_{FH} point is achieved by setting the anchor threshold low. When the anchor is ‘off’, noise alone can trigger a low threshold allowing us to achieve the

same high P_{HI} , high P_{FH} point. The main difference is in the intermediate points. However since the extremal points can be achieved many internal points can be achieved via time sharing.

5.3 Experimental results

In this section we analyze experimental data described in Section 4.2.3 and in Appendix B to provide evidence for the environment facts listed in the introduction and examine the performance of multiband Sensing in the real world (See Tables 4.2, 4.3 and Figure 4.2).

5.3.1 Shadowing Correlation

First we determined the shadowing correlation across channels at different locations (where a location could be considered as the BWRC Lab, the whole of BWRC etc). For each location, the measured signal strength for a given channel was treated as a vector of realizations of a single random variable. Given two vectors corresponding to two different channels the shadowing correlation was calculated. By doing the calculation in this manner, we are making two simplifications: firstly, we are assuming that the average represents the nominal received signal strength at any location. Secondly, we are clubbing shadowing and multipath together into a single random variable (multipath is independent across frequencies and hence lowers the calculated correlation). Both these simplifications will cause us to underestimate the shadowing correlation across frequencies. Another important aspect is ensuring that we capture the shadowing component. If we only take the measurements in the BWRC lab and subtract the mean, we will have removed the shadowing component introduced by the signal penetrating through the walls/roof. This component is very similar across frequencies and hence will be removed in the mean.

The second column in Table 5.3 shows the average measured correlation over all channel pairs at different locations. Shadowing correlation is high for outdoor Non Line-of-Sight (NLOS) locations (Milvia and California streets) as well as for indoor scenarios (BWRC). Unfortunately the correlation coefficient is unable to distinguish between Line-Of-Sight scenarios (the antenna pointing west i.e line-of-sight to Sutro at Prof. Sahai's deck where we see no shadowing and Non Line-Of-Sight scenarios where the correlation calculation process has removed the shadowing component (BWRC lab for example). One of the ways to get around this problem is to designate the average power at a given non-shadowed location as the nominal power and to calculate the correlation with respect to that nominal value. We assumed that the nominal received power in Berkeley was the average signal strength at Prof. Sahai's deck. Using this as the 'universal' mean we can calculate shadowing at various locations. These values are shown in the third column in Table 5.3 which illustrates the 'correct' shadowing correlation across frequencies. The antenna pointing west on Prof.




Figure 5.4: Performance characteristics of a single wideband radio when an anchor node may or may not be on ($P_{on} = .8$). The 20 percent chance that the anchor is not on does not lead to many missed opportunities.

Sahai's deck sees little correlation while the BWRC laboratory measurements see a lot of correlation.

Location	Correlation Coefficient	Correlation Coefficient using Prof Sahai's deck as the nominal value
BWRC Lab Night	0.11	0.99
BWRC Lab Afternoon	0.09	0.99
All BWRC	0.74	0.99
California Street	0.43	0.91
Milvia Street	0.41	0.93
All Outdoors (NLOS)	0.89	0.97
Dr. Sahai's Deck (West only)	0.07	0.14
Dr. Sahai's Deck (All directions)	0.7	0.55
All Berkeley	0.89	0.97

Table 5.3: Shadowing correlation at various locations across Berkeley

5.3.2 Spread of received signal strengths

We measured the spread of the received signal strength across the 12 channels at each measurement point and compared that to the predicted spread of 11dB [77] (For channels only on Sutro this number was 8dB). To accomplish this we determined the maximum received power on any DTV channel at each measurement point. Next we set a detection threshold, a certain number of dB's below the maximum received power and determined the number of channels that we missed by using this new detection threshold (i.e. the Probability of causing Harmful Interference - P_{HI}). Figure 5.5 shows the fraction of locations times channels that were missed as the detection threshold relative to the maximum received power was scaled. The maximum gap between the maximum and the minimum received strength was around 25dB. As per the FCC's database the predicted variation in received signal strength based on the ITU propagation curves [77] was 11dB which means that an additional 14 dB of variation is introduced by shadowing and multipath.

As stated earlier, the predicted variation for channels only on Sutro was 8dB. To determine if 8dB was a typical variation to expect for channels on the same tower, we determined the spread in the received signal strength over all 47 channels at a random sampling of locations in the continental United States for transmitters on the same tower. The FCC database of all high and low power towers was used [36, 37] to determine the locations, HAAT and transmit powers of all transmitters. Figure 5.6 shows the cumulative distributed function of

the predicted received signal strength variation over all locations in the United States. The median variation in signal strength is 5dB.

This analysis confirms the fact that the shadowing environment seen by signals at a given location is highly correlated and leads to bounded variation in received signal strengths. This insight can be used to implement a simple multiband algorithm in which the maximum signal strength of any channel is used to set the detection threshold.

5.3.3 Multipath variation and its correlation across channels

It is well known that multipath across frequencies greater than a few coherence bandwidths apart is independent. How about the multipath spread? Multipath spread is related the 'size' of the room/building through which the signal travels. Signals from two TV transmitters on the same tower reaching a single sensor will travel through similar multipath environments and hence their delay spreads should be positively correlated. Unfortunately we do not have a mechanism to measure the delay spreads of these signals (to measure the delay spread we need a controlled impulse transmitter). To understand how we can determine the positive correlation across delay spreads lets start with the problem of detecting a signal undergoing fading:

$$\begin{aligned}\mathcal{H}_0 : y[n] &= w[n] & n = 0, \dots, N-1 \\ \mathcal{H}_1 : y[n] &= \sum_{l=0}^{L-1} h_l x[n-l] + w[n] & n = 0, \dots, N-1\end{aligned}$$

where $x[n] \sim \mathcal{CN}(0, P_w)$, $(10 \log_{10}(P_w) = P = p_t - (l(r) + S))$ and $h_l \sim \mathcal{CN}(0, \frac{1}{L})$ are the channel filter taps.

For the case when the channels coefficients are unknown, the test statistic ($T(y)$) for this problem can be shown to be $\frac{1}{N} \sum_{n=1}^N |y[n]|^2$.

We can calculate the mean and variance of $T(y)$ where the mean and variance is over the random variables $h_0, \dots, h_{L-1}$ and $x[0], \dots, x[N-1]$. We make two simplifications in deriving these values. First, we assume that we are operating in the high SNR regime (so the impact of $w[n]$ is small and can be ignored). Second, we assume that $N \gg L$.




Figure 5.5: Performance of multiband sensing: the P_{HI} is plotted against the threshold (value below the maximum energy at which detection is performed). No significant effect is observed if Sutro and San Bruno are considered together.




Figure 5.6: Cumulative distribution of the received signal strength spread from all transmitters on a single tower. All transmitters in the FCC’s high and low database were used. The distribution is computed by sampling multiple locations across the continental United States.

$$\begin{aligned}
 \mathbb{E}[T(\mathbf{y})] &= \mathbb{E}[y[n]y^*[n]] \\
 &= \sum_{l=0}^{L-1} \sum_{l'=0}^{L-1} \mathbb{E}[h_l x[n-l] x^*[n-l'] h_{l'}^*] \\
 &= \sum_{l=0}^{L-1} \mathbb{E}[|h_l|^2] \mathbb{E}[|x[n-l]|^2] \\
 &= \frac{1}{L} \sum_{l=0}^{L-1} \mathbb{E}[|x[n-l]|^2] \\
 &= P_w
 \end{aligned}$$

Similarly, we can calculate the variance of $(T(y))$. First we calculate $\mathbb{E}[(T(y))^2]$.

$$\begin{aligned}
 \mathbb{E}[(T(\mathbf{y}))^2] &= \mathbb{E}\left[\frac{1}{N^2} \left(\sum_{n=0}^{N-1} y[n]y^*[n]\right)^2\right] \\
 &= \frac{1}{N^2} \left(\sum_{n=0}^{N-1} \mathbb{E}[|y[n]|^4] + \sum_{\substack{n=0, n'=0 \\ n \neq n'}}^{N-1} \mathbb{E}[|y[n]|^2 |y[n']|^2] \right)
 \end{aligned}$$

Next we calculate $\mathbb{E}[|y[n]|^4]$:

$$\begin{aligned}
 \mathbb{E}[|y[n]|^4] &= \mathbb{E}\left[\left(\sum_{l=0}^{L-1} \sum_{l'=0}^{L-1} h_l h_{l'}^* x[n-l] x^*[n-l']\right)^2\right] \\
 &= \mathbb{E}\left[\sum_{l_1, l'_1, l_2, l'_2=0}^{L-1} h_{l_1} h_{l'_1}^* h_{l_2} h_{l'_2}^* x[n-l_1] x^*[n-l'_1] x[n-l_2] x^*[n-l'_2]\right] \\
 &\stackrel{(a)}{=} \mathbb{E}\left[\sum_{l_1=0}^{L-1} |h_{l_1}|^4 |x[n-l_1]|^4\right] + \mathbb{E}\left[\sum_{\substack{l_1=0, l_2=0 \\ l_1 \neq l_2}}^{L-1} |h_{l_1}|^2 |h_{l_2}|^2 |x[n-l_1]|^2 |x[n-l_2]|^2\right] \\
 &\stackrel{(b)}{=} \frac{K^2 P_w^2}{L} + L(L-1) \frac{P_w^2}{L^2}
 \end{aligned}$$

where the left term in (a) is for the case when $l_1 = l_2 = l'_1 = l'_2$ and the right term is for the case when $l_1 = l'_1, l_2 = l'_2, l_1 \neq l_2$ or $l_1 = l'_2, l_2 = l'_1, l_1 \neq l_2$. All other cross terms produce zero expectation. In (b), $K = \frac{E[|X|^4]}{E[|X|^2]}$, $X \sim \mathcal{CN}(0, 1)$.

Similarly, we can calculate, $\mathbb{E}[|y[n]|^2 |y[n']|^2]$ as:

$$\begin{aligned}
 \mathbb{E}[|y[n]|^2 |y[n']|^2] &= \mathbb{E} \left[\left(\sum_{l_1=0}^{L-1} \sum_{l'_1=0}^{L-1} h_{l_1} h_{l'_1}^* x[n-l_1] x^*[n'-l'_1] \right) \right. \\
 &\quad \left. \left(\sum_{l_2=0}^{L-1} \sum_{l'_2=0}^{L-1} h_{l_2} h_{l'_2}^* x[n-l_2] x^*[n'-l'_2] \right) \right] \\
 &= \mathbb{E} \left[\sum_{l_1, l'_1, l_2, l'_2=0}^{L-1} h_{l_1} h_{l'_1}^* h_{l_2} h_{l'_2}^* x[n-l_1] x^*[n-l'_1] x[n-l_2] x^*[n'-l'_2] \right] \\
 &\stackrel{(a)}{=} \mathbb{E} \left[\sum_{l_1=0}^{L-1} |h_{l_1}|^4 |x[n-l_1]|^2 |x[n'-l_1]|^2 \right] + \\
 &\quad \mathbb{E} \left[\sum_{\substack{l_1=0, l_2=0 \\ l_1 \neq l_2}}^{L-1} |h_{l_1}|^2 |h_{l_2}|^2 |x[n-l_1]|^2 |x[n'-l_2]|^2 \right] \\
 &\stackrel{(b)}{=} \frac{K P_w^2}{L} + L(L-1) \frac{P_w^2}{L^2}
 \end{aligned}$$

where the left term in (a) is for the case when $l_1 = l_2 = l'_1 = l'_2$ and the right term is for the case when $l_1 = l'_1, l_2 = l'_2, l_1 \neq l_2$. The case $l_1 = l'_2, l_2 = l'_1, l_1 \neq l_2$ gives zero expectation. We are ignoring the case when n and n' are close enough that $n - l_1 = n' - l_2$ since the cardinality of such terms is small when compared to N . All other cross terms produce zero expectation. In (b), $K = \frac{E[|X|^4]}{E[|X|^2]}$, $X \sim \mathcal{CN}(0, 1)$.

Finally, we can calculate the variance of the test statistic as:

$$\begin{aligned}
 \mathbb{V}[T(\mathbf{y})] &= \mathbb{E}[(T(\mathbf{y}))^2] - (\mathbb{E}[T(\mathbf{y})])^2 \\
 &= \frac{1}{N^2} (N \mathbb{E}[|y[n]|^4] + N(N-1) \mathbb{E}[|y[n]|^2 |y[n']|^2]) - P_w^2 \\
 &= \mathbb{E}[|y[n]|^2 |y[n']|^2] - P_w^2 + \frac{1}{N} (\mathbb{E}[|y[n]|^2 |y[n']|^2] - \mathbb{E}[|y[n]|^4]) \\
 &= P_w^2 \left(\frac{K}{L} - 1 + \frac{1}{NL} K(K-1) \right)
 \end{aligned}$$

i.e. the variance of the test statistic scales as $\mathcal{O}(\frac{1}{L})$. This means that the *variance in the measured signal power at a given location is inversely proportional to the delay spread*. The reader might be surprised by the fact that the variance does not scale to zero with the number of samples N . This is due to the nature of the experiment we are considering. We fix the number of samples and conduct the experiment multiple times to obtain a sampling of the multipath. For the duration of each experiment, the multipath is assumed constant. The experiment is repeated several times for the each N i.e. we are considering scaling with N after averaging over the multipath.

Unfortunately, in our measurements we obtain a single snap shot of the multipath spread at a given location. However we can estimate the variance of the measured energy over smaller frequency bands. As a concrete example, consider the simple case where $x[n]$ is an impulse and we collect $N \geq L$ samples at a ADC sample rate R and perform the following procedure:

- Take the K -point FFT (where $K > N$) of the received sequence $z[0], \dots, z[K-1] = FFT(y[0], \dots, y[N-1])$.
- Compute the total energy in each frequency block of size B (expressed as a multiple of the bandwidth of a single FFT Bin). There are $Q = \lceil \frac{K}{B} \rceil$ such blocks.

$$E(q) = 10 \log_{10} \left(\sum_{n=qB}^{(q+1)B-1} |z[n]|^2 \right), \quad q = 0, \dots, Q-1$$

- Compute the standard deviation of the computed energies:

$$\hat{S} = StdDev(E(0), \dots, E(Q-1)).$$

In Figure 5.7, $\hat{S}$ is computed for different values of the delay spread (measured in number of taps) and for different averaging bandwidths ($Q = 16, 64, N = 1024, K = 1024$). Two

points stand out: firstly, the standard deviation/variance scales inversely with number of taps. Secondly, different taps perform the same relative to each other irrespective of the averaging bandwidth B being used i.e. examining the variance at low bandwidths is a accurate predictor of the variance the overall bandwidth of the transmission.

We can perform a similar calculation for for the measurement data. Figure 5.8 shows the calculated standard deviation ($\hat{S}$) versus averaging bandwidth for various locations. For each location, the measured standard deviation is calculated over all 12 TV channels. The difference between outdoor and indoor locations is evident. For outdoor locations (Dr Sahai’s deck for example), the multipath reflections come from hills and buildings and hence the delay spread is very large – consequently the calculated standard deviation is small. This can be contrasted to the indoor locations (BWRC lab for example) where the delay spread is small and hence the corresponding standard deviation is large.

Figure 5.9 shows the calculated standard deviation versus averaging bandwidth for two locations but plots all 12 channels separately. While the different frequencies do not fall on top of each other, the plots for each are relatively close together and some positive correlation can be expected among the standard deviations for different channels.

Table 5.4 shows the measured correlation among a few channels. To calculate these correlations, the standard deviation for each channel at each measurement point was computed and strung together into a single vector. Correlation between two channels was then calculated using the standard correlation formula. The average calculated correlation across channels was $\sim .3$.

Channel	19	29	30	44
33	0.24	0.3	0.31	0.21
38	0.28	0.37	0.34	0.31
45	0.25	0.3	0.33	0.24

Table 5.4: Multipath standard deviation between various channels using a averaging bandwidth of 10kHz. The average correlation across all channel pairs is 0.3

5.3.4 Multiband sensing in the real world

In the previous section we saw experimental evidence for the following facts:

- Shadowing is highly correlated across frequencies
- Multipath spread (as measured by the standard deviation of the energy in a frequency block) is correlated across frequencies.




Figure 5.7: Standard deviation ($\hat{S}$) of the energy measured in blocks of varying bandwidths. The standard deviation/variance scales inversely with number of taps. Different taps perform the same relative to each other irrespective of the averaging bandwidth B being used i.e. examining the variance at low bandwidths is an accurate predictor of the variance at full bandwidths.




Figure 5.8: Standard deviation ($\hat{S}$) of the energy measured in blocks of different bandwidths. The standard deviation/variance scales inversely with number of taps. Different taps perform the same relative to each other irrespective of the averaging bandwidth B being used i.e. examining the variance at small bandwidths is an accurate predictor of the multipath variance at larger bandwidths.




Figure 5.9: Standard deviation ($\hat{S}$) of the energy measured in blocks of different bandwidths for two locations. The standard deviation is measured separately across the 12 channels. While the different frequencies do not fall on top of each other, the plots for each are relatively close together and some positive correlation can be expected among the standard deviations for different channels.

- Noise does not exhibit any multipath spread (Standard deviation of the energy in a frequency block occupied by just noise is close to 0).

Based on these observations, we will define two detectors.

Multiband detector based on shadowing correlation only

The first detector is based solely on the fact that shadowing across frequencies is correlated. The procedure for executing this multiband detector is as follows:

- Calculate the energy in channel i as: $T(\mathbf{y}_i) = 10 \log_{10} \left(\sum_{n=0}^{N-1} |y_i[n]|^2 \right)$.
- Set the detection threshold as $\lambda_M = \max_i T(\mathbf{y}_i) - \delta$, where δ is a free parameter to change the tradeoff between P_{HI} and P_{FH} .
- For each channel decide on the channel present/absent ($\mathcal{H}_{1_i}/\mathcal{H}_{0_i}$) as:

$$T(\mathbf{y}_i) = \frac{1}{N} \sum_{n=1}^N |y_i[n]|^2 \underset{\mathcal{H}_{0_i}}{\overset{\mathcal{H}_{1_i}}{\gtrless}} \lambda, \quad (5.13)$$

For this detector we can calculate P_{HI} as the fraction of the 12 ‘on’ channels that we declare as absent ($\mathcal{H}_{0_i}$). Similarly, P_{FH} is the fraction of the remaining 21 channels that we declare as absent. Figure 5.10 compares the performance of a simple energy detector with the performance of the multiband detector discussed above. The multiband detector is able to give P_{FH} values close to 1 for low values of P_{HI} . For the energy detector, we need to raise the threshold high enough to ensure that we do not detect noise/interference. This impacts the number of spectrum holes we can recover.

Multiband detector based on correlation in shadowing and multipath spread across frequencies

Next, we define an enhanced multiband algorithm that also utilizes the facts that the multipath spread across frequencies is correlated and that noise has no multipath.

- Calculate the energy in channel i as: $T(\mathbf{y}_i) = 10 \log_{10} \left(\sum_{n=0}^{N-1} |y_i[n]|^2 \right)$.
- Calculate the standard deviation of multipath in a bandwidth B as follows:
 - Take the K -point FFT (where $K > N$) of the received sequence $z[0], \dots, z[K-1] = FFT(y_i[0], \dots, y_i[N-1])$.




Figure 5.10: Performance of multiband sensing: the probability of finding a hole P_{FH} (where the probability is calculated over the empirical distribution) is plotted against P_{HI} . The Multiband detector based on simple energy estimates in each channel is able to give values close to 1 for low values of P_{HI} . The Multiband detector based on the estimated energy and the multipath variance calculated at 10kHz resolution performs even better.

- Compute the total energy in each frequency block of size B (expressed as a multiple of the bandwidth of a single FFT Bin). There are $Q = \lceil \frac{K}{B} \rceil$ such blocks.

$$E_i(q) = 10 \log_{10} \left(\sum_{n=qB}^{(q+1)B-1} |z[n]|^2 \right), \quad q = 0, \dots, Q-1$$

- Compute the standard deviation of the computed energies:

$$\hat{S} = StdDev(E(0), \dots, E(Q-1)).$$

- Set the detection threshold as $\lambda_M = \max_i \left(T(\mathbf{y}_i) + \hat{S}(\mathbf{y}_i) - \delta \right)$, where δ is a free parameter to change the tradeoff between P_{HI} and P_{FH} .
- For each channel decide on the channel present/absent ($\mathcal{H}_{1_i}/\mathcal{H}_{0_i}$) as:

$$T(\mathbf{y}_i) + \hat{S}(\mathbf{y}_i) = \frac{1}{N} \sum_{n=1}^N |y_i[n]|^2 \underset{\mathcal{H}_{0_i}}{\overset{\mathcal{H}_{1_i}}{\gtrless}} \lambda_M, \quad (5.14)$$

Again, for this detector we can calculate P_{HI} and P_{FH} as a function of δ . Figure 5.10 compares the performance of a singleband detector (this is an enhanced version of the energy detector and uses the test statistic $T(\mathbf{y}_i) + \hat{S}(\mathbf{y}_i)$). The performance of the singleband detector is much better than that of the simple energy detector. This is because the mean of the test statistic $\hat{S}(\mathbf{y}_i)$ when the signal is present is much larger than when only noise is present. The enhanced multiband detector gains from the higher mean of $\hat{S}(\mathbf{y}_i)$ and also from the positive correlation of $\hat{S}(\mathbf{y}_i)$ across frequencies.

5.4 Summary

In this chapter we go beyond single band sensing and use information from nearby frequencies to estimate our shadowing environment. We call this approach multiband sensing. Three environmental ‘truths’ enable multiband sensing: towers are scarce and so a single tower typical houses many transmitters, shadowing and multipath spread is correlated across frequencies and noise exhibits no multipath spread. These environmental truths were validated using more than a thousand Digital TV measurements made at various locations in Berkeley, CA. Multiband detectors based on these insights outperform singleband detectors. Cooperation across multiband radios is robust against channel correlation and provides gains by weeding out radios afflicted by adverse propagation environments.

Chapter 6

Coexistence with small scale primaries

So far we have concentrated on large primaries. DTV transmitters are mounted on large towers, are high power devices (100kW) and have huge coverage areas (50-60km). The FCC had mandated white space devices with a maximum power of 1W with a 6dBi antenna [12] i.e. the secondary user is 5 orders of magnitude ‘smaller’ than the primaries they wish to coexist with. While TV transmitters are the main primary users in VHF/UHF frequency range, wireless microphones (part 74 devices) also need to be protected [84]. These devices are low power devices (50mW-250mW) and have small coverage areas (100m). So what does it take for IEEE 802.22 devices to protect wireless microphones? Or in a general case, what does it take for secondary devices to coexist with similar sized primary users? Coexistence with primaries of different scales is the focus of this chapter.

Three aspects of the problem change when we move from a large scale primary to a small scale primary:

- Radios cannot make individual decisions about their transmission when a primary is not found. An admission control network is needed to decide on which users can transmit.
- Sensing results are valid in a small area around the primary (due to the footprint of the primary). This requires a minimum sensing density to achieve the required diversity during sensing.
- With large primary systems, decisions were made about a primary that was very far away from the secondaries. Hence location uncertainty of the primary/secondary did not impact results. This is not the case for small scale primaries - location uncertainty can lead to very conservative estimates about the area where transmissions can occur.

Analyzing various aspects of the above problem forms the core of this Chapter. In Section 6.1 we provide a summary of the related work in this area. In Section 6.2, we seek to

answer the question of whether primary and secondary users of the same scale can coexist. In this section we also establish the need to decouple sensing from admission control. Admission control is analyzed in detail in Section 6.3. This section also quantifies gains from having location aware devices on each secondary transmitter. We focus our discussion around a toy model of the Part 74 ‘wireless microphones’. This chapter is based on the paper [25].

6.1 Related Work

Sensing wireless microphones requires sensing very weak signals since we are forced to confront the indoor/outdoor scenario. The wireless microphone and/or cognitive radio may be indoors or outdoors. Sensing very weak signals is impossible due to uncertainties in the noise+interference environment [85].

In [86], the authors describe the typical characteristics of wireless microphone signals. Wireless microphones are FM modulated signals. When a speaker is silent only a single the FM carrier and a tone key (a supra audible tone used by the receiver to identify the desired radio signal, so that it can unmute/mute) is seen. When the speaker is soft the transmitted signal transmitted is the FM carrier with moderate amount of deviation. When the speaker is loud this deviation increases substantially. In [87] the authors use Eigen value detection to detect wireless microphone signals. Two sinusoidal detection techniques are proposed in [88] for the detection of pilots and wireless microphones. The first technique depends on the FFT being a partially coherent detector for sinusoidal signals while the second technique runs two simultaneous PLLs. If a sinusoidal input is present the two PLLs lock to the same frequency. Beacons have also been proposed for wireless microphone to make them easy to detect.

When a secondary user transmits, it also needs to know the impact of its transmission on the primary receivers. This task is rendered difficult since the transmitter does not know its channel to the primary receivers; hence in general it needs to budget for the worst case e.g. its transmission has a line of sight to the primary receiver. This worst case budgeting can be avoided by placing a sensor near the primary receiver [89]. The sensor can feedback the channel information to the secondary transmitter which can then estimate the shadowing to the primary receiver.

In [90], the authors have also considered the required density of a sensor network performing distributed sensing. However the key constraint is the power density of the sensor network to relay the information back to the fusion center. In this context, the appropriate question is ‘Is it better to have few high power sensors or many smaller sensors?’.

6.2 Can secondary users protect similar scale primary users?

The key question to answer is, ‘*What should a secondary device do when it does not hear the primary user?*’. If the secondary device chooses to transmit, it must ensure that its transmission will not interfere with the primary receiver. This situation is depicted in Figure 6.1(a) - the secondary sensor must detect the primary transmitter while the secondary transmitter must ensure that the primary receiver is not interfered with. Figure 6.1(b) shows the scenario that we focus on first: the secondary sensor and transmitter are co-located (we will revisit this assumption at the end of this section) while the primary receiver is as close to the secondary as possible. Table 6.1 lists the relevant parameters, their description and their typical values for wireless microphones and IEEE 802.22 Customer Premise Equipment (CPE). These conditions are summarized in equation form as [91]:

$$\frac{P_1 r_s^{-\alpha_{12}} 10^{\frac{-\Delta_{12}}{10}}}{\sigma^2} \geq 10^{\frac{\beta}{10}} \quad (6.1)$$

$$\frac{\sigma^2 + P_2 * (r_s - r_p)^{-\alpha_{21}} 10^{\frac{-\Delta_{12}}{10}}}{P_1 r_p^{-\alpha_{11}} 10^{\frac{-\Delta_{11}}{10}}} \leq \frac{1}{10^{\frac{\delta}{10}}} \quad (6.2)$$

(6.1) states that the received SNR of the primary transmitter at the secondary sensor should be above the SNR wall. (6.2) specifies that the SINR at the primary receiver should be above the Desired/Undesired ratio requirements of the primary receiver.

The sensitivity of a wireless sensor is limited by the noise uncertainty. We would like to determine if the received signal power at the sensor will be above this value. We analyze this question starting with a very simple case where the primary receiver is co-located with the transmitter and complete reciprocity holds between the primary transmitter-secondary sensor channel and the secondary transmitter-primary receiver channels. We develop this story using the wireless microphone as the primary transmitter and the IEEE 802.22 CPE as the secondary sensor and transmitter. We gradually move to more realistic scenarios quantifying the effect of each change. These scenarios are depicted in Figure 6.2 which examines if the received power at the primary receiver can be above the desired sensitivity in each case.

Case I: Complete reciprocity, primary receiver co-located with primary transmitter

At the edge of the sensing region, (6.1) becomes an equality and we can express the




Figure 6.1: (a) Sensing and interference scenario where we must ensure that the received energy at the sensor is above the sensing sensitivity. Further, we must ensure that interference to the primary receiver must be below the SINR requirements of the receiver. (b) Typical sensing scenario where the secondary sensor and transmitter are co-located and the primary receiver is at its worst case location with respect to interference.

Parameter	Description	Typical values
P_1	Power of primary transmitter	100mW in 200kHz
P_2	Power of secondary transmitter	1W (30dBm) in a 6MHz band 15dBm in 200kHz
σ^2	Noise power	-121dBm (200KHz)
r_p	Protected Radius	100m
β	SNR wall	-3dB
Δ_{12}	Fading from primary transmitter to secondary sensor	0 - 50dB
Δ_{21}	Fading from secondary transmitter to primary receiver	0 - 50dB
Δ_{11}	Fading from primary transmitter to primary receiver	0 - 50dB
α_{12}	Path loss exponent from primary transmitter to secondary sensor	2 - 5
α_{21}	Path loss exponent from secondary transmitter to primary receiver	2 - 5
α_{11}	Path loss exponent from primary transmitter to primary receiver	2 - 5
δ	SINR	20dB
D	Transmission Density of the Secondary users (Watts/ km^2)	
N_w	Number of secondary radios around a primary receiver (worst case)	16
R	Length of a secondary cell (assuming a square cell)	10-33km
H	Number of secondary cells between the primary transmitter and the sea of users	5-15 5-15

Table 6.1: Relevant parameters, their description and typical values for wireless microphones as primaries and IEEE 802.22 CPEs as secondaries.




Figure 6.2: (Case I) Complete reciprocity ($\alpha_{21} = \alpha_{12} = \alpha_{11}$ and $\Delta_{12} = \Delta_{21}$); primary transmitter is co-located with the primary receiver ($r_p = 0$ and $\Delta_{11} = 0$). (Case II) Primary receiver is r_p distance away from the primary transmitter. (Case III) $\Delta_{12} - \Delta_{21} = 20\text{dB}$. (Case IV) $\alpha_{12} - \alpha_{21} = 3$.

sensing radius (r_s) as:

$$r_s^{-\alpha_{12}} = \frac{\sigma^2}{P_1} 10^{\frac{\beta + \Delta_{12}}{10}} \quad (6.3)$$

To obtain a bound on the secondary power (P_2) from a secondary transmitter at the edge of sensing region, we assume that the interference term in (6.2) dominates the noise power and that the sensing radius is very large as compared to the protected radius ($r_s \gg r_p$). Furthermore, since the primary transmitter and primary receiver are co-located, $\Delta_{11} = 0$. Our reciprocity assumptions also imply that $\alpha_{12} = \alpha_{21} = \alpha_{11}$ and that $\Delta_{12} = \Delta_{21}$. Under these assumption and substituting (6.3), in (6.2) we get:

$$P_2 \leq \frac{P_1^2}{\sigma^2} 10^{\frac{\beta - \delta}{10}} \quad (6.4)$$

In the log domain, this translates into,

$$P_2(dB) \leq 2 P_1(dB) - \sigma^2(dB) + \beta - \delta \quad (6.5)$$

Figure 6.2 illustrates this case for the wireless microphone. Since the D/U ratio for the microphone is around 20dB, the maximum interference at the primary receiver can be as high as 0dBm. Since the secondary signal only needs to decay by 15dB, complete reciprocity implies that the primary signal power at the secondary sensor will also be very high (20dBm - 15dB = 5dBm). This is must larger than the target sensitivity (-124dBm) and hence the IEEE 802.22 CPE can sense the microphone and does not risk interfering with the receiver.

Case II: Complete reciprocity, primary receiver is 100m away from the primary transmitter

For this case we assume that primary transmitter and receiver pair is separated by a distance of r_p meters. In this case, (6.2) at the edge of the sensing radius (assuming $\frac{r_s}{r_p} \gg 1$) takes the following form,

$$P_2 \leq r_p^{-\alpha_{12}} \frac{P_1^2}{\sigma^2} 10^{\frac{\beta - \delta}{10}} \quad (6.6)$$

Again in the log domain, this translates into,

$$P_2(dB) \leq 2 P_1(dB) - \sigma^2(dB) - 10\alpha_{12} \log_{10}(r_p) + \beta - \delta \quad (6.7)$$

This is the same as (6.5) except for the term $10\alpha_{12} \log_{10} r_p$ which could be as large as 100dB for the wireless microphone ($\alpha_{12} = 5$, $r_p = 100\text{m}$). With this additional term, the interference at the primary receiver has to be limited to -100dBm (20dBm (primary power) - 100dB (distance loss) - 20dB (D/U ratio) = -100dBm). This requires the loss from the secondary transmitter to be larger (115dB as seen in Figure 6.2). With complete reciprocity this implies that the primary power at the secondary sensor is still large enough (20dBm - 115dB) and hence above the sensing sensitivity.

Case III: Primary receiver is 100m away from the primary transmitter, fading is not reciprocal

Since the primary transmitter-receiver pair are more than a wavelength apart, the multipath at both these places is essentially independent. Shadowing on the other hand could be highly correlated since the primary transmitter-receiver pair are very close as compared to the secondary. In this case, (6.2) at the edge of the sensing radius takes the following form,

$$P_2 \leq 10^{\frac{\Delta_{21}-\Delta_{11}-\Delta_{12}}{10}} r_p^{-\alpha_{12}} \frac{P_1^2}{\sigma^2} 10^{\frac{\beta-\delta}{10}} \quad (6.8)$$

In the log domain, the above equation is the same as (6.7) except for the difference between fading margins ($\Delta_{21} - \Delta_{11} - \Delta_{12}$). We are interested in the case (worst case) where the fading between the secondary transmitter and the primary receiver is less than the fading between the primary transmitter and the secondary sensor ($\Delta_{12} > \Delta_{21}$)¹. In [20], it has been shown that the multipath variability can be as high as 20dB for reasonable detection probability. Hence we budget 20dB for this mismatch.

Case IV: Primary receiver is 100m away from the primary transmitter, fading is not reciprocal, path loss exponents are unequal

Inequality between exponents can arise from the physical configuration of the transmitters. In the case of wireless microphones, the wireless microphone transmitter may be on the ground while the wireless microphone receiver may be mounted on a truck. The IEEE 802.22 transmitter may be mounted in the roof top. Here the propagation between the microphone

¹In this case we are assuming that $\Delta_{11} = 0$

and the secondary sensor is worse than the propagation between the secondary transmitter and the primary receiver. For this scenario, (6.2) at the edge of the sensing radius (assuming $\frac{r_s}{r_p} \gg 1$) takes the following form,

$$P_2 \leq \frac{r_p^{(\alpha_{12}-\alpha_{11})}}{r_s^{(-\alpha_{21}+\alpha_{12})}} 10^{\frac{\Delta_{21}-\Delta_{11}-\Delta_{12}}{10}} r_p^{-\alpha_{12}} \frac{P_1^2}{\sigma^2} 10^{\frac{\beta-\delta}{10}} \quad (6.9)$$

Again, this is the same as (6.8) except for the factor $\frac{r_p^{(\alpha_{12}-\alpha_{11})}}{r_s^{(-\alpha_{21}+\alpha_{12})}}$. We assume that $\alpha_{12} = \alpha_{11}$ and would like to budget for the case that $\alpha_{12} > \alpha_{21}$. Hence, the secondary power has to be further scaled down by $10(\alpha_{12} - \alpha_{21}) \log_{10}(r_s)$. Since r_s is at least as large as r_p , this mismatch could be 60dB if $\alpha_{12} - \alpha_{21}$ is 3.

Figure 6.3 shows the maximum secondary power as a function of secondary sensor sensitivity for the four cases discussed above. It is interesting to note that considering all mismatches and a SNR wall of -3dB, the secondary power should be limited to -36dBm; thats 56dB below the primary power in the band of interest. This gap between the primary and secondary power is similar to the gap between a TV transmitter and a IEEE 802.11 CPE (100kW versus 1W $\sim$ 50dB).

Before we continue with our discussion, let us examine two aspects of the above analysis. The first is the choice of exponents and the second questions the impact of diversity on the above analysis.

The case for Path Loss Exponents: Can the loss exponent between the primary transmitter and the secondary sensor be as high as 5? The fixed fading margin that we have assumed, takes into account multipath and the shadowing at the receiver/transmitter. To factor in the effect of distance dependent shadowing, we need to take into use higher path loss exponents. Secondly, there are example scenarios where the path loss exponent for sensing is large while the path loss exponent for interference could be small. Consider the case where the wireless microphone is on the ground while the microphone receiver is mounted on a truck 100 meters away. The secondary transmitter is roof mounted and hence has a clear line of sight to the primary receiver. Due to ground bounce and shadowing local to the transmitter, the path loss exponent for sensing could be high while the exponent for interference could be small.

The impact of multiuser diversity: The 20dB fading mismatch that we budget has two components: The first part is the fading between the primary transmitter and the secondary sensor. This fading can be reduced by multiuser diversity at the sensors. The second part is the fading from the secondary transmitter to the primary receiver. For this, each secondary transmitter must budget on its own - this process does not gain from multiuser diversity.




Figure 6.3: Maximum secondary power as a function of secondary sensor sensitivity. Considering all mismatches and a sensitivity of -124dBm, the secondary power should be limited to -36dBm; that's 56dB below the primary power in the band of interest.

Similarly, during cooperative sensing, a secondary user may have a line of sight to the primary transmitter. In this case, we do not need to budget for the worst case path loss exponent from the primary transmitter to the secondary.

The preceding analysis emphasizes the following fact: if a IEEE 802.22 CPE does not sense a wireless microphone and decides to transmit, it may interfere with a wireless microphone receiver. So what can a secondary device like a IEEE 802.22 CPE do to deal with a primary transmitter which is spatially matched in scale?

One of the solutions to this problem is to decouple sensing and admission control. Sensing answers the question ‘*Is the primary close enough?*’ while admission control seeks to answer the question ‘*Can I transmit?*’. In the analysis so far, the radio that performs sensing also decides whether it can transmit. This approach was valid for large scale primary users [92] where a secondary user could make a transmission decision on its own (with a little cooperation with nearby secondaries to increase robustness). This is not the case for small scale primary users. For small scale primaries, we need a sensor network to localize the position of the primary user² and an admission control network to decide on which radios can transmit. This decoupling is the key to handling small scale primary users.

To calculate a target sensor density, we set a diversity target of 8 users (we would like 8 users to be within the sensing region with the hope that at least one of them will have a good channel to the primary user). This would mean we only need to budget 20dB for the fading margin since cooperation among 8 users causes sensitivity to get within 20dB of the path loss [20]. Table 6.2 shows the sensing density requirements for various path loss exponents. Ideally, we would like the density of active users (users that have data to transmit) to meet the requirements of Table 6.2 so that we do not require inactive users to sense. Unfortunately, the IEEE 802.22 Working Group targets areas where the population density is around 1.25 people/ km^2 [93]. Hence, in the best case this implies a deployment density of 1.25 secondary nodes/ km^2 . In this case, the deployment density (leave alone the active user density), is not adequate if the primary to secondary loss exponent is 4 or greater.

A note about path loss exponent: We have seen that an increase in the primary-to-secondary path loss exponent causes the sensor density to increase. What if we budget for the worst case ($\alpha_{12} = 5$) and the actual loss exponent turned out to be 2; would that impact the system adversely? If the actual loss is 2 then many of the radios/sensors around the primary user will detect the primary. The impact of this would depend on the actual admission control algorithm. If all the radios that sense the primary need to be quiet, then a small loss exponent would render a large area unusable. On the other hand, many sensors sensing the primary could help localize the primary’s location better. With this information, even radios that sense the primary may be able to transmit. The opposite is true for the

²This sensor network may not be separate from the radio network; however sensor density requirements greater than the density of active users would require inactive nodes to sense

Loss Exponent (α_{12})	Area of sensing region (km^2)	Secondary density requirements secondary nodes/ km^2
2	7.89×10^6	10.16×10^{-7}
3	580.52	0.0136
4	4.98	1.68
5	0.29	27.92

Table 6.2: Area of sensing region for various primary-to-secondary path loss exponents. Since we require 8 secondary sensing nodes in each sensing region, this gives the equivalent sensing density requirements. The target IEEE 802.22 deployment density is 1.25 secondary nodes/ km^2 [93]. This means that the deployment density is not adequate if the primary to secondary loss exponent is 4 or greater.

Path Loss Exponent (α_{21})	Pollution radius of a single secondary user (km)
2	1778
3	14
4	1.33

Table 6.3: Pollution area is a region around a secondary user where a primary receiver will face undesired interference. Due to the symmetry in our model, the pollution area is a circle and hence the pollution radius is the right parameter to study. This table shows the pollution radius around a single secondary user for various secondary-to-primary path loss exponents

secondary-to-primary path loss exponent. In this case a large exponent is preferred (see Table 6.3) but we need to budget for a small exponent.

To deal with secondary density requirements, the IEEE 802.22 Working Group must either target suburban/urban environments where the deployment density is large enough to provide the required sensor density, or it must characterize the rural environment to ensure that the path loss exponent is better than 4.

6.3 Admission Control

In the last section we established the fact that an admission control network is required to decide on which users can actually transmit; a single user cannot make this decision on its own.

6.3.1 Secondary Pollution Area

To develop a simple admission control algorithm, we first introduce the concept of a *secondary pollution area*. Pollution area is a region around a secondary user where a primary receiver will face undesired interference. This pollution area is easy to calculate for a single secondary radio, but the presence of multiple secondary users make this pollution radius larger due to cumulative interference.

To start, we first calculate the acceptable interference level at the primary receiver. Assuming a primary-primary loss exponent (α_{11}) of 5 and a fading margin of 20dB, the acceptable interference turns out to be -110dBm.

To obtain the worst case pollution radius we need a model for the placements of the secondary radios. In [91] the power of a secondary radio was smeared over its footprint – this was a viable option since the primary was far away from the secondary and from its perspective, the secondaries seemed like a sea of transmitting users. The main mode of failure was that all the secondary receivers at the edge of the protected radius received undesired interference. For small scale primaries this is not the case - the typical mode of failure is that some secondary radio is close enough to interfere with the primary. In the absence of a trusted model for radio placements, we need to budget for worse case placements. Hence we must calculate the maximum number of secondary users simultaneously transmitting on the uplink in the band of a microphone ($\sim 200\text{kHz}$). The uplink bandwidth per user is targeted to be 384kbps [93]. We arbitrarily assume that 4 secondary users in a single cell can be active simultaneously on the channel occupied by a wireless microphone.

Next we factor in the interference from the other cells. For these cells, in keeping with our modeling philosophy, we assume worst case radio placements for the transmitting radios i.e. we assume that the radios are placed at the corner which is nearest to the microphone




Figure 6.4: The worst case pollution radius (and assuming that $r_{pol} \gg r_p$ is for a radio that is at the corner of a cell (for convenience we assume square cells). Around each microphone, each adjoining cell can have up to 4 radios at the pollution radius boundary. Radios outside the horizon are treated as a seas of secondary users with a power density corresponding to 4 simultaneous transmitting users. For cells within the horizon, we assume worst case radio placements which will impact the microphone adversely.

Number of radios around the pollution circle	Radius of pollution circle (km)
4	1.89
8	2.24
12	2.48
16	2.67

Table 6.4: Pollution radius for various number of radios around the pollution circle. Since the transmission densities in the sea of radios beyond the horizon are very low, the position of the horizon does not impact the pollution radius.

(See Figure 6.4). We make this exact computation for all cells which are within the horizon (the horizon (H) is defined in terms of number of cells and is a variable in our model).

To calculate the influence from the users beyond the horizon, we use Equation (12) in [91]. This equation states that the effect of a sea of secondary users with power density D is the same as a single user of the same power but a loss exponent which is reduced by 2 (for this case the power density will correspond to a density of 4 radios per cell, $D = \frac{4P_2}{R^2}$ where R is the length of each side of the cell and P_2 is the power of a single secondary node).

Hence, the pollution radius is the smallest value of r that satisfies:

$$\begin{aligned}
 & D K(\alpha_{21}) (R H)^{-\alpha_{21}+2} + N_w P_2 r^{-\alpha_{21}} + \\
 & 4 P_2 R^{-\alpha_{21}} \sum_{i=1}^H \sum_{j=1}^H (i^2 + j^2)^{-\alpha_{21}/2} + 8 P_2 \sum_{j=1}^H (R \times j)^{-\alpha_{21}} \\
 & \leq 10^{\frac{I_{max}}{10}}
 \end{aligned}$$

where D is the maximum transmission density of the secondary nodes (as calculated above) (Watts/ m^2), $K(\alpha_{21}) = \frac{1}{\alpha_{21}-2} \int_{-\pi/2}^{\pi/2} \cos(\theta)^{-\alpha_{21}+2} d\theta$, I_{max} is the interference limit (-110dBm), N_w is the number of users around the pollution circle, and H is the number of cells within the horizon.

Our analysis shows that the sea of secondary users beyond the horizon has very little effect if the transmission densities are assumed to be small (4 transmitting radios per cell). For this case, Table 6.4 shows the value of the pollution radius for various worse case configurations (4/8/12/16 radios just around the pollution circle).

To quantify the effect of a microphone on the performance of a 802.22 cell, we introduce the notion of *effective microphone footprint*. This is the area around the microphone where a secondary user cannot transmit. The microphone has an actual footprint that is dependent on the situation on the ground. Since we do have complete knowledge of this, we have to be conservative and need to budget by using worst case exponents, location uncertainty and

worst case placement assumptions. The *effective microphone footprint* captures this loss from conservatism.

Since it is impossible to know if the sensed microphone is near the edge of the cell or near the center we need to budget for the worst case (microphone is in a corner and hence there can be up to 16 users just outside the pollution circle). In this case the pollution radius is 2.67km and the effective microphone footprint is $7.27km^2$ ($\alpha_{21} = 4$).

It is prudent to compare these results to the results from a model which considers all secondary users as a sea of secondaries with a certain power density. The density model yields a pollution radius of 315 meters for the same set of parameter values; this is much more optimistic than the prediction from worst case placement assumptions.

6.3.2 Location uncertainty of the primary

In section 6.3.1, we assumed that the location of the primary and secondary users is known exactly. In this section we shall first relax the assumption that the location of the primary is known. This is important since the the primary is being sensed and hence the location of the primary will only be known approximately.

We first characterize the *location uncertainty* of the primary user as a function of the sensor density. A location uncertainty of x meters implies that the primary could be closer/further by x meters from its true location. It is important to note that this is an example of a case where uncertainty can be reduced by actual measurements on the ground and hence we do not need to stick with worst case assumptions.

In Table 6.4, the maximum sensing radius was calculated by requiring a diversity of 8 sensing radios around the primary. A diversity of 8 assumes that at least one radio will catch the primary and the uncertainty would be equal to the sensing radius³. When the sensing density increases beyond this, multiple radios exist in the sensing circle - in this case the sensing radius is governed by the distance to the radio that sees the best channel. This distance shrinks (in an expected sense) as the density of sensing nodes increases.

Figure 6.6 shows the *effective microphone footprint* for various secondary densities. Increasing the densities of secondary sensors improves the location uncertainty which in turn decreases the effective microphone footprint. At the basic required density for $\alpha_{12} = 4$ (1.68 sensing nodes/ km^2 - see Table 6.2), the effective microphone footprint goes up by a factor of 2.

6.3.3 Location uncertainty of the secondary

There are three different ways to resolve location uncertainty of secondary users. The first is to have people/technicians register these devices. The method suffers from costs related

³This is a worst case assumption - location uncertainty would be better than the sensing radius in most cases




Figure 6.5: Location uncertainty about the primary as the secondary density is scaled. This uncertainty must be added to the pollution radius to ensure that the primary radio is protected.




Figure 6.6: Effective Microphone Footprint with location uncertainty of the primary and secondary users.

to technicians or to trust issues related to people registering their own devices. The second method is to install location aware devices on these radios – this method also incurs costs. The third (and the method analyzed in this paper) is to have the sensor nodes that sense the primary also sense the secondary users in order to estimate their locations.

Since the IEEE 802.22 devices have large power as compared to wireless microphones (1W versus 100mW), they can be sensed by the same infrastructure that is sensing the primary user. By our previous assumption, the location uncertainty related of these secondary devices would be the same as the location uncertainty of the primary for a given secondary sensor density.

Figure 6.6 again shows the effect of the combined location uncertainty of secondary and primary users. At the basic required density for $\alpha_{12} = 4$ (1.68 sensing nodes/ km^2), the effective microphone footprint goes up by a factor of 4.

6.4 Discussion

So what choices do we have for enabling coexistence between primaries and secondaries of similar scale? We can place the burden of coexistence on either the secondary user alone or we can have obtain primary assistance as well.

If the burden lies entirely on the secondaries, they can either increase their sensor densities (as explained in earlier sections) or they can employ sophisticated sensing techniques (like coherent detection)⁴. The SNR wall for energy detection can be reduced by 20dB by using coherent detection [85]. This increases the sensing radius which in turn reduces sensor density requirements. For $\alpha_{12} = 4$, every dB improvement in sensitivity corresponds to a .5dB equivalent increase in the sensing area.

The primary user can enable opportunistic sharing by one of the following ways: Firstly, primaries could introduce beacons which could facilitate coherent detection (as discussed above). Secondary, the primary receivers can issue denials whenever the interference is larger than the desired/undesired ratio⁵. Since denials have low probability of occurrence the entire secondary network could cease transmission once a denial is issued. After this, secondary nodes could gradually reenter the system calibrating their impact on the primary.

6.5 Summary

Real world cognitive systems will need to coexist with multiple classes of primary users each with its own scale of operation. Unfortunately, coexistence with primary users of different

⁴This does require primary users to have beacons/pilots that can be easily detected

⁵A denial is a identifiable signature that can be issued by the receiver whenever interference reaches an unacceptable level

scale leads to different infrastructure requirements on the secondary system. While sensing a primary user of a scale much larger than the secondary user (primary power 5 orders of magnitude greater than secondary power), the secondary user can make the decision about its transmission based on the sensing results in its own neighborhood. This key system assumption breaks down in the presence of small scale primary users. For such primaries, we need to decouple the sensing and admission control decisions. A secondary user that does not sense a primary *cannot* transmit until a separate admission control entity/network allows it to do so. In the case of IEEE 802.22, the sensing network requirements needed to ensure that all microphones can be reliably sensed are much larger than the target deployment density.

A simple admission control algorithm is proposed which only allows a secondary user to transmit if it is beyond the worst case pollution radius from the primary. Even if we assume that the location of the primary and secondary radios is known exactly, the effective microphone footprint is $7.13km^2$. If the location of the primary and secondary users is known approximately (based on having the sensor network sensing the primary and secondary radios), then the effective microphone footprint scales up by a factor of (around) 4.

Chapter 7

Conclusions and Future Work

Cognitive radios is a nascent field — defining the right problems is as important as finding the right solutions. This has been the guiding principle of this research which allowed us to gain many insights and make significant steps towards the understanding of the most pertinent issues in cognitive radios.

7.1 Advances

Most research in cognitive radios today is focussed on improving detection sensitivity. However, as shown in the early part of this thesis, improving detection sensitivity is not always congruent with the goal of recovering spectrum holes. In this thesis we focussed on the problem of recovering spectrum holes – especially spatial spectrum holes. To this end, we analyzed the FCC’s DTV databases and the population data of the United States (as per the Census of 2000), to determine the white space opportunity enabled by the FCC’s ruling of Nov 4th, 2008. This analysis reveals an interesting contrast between the FCC’s geo-location rules (which pick the appropriate political and engineering trade-off point between whitespace availability and broadcast use) and sensing rules (which are extremely conservative and are unable to recover most of the available whitespace opportunity).

Having identified the gap between geo-location and sensing via use of Monte Carlo simulations, we presented a simplified single-tower model that captures this gap fairly accurately. Together with the proposed metrics of Fear of Harmful Interference (F_{HI}) and Weighted Probability of Area Recovered ($WPAR$) this model can be used to compare advanced sensing rules.

Cooperative and multiband sensing aim to use information about the environment (either collected from nearby radios as in the case of cooperation or from nearby frequency bands as in the case of multiband) which is as vital to the detection process as the information in the band of interest at the radio. Both these schemes provide a mechanism to implement

‘dynamic sensing threshold’ where the dynamism allows a close relationship between the sensing threshold and the actual spectral environment of the radio. The median rule for cooperation was shown to be robust to outliers, misbehaving radios and fading uncertainties and works well for both temporal and spatial holes. Unfortunately cooperation suffers from a loss in spatial diversity and from the presence of extremely faded/shadowed radios. Multiband sensing can be used to combat shadowing correlation and weed out faded/shadowed radios. Cooperation across multiband radios is robust against channel correlation and enables the ‘weeding out’ process. A mobile, wideband testbed was designed and used to capture Digital TV signals in the 500-700MHz band at various locations in Berkeley, CA. Analysis of these DTV measurements reveals high shadowing and multipath spread correlation across frequencies. These new insights into the spectral environment forms the basis for multiband sensing — detectors based on multiband sensing are shown to outperform singleband sensing strategies in theory as well as in the real world.

While large scale primaries (e.g. TV transmitters) have been the focus of attention, protecting small scale primaries (e.g. Wireless Microphones) poses a even harder problem — these primaries have no predefined frequency of operation and no distinguishable signature. Furthermore the interference radius of 4W secondaries is much larger than the sensing radius of a 50mW primary. In these cases, we need to decouple sensing from admission control — a sensor network is required to perform the sensing and localize the primary. For small scale primaries, the environment over which the sensing results are valid is small which imposes certain minimum density requirements for sensor nodes. Furthermore, location information of the primary and secondary users is key for such an admission control algorithm to operate successfully.

7.2 Future work

This thesis has opened up an interesting set of issues and insights which provide a rich set of problems for further research. Some of these problems are near-term i.e. they form natural next steps to existing problems (they have either being investigated or have already been investigated) while some other problems arise by taking the ten thousand foot view of this research.

7.2.1 Near Term

White space Capacity: In this thesis we have determined white space availability per person across the United States using the protection and pollution viewpoints. The natural next question is: “What is the available white space capacity per person?”. At each location we need to exclude certain channels due to protection requirements. For the other channels we have to deal with interference from primaries and other secondaries. Based on the char-

acteristics of the secondary base station, we can determine the downlink capacity at each location. This work is presented in [16].

Theoretical basis for people tradeoff plots for all detectors types: In Chapter 2 we presented people-channel tradeoff plots for geo-location using actual population data. Then in Chapter 3 we derived a theoretical model for population density as a function of distance from a tower using the same data. The next step would be to use this model to obtain people tradeoff plots for all detectors without having to perform extensive Monte Carlo simulations.

A people and tower based weighting strategy: The weighting that we finally used for all our plots was based on the Voronoi approach and neglected the population density. We need a weighting strategy that combines the population density model with a ‘power Voronoi’ (where the Voronoi region around each tower is a function of its transmit power and height) to be able to predict the sensing performance relative to geo-location in terms of channels per person.

Time space metrics: The metrics proposed in this section address the problem of recovering spatial holes. We need a combined metric to capture the secondary’s ability to recover the time-space hole of a ‘on-off’ primary.

Shadowing and noise calibration gains in multiband sensing: In Chapter 5 we have looked at the possibility of using noise calibration across multiple bands to overcome noise uncertainty when the uncertainty in the two bands is the same. This was analyzed for the case of infinite samples. The interplay of shadowing correlation and noise calibration needs to be evaluated.

7.2.2 Long term

Combined sensing and geo-location – Map making: In Chapter 2 we assumed that geo-location provides the golden standard to determine exclusion zones. Unfortunately, geo-location is based on ITU propagation curves which may not apply to all deployment types. What we need are ‘anchor’ sensors that can provide received strength measurements for channels at a given location. Given a secondary’s measurements it can then compute its location relative to these anchors. There are two advantages of this approach. Firstly, it eliminates the need for each tower to register in a database. Only the anchors and their measurements are registered — the onus is on the secondaries (and not the primaries) to keep the database consistent. Secondly, it eliminates the problems with using the ITU propagation curves — protection radii are computed dynamically.

A generalized framework to study directional primaries and directional secondaries: The

work in this thesis is based on the assumption of omni-directional primary and secondary users. The next likely candidate bands for opportunistic communication are the long-range microwave links and satellite bands. These bands employ directional antennas and hence secondaries also need to be able to create ‘nulls’ in the direction of the primary. We need a generalized framework to study such primaries and secondaries.

Certification and/or enforcement of systems rather than radios: This thesis proposes cooperative sensing as a way to recover spectrum holes and improve on the performance of a single radio sensor. The advantage of the single radio sensor, however, is actually in the certification process. It is much easier to certify that a radio meets the -114dBm sensing threshold than it is to certify that a system of radios will be able to reliably sense for the presence of the primary under all deployment scenarios. Certification of cooperative sensing is hence of great concern.

Bibliography

- [1] F. C. Commission, “Auction 73: 700 MHz Band,” online, 2008. [Online]. Available: http://wireless.fcc.gov/auctions/default.htm?job=auction_summary&id=73
- [2] T. Erpek, K. Steadman, and D. Jones, “Spectrum Occupancy Measurements: Dublin, Ireland, Collected on april 16-18, 2007,” Shared Spectrum Company, Tech. Rep., 2007. [Online]. Available: http://www.sharespectrum.com/measurements/download/Ireland_Spectrum_Occupancy_Measurements_v2.pdf
- [3] M. A. McHenry and D. McCloskey, “Spectrum Occupancy Measurements: Chicago, Illinois, November 16-18, 2005,” Shared Spectrum Company, Tech. Rep., 2005. [Online]. Available: http://www.sharespectrum.com/measurements/download/NSF_Chicago_2005-11_measurements_v12.pdf
- [4] T. Erpek, M. Lofquist, and K. Patton, “Spectrum Occupancy Measurements: Loring Commerce Centre, Limestone, Maine, september 18-20, 2007,” Shared Spectrum Company, Tech. Rep., 2007. [Online]. Available: http://www.sharespectrum.com/measurements/download/Loring_Spectrum_Occupancy_Measurements_v2_3.pdf
- [5] M. Islam, C. Koh, S. Oh, X. Qing, Y. Lai, C. Wang, Y.-C. Liang, B. Toh, F. Chin, G. Tan, and W. Toh, “Spectrum Survey in Singapore: Occupancy Measurements and Analyses,” in *Cognitive Radio Oriented Wireless Networks and Communications, 2008. CrownCom 2008. 3rd International Conference on*, May 2008, pp. 1–7.
- [6] R. W. Broderson, A. Wolisz, D. Cabric, S. M. Mishra, and D. Willkomm, “White paper: CORVUS: A Cognitive Radio Approach for Usage of Virtual Unlicensed Spectrum,” University of California, Berkeley, Tech. Rep., 2004. [Online]. Available: http://bwrc.eecs.berkeley.edu/Research/MCMA/CR_White_paper_final1.pdf
- [7] J. M. III, “Cognitive Radio An Integrated Agent Architecture for Software Defined Radio,” Ph.D. dissertation, KTH Royal Institute of Technology, Stockholm, Sweden, 2000.

- [8] “Spectrum policy task force report,” Federal Communications Commision, Tech. Rep. 02-135, Nov 2002. [Online]. Available: [http://hraunfoss.fcc.gov/edocs\\$_public/attachmatch/DOC-228542A1.pdf](http://hraunfoss.fcc.gov/edocs$_public/attachmatch/DOC-228542A1.pdf)
- [9] “Unlicensed Operation in the TV Broadcast Bands and Additional Spectrum for Unlicensed Devices Below 900 MHz in the 3 GHz band,” Federal Communications Commision, NOTICE OF PROPOSED RULEMAKING 04-186, May 2004. [Online]. Available: [http://hraunfoss.fcc.gov/edocs\\$_public/attachmatch/FCC-04-113A1.pdf](http://hraunfoss.fcc.gov/edocs$_public/attachmatch/FCC-04-113A1.pdf)
- [10] “The XG vision, rfc v1.0,” Darpa XG working group, Tech. Rep., 2003. [Online]. Available: http://www.darpa.mil/ato/programs/XG/rfc_vision.pdf
- [11] “The XG architectural framework, rfc v1.0,” Darpa XG working group, Tech. Rep., 2003. [Online]. Available: http://www.darpa.mil/ato/programs/XG/rfc_af.pdf
- [12] F. C. Commission, “In the Matter of Unlicensed Operation in the TV Broadcast Bands: Second Report and Order and Memorandum Opinion and Order,” Federal Communications Commision, Tech. Rep. 08-260, Nov. 2008. [Online]. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-260A1.pdf
- [13] S. M. Mishra and A. Sahai, “How much white space is there?” Department of Electrical Engineering and Computer Science, University of California Berkeley, Tech. Rep. EECS-2009-3, Jan. 2009. [Online]. Available: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-3.html>
- [14] —, “How much white space has the fcc opened up?” *To appear in IEEE Communication Letters*, 2010.
- [15] —, “Pollution vs protection in determining spectrum whitespaces: a semi-empirical view,” *Submitted to IEEE Transactions on Wireless Communications*, 2010.
- [16] K. I. Harrison, S. M. Mishra, and A. Sahai, “How much white space capacity is there?” in *New Frontiers in Dynamic Spectrum Access Networks, 2010. DySPAN 2010. 4th IEEE International Symposium on*, Singapore,, 2010.
- [17] S. M. Mishra, “White space code and data, version 0.1.” [Online]. Available: http://www.eecs.berkelet.edu/~smm/white_space_data_and_code.zip
- [18] R. Tandra, S. M. Mishra, and A. Sahai, “What is a spectrum hole and what does it take to recognize one?” *Proc. IEEE*, vol. 97, pp. 824–848, May 2009.
- [19] R. Tandra, “Spectrum sharing by cognitive radios: Opportunities and challenges,” Ph.D. dissertation, University of California, Berkeley, 2009. [Online]. Available: http://www.eecs.berkeley.edu/~sahai/Theses/Rahul_PhDThesis.pdf

- [20] S. M. Mishra, A. Sahai, and R. W. Brodersen, "Cooperative sensing among cognitive radios," in *IEEE International Conference on Communications*, vol. 4, June 2006, pp. 1658–1663.
- [21] S. Mishra and R. W. Brodersen, "Cognitive Technology for improving Ultra-Wideband (UWB) Coexistence," in *Proc. of the IEEE International Conference on Ultra-Wideband (ICUWB'07)*, Sept. 2007, pp. 253 – 258.
- [22] S. M. Mishra and A. Sahai, "Cooperative sensing for recovering spectral holes," *Submitted to IEEE Transactions on Wireless Communications*, 2010.
- [23] S. M. Mishra, R. Tandra, and A. Sahai, "The case for Multiband Sensing," in *Proc. of the Allerton Conference on Communications, Control and Computing*, September 2007.
- [24] S. M. Mishra, A. Sahai, and R. W. Brodersen, "Multiband Sensing: A theoretical and experimental investigation," *Submitted to IEEE Transactions on Wireless Communications*, 2010.
- [25] S. M. Mishra, R. Tandra, and A. Sahai, "Coexistence with primary users of different scales," in *Proc. of 2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Apr. 2007, pp. 158–167.
- [26] P. Kolodsky, "Interference temperature: a metric for dynamic spectrum utilization," *International Journal of Network Management*, vol. 16, no. 2, pp. 103–113, March 2006.
- [27] J. Snider, "The Art of Spectrum Lobbying: America's \$480 Billion Spectrum Giveaway, How it Happened, and How to Prevent it From Recurring," New America Foundation, Tech. Rep., Aug. 2007. [Online]. Available: http://www.newamerica.net/publications/policy/art_spectrum_lobbying
- [28] F. Communication Commission, "DTV Reception Maps." [Online]. Available: <http://www.fcc.gov/mb/engineering/maps/>
- [29] C. Jackson, D. Robyn, and C. Bazelon, "Comments of Charles L. Jackson, Dorothy Robyn and Coleman Bazelon," The Brattle Group, Tech. Rep., June 2008. [Online]. Available: http://fjallfoss.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6520031074
- [30] A. Sahai, S. M. Mishra, R. Tandra, and K. A. Woyach, "DSP Applications: Cognitive radios for spectrum sharing," *IEEE Signal Processing Magazine*, Jan. 2009.
- [31] B. Scott and M. Calabrese, "Measuring the TV 'White Space' Available for Unlicensed Wireless Broadband," New America Foundation, Tech. Rep., Jan. 2006.

- [32] N. Devroye, P. Mitran, and V. Tarokh, "Achievable rates in cognitive radio channels," *IEEE Trans. Inform. Theory*, vol. 52, pp. 1813–1827, May 2006.
- [33] A. Jovicic and P. Viswanath, "Cognitive radio: An information-theoretic perspective," in *IEEE International Symposium on Information Theory*, Seattle, USA, July 2006, pp. 2413–2417.
- [34] P. Grover and A. Sahai, "On the need for knowledge of the phase in exploiting known primary transmissions," in *Proc. of 2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Dublin, Ireland, Apr. 2007, pp. 462–471.
- [35] R. Tandra and A. Sahai, "Is interference like noise when you know its codebook?" in *IEEE International Symposium on Information Theory*, Seattle, USA, July 2006, pp. 2220–2224.
- [36] "Memorandum Opinion and Order on Reconstruction of the Seventh Report and Order and Eighth report and Order," Federal Communications Commision, Tech. Rep. 08-72, Mar. 2008. [Online]. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-72A1.pdf
- [37] "List of All Class A, LPTV, and TV Translator Stations," Federal Communications Commision, Tech. Rep., 2008. [Online]. Available: <http://www.dtv.gov/MasterLowPowerList.xls>
- [38] "Cisco HWIC-AP WLAN Module Receiver sensitvity specifications," CISCO Systems, Inc., Tech. Rep. [Online]. Available: http://www.cisco.com/en/US/prod/collateral/modules/ps5949/ps6246/product_data_sheet0900aecd8028cc7b.html
- [39] "Rural population density in the United States." [Online]. Available: <http://www.demographia.com/db-usa-staterural.htm>
- [40] U. Census Bureau, "US census 2000 Gazetteer files." [Online]. Available: <http://www.census.gov/geo/www/gazetteer/places2k.html>
- [41] —, "US Census Cartographic Boundary files." [Online]. Available: <http://www.census.gov/geo/www/cob/st2000.html#ascii>
- [42] "Method for point-to-area predictions for terrestrial services in the frequency range 30 mhz to 3 000 mhz," International Telecommunications Commission (ITU), RECOMMENDATION ITU-R P.1546-3, 2007.
- [43] "Establishment of an interference temperature metric to quantify and manage interference and to expand available unlicensed operation in certain fixed, mobile and

- satellite frequency bands,” Federal Communications Commission, Tech. Rep. 03-237, Nov. 2003. [Online]. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-03-289A1.pdf
- [44] F. Database, “Summary of TV transmitters.” [Online]. Available: <http://www.fcc.gov/fcc-bin/tvq?state=&call=&arn=&city=&chan=&cha2=69&serv=&type=3&facid=&list=2&dist=&dlat2=&mlat2=&slat2=&dlon2=&mlon2=&slon2=&size=9>
- [45] Wikipedia, “List of all Music Venues in the United States.” [Online]. Available: http://en.wikipedia.org/wiki/Category:Music_venues_in_the_United_States
- [46] G. Chouinard, “Ieee 802.22-08/0118r0: Use of collaborative sensing to reduce false positive results,” IEEE 802.22 Meeting Documents, May 2008. [Online]. Available: <https://mentor.ieee.org/802.22/dcn/08/22-08-0118-00-0000-collaborative-sensing.ppt>
- [47] S. M. Kay, *Fundamentals of statistical signal processing: Detection theory*. Upper Saddle River, New Jersey: Prentice Hall PTR, 1998, vol. 2.
- [48] R. Tandra, “Fundamental limits on detection in low snr,” Master’s thesis, University of California, Berkeley, 2005. [Online]. Available: http://www.eecs.berkeley.edu/~sahai/Theses/Rahul_MSThesis.pdf/
- [49] G. Hufford, “A characterization of the multipath in the HDTV channel,” *Broadcasting, IEEE Transactions on*, vol. 38, no. 4, pp. 252–255, Dec 1992.
- [50] D. Tse and P. Viswanath, *Fundamentals of Wireless Communications*. Cambridge University Press, 2005.
- [51] J. Hirshleifer and J. G. Riley, *The Analytics of Uncertainty and Information*, ser. Cambridge Surveys of Economic Literature Series, 1992.
- [52] H. Niedercorn, “A Negative Exponential Model of Urban land use densities and its implications for Metropolitan Development,” *Journal of Regional Science*, vol. 11, pp. 317–326, May 1971.
- [53] A. Sendonaris, E. Erkip, and B. Aazhang, “User Cooperation Diversity–Part i: System Description,” *IEEE Trans. Commun.*, vol. 51, no. 11, pp. 1927–1938, November 2003.
- [54] P. K. Varshney, *Distributed Detection and Data Fusion*. Springer-Verlag, 1997.
- [55] G. Ganesan and Y. Li, “Cooperative spectrum sensing in cognitive radio, part I: Two user networks,” *IEEE Trans. Wireless Commun.*, pp. 2204–2213, June 2007.
- [56] —, “Cooperative spectrum sensing in cognitive radio, part II: Multiuser networks,” *IEEE Trans. Wireless Commun.*, pp. 2214–2222, June 2007.

- [57] T. Weiss, J. Hillenbrand, and F. Jondral, "A diversity approach for the detection of idle spectral resources in spectrum pooling systems," in *Proc. of the 48th Int. Scientific Colloquium, Ilmenau, Germany*, 2003.
- [58] E. Vistotsky, S. Kuffner, and R. Peterson, "On Collaborative Detection of TV Transmissions in Support of Dynamic Spectrum Sharing," in *Proc. of 1st IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Nov. 2005, pp. 338–345.
- [59] A. Ghasemi and E. S. Sousa, "Collaborative Spectrum Sensing for Opportunistic Access in Fading Environments," in *Proc. of 1st IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*, Nov. 2005, pp. 131–136.
- [60] —, "Spectrum sensing in cognitive radio networks: requirements, challenges and design trade-offs [cognitive radio communications]," *IEEE Commun. Mag.*, no. 4, pp. 32–39, Apr. 2008.
- [61] J. Ma and Y. G. Li, "Soft combination and detection for cooperative spectrum sensing in cognitive radio networks," in *Proc. of the IEEE Global Telecommunications Conference (GLOBECOM)*, Nov. 2007, pp. 3139–3143.
- [62] J. Unnikrishnan and V. Veeravalli, "Cooperative sensing for primary detection in cognitive radios," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, pp. 18–27, Feb. 2008.
- [63] R. Chen, J.-M. Park, and K. Bian, "Robust Distributed Spectrum Sensing in Cognitive Radio Networks," Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University, Technical Report TR-ECE-06-07, July 2006.
- [64] R. Chen, J.-M. Park, and J. Reed, "Defense against Primary User Emulation Attacks in Cognitive Radio Networks," *IEEE J. Select. Areas Commun.*, vol. 26, pp. 25–37, Jan. 2008.
- [65] P. Kaligineedi, M. Khabbazzian, and V. Bhargava, "Secure Cooperative Sensing Techniques for Cognitive Radio Systems," in *Proc. of the IEEE International Conference on Communications (ICC)*, May 2008, pp. 3406–3410.
- [66] C. Song and Q. Zhang, "Sliding-Window Algorithm for Asynchronous Cooperative Sensing in Wireless Cognitive Networks," in *Proc. of the IEEE International Conference on Communications (ICC)*, May 2008, pp. 3432–3436.

- [67] N. Ahmed, D. Hadaller, and S. Keshav, "GUESS: gossiping updates for efficient spectrum sensing," in *Proceedings of the 1st international workshop on Decentralized resource sharing in mobile computing and networking*, 2006, pp. 12–17.
- [68] G. Ganesan, Y. Li, B. Bing, and S. Li, "Spatiotemporal Sensing in Cognitive Radio Networks," *IEEE J. Select. Areas Commun.*, vol. 26, pp. 5–12, Jan. 2008.
- [69] X. Huang, N. Han, G. Zheng, S. Sohn, and J. Kim, "Weighted-Collaborative Spectrum Sensing in Cognitive Radio," in *Second International Conference on Communications and Networking in China, (CHINACOM)*, Aug. 2007, pp. 110–114.
- [70] W. Zhang, R. Mallik, and K. B. Letaief, "Cooperative Spectrum Sensing Optimization in Cognitive Radio Networks," in *Proc. of the IEEE International Conference on Communications (ICC)*, May 2008, pp. 3411–3415.
- [71] F. Visser, G. Janssen, and P. Pawelczak, "Multinode Spectrum Sensing Based on Energy Detection for Dynamic Spectrum Access," in *Proc. of the IEEE Vehicular Technology Conference (VTC)*, May 2008, pp. 1394 – 1398.
- [72] G. Atia, E. Ermiş, and V. Saligrama, "Robust Energy Efficient Cooperative Spectrum Sensing in Cognitive Radios," in *Proc. of the IEEE/SP 14th Workshop on Statistical Signal Processing, (SSP'07)*, Aug. 2007, pp. 502–506.
- [73] D. Cabric, A. Tkachenko, and R. Brodersen, "Experimental study of Spectrum Sensing based on Energy Detection and Network Cooperation," in *ACM TAPAS*, 2006.
- [74] M. Gudmundson, "Correlation model for Shadow fading in Mobile Radio Systems," *Electronic Letters*, vol. 27, no. 23, pp. 2145–2146, 1991.
- [75] J. N. Tsitsiklis, *Advances in Signal Processing, Vol. 2*. JAI Press, 1993, ch. Decentralized Detection, pp. 297–344.
- [76] H. A. David and H. N. Nagaraja, *Order Statistics, 3rd Edition*. Wiley, 2003.
- [77] "Digital TV Website," 2009. [Online]. Available: <http://www.fcc.gov/mb/engineering/maps/>
- [78] J. Yang, R. W. Brodersen, and D. N. Tse, "Addressing the Dynamic Range Problem in Cognitive Radios," in *Proc. of the IEEE International Conference on Communications (ICC)*, June 2007, pp. 5183 – 5188.
- [79] R. Lpez-Valcarce and G. Vazquez-Vilar, "Spectrum Sensing in Cognitive Radio: Joint Estimation of Noise Variance and Multiple Signal Levels," in *Signal Processing Advances for Wireless Communications, 2007. SSP '07. IEEE International Workshop on*, June 2009.

- [80] Z. Quan, S. Cui, A. H. Sayed, and H. V. Poor, "Optimal Multiband Joint Detection for Spectrum Sensing in Cognitive Radio Networks," *IEEE Trans. Signal Processing*, vol. 57, no. 3, pp. 1128–1140, Mar. 2009.
- [81] G. Atia, E. Ermiş, and V. Saligrama, "Robust Energy Efficient Cooperative Spectrum Sensing in Cognitive Radios," in *Statistical Signal Processing, 2007. SSP '07. IEEE/SP 14th Workshop on*, Madison, WI, USA,, Aug. 26–29, 2007, pp. 502–506.
- [82] Y. Zeng and Y.-C. Liang, "Covariance based signal detections for cognitive radio," in *New Frontiers in Dynamic Spectrum Access Networks, 2007. DySPAN 2007. 2nd IEEE International Symposium on*, Dublin,, Apr. 2007, pp. 202–207.
- [83] S. Shellhammer, S. Shankar, R. Tandra, and J. Tomcik, "Performance of Power Detector Sensors of DTV Signals in IEEE 802.22 WRANs," in *Proc. of the First ACM International Workshop on Technology and Policy for Accessing Spectrum (TAPAS)*, Aug. 2006. [Online]. Available: <http://doi.acm.org/10.1145/1234388.1234392>
- [84] C. Cordeiro, K. Challapali, D. Birru, and S. Shankar, "IEEE 802.22: the first worldwide wireless standard based on cognitive radios," in *Proc. of the first IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks, DySPAN*, 2005.
- [85] R. Tandra and A. Sahai, "SNR walls for signal detection," *IEEE Journal on Selected Topics in Signal Processing*, vol. 2, pp. 4 – 17, Feb. 2008.
- [86] C. Clayton, , M. Kenkel, and Y. Tang, "Wireless Microphone Signal Simulation Method," IEEE 802.22 Meeting Documents, Mar. 2007. [Online]. Available: <https://mentor.ieee.org/802.22/dcn/07/22-07-0325-00-0000-i2r-simulation-for-wireless-microphone-detection.ppt>
- [87] Y. Zeng and Y.-C. Liang, "Covariance based sensing algorithms for detection of DTV and wireless microphone signals," IEEE 802.22 Meeting Documents, Nov. 2006. [Online]. Available: http://grouper.ieee.org/groups/802/22/Meeting_documents/2006_Nov/22-06-0187-01-0000_I2R-sensing-2.ppt
- [88] M. Ghosh, V. Gadam, and G. Turkenish, "DTV Signal Sensing Using Pilot Detection," IEEE 802.22 Meeting Documents, Mar. 2007. [Online]. Available: <https://mentor.ieee.org/802.22/dcn/07/22-07-0125-00-0000-philips-pilot-detection-based-sensing.ppt>
- [89] M. McHenry, "The Probe Spectrum Access Method," in *Proc. of the first IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks, DySPAN*, 2005.

- [90] J. F. Chamberland and V. V. Veeravalli, "How Dense Should a Sensor Network Be for Detection With Correlated Observations?" *IEEE Transactions on Information Theory*, vol. 51, no. 11, pp. 5099–5106, November 2006.
- [91] N. Hoven and A. Sahai, "Power Scaling in Cognitive Radios," in *WirelessComm, 2005*, 2005.
- [92] A. Sahai, R. Tandra, S. M. Mishra, and N. Hoven, "Fundamental Design Tradeoffs in Cognitive Radio Systems," in *ACM TAPAS*, 2006.
- [93] C. R. Stevenson, C. Cordeiro, E. Sofer, and G. Chouinard, "RAN requirements," IEEE 802.22 Working Group, Tech. Rep. 22-05-0007-47-0000, January 2006.
- [94] "The Global Land One-km Base Elevation (GLOBE) project." [Online]. Available: <http://www.ngdc.noaa.gov/mgg/topo/globe.html>
- [95] "Longley-rice methodology for evaluating tv coverage and interference," Federal Communications Commission, OET BULLETIN 69, Feb. 2004.
- [96] "Xilinx System Generator," 2009. [Online]. Available: <http://www.xilinx.com/tools/sysgen.htm>

Appendix A

Example for Evaluating Whitespace

In this appendix we shall use the running example of the KCNS transmitter in San Francisco to illustrate the calculations in Chapter 2. The KCNS digital transmitter occupies channel 39 and is housed on Sutro tower at a height above sea level of 459m and a transmit power of 1000kW. To determine the propagation characteristics of this tower we need to calculate the effective height of the tower — called HAAT for height above average terrain. To calculate the HAAT, the elevation of 50 random points are taken at distances between 3km and 16km around the tower at 8 (or more) evenly spaced radials from the transmitter site (These calculations use the Globe Terrain database available from the National Geophysical Data Center [94]). The elevation points along each radial are averaged, then the radial averages are averaged to calculate the average height of the terrain. Subtracting this value from the height of the tower above sea level gives the HAAT of the tower. For the KCNS transmitter this value turns out to be 430m. This is mainly because San Francisco is at an average elevation of ~ 16 meters above sea level.

Before we continue, we should also determine the operational SINR of the KCNS transmitter. According to [12] the Desired Signal to Interference ratio (SIR) ratio at the Grade-B contour is 23dB. With noise power over a 6MHz band at -106dBm we can calculate the required SINR at the Grade-B contour. These values are shown in Table A.1. For channels 39, the required SINR is 15.46dB.

A.1 Calculating the pollution radius

The maximum interference that a secondary can tolerate (γ) is set as a value above the noise level. To calculate the pollution radius, we need to find the distance where the signal power drops to $\gamma + N_0$ dBm for 50% of the locations, 50% of the time (the F(50, 50) curve from the ITU specifications). For the KCNS tower and a $\gamma = 5$ dB, the pollution radius is 220km.

Parameter	Description	Value
Δ	Minimal Operational SINR	LVHF: 19.69dB ULHF: 19.13dB ULHF: 17.33dB - 14.67dB
$E_{r_p}(dBu)$	Electric field at the protected radius (r_p)	
$E_{r_{nl}}(dBu)$	Electric field at the noise limited radius (r_{nl})	
$E_{r_b}(dBu)$	Electric field at the Grade B radius (r_b)	See Table A.2
ψ	Margin eroded	1dB
$I_{r_p}(dBm)$	Interference at the protected radius	
$N_0(dBm)$	Noise in a 6MHz band	-106.22dBm

Table A.1: Parameters to be used in white space calculations.

A.2 Calculating the protection radius r_p

The FCC calculations give rise to a range of erosion margin values (See Figure 2.4(a)) the median of which is 1dB. We shall use this as the target erosion margin ($\psi = 1dB$). Hence r_p is the distance at which the signal strength is greater than $N_0 + \Delta + \psi = -89.75dBm$ for 50% of the locations, 90% of the time. For channel 39, this translates to a electric field of 41.16dBu ($E_{r_p}(dBu)$). $E_{r_p}(dBu)$ was used to calculate the r_p using the following procedure.

- The Effective Radiated Power is converted to Effective Isotropic radiated Power ($EIRP(dBm) = ERP(dBm) + 2.15dB$). For the KCNS transmitter the EIRP (dBm) is 92.5dBm.
- The Electric field at a distance of 1m from the transmitter is calculated as $E_{1m}(dBu) = 104.8 + EIRP(dBm)$, which is 197.3dBu for the KCNS transmitter.
- The required path loss is calculated as $RPL(dB) = E_{1m} - E_{r_p}$ which turns out to be 156.3dB for KCNS.
- The ITU-R recommendations are used to determine the maximum distance (beyond 1m) at which the path loss is less than (or equal to) RPL for 50% of the locations, 90% of the time. (The ITU-R recommendations provide a mechanism to extrapolate the tables for different distances, heights and frequency [42]).

For KCNS, the r_p calculated using this method is 129.3km.

We further assumed that all television signals are ATSC signals *i.e.* the assumption is made that low power and Class A transmitters would switch to ATSC signals with the same power.

<i>Channels</i>	<i>Required Field Strength (dBu)</i>
2 – 6	23
7 – 13	36
14 – 69	41

Table A.2: ATSC (Digital) Field Strength dBu required at the Grade B contour.

<i>Channels</i>	<i>Formula to convert from dBu to dBm</i>
2 – 6	$P(\text{dBm}) = E(\text{dBu}) - 111.8$
7 – 13	$P(\text{dBm}) = E(\text{dBu}) - 120.8$
14 – 69	$P(\text{dBm}) = E(\text{dBu}) - 130.8 + 20 \log_{10}(\frac{615}{f_h + f_l})$

Table A.3: dBm to dBu conversion values for various frequencies. f_l and f_h are the channel's lower and higher frequency limits (in MHz) respectively. [95]

A.3 Calculating the no-talk radius (r_n)

To calculate r_n we first calculated the distance beyond the protected radius where the secondary can transmit ($r_n - r_p$) *i.e.* we need to determine the distance r_n such that a transmission from r_n results in a signal level at r_p of I_{r_p} (See Equation 2.4). After we converted I_{r_p} to an electric field, we used the procedure outlined in Appendix A.2 to determine the value of $r_n - r_p$ with the exception that we used the $F(50, 10)$ propagation curves for predicting the distance. This was to ensure that transmissions from a secondary just outside r_n can cause harmful interference only 10% of the time.

For the KCNS station, I_{r_p} was -112.1dBm which translated into a electric field strength of 18.83dBu. This required the secondary transmitter to be 20.57km outside KCNS's protected radius.

A.4 Calculating the no-talk radius (r_n) for adjacent channels

For adjacent channels, the additional distance beyond the protected radius that we need to budget reduces since the TV receiver can tolerate higher interference from adjacent channels. The FCC specifies that the adjacent channel interference can be ~ 27 dB higher than the desired signal. Using this value, $r_n - r_p$ was calculated for adjacent channels as per the procedure outlined in Appendix A.3 except that I_{r_p} can be as high as -62.1dBm. To create this interference level, the secondary would have to be 1.1km from r_p .

A.5 Calculating r_n using the FCC method

The FCC assumes that the protection radius is the Grade B contour (r_b). The required field strength ($E_{r_b}(dBu)$) at the Grade B contour is defined by the FCC for ATSC signals [12] as shown in Table A.2. Hence for the KCNS transmitter we need to determine the distance at which the received signal is above 41dBu for 50% of the locations, 90% of the time. This distance turns out to be 129.8km.

To calculate $r_n - r_b$ we found the distance beyond r_b such that a transmission from r_n results in a signal level at the Grade-B contour of $E_{r_b} - 23dBu$. Since E_{r_b} is only a function of the frequency (and not the transmit power/height of the primary transmitter) there is a single $r_n - r_p$ value for a given channel. The procedure outlined in Appendix A.2 can be used to determine this value of $r_n - r_b$ with the exception that we use the the $F(50, 10)$ propagation curves for predicting the distance. For the KCNS transmitter, this distance turns out to be 21.42km.

An alternative approach adopted by the FCC also specifies the value of $r_n - r_p$ to be used for all channels (See Section 15.712 [12]).

A.6 Calculating the radius for single-threshold sensing rules

A sensing based approach does not distinguish between different towers. The FCC's -114dBm rule imposes a 20dB margin on the nominal signal level at r_n . To determine the area lost due to such a rule, we first converted the -114dBm power to an equivalent dBu value using the equations in Table A.3. Using the procedure outlined in Appendix A.2 we calculated the equivalent distance. As discussed in the main text we use the $F(50, 50)$ curves for predicting distances. For KCNS the -114dBm rule would wipe out all area to a distance of 255km.

Next, we derived an equivalent rule for adjacent channels. In Appendix A.4 we had determined r_n for the adjacent channel. At this new r_n we calculated the average signal level and added a budget of 20dB to obtain a -110dBm sensing rule for adjacent channels. With such a rule the area around the KCNS station that needs to be excluded for usage is 223km. It should be noted that the FCC specifies a single -114dBm rule for all channels (co-channel and adjacent channel sensing).

A.7 Incorporating multipath into the path loss model

The effect of multipath fading is incorporated in the calculations of r_p , $r_n - r_p$ and the distance corresponding to the -114dBm rule. To remain conservative we neglected multipath for calculating r_p (with multipath, the $F(50, 90)$ point is lower and this leads to the under-

estimation of r_p). For determining $r_n - r_p$ we assumed Rayleigh multipath – this increased the value of $r_n - r_p$ for the KCNS tower from 20.57km to 24km. For the -114dBm rule the opposite was true; the distance excluded shrinks from 255km to 240km.

The generic method to obtain the $F(50, 90)$ tables for a channel that includes Ricean (Note that Rayleigh multipath is a special case of Ricean multipath) multipath is as follows:

Let H be the Ricean distributed multipath random variable with unit total power ($\sigma^2 + A^2 = 1$), then the density of H is,

$$f_H(h) = \frac{h}{\sigma^2} \exp\left(-\left(\frac{r^2}{\sigma^2} + K\right)\right) I_0\left(\frac{\sqrt{2K}h}{\sigma}\right) \quad (\text{A.1})$$

where K is the Ricean factor which measures the ratio of the energy in the line-of-sight path to the multipath variance ($\frac{A^2}{2\sigma^2}$). If the average received power is G^2 , then the instantaneous received power is given by $G^2 h^2$ or:

$$P^I = P^A + 20\log_{10}(h) \quad (\text{A.2})$$

where P^I is the instantaneous received power ($G^2 h^2$) and P^A is the average received power (G^2). Since the conversion from dBm to dBu only involves a constant addition/subtraction at all frequencies, converting this to electric field (in dBu) we get:

$$E^I = E^A + 20\log_{10}(h) \quad (\text{A.3})$$

From the ITU we have the $F(50, q)$ ($q = 50, 10$), tables for 100MHz, 600MHz and 2000MHz for various transmitter heights and distances i.e. for a given frequency, transmitter height and distance, these tables provide the electric field e such $\mathcal{P}(E^A > e) = \frac{q}{100}$ for 50 percent of the time.

Now to generate the appropriate tables all we need to find is the new values of the electric field e' such that $\mathcal{P}(E^I > e') = \frac{q}{100}$. Once we preprocessed these files we can use the new tables in all the subsequent calculations.

Appendix B

Testbed

In this Appendix we describe the equipment and methodology to collect DTV measurements at various locations in Berkeley, CA. We wanted to build a wideband detector that could collect DTV signals over a frequency range of 500-700MHz. Furthermore, we needed this equipment to be mobile so that it could be driven around in a car and transported easily.

B.1 Overall Architecture

Figure B.1 shows the overall architecture of the DTV mobile experimental unit. Since we were interested in the 500-700MHz band, we chose a purely baseband architecture (no RF down-conversion). The analog gain stages provide a variable gain from 0 to 88dB. The analog filtering is implemented to attenuate cellular bands (850-900MHz) and the FM/VHF bands. The analog signals are converted to digital signals using the Atmel/ev2 (AT84AD001B) ADC which is capable of a conversion rate up to 2Gsps (Giga-samples per second) in interleaved mode. The raw ADC resolution is 8bits with a ENOB (effective number of bits) of 7bits. The resulting 42dB of dynamic range is low for sensing weak DTV signals using a wideband receiver but serves the purpose of investigating signals from Mt San Bruno and Mt Sutro. According to the FCC, the received signals from these towers in Berkeley should be around 30dBm [77].

For our experimentation, the ADC runs at 900MHz in interleaved mode yielding a sampling rate of 1.8Gsps. These samples are processed in the Xilinx Virtex-II Pro FPGA and stored in the external 2Mbyte SDRAM. The Power PC on the Xilinx is running Linux and a IP stack which can be used to transfer the captured data to the Laptop (via a UDP connection) for later processing. The 12V battery enables the unit to be mobile and can be used to make measurements while driving.

The output of the ADC is digitally de-modulated to 450MHz and down converted to a sampling rate of 900MHz. This operation is shown in Figure B.2. The reason for this is the

rate mismatch between the ADC and SDRAM rates. The output of the ADC are 8 samples (64bits) at a clock rate of 225MHz. The external SDRAM also runs at 225MHz but has a 32bit interface. Hence the need to down convert to a sample rate of 900MHz. The samples are demodulated to 450MHz using a very simple digital demodulator (this step involved multiplication by the vector $[1 \ 0 \ -1 \ 0 \ 1 \ 0 \ -1 \ 0]$). These samples are filtered by a 32-tap low pass filter (implemented as a 8 channel polyphase filter). The filter has a passband of 0-300MHz. The output of the filter is downsampled by 2 and copied into a ping-pong buffer which feeds the SDRAM. Once the SDRAM is full, the data (it can store data contiguously data for just over 2ms) is moved to the laptop.

The FPGA was programmed using Matlab/Simulink from Mathworks coupled with the Xilinx system generator [96]. The tool chain is augmented with BWRC developed automation tools for mapping high level block diagrams and state machine specifications to FPGA configurations. A set of parameterized library blocks have been developed for communications, control operators, memory interfaces and I/O modules. This flow enabled fast prototyping of the design.

B.2 Antenna

We used a simple indoor 7dBi antenna (the Silver Sensor) from Phillips/Zenith. The measurements were verified using a similar antenna from Terk (HDTV_i). Both these antennas have a relatively flat frequency response in the 500-700MHz and present the right compromise between directionality and maneuverability. Figure B.3 shows the result of making multiple measurements at a given location by rotating the antenna around. Measurements made on Prof. Sahai's deck and the BWRC classroom near the window reveal a total variation of 8dB. This is because in both locations there is a single dominant signal direction (west at Prof. Sahai's deck and facing the window in the BWRC classroom). The measurements that were made in the BWRC laboratory have relatively little variation. The BWRC laboratory has no windows and hence there is dominant signal direction.

B.3 Filtering and Gain

The analog filtering was implemented using discrete components from Mini circuits. The final impact of this filtering is shown in Figure B.4. At 900MHz, the unwanted cellular signals are attenuate to below the dynamic range of the ADC. The digital filter is used to remove the high frequency components after digital de-modulation. Figure B.5 shows the frequency response of the digital filter. The passband is not completely flat – this is to compensate for the attenuation introduced by the ADC at high frequencies. The overall passband response (ADC, analog filtering, gain stages, digital filtering) is shown in Fig B.5. The maximum variation in the magnitude response in the range 500-700MHz is less than 1dB.




Figure B.1: Overall architecture of the mobile baseband TV sensor using a 2Gps ADC.




Figure B.2: Block diagram of the FPGA design. The down conversion operation is needed to match the rates of the ADC and the SDRAM.

The gain stages provide a gain of 0, 22, 44, 66 and 88dB. The calculated overall noise figure of the 4 gain stages is 2.72dB while the measured Noise Figure is 10dB.

The overall linearity of operation is limited by the ENOB of the ADC. This can be seen by inputting a sine wave at different frequencies. The linear operation is limited to around 40dB input range (See Figure B.6).

B.4 Locations and Measurement methodology

We wanted a good sampling of locations in Berkeley to conduct experiments. Finally we chose the following locations:

- BWRC laboratory, BWRC classroom, BWRC library, BWRC cubicles 87 and 90 to provide a rich set of Non Line-of Sight (NLOS), indoor locations.
- BWRC basement to provide a highly shadowed indoor environment.
- Prof. Sahai's deck to provide LOS, outdoor measurements. From the deck we have line-of-sight to Mt Sutro.
- Drive measurements along Milvia, California and Oxford streets to provide NLOS, outdoor measurements.




Figure B.3: Measurements using the 7dBi indoor antenna from Phillips. Measurements made at Prof. Sahai's deck and the BWRC classroom reveal a total variation of 8dB since there is a single dominant signal direction. The measurements made in the BWRC lab have relatively less variation. The BWRC lab has no windows and hence there is no dominant signal direction.






Figure B.4: The gain at each frequency from analog filtering.

Figure B.7 shows the areas within the BWRC which served as the measurement locations. At each location a specific corner was designated as the origin and X/Y-axes were established relative to the walls. At each location multiple measurements were made and indexed by their X and Y co-ordinates relative to the origin at that location. At each measurement point 4 captures were performed with the antenna facing the magnetic North, the magnetic South, the magnetic West and the magnetic East (A standard compass was used to align the antenna to each direction). The measurements were stored in a file named as $\langle locId \rangle - \langle x \rangle - \langle y \rangle - \langle dir \rangle .txt$, where $\langle locId \rangle$ is the location index (for example, BWRC lab has the index 'BLB'), x and y are the distances from the location's origin (measured in inches) to the measurement point and $\langle dir \rangle$ is the direction ('W', 'E', 'N' and 'S') at which the antenna was pointed.

The outdoor measurement locations are shown in Figure B.8. The antenna and a GPS unit were mounted on top of a car. as shown in Figure B.9. The antenna was pointing west throughout the drive experiments. The average spacing between measurements was 10m.



(a)



(b)

Figure B.5: (a) The frequency response of the 32-tap digital filter. (b) The shape of the passband is used to compensate for the frequency response of the ADC.




Figure B.6: The input-output transfer curve of the system for a gain setting of 44dB. The linear operation occurs for a input range of 40dB.




Figure B.7: The BWRC seating map with the various locations marked out. The X and Y axis is used to index and store measurements at each location.




Figure B.8: A map of the outdoor NLOS locations: Milvia, California and Oxford streets. courtesy: Google Maps and <http://www.gpsvisualizer.com>




Figure B.9: The GPS and antenna mounted on atop a car.