Addressing Heterogeneous Wireless Interference in Spectrum Bands with Unlicensed Access

George Nychis

Electrical and Computer Engineering

Carnegie Mellon University

2013

A dissertation submitted for the degree of

Doctor of Philosophy (PhD)

Abstract

In the past two decades, we have seen an unprecedented rise in unlicensed wireless devices and applications of wireless technology. To meet various application constraints, we continually customize the radios and their protocols to the application domain which has led to significant diversity in spectrum use. Unfortunately, this diversity (coupled with increased demand) complicates spectrum sharing, exacerbates interference, and as a result: reduces network performance and capacity. The introduction of the "white spaces" to address the increasing demand for spectrum further complicates this problem. Now, unlicensed devices must also adhere to strict regulations against interference on *spectrum primaries*, i.e., licensed devices in these bands.

In an attempt to address diversity between devices in unlicensed spectrum, our community has focused on developing coexistence techniques, i.e., modifications to the radios and protocols to reduce interference between two specific technologies when operating in the same band. This general approach, however, requires N^2 solutions, that are rarely deployed and often short-lived due to rapid changes in unlicensed technologies. To address diversity with spectrum primaries in the white spaces, spectrum management approaches are being enforced that are effective but extremely spectrum inefficient, threatening the very the goal of the white spaces: additional spectrum.

In this dissertation, we explore alternative approaches. We argue that spectrum management can be a better long-term solution to address diversity between technologies sharing unlicensed spectrum, whereas coexistence techniques can provide spectrum-efficient solutions to protect spectrum primaries. However, one cannot simply apply white space spectrum management techniques to unlicensed spectrum due to a lack of information and algorithms that support the high degree of unlicensed technologies. Similarly, traditional coexistence techniques between unlicensed devices do not meet strict zero-interference policies to be applied to spectrum primaries.

To overcome these challenges and provide more efficient and long term solutions to interference between heterogeneous technologies, we make three key contributions in this dissertation. First, we introduce a novel system on the smartphone which allows it to collect the necessary information about heterogeneous wireless technologies towards proper spectrum management between unlicensed devices. Using this information, we then introduce an efficient and effective spectrum assignment model and algorithm, capable of supporting various unlicensed technologies (even as they evolve). Finally, we switch our focus to the white spaces and show how spectrum management is spectrum inefficient to be effective, and introduce a novel coexistence protocol between unlicensed devices and spectrum primaries that allows spectrum-efficient interference-free coexistence.

To the two amazing people that I have lost in my life since starting the Ph.D. program: my “little” γιαγιὰ, and ϑεία. Γιαγιά, I used to joke that I was dedicating my Ph.D. to you. I'm not sure that I can dedicate my Ph.D., but I can definitely dedicate my thesis...

Αιὼνια η μνὴμη της.

Acknowledgements

It is extremely difficult (if not impossible) to acknowledge all of the help, support, and advice that I have received from such a large number of people throughout graduate school and the Ph.D. program. In a seemingly short amount of space (compared to the vast amount of help I have received), I will do my best to not leave anyone out.

First and foremost, I extremely grateful for having the amazing pair of advisors that I did: Peter Steenkiste and Srinivasan Seshan. I do not only owe them my sincere gratitude for guiding me throughout the Ph.D. program, but also for inspiring me to be here in the first place. It was their enthusiasm, insight, and general love for our field and research that led me to apply to the Ph.D. program. Over the past six years that I have worked with them, I have learned a tremendous amount about our field from them, including how to conduct research, approach hard problems, and focus on key issues. I feel extremely privileged to have had Srini and Peter as advisors, and I hope to pass on what I have learned from them.

I would also like to thank Raj Rajkumar and Ranveer Chandra for being members on my thesis committee. Their feedback from the introduction of the thesis proposal, to the conclusions section of this dissertation has been extremely valuable. In particular, I owe Ranveer Chandra many thanks for his guidance and support throughout my entire career. His trust and belief in me as a researcher helped me push myself in many ways to take on hard problems, and push every project we worked on together further. He was an excellent mentor for two of my internships at Microsoft Research, and I am extremely thankful for him taking me on as an intern (with Thomas Moscibroda) and guiding me well beyond the scope of my internships. Your hard work has always inspired me, Ranveer.

There are many other colleagues that I would like to thank for their support, collaboration, and guidance. Thomas Moscibroda, Onur Mutlu, Hui Zhang, Victor Bahl, and David Andersen were all amazing collaborators to work with on projects outside the scope of my thesis. Although a lot of time has passed since I first entered CMU, Vyas Sekar and Michael Bigrigg played very special roles in the early days of graduate career, mentoring me, and teaching me a lot about how to write research papers and conduct research. Additionally, I owe thanks to many other wonderful colleagues: Dongsu Han, Dan Halperin, Rohan Murty, Chris Fallin, Elie Krevat, Dan Wendlandt, Xi Liu, Eric Rozner, Colin Dixon, Jeffery Pang, Fahad Dogar, Kaushik Lakshminarayanan, Athula Balachandran, Bin Fan, David Naylor, Matthew Mukerjee, Eduardo Cuervo, Souvik Sen, and

Justin Manweiler. Special thanks to Brian Padalino, Matt Ettus, Eric Blossom, and many others that spent hours and hours of their time helping me grasp (and debug) the many complexities of software-defined radios. Many thanks to Dena Haritos Tsamitis and her faith in me also, helping me to be where I am today.

Finally, I have the most *amazing* family and set of friends that helped me simply make it through the often grueling process of the Ph.D. program. Thank you for all of your emotional support. First, my family: Vicki Nychis (mom), Peter Nychis (dad), Stacey Nychis (sister), and Nicole and Deanna Yialouris (cousins). Second, my (often crazy) friends: Brock Harrison, Shawn Robinson, Evan Easton, Adele Ferris, JJ Stamatelos, Thomas Joseph Slosky III, Carine Bou-Abboud, Hilary Michel, Elie Krevat, Theodore Andrianos, Lauren Brilli, Melissa Jankowski, Mary-Kaylor Hanger, Michael Gregory, Lauren Easton, Dean Leventopoulos, Mary Doreza, Alexis Ellis, Andie Kyros, CB, Jon New, Elisa Bosilovik, Korina Loumidi, Iryna Zenyuk, Siddharth Komini Babu, Alex Styler, Amy Nibert, Udhyakumar Nallasamy, Bill Ray, Chad Michael, Schlou, and ET.

My graduate work was supported by the National Science Foundation (CNS-0855137 and CNS-0520192), the Nick Vlahakis Fellowship, and through internships at Microsoft Research.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In the past two decades, we have seen a transition in the use of the unlicensed wireless spectrum from being underutilized by a few dominant technologies technologies (e.g., Wi-Fi), to being over-utilized by many competing and heterogeneous technologies (now *many* protocols and devices using the same spectrum). Behind this trend are several key factors that are important to this dissertation. First, there is an increasing number of unlicensed wireless devices and applications for wireless technology. This trend continues to strain spectrum availability. As put by the wireless association, the CTIA:

> *"... devices are then developed to take advantage of next generation networks, application and content developers then create new content to take advantage of new capabilities, and ultimately, consumers demand more. It's a cycle that never ends as long as spectrum is available."* [1]

Second, to meet the specific needs and constraints of each application domain for wireless technology, we are catering the radios and their protocols to the application domain. This is the primary driver of heterogeneous technologies where varying application constraints (e.g., in latency, bandwidth, or reliability) has led to the development of a number of wireless protocols. ZigBee, 802.11, Bluetooth, DECT, and Z-Wave are just a few examples wireless protocols used in spectrum bands with unlicensed access, that are designed to address the constraints of different types of applications [2,3,4,5,6]. Further meshing these unlicensed devices and protocols with others in the spectrum, the FCC made a historic ruling in 2008 that allows "white space" spectrum access in licensed (but idle) parts of the spectrum where these secondary (unlicensed) devices operate around FM microphones, TVBD broadcasts, and potentially many other licensed devices and protocols in the future.

This increasing demand on the spectrum from such a diverse set of devices and wireless technologies (both unlicensed and licensed), has created ***2 first-order problems*** in current and future access of the wireless spectrum:

1. *General Heterogeneous Interference:* The sharing of any broadcast medium is based on the ability for the endpoints to coordinate and detect when the medium is idle. In historically homogeneous environments, this was guaranteed by a common or dominant PHY and MAC. However, this most basic principal of the wireless medium breaks down in current environments with a high degree of heterogeneous technologies, where diversity in the PHY and MAC affects devices' ability to coordinate and properly sense the medium. As a result, significant amounts of cross-technology interference have been observed [7,8]. This general problem of heterogeneity is, in particular, a problem between unlicensed devices whose protocols and standards evolve quickly. Being able to overcome cross-technology interference with a long-term and robust solution, despite this quick evolution, is a critical challenge.

2. *Efficient Interference Avoidance with Spectrum Primaries:* Spectrum bands with white space access create new challenges in interference avoidance. In particular, with licensed users (i.e., the spectrum primaries) that are also diverse in their technologies and have strict rules prohibiting *any* interference generated on them. Given that the goal of the white spaces is to provide additional spectrum, the mechanisms to avoid interference with the primaries should be spectrum-efficient. That is, they should be designed to allow the secondary (i.e., unlicensed) devices to use as much spectrum around the primaries as possible without interfering with them.

In this dissertation, we address these two first-order challenges. By doing so, we show the potential to significantly reduce cross-technology interference and improve spectrum efficiency. At the center of this dissertation, we argue that a long term solution to the general problem of heterogeneity in unlicensed bands is not through the development coexistence techniques. These techniques are often short-lived and rarely deployed due to rapid changes in the technologies of unlicensed devices. Instead, we argue that better heterogeneous monitoring and spectrum management are the keys to addressing the problem of general heterogeneity. To the contrary, we argue that coexistence techniques are better suited for dealing with the coexistence between primaries and secondaries in the white spaces. Licensed devices rarely evolve (if ever), making their signals predictable and coexistence techniques longer-lived and more highly spectrum efficient for the white spaces. Towards these goals, this dissertation provides two case studies that provide positive results which support these claims: one in the TV white spaces where a coexistence technique is shown to provide a more spectrum-efficient solution than current spectrum management based regulations, and one in the unlicensed bands where we show spectrum management can provide an effective single solution to general heterogeneity in home environments (where diversity is significantly high).

1.1 Spectrum Bands, Technologies, and Access Schemes

Before we discuss the problem this dissertation addresses in more detail, it is important to first have an understanding of the use of unlicensed and white space spectrum bands, as well

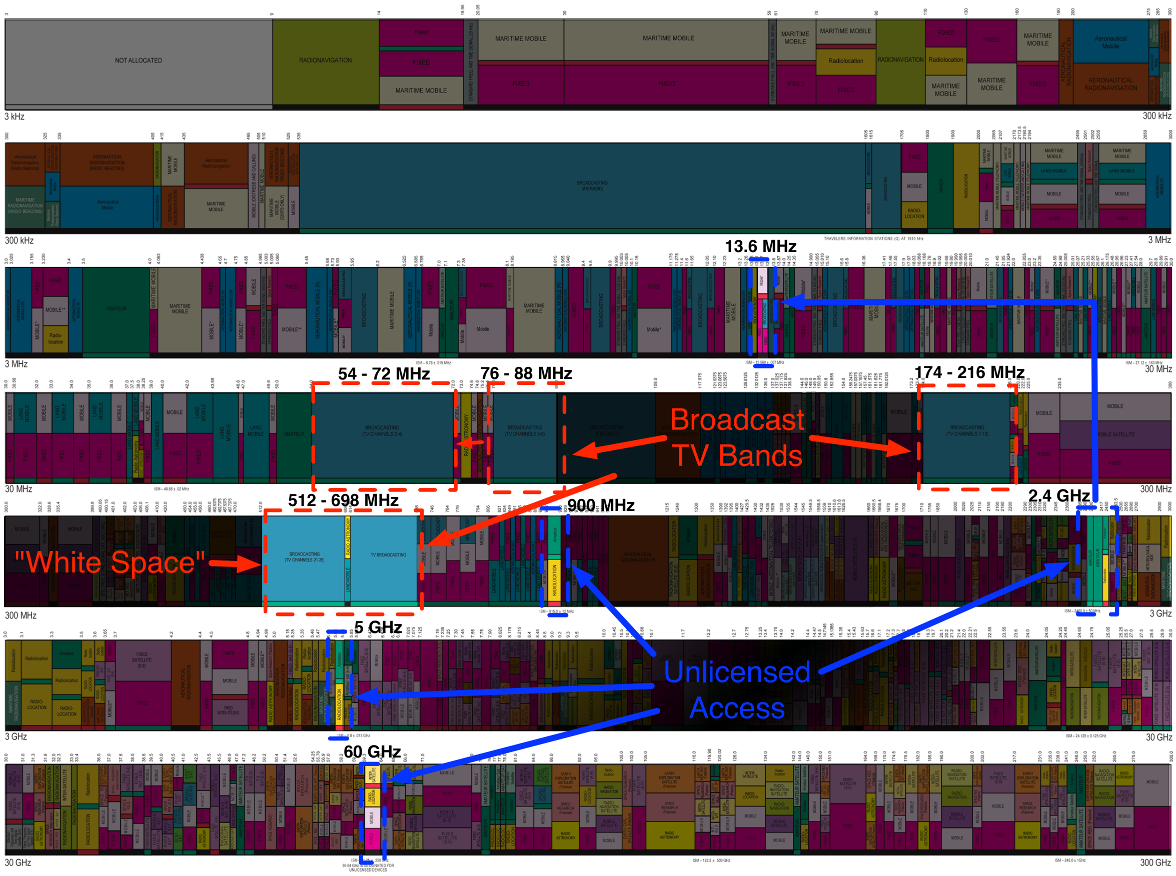

Figure 1.1: A map of the wireless spectrum from 3KHz to 300GHz. The map shows which bands are licensed to which entities, and we highlight examples of spectrum bands with unlicensed access, and white space access.

as the regulations, trends, and access schemes that lead to the 2 first-order problems we have described.

In the remainder of this section, we first describe the two primary types of spectrum bands that are the focus of this dissertation: 1) Spectrum bands with unlicensed access, and 2) White space spectrum bands, and in particular: TV white spaces. Then, we describe the technologies and access schemes used by devices that access the spectrum. Understanding the diversity between these technologies and access schemes is critical to understanding the interference created as a result, the main problem this dissertation.

1.1.1 Spectrum Bands with Unlicensed Access

There are many different fragments of spectrum that are permitted access from unlicensed devices, as shown with a few examples in Figure 1.1. These bands, and in particular the

ISM bands (e.g., at 900MHz, 2.4GHz, and 5GHz), contains the majority of wireless network traffic from unlicensed technologies today [9]. Wi-Fi, ZigBee, Bluetooth, NFC, and many other common technologies utilize these bands due to their availability of spectrum, lax set of rules, and relatively good propagation qualities (e.g., at 2.4GHz and 5GHz).

Spectrum Details: There are several bands with unlicensed access highlighted in Figure 1.1. The lower frequencies around 13.6 MHz are commonly used by RFID and NFC, whereas the higher frequencies support the majority of the unlicensed networking equipment. Wi-Fi, ZigBee, Bluetooth and many others operate in 2.4 GHz, in addition to dual-band Wi-Fi devices operating in 5 GHz where they can avoid high levels of interference in 2.4 GHz due to crowding and heterogeneity. The 60 GHz "extremely high frequency" bands are typically used for high-bandwidth applications where the radios involved can be assumed to have line-of-sight, since signals at this frequency typically will not penetrate walls. Current applications for 60 GHz frequencies, for example, have been to augment data center racks with high throughput wireless links [10,11] and wireless HD/HDMI [12].

In this dissertation, the focus of our work is in the 2.4 GHz, and 5 GHz ISM bands where heterogeneity and interference are high. 60 GHz currently has few deployed applications, and receives little interference due to signals in this band being unable to pass through walls and being predominately line of sight. The 13.6 MHz ISM band is low in bandwidth and heterogeneity, dominated by RFID and NFC which are compatible technologies.

Rules and Regulations: Within the United States, unlicensed devices operate under the FCC Part 15 rules of Title 47 [13]. Unlike licensed transmitters where the FCC guarantees freedom from harmful interference, unlicensed devices are granted no protective rights. In particular, unlicensed devices operating under Part 15 rules cannot cause interference on any licensed devices within range, and must accept all interference from other sources (licensed and unlicensed). Additional cardinal rules include: a maximum transmission output power of 1 watt (30 dBm) at the antenna, ceasing of operation if notified by the FCC of interference, and all equipment must be certified and shown to comply with FCC standards before it is marketed [14]. Note that, importantly, there are no rules that regulate how, when, and for how long the radios can access the spectrum i.e., there are no general rules that regulate the protocols used by the radios.

Trends: There are several important trends in the use of spectrum bands with unlicensed access that are important to this thesis. Historically, these bands have served the majority of unlicensed technologies and their demand. Today, nearly every wireless device that consumers own has at least 1 radio that operates in the 2.4GHz and 5GHz bands. Cell phones, laptops, gaming controllers, cordless phones, and wireless headsets are just a few examples. Given the increasing number of applications and demand on the spectrum, many of these heterogeneous (and incompatible) technologies have been forced to share channels in these bands, exacerbating interference [7,8,15,16]. Addressing these trends and concerns has become of critical importance moving forward in the wireless spectrum since it is likely demand and density will continue increasing. Mitigating interference between these many unlicensed heterogeneous technologies is a primary focus of this dissertation.

1.1.2 White Space Spectrum Bands

Addressing the "impending spectrum crisis" and the growing demand for spectrum (e.g., in the ISM bands), has become a critical challenge moving forward, as we have discussed. To keep up with this demand, many have petitioned the FCC to simply "unlicense" additional spectrum, e.g., the CTIA in petitioning the FCC to free up an additional 800 MHz over the next 6 years to address these concerns [1]. While one might think that there should be additional spectrum in the 300 GHz pictured in Figure 1.1, all of the spectrum shown in the figure has already been licensed.

Understanding the need for additional spectrum, yet a lack of "free spectrum," the FCC made a historic ruling in 2008 that allows "secondary" devices (i.e., the non-licensed users) to access licensed (but idle) spectrum for the first time in history [17,18,19,20]. This spectrum exists in the TV broadcast bands (highlighted in Figure 1.1), and provides up to an additional 180 MHz with 4x the range compared to 2.4 GHz due to the different propagation characteristics at the lower frequencies. The transition from analogue to digital TV broadcasts not only quelled concerns around low quality video and audio broadcasts, but it also addressed the concern about inefficient analogue broadcasts. That is, the digital conversion introduced higher quality video while using less spectrum. The FCC saw this as an opportunity to open up additional spectrum in the "white spaces" or idle gaps in-between digital TV broadcasts.

This "white space" style of spectrum access introduces many new challenges, however [21]. Secondary devices now need to locate the idle parts of the spectrum, which are dependent on their particular location in the US, and avoid parts of the spectrum being occupied by the TV broadcasts and another licensed user in the spectrum: wireless microphones. This means that unlicensed devices in this spectrum not only have to worry about interference from each other, but also about strict rules against creating interference on primary users, as we briefly discuss below.

Spectrum Details: The TV band spectrum that the FCC has allowed unlicensed devices to access through their ruling is in the upper UHF TV band from 512 - 698 MHz, highlighted in Figure 1.1. This band can provide up to 180 MHz of additional spectrum, subject to the active TV broadcasts and wireless microphones within range of the particular unlicensed white space device (WSD). This spectrum is broken up in to 30 channels, each 6 MHz wide, with one of these channels being restricted from unlicensed access.

Rules and Regulations: Since the initial announcement of the TV white spaces, many rules and regulations have been both introduced and changed. With the TV white spaces being the first band of their type, it is expected that changes will continue to be made to ensure safe usage around spectrum primaries, and to ensure unlicensed devices can capably and practically access the spectrum. We summarize the rules that are important to this thesis below, and refer the reader to the rulings for a more complete view [17,18,19,20].

The particular white space rules that pertain to this dissertation revolve around unlicensed devices operating in the spectrum without interfering with the spectrum primaries. Ultimately, secondary devices can create absolutely zero interference on the spectrum primaries. In the FCC's original ruling, unlicensed white space devices (WSDs) need to sense for active TV broadcasts and wireless microphones in the spectrum and avoid channels with either present [18,19]. Due to complexity and inaccuracies involved in sensing, however, petitions were received from many manufacturers against sensing requirements [20]. Instead, to provide protection, it was proposed that WSDs could access a database using their known location and query for active TV broadcasts and wireless microphones within range to perform spectrum management and avoid interference with primary users [22]. The known and stable locations of TV broadcast antennas would be used by the database in conjunction with the WSD's location to estimate if it is within range. Likewise, microphones register their location with the database. After testbeds and experimental results showed that the database could provide complete protection [22,23], the FCC dropped the sensing requirement in favor of a database in the Second Order [20]. In addition to this change, microphones have two reserved channels, yet are still allowed to operate in all possible channels.

Trends: Given the recent change in rulings towards a database approach, there have also been many proposed, and several accepted, FCC registered databases i.e., officially accepted for WSD usage [22,24,25]. Additionally, there is movement in the market to treat the spectrum as capable of housing "Super Wi-Fi" i.e., networks that provide connectivity like Wi-Fi with further range due to the lower frequencies of the TV bands [21,26,27,28]. If successful, we could see many wideband WSD in the spectrum that need to ensure interference-free avoidance or coexistence with the spectrum primaries. While early uses of the spectrum may be dominated by a few protocols, as new technology is developed to use the white space, a higher degree of heterogeneity is expected over time. This behavior was seen in the ISM bands, which was dominated early on by 802.11 traffic, but is now used by many technologies.

1.1.3 Wireless Technologies and Spectrum Access Schemes

Just like any broadcast medium, efficient and interference-free access of the wireless spectrum depends on proper coordination of its participants (i.e., radios) within range. Improper coordination can lead to two or more radios accessing the medium at the same time, causing a collision and failure in communication. Here, we describe the common and dominant access control schemes used in the bands that are important to this thesis.

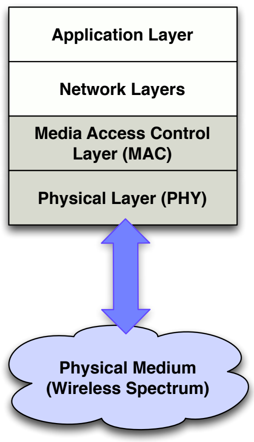

Figure 1.2: Layers in the wireless networking stack.

Basics of Communication and Coordination: Referring the reader to the wireless networking stack shown in Figure 1.2, communication is implemented at the physical (PHY) layer which allows to radios to communicate directly through modulation schemes (analog or digital). Two radios must share the same physical layer (i.e., modulation) to directly communicate with each other. The MAC layer coordinates control of the medium, often establishing coordination between radios via information from the PHY layer. Commonly, the PHY will

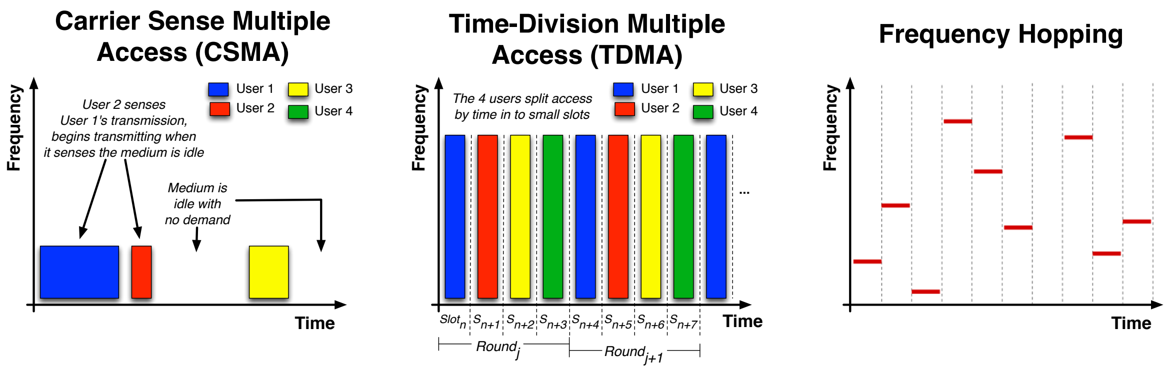

Figure 1.3: Common types of spectrum coordination/access used by unlicensed devices in the spectrum bands we consider in this dissertation.

decode transmissions (e.g., in to packets) that allow nearby radios to establish coordination. Additionally the PHY can provide analog information about the medium (e.g., the current level of power in it) that can be used to coordinate spectrum access.

Common Coordination Schemes: Based on the information provided by the PHY, there are several common coordination schemes that MAC layers implement. These schemes do not only address coordination, but their differences are also meant to handle different styles of traffic (e.g., bursty vs. continuous) and latency of the communication. We highlight a few of these schemes in Figure 1.3.

- *CSMA*: The basic principle of CSMA is to "listen before you talk," as illustrated by the 4 user example in Figure 1.3. This is a distributed means of coordination, allowing a radio to transmit only when it needs to (suitable for bursty application traffic), and only when there is no other active transmission in range (i.e., it is distributed). The "listening" done before each transmission typically happens in 2 ways: 1) Through analog power-based sensing, and 2) Through digital announcements of a transmission. The power-based method allows coordination between radios when they are within range yet do not share the same physical layer (i.e., modulation). The digital method is more robust in prevent deferral to spurious RF interference (e.g., from a microwave), but requires direct coordination. CSMA-based networks use one, and sometimes both, of these methods. Collisions can still occur in CSMA. For example, if two transmitters are out of range of sensing each other, but their receivers are within interference range of the opposing transmitter (i.e., hidden terminals). Additionally, when two or more radios sense an idle medium at the same time and begin to transmit, therefore random back-off before access is typically introduce which also improves fairness [29].
- *TDMA*: Illustrated in Figure 1.3, TDMA divides access to the spectrum in to pre-defined and fixed length time slots, allowing each user alternating access to the spectrum one after another in to these slots. Users alternate turns, and one series of transmissions

from each user is referred to a *round*. Time slots in TDMA-networks are typically short, e.g., on the order of microseconds, to bound the delay any single user has in accessing the medium. To establish the slots, synchronization, and ordering, one user is typically denoted a *master* by which all other users synchronize a logical clock to establish the beginning of rounds and slots. The benefit of TDMA over CSMA is that it avoids collisions between radios within a given network since access to the spectrum is broken in to pre-defined slots that are globally coordinated to the radios in the network. To contrast to CSMA, since time access is broken in to pre-defined slots and rounds, a bound on the latency any single radio has to access the spectrum can be made. This can provide guarantees of low latency in spectrum access (with small slot times), making TDMA a common protocol for providing wireless voice communication [30].

- *TDMA + Frequency Hopping*: Finally, *frequency hopping* is another common access scheme used in conjunction with TDMA. Radios coordinate a set of frequencies that are acceptable for use and quickly hop across these frequencies using a pre-established pseudo-random sequence. The scheme can ensure that narrow and high levels of interference in the spectrum do not adversely affect performance. The network may briefly hop into a part of the spectrum with high interference, but will quickly hop out of it.

1.2 The Problem: Heterogeneity breaking spectrum sharing

What should be clear from our overview of the spectrum and access schemes is that coordination between radios relies on shared MAC and PHY properties. Fundamentally, these layers are meant to work together across radios sharing the same frequencies by detecting when the medium is free before transmitting (e.g., by sensing), and/or by announcing that a radio will take control of it at a given time for a certain period (e.g., digitally). These principles ensure proper and efficient coordination of the wireless spectrum, and prevent two or more users from accessing the medium at the same time.

Today, diverse protocols and radios are being developed without these principles in mind which leads to the problem of interference between heterogeneous technologies. Instead of ensuring proper coordination of the spectrum, these layers are being customized on a per-application basis to improve performance. Diverse PHY layers vary the communication robustness and bitrate, differing MAC layers consider the application's traffic (e.g., CSMA – bursty, TDMA – continuous), while different spectral bandwidths and transmission powers are used to achieve the needed throughput and power constraints. Although this per-application optimization shows benefits in the short term (e.g., when spectrum utilization and heterogeneity are low), it is destructive in the long term as the numbers of applications, protocols, and spectrum utilization increase – exacerbating interference from a lack of coordination.

In fact, this destructive trend and behavior has been observed and advised against in other fields (e.g., in economics and ecology), referring to it as the tragedy of the commons:

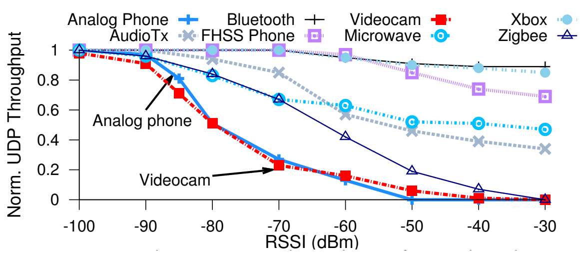

Figure 1.4: Heterogeneity can lead to significant reductions in performance from various types of devices and protocols.

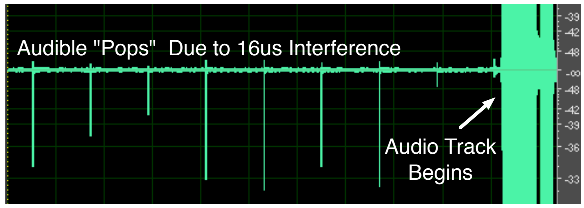

Figure 1.5: Even minimal heterogeneous interference on spectrum primaries in the white spaces can lead to notable problems.

> **Tragedy of the Commons:** *the depletion of a shared resource by individuals, acting independently and rationally according to one's self-interest, despite their understanding that depleting the common resource is contrary to the group's long-term best interests.* [31]

Within our discussion, the common resource is the spectrum and interference depletes it under high utilization and density due to diversity across technologies.

This destructive behavior is already evident. Heterogeneous interference breaking down spectrum efficiency and application performance has been widely reported across many spectrum bands, technologies, and environments (e.g., home and enterprise) [7,8,16,32].

We briefly refer the reader to Figure 1.4 to illustrate the impact of diverse technologies sharing overlapping channels. The figure shows an 802.11 network's performance as it comes within range of networks and devices using other technologies (provided by Rayanchu et al. [7]). As the interfering technologies are moved closer to the 802.11 network (towards the right on the x-axis), the 802.11 network's throughput can drop far below its fair share, potentially down to a complete loss (depending on the opposing device). Considering 802.11 as the interferer, it has been shown to increase the loss rates of Bluetooth and ZigBee to 40-90% [16,33]. Of course, heterogeneous interference is not limited to the ISM band. In Figure 1.5 we show the impact that heterogeneous devices can have on white space primary users: wireless microphones. Even a data packet as short as $16\mu s$ will cause audible pops in an audio recording from the microphone. This interference on the spectrum primary, no matter how small, is strictly against FCC regulation.

The interference highlighted in Figures 1.4 and 1.5 are created by various different types of conflicts driven by diversity. Understanding these technologies and the various types of conflicts they create sis important to this dissertation. In the remainder of this section, we provide details of unlicensed technologies relevant to coordination and interference avoidance.

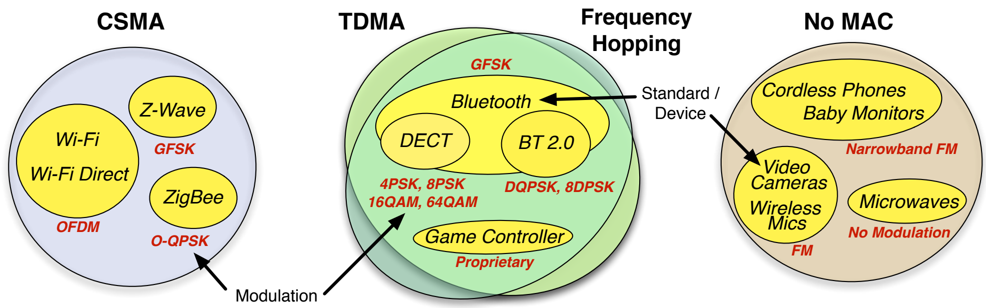

Figure 1.6: Taxonomy of protocols/devices in the unlicensed bands & their modulation schemes.

Heterogeneous Conflict Classification	Conflict Description
MAC Conflict	• Incompatible access techniques (e.g., CSMA vs. TDMA) • Incompatible sensing of the carrier (e.g., digital vs. power-based) • Varying techniques to determine time-based sharing schedules
PHY Conflict	• Incompatible modulation, preventing digital coordination • Varying transmission powers causing asymmetric hidden terminals • Differing spectrum usage (e.g., bandwidth), that can also vary spectral power

Table 1.1: A summary of common heterogeneous conflict types, broken down by MAC & PHY.

1.2.1 Conflicts created by heterogeneous technologies

The types of conflicts that two heterogeneous technologies can experience can be classified in to two main categories: heterogeneous MAC conflict, or heterogeneous PHY conflict. We summarize these categories and conflicts within them in Table 1.1 and describe the conflicts in detail below.

Heterogeneous MAC Conflicts: Starting at the MAC layer, the most basic conflict is from incompatible media access techniques. That is, for example, a network that assumes all radios have coordinated time-division based access, while sharing the spectrum with another network that assumes all radios will sense the medium before they transmit. This is in addition to operating with wireless devices that have no MAC at all, i.e., they continuously transmit without using a coordination mechanism. As a result of this type of conflict, studies have shown significant interference between CSMA and TDMA networks (e.g., 60-70% between 802.11 and Bluetooth [33]), particularly when a frequency hopping TDMA network hops in to the active frequency of a CSMA network. The impact of conflicts that include devices with no MAC have been shown to be even more severe. For example, Figure 1.7

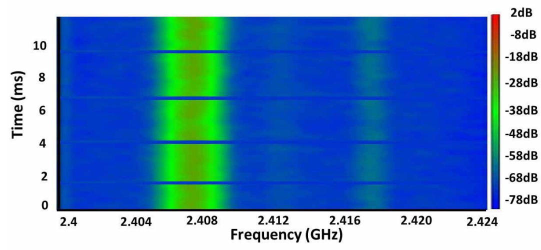

Figure 1.7: Continuous transmission without back-off or sensing from a cordless phone is shown in this spectrogram.

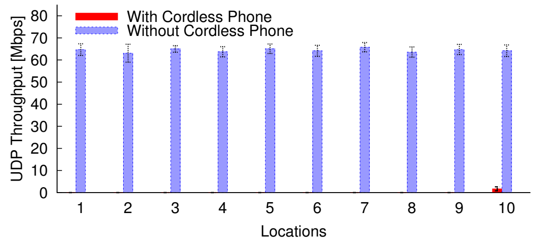

Figure 1.8: The continuous transmission from the cordless phone causes 802.11 to continually back off, driving throughput to 0.

shows the spectrum usage of a cordless phone in 2.4 GHz: a continuous transmission with no back-off or sensing (provided by Gollakota et al. [8]). When the cordless phone is in range of an 802.11 device, the 802.11 device can continuously back-off, driving throughput to 0 Mbps (shown in Figure 1.8).

The next two MAC conflict types in Table 1.1 describe conflicts that exist within a given style of MAC, i.e., between protocols/devices within the same MAC category shown in Figure 1.6. The first is within CSMA-based standards where the style of sensing conflicts, commonly digital sensing vs. analog sensing. For example, to prevent the constant back-off conflict with the cordless phone, 802.11 networks can be configured to use digital sensing, i.e., only back-off when they decode a transmission header from another network within range. Transmissions from the cordless phone would not be decodable, ensuring the 802.11 device would not back-off to it. While this can prevent the constant back-off conflict, it can create a conflict where the digital sensing network will not back-off to CSMA networks that use a different PHY layer. As a result, for example, studies have shown conflicts between 802.11 and ZigBee networks whose transmissions (and headers) are modulated differently, preventing them from digitally sensing each other [16,34].

TDMA networks that use different scheduling algorithms can, of course, conflict. While possible, the majority of TDMA-based standards use narrow channels (e.g., 1 MHz wide) and hop across 100+ MHz of spectrum quickly (on the order of 100s of μs) to avoid conflicting with each other and CSMA networks. As a result, he probability of two frequency hopping TDMA networks conflicting is low. While possible, this type is not common in practice and not a focus in our work.

Heterogeneous PHY Conflicts: At the most basic level, different modulation schemes prevent digital coordination. Therefore, for example, ZigBee networks cannot digitally coordinate with 802.11 networks (also briefly described in the MAC section). Failure to digitally coordinate due to differences in PHY layers can even happen within standards. For example, the different modulations schemes used between 802.11n and legacy 802.11 networks prevents legacy networks from detecting newer "greenfield" 802.11n pre-ambles, causing interference due to improper back-off [35].

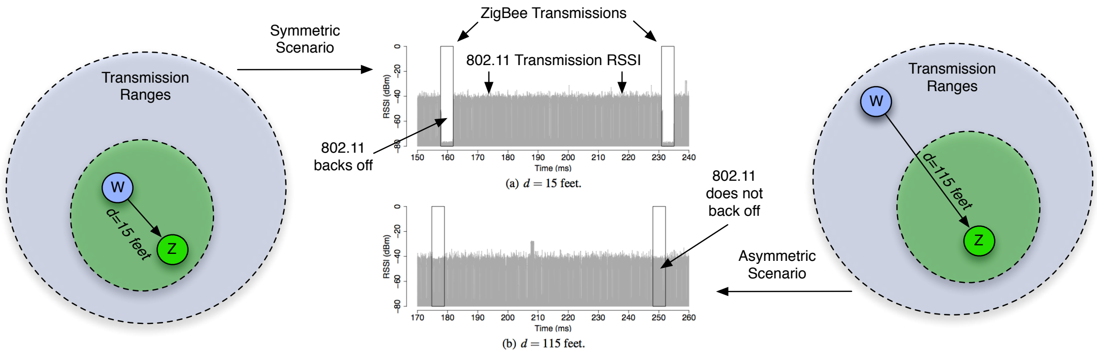

Figure 1.9: Asymmetric interference on ZigBee due to differing transmission powers.

The second common heterogeneous PHY conflict is asymmetric hidden terminals. This conflict is driven by differences in transmission powers across heterogeneous radios. While many 802.11 radios use similar transmission powers and tune their carrier sense threshold to these powers, other devices such as ZigBee or Z-Wave devices transmit at much lower powers due to power constraints (e.g., ZigBee sensors that are battery powered). As a result, this disparity at the PHY layer can cause the lower power transmitter to back-off to the higher power one at further distances, but not visa-versa leading to interference. To provide an example of this, we refer the reader to Figure 1.9 provided by Liang et al. [16]. Their results show that when the ZigBee transmitter is within close range of the Wi-Fi transmitter (e.g., shown at 15ft), the 802.11 transmitter will back-off. However, when the distance is increased (e.g., to 115ft), the 802.11 transmitter will not back-off, resulting in 40% or higher loss rate. Note that it may not reach a higher loss rates in the asymmetric scenario because, as the figure shows, the Wi-Fi signal is still strong at 115ft for the ZigBee network to back off.

Finally, the many different bandwidths used by the many standards creates conflict. 802.11 can use 20 MHz (legacy), 40 or 80 MHz (802.11n), and now up to 160 MHz (802.11ac). Bluetooth and DECT commonly use 1 MHz channels, ZigBee uses 5 MHz channels, whereas the analog transmitters can vary (e.g., between 200 KHz and 4 MHz). This creates two key conflicts: 1) Different bandwidths, even if the same modulation is used, can prevent digital coordination and sensing, and 2) The differing bandwidths lead to different spectral powers which complicates basic power-sensing. Sensing for the same amount of power across 20 MHz is different than 40 MHz in terms of spectral density. This often prevents wider band networks (e.g., 802.11) from backing off to more narrowband networks, e.g., 1 MHz used by Bluetooth [33].

1.2.2 Will this heterogenous trend continue?

What we have presented are various types of conflicts that degrade performance across heterogeneous technologies. It should be clear that, as a community, we understand these many types of conflicts and what creates them. Therefore, one might ask: can we reverse this trend, or will this trend continue? Why does the FCC not recognize this and take action? In fact, the FCC *does* recognize it, but has taken a firm stance against rules and regulations that would force some form of unifying behavior at the PHY or MAC. Unfortunately, they state that rules and regulations take longer for them to pass and modify than the "Internet time" by which technology changes [36]. A rule they introduce could prevent a new and breakthrough technology from being deployed, since it may take a year or more to make it compliant. This hinders innovation, consumer demand, and the marketplace. Additionally, there is simply no single PHY or MAC protocol that can satisfy the needs and constraints of all applications for wireless technology [37]. Going back to the 2 first-order problems we presented at the beginning of this thesis, these arguments suggest that the first challenge – increasing heterogeneity between unlicensed devices that quickly evolve – will continue, leading to the performance degradations we highlighted in Figure 1.4.

Second, as the demand on the spectrum increases through these many applications of wireless technology, we will continue to look for additional spectrum to meet the growing demand [1]. With increasing demand on the wireless spectrum, we will likely continue to take the white space approach to provide additional spectrum to meet this demand [17]. This means that the second challenge we presented at the beginning of this thesis – heterogeneity with spectrum primaries – will also continue.

Given: (a) The fundamental trend to support diverse applications with wireless technology, (b) The inability to develop a single protocol to meet the needs and constraints of each application, and (c) The stance by regulation authorities against unifying rules, and their push for additional spectrum through white space access – we conclude that this movement is very likely to continue and that new technologies are needed to address heterogeneous interference, which includes both general heterogeneous interference between many unlicensed devices (that quickly evolve), and heterogeneity in the presence of licensed devices (i.e., spectrum primaries with strict rules against interference).

1.3 Proposed Solutions and Their Predominant Application

This on-going trend of heterogeneity and its negative effects has led to many proposed solutions over time. The most common of these solutions can be placed in to 2 main categories: 1) *Coexistence techniques*, and 2) *Spectrum management-based approaches*. Below, we provide a brief overview of these approaches and argue that *each is being applied to address the wrong first-order challenge in addressing heterogeneous technologies*. Ultimately, we believe that the focus of their application in addressing the 2 first-order problems is reversed. This observation motivates our thesis statement and the research presented in the rest of this thesis.

Coexistence techniques are pair-wise solutions that address heterogeneous interference between two heterogeneous technologies sharing a channel. These techniques can allow *Technology A* to avoid generating interference on *Technology B* (e.g., through modifications to *A*'s

MAC and/or PHY), as well as allow *Technology B* to avoid interference from *Technology A* through changes to *B*'s protocols, making both *A* and *B* more robust with respect to interference from the other technology. For example, Liang et al. propose a coexistence technique that allows ZigBee to become more robust to 802.11 interference, through a modification to the MAC/PHY where the ZigBee header is transmitted twice with every packet (increasing probability of reception) [16].

Today, coexistence techniques are the most popular approach in our community in addressing the general problem of heterogeneity between the *many* technologies used by unlicensed devices (first-order challenge 1). For example, proposals between 802.11 and ZigBee [16,38,39], Bluetooth and 802.11 [33,40,41], 802.11 and cordless phones (or baby monitors) [8,42], as well as various other combinations of techniques (e.g., [43,44,45]). Very few proposals have been made towards solving unlicensed heterogeneity through spectrum management, and those that have are either Wi-Fi centric, i.e., they only attempt to manage the 802.11 network to avoid heterogeneous interference [15], or general architectures with no detail on the actual management algorithms [46,47].

The fundamental problem with the development of coexistence protocols to solve the problem of heterogeneity is that it is an N^2 solutions-based approach. That is, we are trying to find solutions between all pairs of technologies, and these technologies change so rapidly that the solutions (if they are actually deployed) are extremely short-lived. As technologies evolve, new techniques are needed (e.g., with the introduction of MIMO in 802.11n [8]).

Spectrum management based approaches attempt to avoid interference between incompatible technologies by isolating them in the frequency domain. If complete separation is not possible, the management algorithm attempts to place the network(s) in the spectrum where they will receive the least interference and best performance.

Currently, spectrum management is the most popular approach in addressing heterogeneous environments with spectrum primaries (first-order challenge 2). For example, to avoid interference from unlicensed devices on licensed users in the white spaces, the FCC has mandated a database that has licensed user's location and frequency, allowing unlicensed devices access to choose a channel without a spectrum primary in it [17,18]. Additionally, basic channel avoidance methods exist between 802.11 and radar systems in the 5 GHz band. Fewer coexistence techniques are developed and pushed to deal with spectrum primaries, possibly due to fear that they cannot guarantee interference-free coexistence.

While spectrum management has been shown to avoid interference with spectrum primaries in the white spaces [22], it can be *extremely* spectrum inefficient. As we will show later in this dissertation, spectrum management in the white spaces can waste significant amounts of spectrum, reducing availability to close to 0 in major cities. This is counter to the entire goal of the white spaces: to provide additional spectrum.

1.3.1 Reconsidering the General Applications of these Proposed Solutions

Clearly, coexistence techniques can provide spectrum-efficient solutions to cross-technology interference by allowing the diverse technologies to operate on the same channel. The downside to the general approach of applying coexistence techniques is that it requires N^2 solutions between N heterogeneous technologies. If new technologies are introduced often, or existing technologies are changed rapidly, then N will increase over time. There will be new types of conflicts, and new techniques are required to avoid the specific interference sustained between two particular technologies. This increases the complexity of the solution space given a coexistence protocol based approach, making such techniques difficult to develop and deploy under these conditions.

These challenging conditions have observed in the unlicensed bands with the technologies in them. Unlicensed technologies change rapidly (e.g., 6 revisions to 802.11 in a decade), and new technologies are introduced in the unlicensed spectrum continuously (e.g., ZigBee, Bluetooth, ZWave, W-Fi Direct, DECT, etc.). This makes coexistence techniques less desirable as the general approach to interference between heterogeneous and unlicensed technologies. However, our community has continued to focus on the general approach of coexistence techniques for unlicensed technologies (e.g., [8,16,39,41,43,44,48]). Despite these efforts, few of these techniques have been deployed due to these rapid changes, as well as the modifications needed at the radio and its protocols. Coexistence does not appear to be a long term solution for the unlicensed spectrum.

On the other hand, spectrum management can provide a "single" solution to interference between heterogeneous technologies through the development of an algorithm that considers the various heterogeneous conflicts (e.g., those we presented in the previous subsection), and assigns spectrum in a way that reduces the number of active conflicts. If designed properly, the spectrum management system can consider new conflicts that are introduced over time in a reasonable way, not requiring a new solution to address it, but rather to be cognizant that the conflict exists and that it should be accounted for.

In the white spaces, the opposite is true: spectrum management has been the primary focus in avoiding interference between the primary and secondary users in the spectrum [21,22, 49,50,51,52,53]. Considering the weaknesses of spectrum management, it can be less spectrum efficient since it does not allow heterogeneous technologies to share a channel. Its focus is separating the incompatible technologies in the frequency domain. When one considers the very goal of the white spaces: to achieve additional spectrum, it becomes unclear why the primary focus has been on spectrum management and not on coexistence techniques that can allow heterogeneous technologies to share the same channel.

Although the general approach of coexistence techniques can require N^2 solutions regardless of the spectrum band (unlicensed, or white space), N is currently small in the white spaces and unlikely to grow significantly since the number of licensed users typically does not increase. Although the number of secondary users can increase significantly, there are restrictions on secondary devices to be wideband; likely leading to many OFDM-based devices, further limiting the potential size of N. This means that few coexistence techniques are needed to achieve interference-free coexistence between primary and secondary users, that can be more spectrum efficient than spectrum management. but, again, the majority of work in this domain has focused on spectrum management techniques.

Given these observations and understanding of the strengths and weaknesses of both approaches, we strongly believe that we need to reconsider these general approaches and where they are applied in addressing heterogeneous technologies (licensed and unlicensed). In particular, we believe that we need to reconsider spectrum management as a solution to the problem of general heterogeneity between unlicensed devices, and coexistence techniques to address interference avoidance with spectrum primaries. Reversing the roles of these proposed solutions and their applications, however, is non-trivial. There are many challenges to applying spectrum management as the appropriate solution to general heterogeneity between unlicensed devices, and coexistence techniques towards addressing interference with spectrum primaries.

The resulting challenges to reverse the application of these general solutions.

Reconsidering the application of these general approaches and reversing them is not as simple as taking spectrum management techniques from the white spaces and applying them to address interference in unlicensed spectrum, and taking coexistence techniques that have been applied in unlicensed spectrum and applying them between unlicensed and licensed technologies. Here, we briefly describe why this is not a simple task, as well as the resulting challenges of reconsidering the general applications of spectrum management and coexistence techniques.

First, performing proper spectrum management requires knowing what signals are present in an environment and what they interfere with. The location stability of TV broadcast towers and wireless microphones made it possible to provide this information through a geo-location database in the white spaces. It is impractical to believe this approach could translate between the many unlicensed devices, their mobility, and their significantly lower power. Second, coexistence techniques used to address heterogeneous interference between unlicensed devices will not directly translate to addressing spectrum primaries. For example, SWIFT enables coexistence of wideband and narrowband unlicensed devices by requiring the wideband device to "poke" the narrowband device with interference to see if it reacts (i.e., determining if it it is within interference range) [42]. Clearly, one cannot "poke" spectrum primaries in the white spaces, further illustrating that techniques will not directly translate.

In particular, we identify 3 key challenges moving forward to apply spectrum management and coexistence techniques in their respective roles:

1. We lack the proper monitoring tools to determine where signals go in an environment, and to determine what they interfere with. This information is critical to proper spectrum management. Particularly, this information should be collected without requiring overly complex and costly multi-sensor deployments that history has shown are rarely deployed (e.g., [15,54,55]).

2. Even with the necessary information, current spectrum management models and algorithms are primarily homogeneous or Wi-Fi centric (i.e., they are only concerned about the interference on a Wi-Fi network) [15,56,57,58]. As a result, the cannot solve the more general problem of allocating the wireless medium across many heterogeneous technologies.

3. Developing efficient coexistence between unlicensed devices and spectrum primaries, at the least, requires an in-depth understanding of interference between these users. Unfortunately, there has been no comprehensive study to at least understand how data transmissions for unlicensed devices impact spectrum primaries in the white spaces. Without this kind of study, we are uninformed about how to avoid interference and establish coexistence.

Addressing these key challenges is critical to addressing the general problem of heterogeneity, and interference avoidance with spectrum primaries and, in particular, reversing the roles of spectrum management and coexistence techniques. If we can overcome these challenges, we can provide practical, deployable, and long-term solutions to heterogeneity.

1.4 Thesis Statement and Approach

In summary, this dissertation shows how we can address the 2-first order challenges in spectrum access driven by heterogeneous technologies in an efficient, practical, and robust manner to reduce (or eliminate) interference between licensed and unlicensed devices in the spectrum.

This dissertation claims that *better monitoring and spectrum management may provide a "single" and more long-term solution to interference between heterogeneous technologies in unlicensed spectrum, whereas coexistence protocols may be more suitable in providing spectrum-efficient interference avoidance with spectrum primaries.*

Approach: To show the potential of this general approach, this thesis presents two case studies: one that focuses on interference between heterogeneous devices in unlicensed spectrum, and another that focuses on interference between primary and secondary devices in white space spectrum bands. One of our goals in these case studies is to focus on environments that we believe are most critical and challenging (to justify our thesis statement). The sheer degree of heterogeneous technologies in the home with the many unlicensed wireless devices, yet a complete lack of expertise and monitoring, makes the home environment the most challenge for addressing general heterogeneity. Towards adoption of the white spaces, urban environments are critical since spectrum primaries are the most relevant. TV broadcasts and wireless microphones are at their highest density in urban environments, making it challenging.

Clearly, these case studies are but 2 points in a much larger problem space that we outline (a portion of) in Figure 1.10. There are various spectrum bands with unlicensed access that have the problems of interference between heterogeneous technologies that we

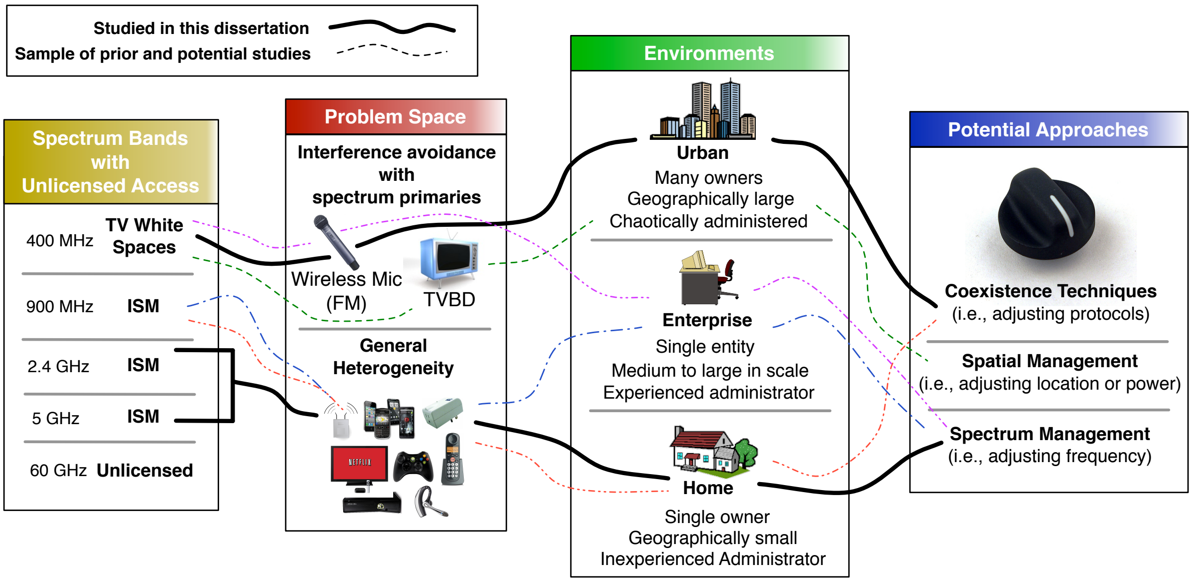

Figure 1.10: The general problem space for interference between heterogeneous technologies with the combination of environments and approaches we focus on.

have discussed. These problems have been shown to exist across 3 diverse environments: urban, enterprise, and the home, all of which have different properties in terms of their geographic size and potential administration. When attempting to address these problems, there are the 2 most common approaches that we discussed: 1) Coexistence, and 2) Spectrum management, in addition to a third that we do not address in this dissertation: 3) Forms of spatial management.

These different bands, problems, environments, and potential approaches lead to a large potential design space that has been the focus of prior studies (e.g., [7,15,16]), and will continue to be studied in the future. Additionally, there are other dimensions to this problem that we have simplified in Figure 1.10 such as specific technologies (e.g., CSMA vs. TDMA) and the degree of information known about the environment and the networks within it. The degree of information about the environment and its accuracy could suggest different approaches. As mentioned, however, we focus on general heterogeneity within the home due to the majority of these diverse devices existing in it, and spectrum primaries in urban areas where additional spectrum from white spaces is needed the most. We study spectrum management and coexistence to begin to address these problems, respectively, with the hope of showing the potential and longevity of the approach. There are challenges in taking our suggested approaches, however, due to a lack of tools, information, and techniques that warrant novel solutions and studies to understand the potential of our suggested approach.

Primary Contributions: This dissertation explores the proposed design spaces through the two case studies we have mentioned, and addresses challenges in 3 key areas that prior work has not been able to.

1. *Better Monitoring of Heterogeneous Technologies:* The first key challenge in spectrum management has been the lack of information about heterogeneous environments, i.e., where heterogeneous signals go and what they interfere with, to perform proper management. We believe that the home is the most critical environment due to the degree of heterogeneity within it, the lack of expertise, and the lack of proper monitoring. Centralized infrastructures with deployments of multiple sensors are too costly and complex for the home [15,54,55]. To overcome this, we introduce a novel single sensor design based on the smartphone. Our system is user friendly, not Wi-Fi centric, and able to collect the necessary information from the environment requiring only simple user interactions. This system provides the information needed for proper heterogeneous spectrum management.

2. *Spectrum Management for Heterogeneous Technologies:* With the proper information, we can now begin to efficiently address the general problem of heterogeneity. The next challenge that we must overcome is the lack of a general heterogeneous spectrum assignment models and algorithms. As discussed, current models target homogeneous environments, or are Wi-Fi centric (§1.3.1). We introduce a heterogeneous spectrum assignment model and algorithm that is generic, i.e., it can support current and future protocols, as well as multiple spectrum bands, yet descriptive and accurate: it is able to represent many different types of heterogeneous conflicts and resolve them. In particular, our approach leverages hypergraphs as a means to accurately represent the environment, given the information from the monitor, and a novel mixed integer program to determine spectrum assignments with minimal interference.

3. *Coexistence with Spectrum Primaries:* Shifting our focus to interference between primary and secondary users in the white spaces, we present results that show spectrum management can be extremely spectrum inefficient at guaranteting interference-free coexistence with primary users. Instead, we explore whether coexistence techniques are better suited for spectrum-efficient interference-free coexistence with primary users. Since primary users are well defined and stable (i.e., they do not evolve quickly), coexistence-based solutions in the white spaces can be long-lived. We perform the first in-depth study of data transmissions from unlicensed devices on wireless microphones (spectrum primaries in the TV bands), to understand how to develop coexistence between them. We then develop a novel coexistence technique with white space primary users (wireless microphones) to avoid complete interference, while significantly increasing white space spectrum availability.

These contributions highlight the potential of the approach presented in our thesis through two case studies that focus in two specific environments: one between unlicensed devices in the home, and one between primary and secondary devices in the TV white spaces and urban areas. Later in this dissertation, we discuss implications of these case studies and our contributions to further validate our proposal to explore spectrum management between unlicensed devices, and coexistence between primary and secondary users.

1.5 Organization of this Thesis

Below, we present the organization for the rest of this dissertation.

- In Chapter 3, we present our novel smartphone-based wireless monitor for heterogeneous networks that is designed to overcome the complete lack of a monitoring system for the home towards proper spectrum management, despite heterogeneity being a first-order problem in the environment.
- Chapter 4 presents our spectrum assignment system for heterogeneous networks in the unlicensed spectrum. This includes our algorithm and interference prediction metric to better organize the spectrum towards minimal interference. We show that this approach is promising towards addressing the general problem of heterogeneous interference without the need for N^2 coexistence protocols in the band.
- Chapter 5 addresses inefficiencies in avoiding interference between primary and secondary users in the white spaces that have relied on spectrum management. We show that spectrum management often requires organizations that are spectrum inefficient to avoid interference. Instead, we introduce a novel coexistence mechanism to overcome this limitation and evaluate its ability to avoid primary interference, as well as its spectrum efficiency.
- We conclude this dissertation in Chapter 6 with the implications and outlook of our work, including future work and a summary of our contributions.

Chapter 2

Smartphone-based Heterogenous Wireless Network Monitor

As we find more applications for wireless communication, we have turned to a diverse set of technologies to meet the specific needs and constraints of each application. This has resulted in the significant increase of heterogeneity in protocols used in spectrum bands with unlicensed access. Unfortunately, this diversity across the PHY & MAC often prevents coordination, exacerbates interference, and ultimately reduces network performance and capacity. This interference has been widely reported through various studies [7,8,16,33,35]. For example, cordless phones can decrease 802.11's performance by 90% or more [7,8], while 802.11 causes the same degradation for ZigBee networks [16].

When considering where heterogeneous devices are the most dense and susceptible to interference, it should be clear that the *home environment* is a major concern. Cordless phones, gaming controllers, baby monitors, wireless speaker systems, and "Smarthome" devices make interference in the home unique and challenging. Proposed solutions to this problem attempt to isolate incompatible technologies (i.e., spectrum management), to modify their protocols to reduce interference when sharing a channel [8,16], and to adapt their frequency usage (e.g., subcarrier suppression) and transmission power over time. To apply any of these solutions, it is critical to know more than whether a signal exists, but also where a signal goes and what it interferes with (i.e., its strength at various locations).

Unfortunately, it is nearly impossible to gather this information in the home due to lack of equipment and expertise. This makes spectrum management difficult to apply in the home. Additionally, bringing the information to a level the home user can understand to address issues is a challenge little work has attempted to address. The average home user has no knowledge of dBm, differences in technologies, or even MAC addresses (i.e., what devices have what address).

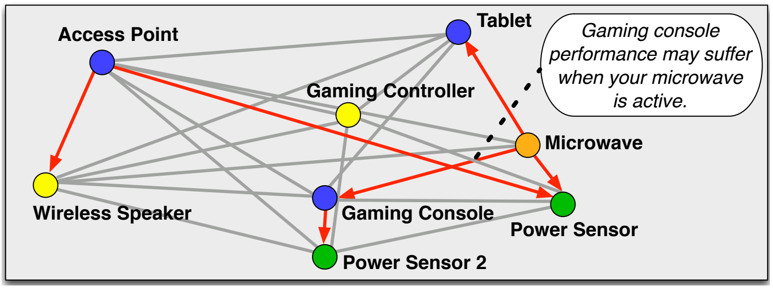

Figure 2.1: An example depiction of a useful RF environmental map.

Motivated by these concerns, the goals of our work are threefold: 1) To make it easy to collect (and update) an accurate view of the home's heterogeneous RF environment (i.e., where signals go and what they interfere with), 2) To make this information accessible to the average home user, and bring it to a level the average user in the home can understand, and 3) To make the information accessible to diagnostics, spectrum management, and coexistence techniques to improve connectivity in the home.

In this chapter, we make the case for a novel and practical home monitoring system that can achieve these goals *based on the smartphone*. The smartphone-based monitor we propose allows the average home user to collect meaningful heterogeneous RF information about the environment, and the system can bring it to a level the user understands. Our monitor differs from many of today's heterogeneous monitors and systems in several key ways.

First, while many heterogeneous monitoring systems rely on dense monitoring infrastructures with multiple monitors to determine where signals go in the environment and what they interfere with (e.g., JigSaw [55], DAIR [54], WifiNet [15]), our system only requires a single monitor: the smartphone. In fact, this single mobile monitor can likely gain more accurate information about the environment, than what is possible from various fixed points that happen to be equipped with a monitor. To collect this information without in-door localization, we leverage the user's natural movement with the phone that provides us with "close monitoring encounters" near each device. Although we do not know the true location of the nearby device, it still allows us to take a measurement near it and derive what is most important: where signals go and what they interfere with.

Second, whereas current heterogeneous monitors (RFDump [59], Airshark [7], DOF [60]) are only able to state that a certain signal classification is present at a certain strength (e.g., a Bluetooth signal @ -52dBm), our smartphone-based system is able to provide more detailed and user-friendly information, e.g., the nearby device named "Bill's iPad" is generating specific Bluetooth and Wi-Fi signals at -52 and -40 dBm. This device-level information is more user friendly and provides insight in to heterogeneous signals that other systems cannot generate. For example, that two heterogeneous signals in the environment coordinate since we know they are generated by the same device. To create these user-friendly device abstractions, we apply novel heuristics to cross-layer information collected by the phone.

Since our system is interacting with the average home user (unlike other systems that are designed for the enterprise [15]), its interactions with the user must be simple. Despite this, the monitoring system must still be able to collect the necessary information. To ensure the proper information is collected without sacrificing user-friendliness, we introduce a novel 3-phase approach. The first phase is *training* which is the most user-intense, but designed to be simple. The main task of the user is to tell the system which wireless devices in range belong to them, and then to bootstrap the system by holding their phone nearby each device. The second phase is *monitoring* which requires no user input and is made up entirely of system tasks: updating information about the environment as the user walks around. The third phase is *diagnostics & management* by which applications use the information collected to improve the environment (such as the one we introduce below).

Finally, the familiar and flexible interface on the smartphone makes it possible to leverage the information collected and present it to the user in a meaningful way to improve the environment. To illustrate this, we introduce a smartphone-based application that uses information collected by the phone and *force-directed graphs* to draw an easy-to-understand environmental map of the home, illustrated in Figure 2.1. The map abstracts complexity of the RF environment away from the user by abstracting signals and radios in to devices (with user-recognizable names). By applying logic to the information, we overlay diagnostics that can be useful to the average home user (e.g., conflict edges, coverage ranges, and interference ranges).

In this chapter, we present our smartphone-based home monitoring system design that overcomes the challenges of monitoring with a single sensor to achieve the goals we have discussed. In addition to the design, we present a 10-home user study which shows that the system is usable, and a heterogeneous testbed evaluation to show that the information collected is accurate.

Chapter Outline: We summarize the requirements, potential, and challenges related to a smartphone-based monitor in Section 2.1. With an understanding of these requirements and challenges, we present our high-level vision for the system's design in Section 2.2. Building on this design, we present the details of each component and our technical contributions in Section 2.3. To show the system is usable and accurate, we present our 10-home user study with heterogeneous testbed evaluation in Section 2.4. Then, to show the potential to build meaningful applications from this information, we present our heterogeneous home RF map with diagnostic overlays in Section 2.5. We present related work in Section 2.6, and then conclude with discussion and future work in Section 2.7.

2.1 Towards a Practical Home Wireless Monitoring System

Before we discuss our design of a smartphone-based home monitor, it is important to understand the requirements and goals of such a system (Section 2.1.1). With this understanding, inherent properties of the phone that make it an attractive platform to build a monitor become clear, in addition to challenges driven by the phone's limitations (Section 2.1.2).

2.1.1 Requirements of a Home Monitor

We believe there are 4 key requirements to a home monitor: 1) It must be heterogeneous in its monitoring capabilities, 2) It must not be Wifi-centric, and it must be 3) Comprehensive, and 4) User-friendly. We motivate these below.

There are many multi-radio and heterogeneous devices in the home, meaning the system must be *heterogeneous* and it *cannot be Wi-Fi centric*. Wi-Fi centricity would be considering all non Wi-Fi transmitters as in-configurable and interferers, and all Wi-Fi transmitters as non-interferers. Unlike the enterprise, there is a lack of priority among technologies in the home, e.g., the Bluetooth game controller may be as important as the iPad connected via Wi-Fi. Additionally, one must be able to differentiate all signals from internal, prioritized, and configurable (i.e., generated by the user's devices), from external (i.e., their neighbor's: contending and unconfigurable devices).

The system must be *comprehensive* in characterizing the environment, just like any other monitoring system. This can be decomposed in to spatial and temporal dynamics. Spatial dynamics include where signals go, what they interfere with, and detecting physical changes (e.g., a device moving). Temporal dynamics are important in accounting for time conflicts, e.g., informing ignoring spatial conflict if two devices never overlap in time, or prioritizing it if their activity overlaps heavily. This often also requires an understanding of communication patterns (i.e., who talks to who).

Finally, to be *user-friendly* the system must be low in cost by leveraging inexpensive and commodity hardware that users can easily install or access. Any interactions with the monitor must be simple, such as requesting input or tasks from the user, in addition to information being presented back to the user at a level they can understand. This means device abstractions are critical to usability: keeping the level of information at a level the user can understand, rather than signal-level information that users rarely understand.

2.1.2 The Potential of a Phone-based Monitor

With these requirements in mind, we present properties of the phone that make it an attractive, yet challenging, platform to develop a home monitoring system.

Beneficial Properties: The smartphone is now commodity, it has a very familiar and flexible interface to aid usability, and it is equipped with multiple heterogeneous radios for low-power sensing of the currently dominant heterogeneous technologies. For example, although ZigBee support is not common now, radio manufacturers are beginning to integrate it to address the rising popularity of the "Smarthome" [61]. This enables simple (and low power) sensing with the ability to passively decode, or proactively query, devices for rich and human readable information.

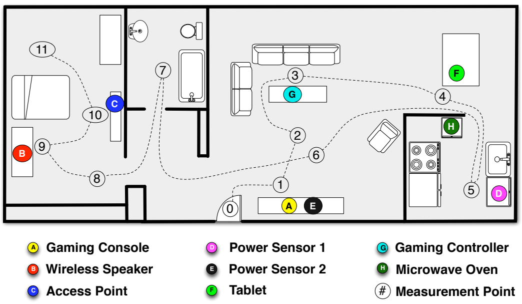

Figure 2.2: Our heterogeneous home environment and testbed.

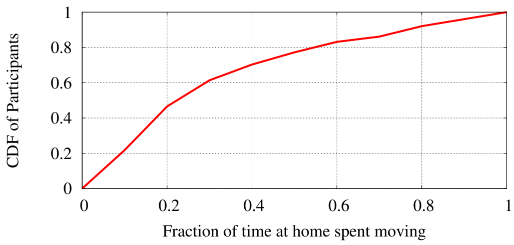

Figure 2.3: Phone movement creates ample opportunities to measure.

Even though the phone is considered a single sensor, its movement through the home allows it to achieve significant spatial diversity. In fact, a mobile can gain more accurate information about the environment than a set of sensors in fixed locations. As anecdotal evidence, we highlight a path that could be considered typical movement in a home, pictured in Figure 2.2. The user walks into his or her home, phone in pocket, and the phone gains several opportunities to take measurements near devices to learn where signals go (i.e., points 1,3,4,5,9,10).

To further validate this claim, we conducted a (accelerometer based) study involving over 100 users. On average, users' phones were home at least 5 to 6 hours a day (not including time when their owners were asleep). Figure 2.3 shows the fraction of this time each phone spends at home moving: half of the phones were mobile more than 20% of this time and one fifth for more than 50% of this time.

Questions and Challenges: A first-order concern of a smartphone-based monitor is whether the phone can provide accurate signal measurements, i.e., given its orientation and levels of obstruction (e.g., being attenuated in a pocket). In addition, the location of the phone at any point in time within the home is generally unknown (*assuming no in-door localization*). This makes deriving spatial dynamics (i.e., where signals go and what they interfere with) a challenge.

Temporal dynamics (e.g., how frequently devices are used), typically learned through continuous monitoring, must instead be learned through periodic monitoring on top of a monitor whose availability in the environment (i.e., home) is unpredictable. Smartphone power usage is a first-order concern for manufacturers and users: when the phone is home we need to decide when to monitor and to what degree.

Finally, there are many system's design questions involved in designing a monitor on the phone. For example, the level of involvement of the user, and the complexity of their tasks.

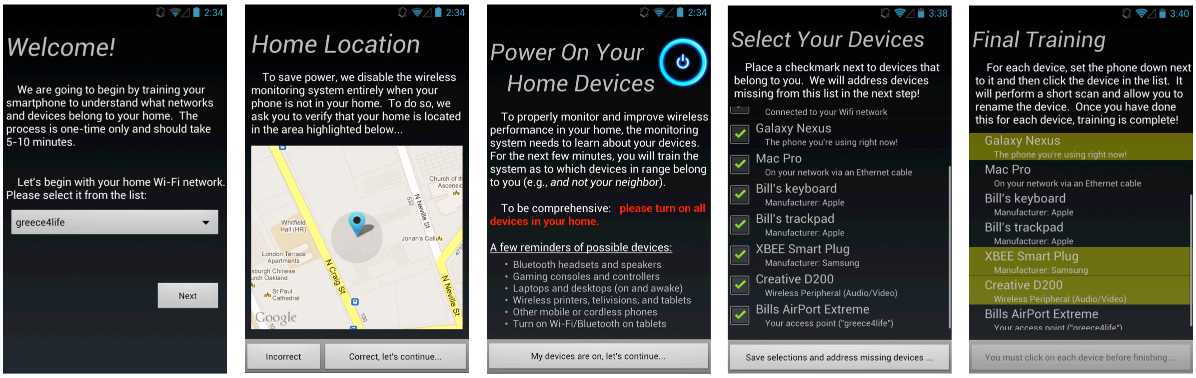

Figure 2.4: Screenshots of our system's interface, highlighting its simplicity, as well as usability through user recognizable identifiers.

2.2 Smartphone-based Home Monitoring System Design

Designing a smartphone-based home monitor to achieve the requirements we have identified (Section 2.1.1) is an open question that, to the best of our knowledge, our work is the first to address. In fact, there are various possible designs, many of which we have considered, that typically trade-off usability, complexity, power efficiency, and effectiveness. The majority of designs can be decomposed to the level of user involvement (tasks/input), and the system's degree of proactivity.

User Involvement: One could imagine requiring the user to walk around the home, trigger measurements, and manually label the physical location of each measurement. However, minimizing information and tasks required from the user make it more usable. The trade-off is that each piece of information needed to properly monitor and manage the environment that is not provided increases system complexity by requiring our power-constrained monitor to derive it (e.g., regarding *spatial diversity:* has a device moved?).

Degree of Proactivity: One could design an entirely reactive system that would only monitor when notified by the user of a problem which would be power efficient, but likely to require more time and user involvement to resolve issues (requiring the user to take the phone to various locations). A highly proactive monitor would continuously monitor and learn the environment to prevent / diagnose issues faster, but requires careful design to be power efficient.

2.2.1 Our High Level Design & Vision

It is hard to argue that any design is perfect in addressing the requirements of the monitoring system. However, we present a 3 phase approach that we believe is a balanced and practical initial design, only requiring simple user involvement to still meet the key requirements of a home monitor.

Phase 1 – Training: The initial training phase is invoked when the user is in their home with our system for the first time. It is the most user-involved phase, but designed to be simple, illustrated by the screenshots in Figure 2.4. First, the user trains the system of their home location, so that for power efficiency reasons (Section 2.3.1.5), our system can use coarse location information and disable itself when not in the home.

Next, the user is asked to turn on all devices in their home. Using the phone's heterogeneous radios and our techniques to derive device abstractions (Section 2.3.1.2), we scan for devices in the area and ask the user to select which devices are theirs (Section 2.3.1.3). This performs signal differentiation (key to avoiding Wi-Fi centricity) *at the device level*, where it is easiest for the user (especially with user recognizable names). Our system then correlates which devices transmit which signals.

Finally, we ask the user to take the phone and set it down next to each device. By doing so, we are able to: 1) Derive initial spatial dynamics and a conflict graph by creating device-centric views (Section 2.3.1.1), and 2) Derive signal strength thresholds indicative of being 1-2ft from the device, used by our system to opportunistically update measurements as the user walks around their home with the phone in their pocket.

Phase 2 – Monitoring: At this point, our system performs background monitoring tasks whose details and operations are largely transparent to the home user. In this phase, the system (*not user*) tasks are: 1) Ensuring an up-to-date and complete list of internal networks by periodically scanning for new devices we believe may be the users, 2) Opportunistically updating where signals go as the user walks past devices, and 3) Learning of spatial and temporal dynamics (e.g., when devices move / how active devices are). These periodic tasks keep an up-to-date view of the environment *without* heavy user involvement or in-door localization.

Phase 3 – Diagnostics & Management: The information collected in the first two phases is useful and important to various applications. This information can be used to perform diagnostics, implement coexistence techniques, or even draw an easy-to-understand environmental map with overlaid connectivity information, as we will show in Section 2.5.

2.3 System Components

An overview of our system is shown in Figure 2.5. The high level design includes a *RF / radio scanner* and a *management system*. The management system deals with the strategy of when to take measurements in the environment to comprehensively map the environment, yet to be power efficient. The RF / radio scanner is responsible for the more low-level scans, including being able to use cross-layer information and heuristics to create device abstractions. These two components access and update the database.

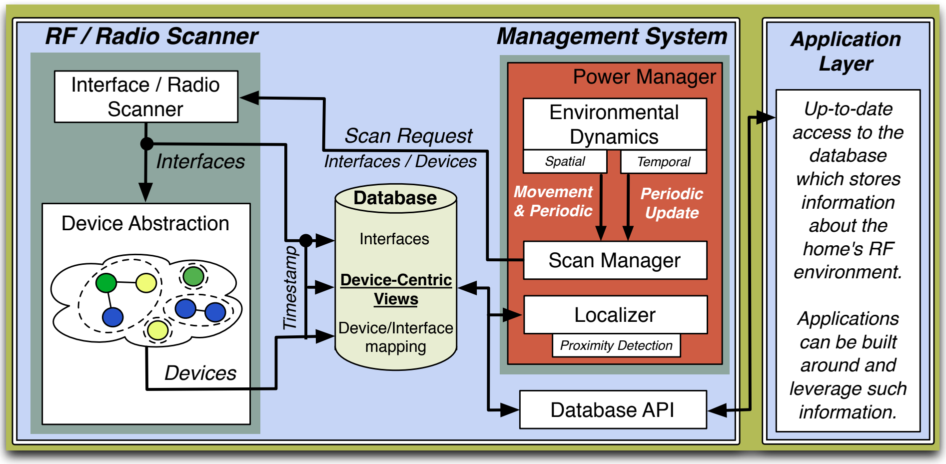

Figure 2.5: Overview of our system design and its components.

2.3.1 System Components

In this section, we present our system components in this design. This includes device-centric views (Section 2.3.1.1), device abstractions (Section 2.3.1.2), differentiation (Section 2.3.1.3), spatial and temporal dynamics (Section 2.3.1.4), and power-efficiency (Section 2.3.1.5).

2.3.1.1 Device-Centric Views

Device-centric views are measurements taken in close proximity to an internal device, allowing our system to estimate what signals reach each device. These are taken during training when the user places the phone next to each device, and opportunistically as the user walks throughout their home. In this section, we address the accuracy of such measurements and how to localize/trigger them when near a device.

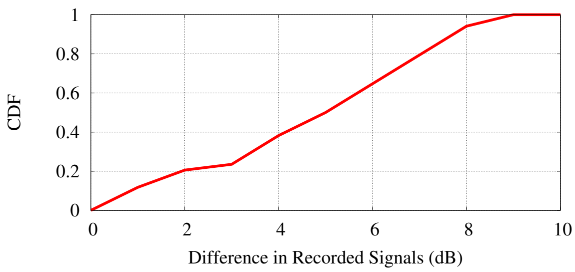

Figure 2.6: The phone can be calibrated to report signals with an accuracy of ±5 dB.

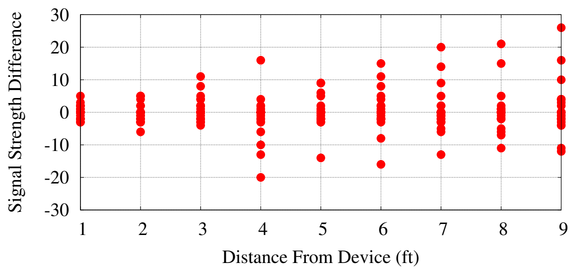

Figure 2.7: When the phone is within 2ft of a device, the view it gains is accurate.

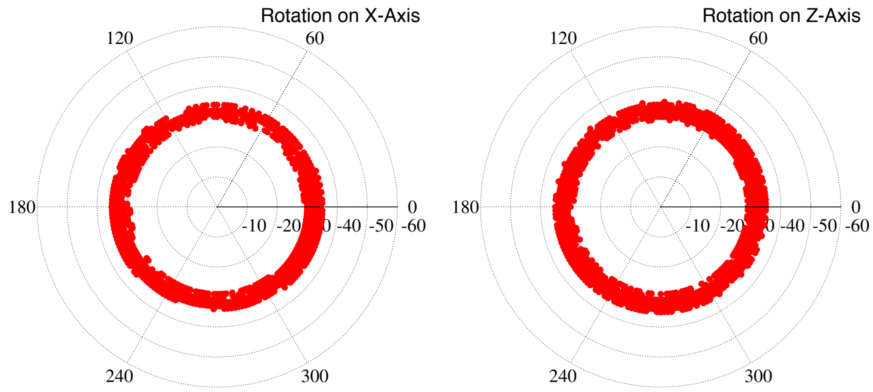

Figure 2.8: Signal strengths are stable across various orientations; at no angle are significantly stronger or weaker strengths observed

.

Accuracy: A first-order concern is how accurate our device-centric views can be. While external monitors have been used in dense monitoring infrastructures (e.g., placed directly next to APs), such infrastructures assume the monitor is stable and unobstructed (unlike the phone), and prior work has assumed that the signal strengths observed by the monitor are the same as at the device. Here, we quantify the accuracy of an "external" device-centric view taken by a phone.

- Accuracy of an External View: In Figure 2.6, we quantify the difference in signal strengths observed by the phone, as compared to the signals observed by a laptop when placed directly next to it. A negative value indicates the phone underestimating a signal's strength, whereas positive values indicate overestimation. As shown, our particular phone (Galaxy Nexus) consistently overestimates signals by approximately 4-5dB. Such error could be corrected through analysis of current generation phones and adjusted based on the hardware model observed by the software. Another method could be to ask the user to install a corresponding App on their laptop that, during the training phase, would measure differences observed in signals for calibration. Our results show that the phone could accurately report signals within $\pm 5dB$ after calibration for this offset.

- Impact of Distance: Beyond the training phase, we want to opportunistically update the signals observed at each device when the user walks near the device. Therefore, it is important to understand the accuracy at various distances. This allows us to derive an acceptable distance to take opportunistic measurements, and when to avoid measurement. Starting directly next to the device (i.e., a distance of 0), we use our monitor to record the signal strengths of all heterogeneous signals in range, recording variation in their strength as we walk away with the phone. From performing such experiments, we find that an accuracy of ±3-5dB is maintained within 1-2ft and then tapers off, as shown by two examples presented in Figure 2.7. Accuracy can tapper off for several reasons at greater distances. For example, the monitor can begin to move into areas where it achieves or loses line-of-sight significantly varies a signal. However, these results motivate only triggering updates near devices, within 1-2ft.

- Impact of the Phone's Orientation: The next concern is whether the orientation of the phone impacts accuracy. If it does, at certain angles it will report higher or lower signal strengths. To measure the impact of orientation, we continuously poll the phone's sensors to record: 1) Its orientation as we rotate it on its X and Z axis and, 2) Its signal strength of a device at that orientation. In Figure 2.8, we show that independent of the phone's orientation, it reports the uniform signal strengths. At no angle does it report different behavior.

- Impact of Pocket and Body Attenuation: Finally, we measure the impact of obstruction from the user's body (e.g., with phone in pocket). To conduct the experiment, we place the phone in our pocket and face in four different directions, obstructing and un-obstructing a nearby device and its "view" we are capturing. Our measurements (not pictured due to brevity), show that the impact of the body on signal attenuation is relatively negligible: approximately 1dB. These results are consistent with other work that found strong multi-path in the home negates the impact of body attenuation [62].

Summary of Accuracy: Our results show signal estimation to be within $\pm 5dB$ when placed directly next to the device, and ±7-10dB when at distances of 1-2ft. While this may seem large, there are still useful applications (Section 2.5).

Localizing/Triggering Measurements: As discussed, our goal is to take opportunistic measurements when near a device. Although we do not know true physical location, our goal is to detect proximity, measure, and *anchor* the measurement to the relative location of the internal device. A history of measurements of these measurements are then stored in the database as being at the device's relative location.

There are two challenges in this, however. The first is illustrated in Figure 2.9. We placed the phone directly next to several devices in the home and record the signal strength of the device observed. As shown, signal strengths when near a device are truly device dependent, even within a manufacturer. Therefore, signal strengths indicative of being near a device needs to be learned per device. As part of our system design, when the user trains the system (i.e., phase 1), the user is asked to set the phone next to each device to empirically learn proximity thresholds for each device. To detect proximity, trigger a measurement, and anchor it we monitor the accelerometer on the phone to detect movement and then continuously monitor signal strengths of internal devices to be within 2dB of the strength learned in training.

The second challenge is that the physical location of a device can change, leading to device-centric views being anchored at two different relative locations. To overcome this, our

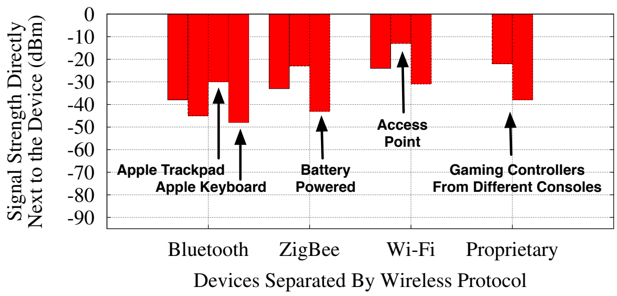

Figure 2.9: No single threshold will work to localize measurements to a device's relative location. It must be learned to detect proximity.

Name extraction method	Derives	Protocols supported	Applicable networks, devices and manufacturers
Passive Packet Inspect	Network Names,	Wi-Fi, Bluetooth, ZigBee	Network names and associations to derive connectivity
Service Requests	Names & Services,	Bluetooth, ZigBee	Wireless speaker/headset manufacturers, Smart-home
SSDP	Names & Services,	Ethernet, Wi-Fi	Cisco (e.g., "Linksys N900"), Microsoft, HP, Dell
Zeroconf / mDNS	Names & Services	Ethernet, Wi-Fi	Apple devices, Linux-based devices
Reverse DNS Lookup	Names	Ethernet, Wi-Fi	Many devices register a hostname that is informative
IEEE OUI Database	Device Manufacturer	Wi-Fi, Bluetooth, ZigBee... Ethernet	Any device with an IEEE registered MAC address

Table 2.1: A summary of the naming methods and protocols used to extract user recognizable identifiers.

system employs two techniques. First, during the training phase when the user sets their phone down next to each device we ask them if it is mobile: "Is this device always in this location?" If not, we only keep a measurement history of depth 1 for the device: its most current measurement is the only one we can trust. On the other hand, fixed devices can still move (e.g., the Xbox changing rooms). In Section 2.3.1.4, we show how to detect the movement of fixed devices and invalidate device-centric views in the database once moved.

2.3.1.2 Creating Device Abstractions

In this section, we describe how to create device abstractions by exploiting cross-layer information (PHY, MAC, network), and applying heuristics to merge multiple radios and/or signals in to the same device. We begin with how to extract user-recognizable names which are critical for usability and used for merging. Then, we describe our merging approach.

Radio / Interface Names: We leverage the phone's packet capabilities to extract user recognizable identifiers using several methods, summarized in Table 2.1. First, we use *passive packet inspection* to extract network names, e.g., "The Smith's Wifi" included in broadcast traffic such as beacons. *Service discovery protocols* (SDPs) such as Bonjour, UPnP, and SSDP provide user specified names (e.g., "Jack's PC"), hardware specifics (e.g., "4th Gen MacBook Air"), and operating system specifics (e.g., "Ubuntu"). SDP use is nearly a standard practice now and *not* limited to 802.11. Bluetooth devices respond with a "Major Service Class" (e.g.,

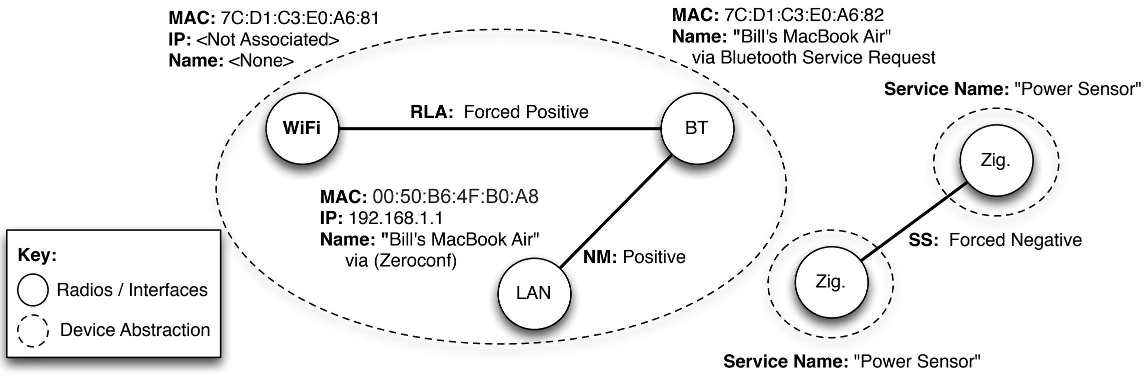

Figure 2.10: To derive radio/interface relationships, we apply heuristics to connect them in a graph, and then extract device abstractions.

networking, audio), as well as a "Major Device Class" (e.g., phone, speaker, toy). ZigBee devices also respond with available services (e.g., "Power Monitor").

As a third method, the smartphone queries *IEEE's OUI* database [63] which maps the first 24-bits of a MAC address (Wifi, ZigBee, Bluetooth, and others) to the organization who assigned the address. This returns names such as: Dell Computer, Microsoft Corporation, and Logitech, helping a user narrow down the device and rename it. Note that OUI typically does *not* return the radio manufacturer (like Atheros).

For devices the phone has no supported radio (e.g., cordless phones), we fall back on the ability of modern 802.11 radios to detect such devices and classify them [7]. We find this to still be usable since "Bluetooth Signal" will not help separate devices, but "Microwave" and "Cordless Phone" are more descriptive. Device activity can also be tied together. For example, the Xbox 360 associating to the Wifi network when a game controller is observed.

Merging Radios / Interfaces in to Devices: To create device abstractions, we leverage hints across the PHY, MAC, and network layers that radios/interfaces belong to the same device. Systematically, we create a weighted graph where radios/interfaces are nodes, and begin by fully connecting the graph with each edge weighted to 0. Then, we apply a set of heuristics on the cross-layer information which suggests each pair of nodes (i.e., radios) belong to the same device (or not), increasing (decreasing) the edge weights between each pair of nodes. After applying each heuristic, we prune edges in the graph whose weight is ≤ 0 and consider the remaining connected nodes as belonging to the same device.

We refer the reader to our example graph in Figure 2.10, and our set of 5 heuristics summarized Table 2.2. Each heuristic will modify an edge weight positively (+1), neutrally (0), or negatively (-1). Some heuristics provide absolute confidence that nodes or interfaces are related or not (forced -, forced +); they override the weight on the edge.[1]

[1] In practice, we have not seen a conflict where a forced negative and positive are applied to the same edge based on our heuristics.

Information / Heuristic	Abbr.	Possible edge weights
Network Level Addressing	NLA	Positive, Neutral
Related Link-layer Addresses	RLA	Forced Positive, Neutral
Radio / Interface Names	NM	Positive, Neutral
Technology Type	TT	Neutral, Negative
Signal Strengths	SS	Neutral, Forced Negative
Timing Analysis	TA	Neutral, Positive, Negative

Table 2.2: Heuristic summary to derive radio/interface relationships.

Briefly, Network Level Addressing (NLA) looks at network-level identifiers (e.g., Wifi network names and IP addresses) to merge radios and interfaces together. Related Link-layer Addresses (RLA) often looks at the relationship between MAC addresses, e.g., it is common for manufacturers to assign adjacent MAC addresses to interfaces that belong to the same device. This creates a forced positive between the Wi-Fi and Bluetooth radios in our example (Figure 2.10). The Radio/interface Name (NM) heuristic examines the user-recognizable identifier names. In our example, the Bluetooth and Wi-Fi interfaces respond with the same name ("Bill's MacBook Air"), causing NM to apply a positive value on the respective edge. Technology type can be used to break merging (e.g., it is uncommon for a ZigBee and Wifi radio to belong to the same device), and finally we use signal strengths to apply forced-negatives by recognizing two radios observe highly different RF environments. We do not use signal strengths to merge radios, because two closely placed devices can observe similar RF environments (i.e., it can be a good negative indicator, but not positive indicator).

We later show the success of these heuristics in our testbed environment and 10 home user-study (Section 2.4). Note our graph-based approach is also flexible to applying other heuristics.

2.3.1.3 User-Friendly Differentiation

Going back to the requirements of a wireless monitor for the home (§2.1.1), being able to differentiate signals as being internal (generated from the user's devices) and external is critical. This is a key piece of information that any spectrum management or coexistence technique would need to know (i.e., whether the signal should be treated as interference).

Our system makes differentiation simple to the average home user after creating device abstractions and performing basic filtering that can make the process more simple. Since we now have a mapping of what devices generate what signals (i.e., §2.3.1.2), we can easily differentiate by reverse-mapping, i.e., presenting a list of devices to the end-user (with user recognizable names) and asking them to select which devices are theirs. We filter out various devices by the networks that they are associated to e.g., all Wi-Fi radios that are not associated to the user's home network. Natural filtering will occur with lower-power devices when

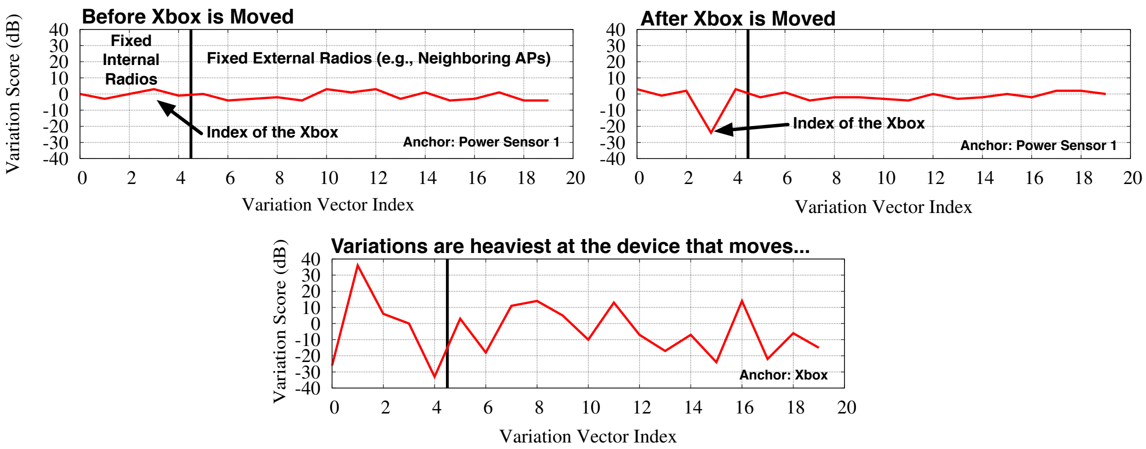

Figure 2.11: Physical movements in "fixed" devices (e.g., the Xbox changing locations) can be tracked by observing variations in signal strengths.

neighboring homes have separation. Results in our evaluation show filtering plays a more key role in dense (apartment) areas (§2.4).

User-recognizable names, device abstractions, and filtering provide the user-friendly list of devices to perform differentiation shown in Figure 2.4. Our user study (§2.4) verifies this simplicity by showing that users only spend a small fraction of time differentiating: < 2 minutes on average.

2.3.1.4 Environmental Dynamics

The two key dynamics in the environment are *spatial* and *temporal* dynamics. Spatial dynamics deal with detecting and then accounting for physical changes (e.g., a device moving locations), and temporal dynamics deal with changes (e.g., network loads) which are time-dependent.

Spatial Dynamics: Detecting spatial changes are critical to: 1) Notify higher layer diagnostics which may need to account for the change, and 2) Invalidate device-centric views that were "anchored" to an old relative location.

We consider differentiating fixed and mobile devices as important, since the movement of fixed (as opposed to mobile) devices is a more critical change. The reason for this is that spectrum management algorithms can prioritize spectrum configuration based on fixed devices that are frequently used, whose interference is stable. In contrast, mobile device interference is harder to account/configure for. We believe the average consumer is less tolerable of their fixed devices performing improperly (e.g., their AP or Xbox), whereas users expect mobile device performance to be variable.

To detect fixed device movement, notifying higher-level applications (e.g., spectrum management), and invalidating its older anchored measurements, *we track the variations of signals* *between fixed devices in the environment, from other fixed devices.* The key information we are leveraging is that signal strengths are relatively stable between fixed devices. However, if a fixed device moves: 1) Its signal strengths from other fixed devices will vary greatly, and 2) Its strength observed at other devices will also diverge.

Systematically, we use the history of (valid) anchored measurements at each fixed device and compute a *variation score* of their signal strengths across their history. Each fixed device maintains a stable variation vector index, and the value at the index is that device's maximum (but stable) variation in signal strength to the device in question. We illustrate this behavior in Figure 2.11 by moving the Xbox in our heterogeneous environment (Figure 2.2) from its pictured location, to the desk next to the AP in the bedroom. The top 3 graphs in Figure 2.11 show, before we move the Xbox: its signals observed from other fixed devices are stable over time, where each index in the vector is the signal from another fixed device. When we move it, we see variations in the Xbox signal at other devices, and the variations in the signals at the Xbox change greatly. We use measurements over time and this understanding to detect when user's "fixed" devices move.

In §2.4, we will show that as long as the device moves more than 5ft, we can accurately detect its change of location.

Temporal Dynamics: Temporal dynamics are largely spatially independent, meaning that as long as the phone is in the home, it can observe temporal network and device usage. For this reason, we can periodically monitor for the usage of devices (once every 30 minutes in our design), and measure their airtime. This usage is kept as a history within our system's database which, again, is queryable by applications which are built on top of our monitor. Over days with our monitor active, we can collect a significant amount of data based on when devices/networks are most active.

2.3.1.5 Energy Efficiency Manager

We make several key design decisions to improve the energy efficiency of our monitor. Importantly, we *only* enable monitoring functionality when the phone is in the home. During the training phase (§2.2.1), we record the phone's coordinates and use them to enable/disable monitoring. This only requires a coarse grained location, where the more power efficient cellular network-level location values are sufficient [64]. Next, we also only enable monitoring when the phone is not in use. The goal is to avoid draining the battery.

For energy efficiency in capturing spatial dynamics, we make the system cognizant of the fact that when the phone is not moving, spatial dynamics are likely the same. Therefore, we only enable proximity detection and device-centric views when the phone begins moving. This can be implemented in a power-efficient manner since major smartphone operating systems provide callbacks on motion sensor changes (e.g., the accelerometer). We find these callbacks are provided even when the phone is in a deep sleep after inactivity.

Since learning of temporal dynamics require more frequent monitoring, we only enable the learning of such dynamics when the battery level is high or the phone is charging. Additionally, we make the system aware of the fact that some radios are coupled to the same hardware and behavior. For example, observing the Xbox's diurnal usage is sufficient to learn of its controller's usage.

2.4 Prototype and Evaluation

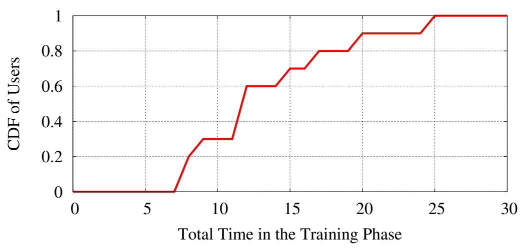

Figure 2.12: The total time spent by users in the training phase.

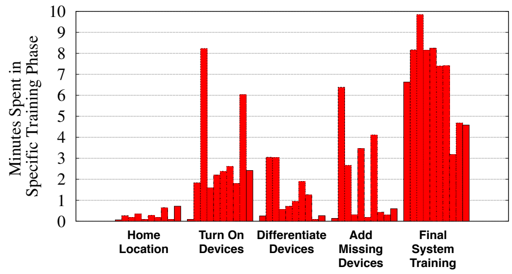

Figure 2.13: The amount of time each user spent in the various parts of the training phase.

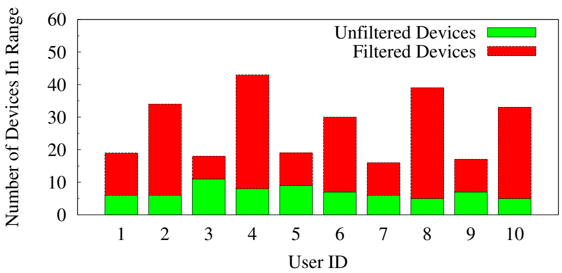

Figure 2.14: The total number of devices within range and filtered out.

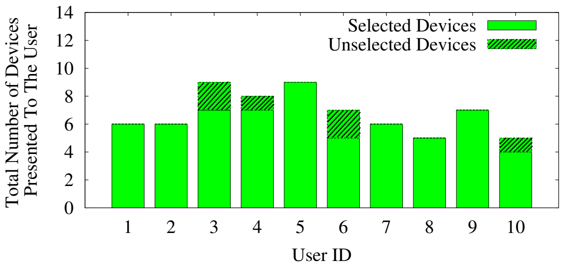

Figure 2.15: The number of devices selected and unselected in the differentiation stage.

To evaluate our system, we build a full prototype on Android OS (Galaxy Nexus hardware), conduct a study of its usability through a 10 user home study, and then details of its accuracy in a controlled testbed. Our prototype integrates a ZigBee radio (given manufacturer's recent announcements of ZigBee support in future smartphones [61]), and Airshark-like functionality using a WiSpy device. Airshark was not publicly available at our submission.

2.4.1 User Study of Training Phase / Interface

We conduct a user study across 10 homes to evaluate the ability for home users to train our monitor in their environment using our interface and ability to bring information up to a level they can understand (e.g., allowing them to differentiate at the device level, not signal level). Users were given no knowledge of the tasks they would have to complete and we

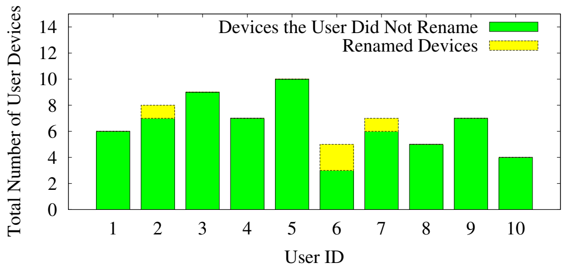

Figure 2.16: The number of devices the user chose to rename and not.

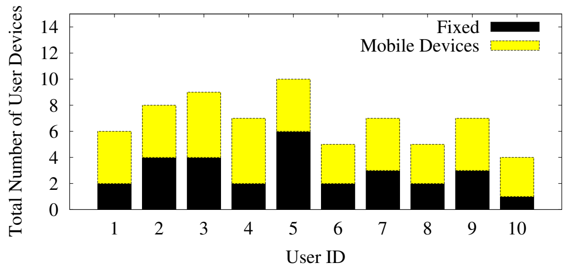

Figure 2.17: By user, a break down of the number of fixed and mobile devices.

record information at each step. Later, we will evaluate the monitors accuracy in a controlled environment (Section 2.4.2).

The training phase has 5 main steps (most of which are depicted by the screenshots in Figure 2.4):. They are: 1) Verifying the user's home location, 2) Turning on all devices in the area, 3) Asking the user to differentiate their devices from external devices, 4) Adding any devices our system missed in the initial scan, and 5) Training the system, which asks the user to set the phone down next to each device in the home, renaming the device if needed, and specifying if it is mobile.

First, we present a summary of the total time the users spent to complete the entire training phase in Figure 2.12. Our system required an average of 14 minutes to complete the training phase, with 3 users only requiring around 10 minutes to complete it, and 2 users 20 minutes to complete.

In Figure 2.13, we break down the total time each user spent in the training phase in the 5 main steps. Verifying the home location took less than 30 seconds for all users, and the result was always correct. The time it took users to turn on their devices greatly varied. In one case (User 1), the user quickly scanned through the list of reminders and claimed "all of my devices are always on," whereas User 3 spent a significant time turning on 3 gaming systems and 2 desktops.

Users were able to differentiate signals in the environment (e.g., *at the device level*) within a maximum of 3 minutes. From speaking to the users, most claimed this was made easy by the names we were able to associate with each device. Table 2.3 shows a summary of our naming techniques used to achieve this. Service-discovery based protocols provided a significant fraction of names, and rarely did we fall back on manufacturer. Microwaves, gaming controllers, and cordless phones were the only devices we used signal-classification.

This differentiation was also made simple through our filtering (Section 2.3.1.3). Figure 2.14 shows, per-environment, how many devices were in range and the number filtered (e.g., due to confidence they were not the user's). The amount of filtering done was highly dependent on the home environment: users 4,8 lived in dense apartment buildings whereas users 1,3,7,9 lived in more separated home environments. From the devices we chose not to filter (presented to the user), the majority ended up being the user's devices, as shown in

Name Resolver	Internally Used
Passive Packet Inspection	13% (10 / 78)
Service Discovery-based	71% (55 / 78)
IEEE OUI Manufacturer	4% (4 / 78)
Signal-Classification Fallback	12% (9 / 78)

Table 2.3: Summary of name resolution from user study.

Figure 2.15: 60% of the lists presented to the users *only* contained their devices (i.e., they selected all of the items). The maximum number of devices that the user had to manually filter out was 2.

The time spent adding missing devices also varied: 60% of the users spent less than a minute on this step, suggesting that they were able to quickly verify all of their devices were listed. Also, 60% of users reported that no devices were missing from their list (i.e., they did not add anything in the "Missing Devices" step). Of the remaining 40% of users, 20% added 1 missing device, and 20% added 2 missing devices. Through inspection of our filtered devices and the devices they added themselves, none of these devices were missing from the initial list due to improper filtering. These were mainly Bluetooth devices out of range of the initial scan, or accidentally not turned on.

The last step, where the user is instructed to place their phone next to each device, was mainly driven by how sparsely the devices were placed and how many devices the user owned. User 3 spent a significant amount of time since their devices were scattered across several rooms in a 3 floor house. We also prompt (and encourage them) the user to rename each device if needed. For example, if our system fell back on presenting the manufacturer of the device such as: "HTC Device" instead of "Bob's Phone" we encourage the user to rename it. In Figure 2.16, we break down the number of devices each user in our study renamed. As shown, 70% of users felt the names we derived for their devices were useful, recognizable, and sufficient.

Finally, we present the fraction of devices that are fixed and mobile in each user's home in Figure 2.17. This includes devices that may be mobile, but the user marked as "fixed" because it never changes locations. This was true of User 6 who, based on the wording of our button as "Fixed: This device is always in this location" was able to properly choose fixed for their laptop, verbally mentioning "this laptop is always in this location." The homes we studied had enough fixed devices to track internal device movement (Section 2.3.1.4).

Summary: Our results show that our system provides a reasonable training phase that can allow our monitor to collect and bootstrap its information about the users home.

2.4.2 Value of the Information Collected

Next, we validate the accuracy of our smartphone-based monitor in a heterogeneous and controlled testbed (pictured in Figure 2.2). Note that, to some extent we have already

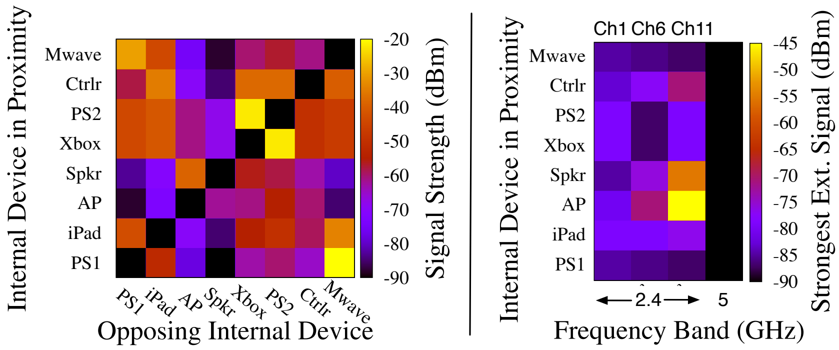

Figure 2.18: We are able to easily map where signals go in the home, providing insight in to signal strengths between devices.

evaluated parts of the monitors accuracy in Section 2.3.1.1.

First, we evaluate our system's ability to collect meaningful and useful information in the training phase. In Figure 2.18, we show a heatmap of the signals at each internal device, from every other device which is collected, stored in our database. It is meaningful and useful: e.g., *Power Sensor 1 (PS1)* is hidden to the *AP*, but not visa-versa; *PS1* is also in strong interference range of the microwave. In our example application (Section 2.5), we will show that this information can be used to predict performance issues, overlaid on a real map to the user, and verified to predict true packet loss. [1]

After training our system, we walk through the environment along the path highlighted in Figure 2.2 with the prototype in our pocket. Through this simple movement with the phone, a set of measurements should be triggered which "update" where signals go and what they interfere with, which (since we did not physically move anything) should match the signals observed in the training phase. Through this movement (which we repeat several times) and the proximity thresholds we derived in the training phase (Section 2.3.1.1), we find that our system consistently triggers measurements when nearby each device to update the device-centric views.

We verify the accuracy of these device-centric views to match our expectation given the phone's orientation, body attenuation, and movement when taking these close-encounter based measurements. For each internal device, we calculate the max, average, and minimum observed error between the device-centric view captured in training, compared to the device-centric view updated when walking. We present the results in Figure 2.19, shown to be bounded within $\pm 5dB$.

In our testbed, 60% of devices obtained user-recognizable names through service discovery protocols, and the 802.11 and ZigBee networks were named using passive packet inspection.

[1] Also note that an external network with high signal strength is observed on channel 11 at the AP, not as strongly observed elsewhere. Had the AP been located elsewhere, an internal channel selection mechanism may have chosen this channel and its clients located in this area (bedroom) could experience interference.

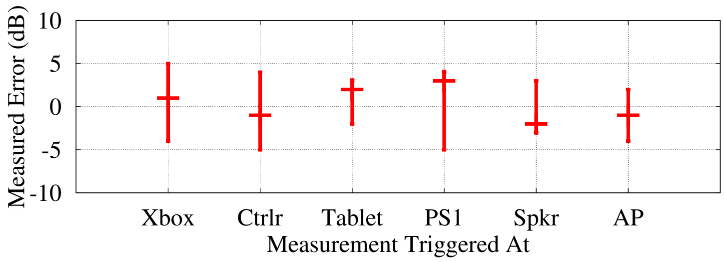

Figure 2.19: Device-centric views are within ±5dB error.

Merge Heuristic	**+ / - Applications**	**Applied To**
Network Level Addressing	+1	AP
Related Link-layer Addresses	+2	AP, Tablet
Radio / Interface Names	+2	AP, Tablet
Technology Type	-1	Power Sensors
Signal Strengths	-1	Power Sensors

Table 2.4: Device abstraction heuristic usage in heterogeneous testbed.

For only 2 devices did we fall back on signal-classification based naming: the microwave and the gaming controller.

We summarize the heuristics used to create the device abstractions in our environment in Table 2.4, which also shows which device types contributed to abstracting. Network-level and link-layer addressing merged the dual-band radios on the AP together, as well as its wired interface, and adjacent MAC layer addresses of the Wi-Fi and Bluetooth of the tablet's radios merged them together. Despite the ZigBee power sensors responding with the same name, we used signal strengths to force them apart.

Finally, we evaluate our system's ability to detect device movement in the environment (§2.3.1.4). To do so, we move each of our fixed devices (i.e., AP, power sensors, gaming console, and speaker) by 1, 2, 5, 10, and 20 feet, walking along the path highlighted in Figure 2.2 after each move. Table 2.5 summarizes the results that show, as long as the device moves more 5 feet or greater, we can detect it has moved with 100% accuracy. Below this distance, the average variation score of signal strengths from other fixed devices does not provide us enough confidence of the move.

2.5 Applying the Information Collected

In this section, we discuss applications of the information collected by our monitor and implement one of these applications: an RF environmental map which reflects device layout that diagnostic information can be overlaid.

Distance Moved (ft)	Accuracy - [FP,FN]	Avg. Variation Score
1	12% - [0%,88%]	2dB
2	25% - [0%,75%]	4dB
5	100% - [0%,0%]	8dB
10	100% - [0%,0%]	17dB
20	100% - [0%,0%]	25dB

Table 2.5: Accuracy of detecting device movement.

2.5.1 Environmental Map with Overlays

To further provide confidence in our claims of possible applications of this information, we build a simple *diagnostic system* with *conflict graphs*, and then overlay this information on a map we generate of the environment to create a *visual aid*. Ultimately, our goals are to: 1) Create an environmental map that reflects true layout of devices in the home, and 2) Overlay information about the RF environment at a level the user can understand on to this map. This includes information about network's coverage, interference, and potential conflicts.

Ensuring the application is suitable and usable for the average home consumer would require a significant user study with a lot more detail and time spent in the human-computer interaction aspects. Our goal is not to claim this, but to simply illustrate the power of our monitor and the information it collects through a sample application, which others could hopefully leverage to make something more user-friendly.

2.5.1.1 Force Directed Environmental Map

Orthogonal to a large amount of work which has attempted to "map" layouts based on relative data (e.g., mapping the Internet using RTT times [65]), our goal is to generate a layout of devices in the home using their relative signal strengths to each other. Like these works, we leverage *force-directed graphing*, where spring-like forces are placed between "connected" nodes (i.e., radios within range) to attract/repulse them towards or away from each other proportional to a distance. Consider the force between nodes X and Z as being denoted $F\{X, Z\}$. If the distance between X and Y is two times the distance of X and Z, then proportionally: $F\{X, Y\} = 2 \times F\{X, Z\}$. In the graph, Y would then be pushed at a distance 2x that of X from Z using such forces.

Challenge in Determining Forces: The key to leveraging the force-directed graph is determining appropriate forces. Assume that a device-centric view of a radio Z shows signals received from two radios, X and Y, at the same strength: -75dBm. One might immediately consider using such received strengths as the forces: $F\{X, Z\} \leftarrow 75$ and $F\{Y, Z\} \leftarrow 75$. This would place X and Y at the same relative distance from Z. However, for two key reasons, *such a simple assignment of forces does not work in heterogeneous environments.*

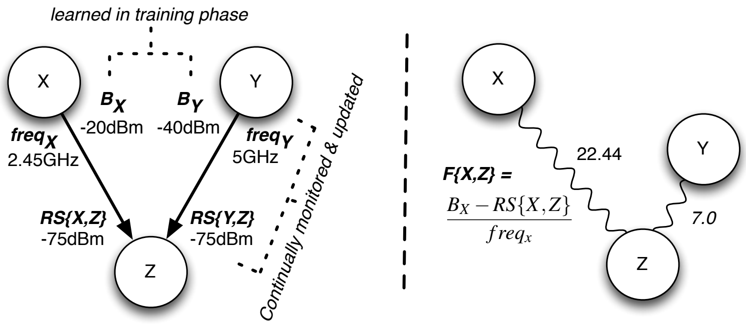

Figure 2.20: Deriving spring force towards RF environmental map.

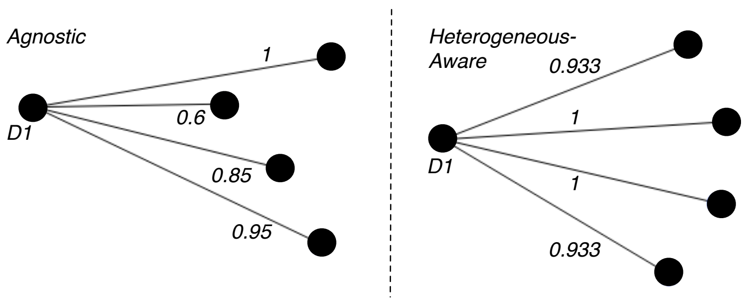

Figure 2.21: Measurement-based force-directed examples showing the importance of accounting for heterogeneous properties in springs.

First, this simple approach is agnostic to differences in transmission powers. If X has a transmission power that is *2x* that of Y, then there is a strong probability its distance is significantly greater than Y to Z. Second, the approach ignores differences in propagation at different frequencies. If X is operating at $2.45GHz$, whereas Y is operating at $5GHz$, their relative distances are likely to be even further.

We refer the reader to Figure 2.21 which illustrates the impact of not accounting for these properties (*agnostic*) vs. accounting for them (*heterogeneous-aware* – which we will describe next). We place 4 devices *at the same distance* from a device *D1*, meaning their spring forces should be equal from *D1*. We normalize the forces to the greatest calculated force, and generate the two force-directed graphs. As shown, ignoring transmission powers and frequency can place the devices at very different distances (e.g., at 60% its distance), accounting for them more accurately reflects distance.

Deriving Appropriate Forces: To account for differences in transmission powers, we can leverage the signal strengths observed when the phone is placed directly next to each internal device in the training phase (§2.3.1.1). These measurements provide a baseline power

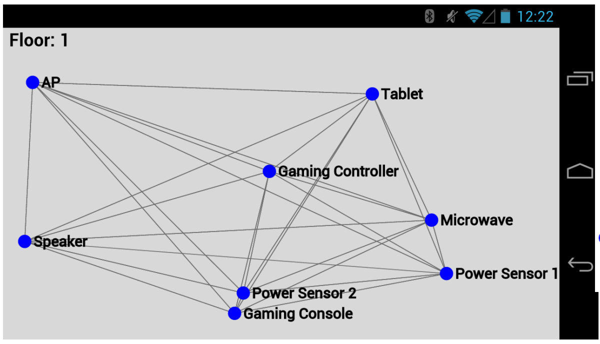

Figure 2.22: Our force-directed model and spring forces reflect true layout.

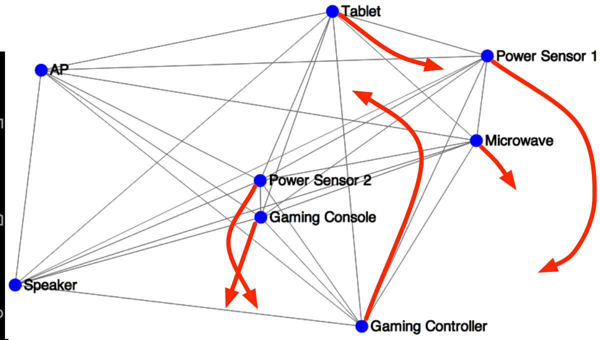

Figure 2.23: Many devices end up out of place when using basic and agnostic spring forces.

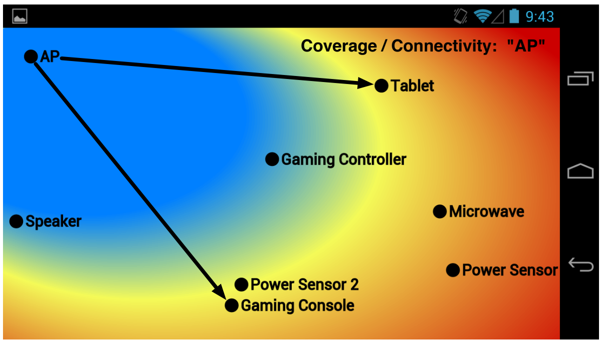

Figure 2.24: Overlays such as this communication/coverage overlay can be easily understood.

as observed directly at the device, and is information unlikely to be collectable by a sparse deployments of fixed monitors.

Therefore, temporarily ignoring difference in propagations due to operational frequency, we can account for transmission power and derive an initial proportional (one-way) force from radio X to Z as: $F'\{X, Z\} = B_X - RS\{X, Z\}$, where B_X denotes the baseline strength observed for a radio X. Referring the reader to Figure 2.20, this would apply the following relative forces: $F'\{X, Z\} = 55$ and $F'\{Y, Z\} = 35$, making X's force greater due to its higher transmission power.

To account for propagation at different frequencies, we can apply operational frequency as an inverse scaling factor to the spring force F'. It is inverse since signals at lower frequencies will degrade less over the same distance as a higher frequency signal (meaning lower frequencies should be pushed further given the same path loss as a higher frequency transmitter). From this we can derive the spring force between two radios X and Z as: $F\{X, Z\} = \frac{B_X - RS\{X,Z\}}{freq_x}$.

Using this equation, we illustrate an example derivation of the forces between the radios shown in Figure 2.20. We will later show that this provides a reasonable proportional force between heterogeneous radios that will reflect true layout.

2.5.1.2 Overlaying Information & Diagnostics

Using the RF environmental map that we generate, there are many types of diagnostics and general information that can be overlaid to inform the user. This mainly involves an API (available on our prototype) to draw lines between devices on the map and shade regions of it. Using this basic functionality, we briefly describe several functions we implement to present environmental information to the user:

Using the signal strengths, knowledge of their heterogeneity, and their operational frequency we can predict various types of *connectivities issues*, drawing arrows to notify of a potential issue. By selecting a device on the map, we show *coverage ranges* by shading areas of the map and/or the colors of each device to show that device receives a strong or weak signal from the device selected. Similar to coverage ranges, but using inverse colors to show *interference ranges*. Using network-level information, we also draw *communication patterns* between devices using lines, where line thickness can reflect load on each network link.

2.5.1.3 Illustration of Our Map & Overlays

We use our testbed environment to illustrate our environmental map and a few of our overlays:

<u>Environmental Map:</u> Using our heterogeneous-aware force-directed graph based model which uses information collected in our database, we generated the following layout of our environment: Figure 2.22. We ask the user to compare this layout to the map of our environment in Figure 2.2. As shown, *our force-directed graph based model is able to reflect the true physical layout of devices in our testbed.*

For comparison to show the importance of taking heterogeneous properties in to account, we use the same force-directed model and information, but calculate the forces using the basic agnostic model. We illustrate the result in Figure 2.23. Our red arrows illustrate where the devices should be placed to reflect true layout, showing significant error and the importance of our heterogeneous-aware derived forces.

<u>Coverage Overlay:</u> Using the environmental map as a basis to overlay information, we illustrate a coverage/connectivity overlay for the AP, illustrated in Figure 2.24. This was generated purely using signal strength information in the database from our device-centric views. We verify and illustrate this overlay is accurate by taking throughput measurements in various locations, presented in Table 2.6. These measurement results show that the coverage map reflects true performance e.g., when taken in a yellow area, performance is mediocore (23% of the potential throughput). This information could guide the user on how to accurately move devices (e.g., the AP) to improve connectivity / performance in certain areas.

<u>Connectivity Issues Overlay:</u> Using information from the database such as device's active frequencies, their spectrum usage, and their signal strengths to and from each other, we are

Measurement Location	Coverage Color	Observed Perf.
Nearby: *AP*	Blue	170Mbps
Nearby: *Gaming Controller*	Light Blue	80Mbps
Nearby: *Tablet*	Yellow	40Mbps
Nearby: *Power Sensor 1*	Red	5Mbps

Table 2.6: Coverage overlay estimates performance can guide the user towards better placement.

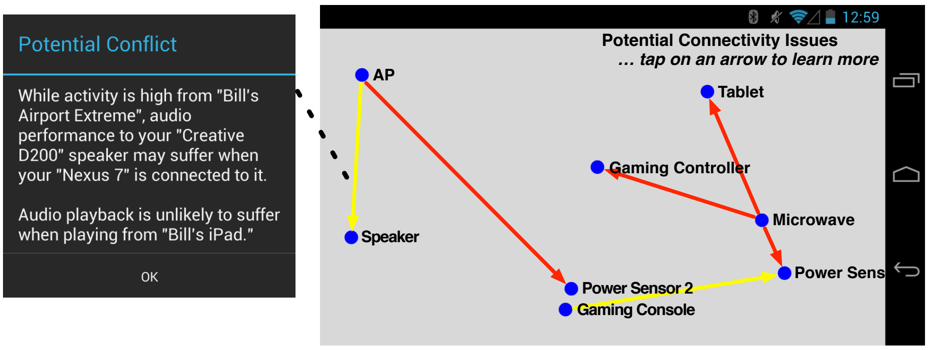

Figure 2.25: Potential connectivity issue graph derived by our heterogeneous home monitor, overlaid on our force-directed graph layout.

able to build a simple diagnostic tool and overlay to detect / display potential connectivity issues. We leverage measurements taken in prior work (e.g., in [7,8,15,16]) to create a "look-up" table within our application of expected (and potentially destructive) behavior between pairs of technologies given their interference levels. Arrows are drawn in the overlay to denote potential connectivity issues, and the user can tap on any arrow to receive meaningful information.

We apply this overlay to our environment and illustrate it in Figure 2.25. Our system detected several potential issues, which we validate through measurement. First, our lookups on the ZigBee and 802.11 information (provided by [16]) suggested that *Power Sensor 2* and the *Gaming Console* were so close that both would defer to each other and therefore, in our overlay, no connectivity issue arrow was drawn. We confirmed this to be true through measurement. The gaming console, however, was flagged to potentially interfere with power sensor 2, which through measurement, we found 23% loss rate when the console was active (e.g., streaming).

The interference of the *Microwave* was also found to be high in the device-centric view of *Power Sensor 1* stored in our database, creating an arrow between them (in addition to the tablet and gaming controller). Since each device-centric view captures the spectrum usage of each device, our database has information about where the microwave leaks power the most, illustrated in Figure 2.26. This information guides our application to flag nearby devices

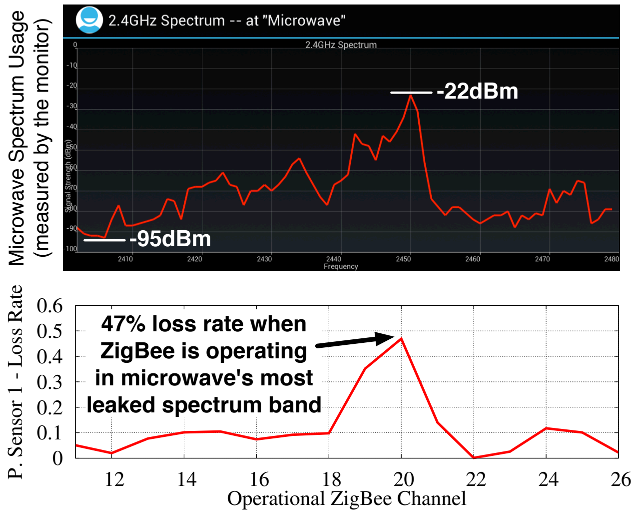

Figure 2.26: Measurements of spectrum usage and interference levels by our monitor reflect observed loss rates and can guide management.

to the microwave if they operate in the areas the microwave leaks most. We confirm this behavior by operating the ZigBee network on all possible channels and plotting its loss rate in Figure 2.26 (aligned with the microwave's spectrum usage).

Finally, our application flags a potential connectivity issue between the AP and Bluetooth speaker by observing different coexistence behavior of devices when paired to the speaker. Our database showed that when an Apple device was paired to the speaker it actively avoided the internal 802.11 network's channel, whereas a Samsung device paired to that same speaker did not – illustrated in Figure 2.27. When the AP is active we observed holes and pops in the audio playback from the Samsung device, not the Apple device.[1]

2.5.2 Applications

There other potential applications to leverage the information collected by our monitor, stored on the phone. However, we remind the reader that the accuracy of our measurements will be within $\pm 5dB$. While this makes some applications infeasible due to the accuracy, there are still many possible applications of this information that we summarize:

- *Heterogeneous Conflict Graphs*: As illustrated by the heatmap generated by our monitor (Figure 2.18), the information can provide several key insights in towards detecting conflicts and generating a conflict graph. First, the information shows which heterogeneous devices are within interference range of each other that do not coordinate (suggesting a MAC conflict). The information also shows hidden terminal situations,

[1]We provide sound samples here: http://tinyurl.com/asr5qep and http://tinyurl.com/apdt5nu

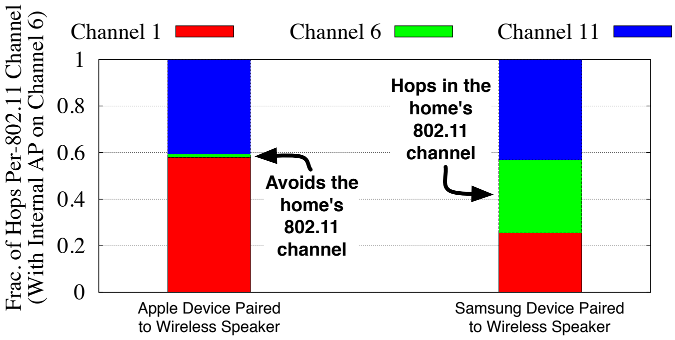

Figure 2.27: Coexistence capabilities can be observed by our system (e.g., between Bluetooth and Wifi); shown to vary by device.

particularly asymmetric hidden terminals created by disparities in transmission powers. Conflicts can be drawn between links that are comprised of these terminals. Finally, digital information collected about heterogeneous networks can provide insight in to their parameters, for example to determine one CSMA network is digitally sensing whereas one is using power-based sensing.

- *Spectrum management:* After creating a heterogeneous conflict graph with this information, spectrum management can be performed improve the heterogeneous environment. First, conflicting heterogeneous networks can specifically be separated in the frequency domain to eliminate conflicts. In more dense scenarios where all networks can be isolated, spectrum management can suggest other frequencies for the networks that will reduce their amount of interference received. For example, by placing them with another heterogeneous network that is active less frequently. TDMA networks could also use this information to know which channels they should avoid hopping to. This general approach of spectrum management from this information is the goal of our work presented in Chapter 3.

- *Coexistence techniques* that do not require extremely fine-grained signal-level information can use the data collected by the monitor. For example, BuzzBuzz's coexistence technique between 802.11 and ZigBee could use this information to enable their transmission of the ZigBee header twice only when needed [16]. Subcarrier suppression-based coexistence in the ISM bands could also be achieved without the need for "poking" or deliberately interfering with other networks [42], instead using the information to better estimate the amount of suppression needed. In general, the information coupled with generated conflict graphs can dictate when to enable or disable coexistence techniques, a key challenge in today's heterogeneous environments where the heterogeneous networks cannot sense each other to know when coexistence should be enabled or disabled.

- *Diagnostics:* Similar to using the information to generate conflict graphs, the information can be used to perform diagnostics (e.g., like WifiNet [15]). That is, to periodically use the information to determine where specific faults or inefficiencies are coming from in the network that pertain to heterogeneous interference. For example, Jupiter's research finds that 67% of all residential Wi-Fi problems are linked to interfering devices, such as cordless phones, baby monitors, and microwave ovens [66]. It is with this information and our tool that diagnostics can be done to determine where sources of interference come from to guide the user.

- *Visual aids:* In line with application of diagnostics, the information taken from the environment such as what devices are causing interference problems from each other can be taken and visually presented to users. That is, the meta-information generated from the data our monitor collects can be used to build tools to help bring the relatively unknown RF environment up to a level the user can understand (e.g., with device abstractions). In particular, this is possible given our monitoring system built on the smartphone which already provides a familiar and flexible interface. This can be used to draw maps of the home environment and overlay information about interference or conflicts within it.

2.6 Related Work

In this section, we provide an overview of related work to our home monitoring system and evaluation. This work is broken down in to 4 main categories: 1) Heterogeneous Monitors, 2) Monitoring systems, and 3) Heterogeneous interference studies and diagnostics, which are related to our interference study and home RF map with diagnostics.

Heterogeneous Monitors: At the base of the problem of heterogeneous networks has been the inability to collect detailed information about heterogeneous networks and signals in the environment. This motivated many commercial products that provide spectrum sensing such as Wi-Spy [67], AirMagnet [68], Bandspeed AirMaestro [69], and Cisco's Spectrum Expert [70]. These devices can provide spectral-power based views that allow users to visually classify signals without logic or automated scanning or signal classification techniques. Therefore, these systems require significant manual work to walk through an environment and determine where signals go and their strengths. Additionally, without any form of low-level signal analysis, it is possible to miss lower power signals within range. Ultimately, acting like heterogeneous spectrum analyzers with interfaces that guide the user to locate various types of devices in the environment.

With the rise of low-level signal access available through software-defined radios, RFDump [59] and DOF [60], and others [71,72,73,74] were able to introduce more complex signal analysis techniques towards accurate and automated scanning and detection. In particular, RFDump [59] and UCS [74] analyze a signal's power, timing, phase, and frequency usage to classify it in the spectrum. DOF [60] improved accuracy through cyclostationary signal analysis, i.e., it builds on hidden repeating patterns in signals that can be used to construct unique signatures. This allows DOF to accurately estimate signal types and their spectral and spatial parameters. While powerful, these approaches require expensive software-defined radio equipment ($1000+) which make them difficult to deploy in practice (motivation for our smartphone-based monitor).

More recently, AirShark [7] has shown that low-level signal information can be extracted from newer 802.11 hardware radios to perform similar signal classification techniques as RF-Dump and others. Through their work, a commodity 802.11 radio can now detect microwaves, cordless phones, gaming controllers, ZigBee networks, and many other types of signals. Note that AirShark is orthogonal to our work. It can be built in to our smartphone-based monitor to allow it to collect information about many other types of signals. Unfortunately, AirShark was not publicly available at the submission of this thesis.

Monitoring Systems: While the heterogeneous monitors we have just described can detect signals in the environment, they require significant manual labor to map where signals go in an environment and what they interfere with over time. To overcome this manual and time-consuming method, popular approaches have been to deploy dense centralized monitoring infrastructures with multiple fixed monitors [15,54,55,75], and robots that traverse the environment (e.g., Roombas) with monitors attached to them [76,77].

In particular, early spatially-aware monitoring systems such as JigSaw [55] and DAIR [54] map where signals go in the environment through a centralized deployment of multiple fixed monitors/sensors. These works targeted the enterprise environment where cost and complexity were not concerns. By deploying multiple monitors in the environment, every event (i.e., a transmission) could hopefully be observed by at least one sensor (an 802.11 packet-level radio). By co-locating the sensors with APs or placing them in specific office locations, events are localized to the static and known location of the sensor that received it strongest. By localizing, collecting, and intelligently synchronizing the events at a central location, one can obtain a global and spatial view of the environment. From this, coverage and interference ranges can be inferred to generate a conflict graph and plan the environment.

These earlier 802.11-based dense monitoring infrastructures motivated WifiNet [15], a heterogeneous monitoring system for enterprise environments that could detect heterogeneous devices in the environment and their impact on the 802.11 network. Similar to DAIR and JigSaw, WifiNet deployed multiple 802.11 radios throughout the environment that act as monitors, but are equipped with AirShark [7] functionality to detect heterogeneous signals and their impact on the enterprise 802.11 network. This approach, however, is Wi-FI centric, i.e., it only monitors heterogeneous signals at the 802.11 access points. It does not attempt to measure heterogeneous signals comprehensively between all devices.

Heterogeneous Interference Studies and Diagnostics: Finally, there have been many studies of heterogeneous interference, and subsequently diagnostic mechanisms and coexistence techniques driven by them. For example, interference studies between 802.11 and ZigBee interference [7,16,34], ZigBee and Bluetooth interference [38], cordless phones and 802.11 [8,15], 802.11 and Bluetooth [33,40], and other general studies (e.g., [78,79]). Many of these studies show the severity of heterogeneous interference in different environments. As a result, there have been many diagnostic techniques and heuristics to detect when heterogeneous interference will be most severe (e.g., [15,79]). For example, Spectrum MRI [79] proposes a multi-radio interference diagnosis framework with the aim of isolating and classifying multi-radio interference problems using heuristic and model-based methods.

Related Work Summary: While there have been many proposed heterogeneous monitors and systems proposed, our work is the first the provide a practical and deployable solution for the home. In particular, it overcomes the cost and complexity of dense monitoring infrastructures. Additionally, our monitor is comprehensive: it determines the signal strengths between all heterogeneous devices rather than being Wi-Fi centric like WifiNet [15].

2.7 Chapter Summary

In this chapter, we explored the possiblity of leveraging the smartphone to create a heterogeneous monitor for the home. We argued that the phone had many beneficial properties that made it an attractive platform to develop a heterogeneous on top of, including its multiple heterogeneous radios, access to packet-level information, and already familiar and flexible interface. We then presented the design of a practical and usable home monitoring system, based on the smartphone.

Our monitoring system is able derive where signals go in the home (i.e., their strength at various locations) with little user involvement, and without the cost and complexity of multi-sensor deployments. To do so, we introduced a novel and initial design that was practical by balancing the level of information requested from the user, with the amount of complexity introduced in the system to infer information that the user does not provide (e.g., has a device moved?). The design introduced 3 main phases to achieve this balance: 1) Training, 2) Monitoring, and 3) Diagnostics.

Our system introduced useful heuristics to take the heterogeneous RF signals in the environment and abstract them in to user-friendly device abstractions. These abstractions are also useful for more accurately depicting the environment (i.e., which heterogeneous signals come from the same device, and therefore coordinate). The information our monitor collects can be used to implement various applications, as illustrated by our force-directed graph based map of the home with diagnostic overlays of the heterogeneous RF environment. example with overlays.

Limitations and Future Research: Today's heterogeneous home environments are becoming more and more dense, with new devices that create new types of conflicts and introduce new signals in to the environment. Being able to monitor and address these environments is critical to the future of their efficiency and performance. This chapter presented work to properly monitor the environment and provide the necessary information to address connectivity and interference issues within it.

Although we presented a single design of a smartphone-based home monitoring system, there are many other possible designs based on the smartphone, as well as limitations introduced by our particular design and work. From this, we outline a few major areas that need to be studied in the future to potential make our system more accurate and usable.

- *Multiple-phone Design*: In our work, we presented the design of a system that was based on a single smartphone collecting information about the home's heterogeneous RF environment. The information collected is therefore based on the interactions of this single phone in the environment (i.e., where that phone goes, and particularly what it comes in close contact with). There are, however, *many* potential smartphones in a single home environment and, in particular, those phones go to different locations of the home and may interact more closely with different devices. For example, phones that belong to particular members of the family will likely be more active in their particular bedrooms. This opens up the possibility of a multi-phone design that introduces new challenges, and potentially more rich information. Now, multiple monitors must be able to coordinate and share information between each other. Is there a "master" monitor in the environment, or is it more distributed? Does correlating events across the phones help? Clearly, new challenges and design questions will arise.

- *User Involvement vs. Complexity*: Although we presented a design that requires a rather significant amount of user involvement in the training phase, many will argue that designs are possible that require absolutely no user involvement. This would be equivalent to a service running on the phone that collects information about the environment over time without any user involvement. While seemingly challenging, such a design is likely possible and considered future work. For example, how does the monitoring system learn which wireless devices belong to the home user? It may be possible that the training phase instead takes days instead of minutes, monitoring which devices it comes in close contact with multiple times, and assumes that these devices belong to the user. There may also be ways to diagnose and reconfigure the environment without involving the user, also. For example, by having the phone generate interference on specific devices to exploit internal coexistence mechanisms in them and get them to change frequencies.

- *Other Spectrum Bands*: Given our use of the smartphone as the base of our monitoring system, we are limited to only being able to monitor certain spectrum bands. While this may change over time (e.g., as phones now have 5 GHz support on top of 2.4 GHz), there are still a few spectrum bands that wireless devices use in the home that may not be supported by smartphones like the 900 MHz band or 60 GHz band. Additionally, it may take time for phones to adopt white space spectrum bands. It may be possible, instead, to augment the phone with small attachable devices that enable their support of other bands when trying to diagnose specific problems. That is, not always having access to these bands, but giving the user an opportunity to monitor them specifically when problems are detected.

Chapter 3

Spectrum Planning for Heterogeneous Wireless Networks

As spectrum use becomes increasingly heterogeneous, wireless network performance can suffer greatly due to the destructive interactions of heterogeneous networks and the interference between them. In the previous chapter, we presented a novel smartphone-based heterogeneous monitoring system that is able to collect the necessary information about an environment to detect heterogeneous conflicts within it. While this information can be useful to both coexistence techniques and spectrum management, we have argued throughout this dissertation that the general coexistence-based approach requires N^2 solutions, and these solutions are difficult to deploy due to their complexity, overhead, and the rate at which each technology changes (requiring new solutions over time). Instead, spectrum management can provide an efficient and general solution that does not require changes with technologies, changes to the endpoints (e.g., at the MAC or PHY), or per-packet overhead.

While spectrum management is a promising approach for dealing with heterogeneity, there are still critical challenges involved in performing spectrum management for heterogeneous networks. These challenges are driven by limitations of current work on spectrum management in what we consider to be 3 key components:

1. *RF environmental models* that represent the networks, radios, and links lack the rich structure needed to represent the diverse properties of heterogeneous networks in the environment needed to perform proper spectrum assignment. For example, what protocols are used by each radio, whether two radios coordinate, or what bandwidths and frequencies each radio supports (e.g., [80,81]).

2. *Spectrum assignment algorithms* that determine the frequencies for each network using the RF environmental model (or, a conflict graph derived from it), which make assumptions about the networks and devices within range sharing a common technology

or standard. For example, they assume that two radios on the same channel and in spatial range will at least coordinate (sharing airtime), and that all radios have the same possible set of channels. These assumptions lend well to coloring algorithms that have been historically popular in wireless spectrum assignment (e.g., [56,81,82]), however, graph coloring does not lend well to heterogeneous environments, as we will later explain in detail.

3. *Predictive channel quality metrics* are at the base of many spectrum assignment algorithms to estimate the performance of a network if operating on a specific channel (e.g., [15,21,56,58,81]). This allows the assignment algorithm to estimate the performance of *many* different configurations without needing to deploy and test them. Unfortunately, current metrics either assume all radios and networks are homogeneous (i.e., they all share the same protocols) [21,56,58,81] or the metrics are Wi-Fi centric [15]. That is, the majority of predictive channel quality metrics will either assume fair-sharing of networks on the same frequency in spatial range (ignoring heterogeneity), or they will be Wi-Fi centric in the sense that the metric and algorithm only predict interference from heterogeneous networks on a Wi-Fi network and reconfigure it.

In this chapter, we present novel contributions in all 3 of these key areas, leading to proper and efficient spectrum management for heterogeneous networks. In particular, our contributions are not *not* Wi-Fi centric, and their structure is meant to be general enough to support various heterogeneous technologies as they evolve over time. The details of our **contributions** in this chapter and these 3 areas are as follows:

First, we introduce a novel *hypergraph-based RF environmental model* that is able to represent rich information about the differences between heterogeneous networks and technologies. In the graph, radios are vertices and they can be connected by 3 unique edge-types: 1) Link edges, 2) Spatial edges, and 3) Hyperedges. Link edges and spatial edges are uni-directional edges that denote communication between two radios, and that one radio is within communication/interference range of another (respectively). Hyperedges connect all radios that belong to a network to inform the assignment algorithm that the radios must be configured to operate on the same frequency. Finally, radios, links, and spatial edges have associated meta-data provided by the monitoring infrastructure (e.g., our smartphone-based monitor) to provide the possible frequencies for each radio, signal strengths at each device, and whether radios coordinate.

Second, we present a new predictive channel quality metric that considers heterogeneity between networks and devices. The metric estimates the expected airtime of a radio on a particular channel by: 1) Accounting for its fair share of airtime from networks it coordinates with, and 2) Degrading this expected airtime due to interference from heterogeneous networks. The degradation is calculated using fundamental properties of the radios such as their airtimes, and whether both radios are unable to coordinate with each other, or whether at least one is able to coordinate (i.e., an asymmetric scenario).

Finally, we feed the hypergraph-based model of the particular environment in to *mixed-integer program (MIP) based spectrum assignment algorithm*. The algorithm uses the constraints given in the model (e.g., the possible frequencies of each radio), decomposes the hypergraph in to a series of conflicts (similar to a conflict graph), and uses HCE to find efficient organizations that reduce interference from heterogeneous networks.

Note that while our previous chapter focused on the home, there is *nothing* specific about the model, metric, or algorithm to the home environment. One could use what we propose in this chapter within other environments, also.

Chapter Outline: We begin this chapter with a brief background on spectrum management and the limitations of current practices in assignment with regards to heterogeneous environments (Section 3.1). We use these limitations to motivate the requirements in spectrum management for today's environments, and then present our principles of design and approach to meet these requirements (Section 3.2). With these principles in mind, we present our system design and components (Section 3.3), followed by an evaluation of our system (Section 3.4). We then conclude with a summary of our contributions, limitations, and future work (Section 3.5).

3.1 Background and Limitations of Current Practice

In this section, we provide a brief overview of spectrum management and its key components (Section 3.1.1). With an understanding of these components, we then present the current practice in spectrum management and its limitations towards supporting general heterogeneity (Section 3.1.2).

3.1.1 Spectrum Management and its Key Components

The goal of spectrum management is to (re)organizing networks in the frequency domain to achieve some objective function. For example, to minimize the overlap of networks in the spatial and frequency domains [83], or to prioritize frequency assignments based on the traffic loads of networks in the spectrum [57,84,85]. The goal of our work is to organize the spectrum in a way that improves performance by minimizing interference between heterogeneous networks and devices.

From studying prior work (e.g., [56,57,81,86,87]), we have found that many spectrum management systems contain a similar structure, shown in Figure 3.1. First, there is an underlying monitoring system that provides information about the environment (e.g., wireless devices in range, and their signal strengths). Then, within the management system, there are typically **3 main components**:

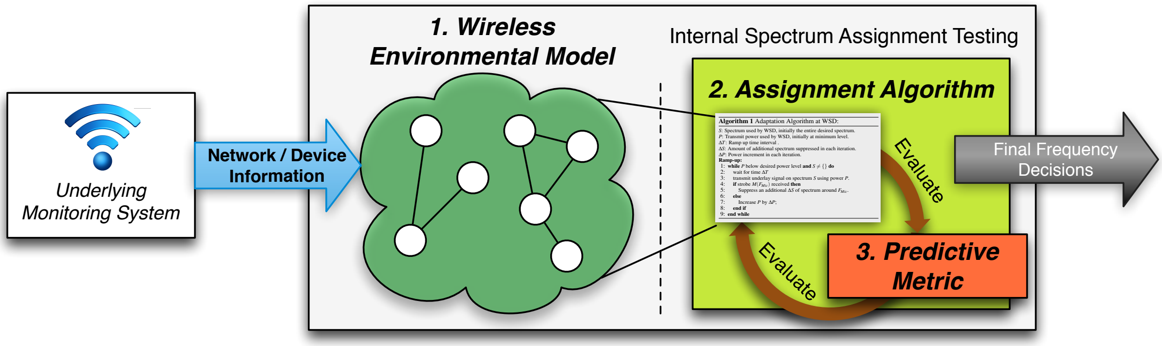

Figure 3.1: Many works in spectrum management contain a similar structure, as shown here with an underlying monitoring system and 3 key components to analyze and optimize the spectrum assignment.

1. An *RF environmental model* that takes the raw information from the monitor, and provides structure to it. This model typically provides constraints in the system, as well as meaningful information about the interactions of the devices through analysis. Examples include a conflict graph that reprints radios that my interfere [80,88,89,90], a weighted graph to convey traffic load and spatial overlap [81,91], network topologies using cliques [92], and graph-based topology including interference and coverage ranges [55,93].

2. An *assignment algorithm* evaluates different frequency-domain configurations of the networks and devices, given the constraints and information provided by the environmental model. The goal of the algorithm is to search the space of potential configurations (an NP-hard problem) to find configuration that reduce conflicts, contention for the medium given traffic loads, and reduce loss rate due to interference. To do so, many approaches have used graph coloring algorithms [56,83,92,92,94], weighted graph coloring algorithms (e.g., to consider traffic load or loss rates) [81,91], simulated annealing [57,95], (mixed) integer programs [96,97,98], and even genetic algorithms [86].

3. To evaluate each of the potential configurations without actually reconfiguring all of the networks (impractical), a *predictive quality metric* is typically used. This metric provides the algorithm with a predicted outcome of when assigning the networks a particular set of frequencies. Examples of such metrics include estimating airtime for each network given fairshare and the residual [21], predicting sustained interference [58, 81,91], and resulting throughput given the traffic loads [57].

Using these 3 key components, the spectrum assignment system provides a set of frequencies that it predicts will provide the best possible outcome in terms of performance. It then provides this final set of frequencies to the network administrator or remote management system to reconfigure the end devices.

3.1.2 Limitations of Current Practices

As we briefly discussed in the introduction, many environmental models, spectrum assignment algorithms, and predictive metrics fail to meet the heterogeneous requirements of the spectrum today. We quickly highlight three key shortcomings of prior work here, and provide a more comprehensive overview of related work in Section **??**.

First, a significant amount of prior work in spectrum assignment is *homogeneous* and predominantly 802.11-based. For example, work by Rozner et al. [57], Akl et al. [99], and Murty et al. [84,85] all assume homogeneous properties while assigning spectrum. That is, they assume overlapping networks will at least coordinate, and that they can all be assigned in the same manner (i.e., they have the same possible set of channels). Even the recent channel changes in 802.11n and 802.11ac would likely require significant changes to these solutions to support the newer standards. However, some of this work does consider different traffic loads (i.e., application-layer requirements) [57,84,85].

Second, more recent work that considers heterogeneous technologies is Wi-Fi centric, i.e., the goal is reconfiguring and optimizing a Wi-Fi network to avoid interference from networks using other technologies (e.g., WifiNet [15]). Their predictive quality metric and framework are only meant to estimate heterogeneous interference on an 802.11 network and reconfigure it. It is not comprehensive in the sense that it does consider or predict interference between all possible heterogeneous radios in the environment. In addition to being Wi-Fi centric, their work also does not provide a concrete algorithm on how to reorganize the spectrum (even the 802.11 network) to avoid the heterogeneous interference estimated.

Third, work that comprehensively considers heterogeneous networks (i.e., it is not Wi-Fi centric) continues to make *overly simplified assumptions* about the networks and the RF environments. For example, work by Peng et al. [56] and Sooyeul et al. [97] make similar critical assumptions untrue of environments with heterogeneous technologies. Both assume that conflicts all have the same weight, i.e., interference from one device is just as severe as from another device (i.e., it is a binary conflict graph). Clearly, this is not true in practice. For example, cordless phones have been shown to reduce 802.11's throughput to near zero [8], whereas ZigBee networks have a lesser impact (e.g., around 60-70% [16]). Given the density of the spectrum, binary conflicts will likely not lend well to efficient configurations. One must know not only that a conflict exists, but also how severe the conflict is to efficiently organize the spectrum of environments with heterogeneous networks.

Additionally, this work and others (e.g., [83,86]) incorrectly assume that all radios use the same channels, i.e., center frequencies and bandwidths. This is a critical assumption that is not true in heterogeneous environments: they have different center frequencies, different bandwidths, and even different spectrum band capabilities (e.g., 2.4GHz vs. 5GHz). We believe that this assumption is made to allow the problem to more easily be reduced to a variant of graph coloring (one of the most popular assignment algorithms). We find that without these simplifying assumptions (e.g., of unified channels), heterogeneous environments will likely be more difficult to reduce to basic variants of graph coloring.

More importantly, we found graph coloring to be overly restrictive when trying to formulate and model an environment and spectrum with heterogeneous networks. For example, the most basic form of graph coloring will consider conflicts to be binary in weight. Weighted graph coloring has been proposed and used to reflect different amounts of interference from partially overlapping channels (e.g., by Mishra [81] et al.), but we found it difficult to try and capture the various degrees of back-off, interference, and asymmetry in a single metric. As we will show in our algorithm, there are various estimates of coordination and interference that are considered to efficiently and properly reconfiguring the spectrum. The mixed integer program representation also easily allows one to introduce multiple types of constraints (e.g., finding a solution where a particular network has no interference, or a solution where a network has an expected level of performance).

In summary, prior work has been insufficient in characterizing and organizing true heterogeneous environments. The majority of this work is homogeneous, Wi-Fi focused, Wi-Fi centric, or overly simplified in assumptions about heterogeneous environments. The goal of our work is to overcome these limitations and better organize heterogeneous environments.

3.2 Requirements and Principles of Design

With a better understanding of spectrum assignment and the limitations in current practice, we present the design requirements needed to support general heterogeneity and the trends of diversity in today's spectrum (Section 2.1.1), followed by the key principles of our design to satisfy these requirements (Section 3.2.2).

3.2.1 Design Requirements To Support Spectrum Trends

To support the trends in the spectrum that we have discussed throughout this dissertation, the spectrum management system must **support general heterogeneity** between networks and devices, and it must **accommodate evolution** of the protocols and bands over time. Given the limitations we discussed in the current practice, it is important to ensure that the design is *not* Wi-Fi centric, and that it does not make overly simplifying assumptions about the bands, protocols, or channels.

Supporting general heterogeneity between networks and devices is the key to properly organizing the spectrum today. This means that the components in the spectrum management system must represent and account for aspects of diversity across the PHY, MAC, and application-layer:

- PHY Layer: The system must support diversity in terms of the potential bands supported by each radio's physical layer, their center frequencies and bandwidths, and the propagation characteristics based on different transmission powers (i.e., the components must support asymmetric spatial properties).

- MAC Layer: Different access schemes used to coordinate the spectrum must be supported, and importantly: it must be possible to represent each pair of radios in the environment and whether they coordinate based on their MAC properties. Like the need to support asymmetric spatial properties at the PHY layer due to varying transmission powers, it is important to support asymmetric coordination at the MAC layer. Different MAC layers and/or settings can lead to one radio coordinating with another, but not visa-versa.

- Application Layer: Like prior (albeit homogeneous) work has shown, it is important to consider application layer properties when organizing the spectrum, e.g., traffic load, desired throughput, or a tolerance to loss when organizing the spectrum [84,85].

Accommodating evolution ensures that, as the protocols, standards, and spectrum bands change over time, the spectrum management system does not require significant changes. The system should be able to support new protocols and new potentially bands without major changes to the model, algorithm, or predictive metric. It would be hard to argue that any single design could support complete evolution of such a complex system, however, we can certainly study how the protocols and bands have evolved up to today, and ensure that similar evolution can be supported.

3.2.2 Our Principle of Design and Approach

The requirements that we have presented highlight the many diverse properties our components must support to accommodate heterogeneity and evolution. Our basic principle of design to meet these requirements is to *describe, represent, and organize the environment* ***using fundamental properties*** *of the spectrum and protocols,* ***remaining protocol independent*** *where possible.*

Using fundamental properties: To accommodate heterogeneity and evolution, we design each component of the spectrum assignment system based on fundamental properties of the spectrum and its protocols, *not* specifics of protocols, standards, or spectrum bands. This ensures that many different protocols and bands can be described using the same basic properties, and that new protocols can be supported by using these fundamental properties (evident in the majority of protocols).

The best way to describe this and its importance is through example. For example, the components make no assumptions about "channels" which are specific to standards and technologies. Simplifying assumptions about channels lead to some of the problems in prior work.

Instead, our work breaks this PHY-layer property (i.e., channels) in to two fundamentals: 1) A set of center frequencies each radio supports, and 2) An associated bandwidth with these frequencies. This removes specifics of protocols and spectrum bands, while maintaing the system's support of heterogeneity and evolution. If two radios support different spectrum bands, this is represented by their possible set of center frequencies. If in the future a new spectrum band supports unlicensed access, the new frequencies only need to be added to the description of the radios that support it. No changes to the model, algorithm, or metric are needed with this evolution.

Likewise, our system does not describe the many specific details of the PHY and MACs. For example, the details of modulations or whether an 802.11n network is operating in green-field mode or not (i.e., a specific to a standard). Instead, as another example of a fundamental property, it breaks these properties down in to the fundamentals that matter towards spectrum management: does network X coordinate with network Y given its properties at the PHY and MAC? The model and algorithm are built on this simple fundamental property which many underlying monitoring systems provide (including our smartphone-based monitor and others [15,55]).

These examples should have provided the reader with a better understanding of how basing our design supports heterogeneity and evolution. There are many other fundamental properties that our system is built on, however, described in the next section.

Remaining protocol independent *where possible*: While we design the majority of our components to only use fundamental properties in support of heterogeneity and evolution, we believe that there are some particular areas of components where being overly generic sacrifices accuracy or efficiency in assignment. In these areas, we believe that specifics should be used to improve accuracy. If the specifics are not available: the system should provide a generic and/or reasonable mechanism to estimate and represent them.

A key example of this is in our predictive channel quality metric, which we will later describe further in detail. However, this metric is meant to estimate how a device will perform in the presence of other heterogeneous devices (e.g., in another channel). Clearly, how often two heterogeneous devices' transmissions will overlap in time will be dependent on the characteristics of their traffic, and how often these overlaps lead to a loss will be dependent on *many* factors. For example, what modulations are used by the specific technologies, which results in different SINR properties. Ignoring these specifics will only lead to inaccuracies and inefficient assignments. Therefore, in these particular areas, we require specifics but make sure that these specifics do not sacrifice evolution and support of general heterogeneity. Our system provides reasonable estimates of these specifics and flexibility to add additional (or more accurate) estimates. Additionally, we allow the monitoring system to directly provide the information if it is available, e.g., if it knows exactly how often two heterogeneous devices attempt to transmit at the same time. Note that this is an improve over prior systems which choose to ignore these specifics (e.g., [56]), or do not comprehensively provide them (or mechanisms to estimate them) between all heterogeneous technologies (e.g., [15]).

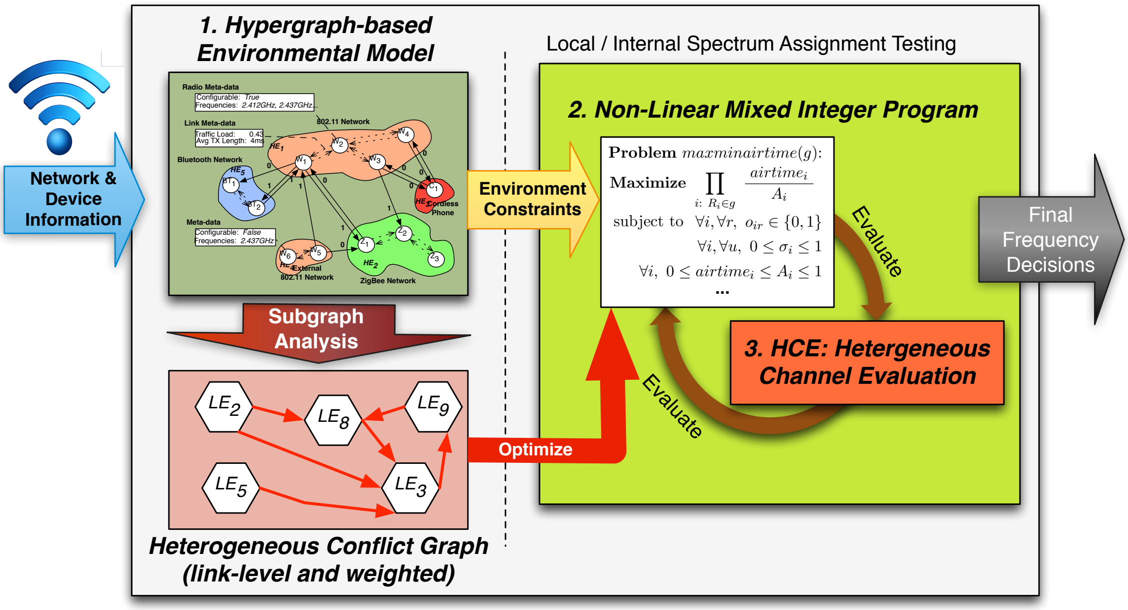

Figure 3.2: An overview of our spectrum management system design, showing each component and the interactions between the components to perform heterogeneous spectrum management.

3.3 Heterogeneous Spectrum Management Design

In this section, we present our novel spectrum management design to accommodate heterogeneity and evolution in the spectrum. This design meets requirements we have discussed in Section 3.2.1, and follows our principles given in Section 3.2.2. We first provide a high-level overview of our system in Section 3.3.1, and then provide the details of each component in Sections 3.3.2 through 3.3.5.

3.3.1 High-level System Overview

The high-level design of our system and its components is illustrated in Figure 3.2. First, the system takes the information provided by the underlying monitoring system and constructs a *hypergraph-based* environmental model that we will describe in detail within Section 3.3.2. The purpose and benefit of the hypergraph-based RF environmental model is three fold: 1) To have a structured representation of the RF environment that unifies the information from diverse underlying monitoring infrastructures, guiding them to collect and structure the necessary information such that our system can organize their environment, 2) To have a rich structure that supports various types of constraints in the RF environment occupied by heterogeneous networks, and 3) To structure the RF environment and the behavior within it in a way that is easily searchable for behavior between heterogeneous networks that match conflict types (to optimize based on). The hypergraph-based model follows our design principles of using fundamentals to represent the many diverse properties of radios and links in the environment, as well as their interactions with each other. For example, using uni-directional edges to represent whether one radio is within spatial range of another, and an indicator on the edge to represent whether the two (potentially heterogeneous) radios coordinate.

Once the environment is represented by our hypergraph-based model, we leverage a key benefit of using such a structure: the ability to search it for specific relationships between its entities (which in our case are radios). Given the various types of heterogeneous conflicts (e.g., one where both transmitters do not coordinate, vs. an asymmetric situation), and the fact that heterogeneous networks do not always lead to conflicts (e.g., a close range Wi-Fi and ZigBee network), detecting conflicts in the RF environment between heterogeneous networks and devices is a more complex task than simply searching for two heterogeneous devices within spatial range of each other. The hypergraph-based structure allows us to leverage *subgraph isomorphism* (also known as *subgraph matching*) to search the larger graph for various signatures that are indicative of conflicts. This allows us to transform the hypergraph (our unified view of RF environments from the underlying monitor) in to a more traditional conflict graph that we can optimize based on.

Finally, as part of the assignment process, we introduce: 1) A *channel evaluation metric* for heterogeneous networks that is able to evaluate the impact of different spectrum assignments given coordinate from homogeneous networks, and a degradation due to conflicts from heterogeneous networks, and 2) A non-linear *mixed integer program* (MIP) based optimization that provides a sufficient amount of flexibility to represent the many different constraints of heterogeneous wireless networks and their interactions to provide accurate and efficient spectrum assignments. In particular, the hypergraph provides the constraints to the optimization, the conflict graph provides a basis to optimize, and the channel evaluation metric provides estimates of interference and performance when networks share a channel. These final components are described in Sections 3.3.4 and 3.3.5.

3.3.2 Heterogeneous RF Environmental Model

Here, we present a highly descriptive hypergraph-based model that represent the key fundamentals in heterogeneous environments. As we briefly described, the purpose of the hypergraph-based RF environmental model is to provide a rich structure and define the input of the information from the underlying monitoring system, used to search for conflicts and provide constraints to the optimization. More generally speaking, hypergraphs are a generalization of a graph in which a *hyperedge* (which we abbreviate *HE*) can connect any number of vertices. This is useful in environments with heterogeneous networks because it can group radios with network and technology dependencies to be configured similarly, or to avoid a similar conflict. In the remainder of this section, we will describe the components in our hypergraph and what they represent, used as input to our system. We refer the reader to an example of our hypergraph-based RF environmental model in Figure 3.3 for discussion throughout the section.

Hypergraph Components & Representation

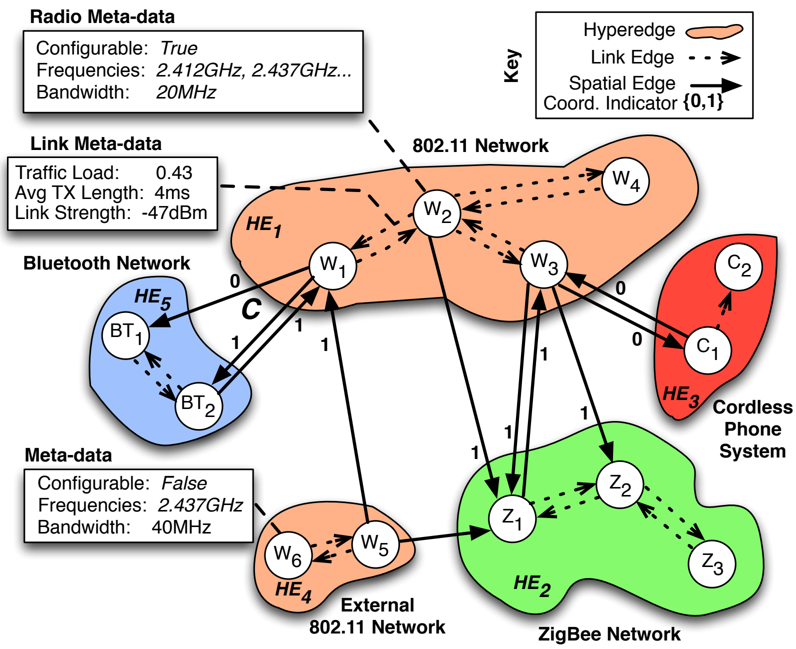

Figure 3.3: Our hypergraph representation of a heterogeneous environment with examples of associated meta-data.

Vertices: At the base of our hypergraph-model is a set of vertices that represent a wireless *radio*. In todays environments, we believe that it is important to make the "base unit" a *radio* rather than network (or device) for several reasons. First, networks can span larger areas in which different radios receive different levels of interference. One cannot assume uniform interference across all radios in a network. Using a network as the base unit will not resolve A level lower, devices can have multiple heterogeneous radios (e.g., a laptop with a Bluetooth and Wifi radio). This also makes devices too coarse-grained. *Radios* truly represent the base unit in today's environments for these reasons.

Edges: Our model has 3 edge types that represent different constraints and properties of the environment:

- Hyperedges: A hyperedge in our model represents a network dependency between specific radios. For example, the network dependency between radios W_5 and W_6, represented by hyperedge HE_4. In terms of spectrum assignment, this is a set of constraints provided for the algorithm to ensure that radios within the same network have uniformly chosen frequencies. Therefore, although our base unit of a radio allows our assignment algorithm to consider interference at the level of a radio or link, the algorithm must consider the impact of configuring all radios uniformly in a network.

- Link Edges: A link edge represents One-way communication between two radios in our

Entity	Meta-data	Metric / Values	Description
Radio	Configurability	True, False	Whether the radio's frequency is reconfigurable.
	Frequencies	List of Frequencies (MHz)	The potential frequencies of the radio.
	Bandwidth	Radio's bandwidth (MHz)	The desired operating bandwidth of the radio.
Link Edge	Traffic Load	Airtime Fraction	Traffic load as a fraction of airtime.
	TX Length	Milliseconds	The average transmission time on the link.
	Link Strength	dBm	The average strength of the link's received signal.
Spatial Edge	Coordination	Binary indicator	Whether the two radios coordinate.
	Coord. Type	Digital, Analog	Whether coordination is analog (sensing) or digital.
	Signal Strength	dBm	The received strength of radio within spatial range.

Table 3.1: Heuristic summary to derive radio/interface relationships.

model, denoted $LE\{X, Y\}$ for communication from radio X to radio Y. These edges imply spatial overlap, and link edges can only exist between radios that are connected by a hyperedge (i.e., communication happens within a network only).

- Spatial Edges: A spatial edge explicitly models a radio Y being within range of a radio X, denoted $SE\{X, Y\}$. This edge is also uni-directional, an important characteristic as we discussed in our requirements to not assume symmetry between radios due to the different transmission powers of heterogeneous radios. Spatial edges also have a binary indicator that indicate whether the radio Y being within range of radio X causes it to back off from X's transmissions. Not that this is *not* an indicator of a conflict.

Meta-data: Finally, radios, links edges, and spatial edges all carry forms of meta-data. Radios carry meta-data that provides additional constraints in assignment. For example, its possible frequencies, bandwidth, and whether the readio is configurable. Not all radios within range will be configurable, like a neighbor's radios that are within range of a user's home. Links carry meta-data that is used by the assignment algorithm and predictive channel metric. This metadata helps estimate fairshare of airtimes between links that coordinate, and predict heterogeneous interference between links that do not. Lastly, spatial edges that denote radio Y is within range of radio X carry meta-data that includes the signal strength of X received at Y. This is also used for predicting heterogeneous interference. We provide a comprehensive list of the meta-data for each entity in the hypergraph in Table 3.1.

Note the difference between the *link strength* on a link edge and a *signal strength* on a spatial edge. The strength on the link edge denotes the reception strength on a communication link, and the signal strength on the spatial edge denotes the reception strength at a radio from another radio. These two pieces of information are used to calculate potential interference and SINR on a communication link due to potentially interfering radios and conflicts between links within range.

Examples of hypergraph flexibility

The hypergraph and its associated meta-data can represent many different types of conflicts in the environment. We will talk about many of these in the next subsection when we discuss how to decompose it to a heterogeneous conflict graph. However, since the graph models spatial overlap uni-directionally, it can capture many types of asymmetry in the environment. This includes two radios being hidden to each other (no spatial edges between them), or just one radio being able to sense and back-off to the other. This latter scenario is shown between W_2 and Z_1 in our hypergraph example (Figure 3.3). The Wi-Fi radio is powerful enough to overlap with the ZigBee radio causing it to back off, but not visa-versa.

Additionally, what our hypergraph shows is that we can model a heterogeneous environment using fundamental properties of the spectrum and its protocols without specifics to standards. This provides flexibility that accommodates evolution. As new protocols enter the spectrum, our model only requires the above fundamental properties about its interactions. If additional fundamental properties become relevant, they can also be flexibly added to the meta-data of each component. In the next subsection, we will show the benefit of its structure in searching for conflicts.

3.3.3 Deriving a Heterogeneous Conflict Graph

The benefit of having a graph-based model to represent the environment and the interactions within it, is the ability to flexibly search the graph for various specific relationships between components. In our case, to search for various types of relationships between radios and links that are indicative of conflicts. This allows us to derive a traditional conflict graph from the hypergraph given to our system as input. This is used as a basis for optimizing spectrum assignment in the environment. However, the hypergraph and its meta-data are still used in conjunction with the conflict graph to provide important constraints in the optimization around the conflicts. As we will show, this approach is flexible and accommodates various conflicts and evolution by searching with "conflict templates" and *subgraph isomorphism.*

Building a conflict graph using *subgraph isomorphism* and conflict templates

Searching for relationships between components in a graph is a common practice used in many fields such as social networking [100,101], data mining [102,103,104], and anomaly detection [105]. These works leverage *subgraph isomorphism* (also known as subgraph matching/analysis) to search a larger graph G for subsets of nodes with specific labels, attributes, and relationships that match a subgraph template H. Therefore, subgraphs in G that match the template H are isomorphic to it.

We use subgraph isomorphism and our hypergraph-based model to provide a flexible (and generic) way to search for conflicts and create a heterogeneous conflict graph. We do so by first creating a set H' that contains subgraph templates with labels, attributes, and relationships between pairs of wireless links that are indicative of conflict behavior. Then, we search the hypergraph G for each $H_i \in H'$. Each time the relationship between a pair of links in G matches subgraph H_i (indicating conflict behavior), we: 1) Create a node in the conflict graph for each wireless link in the conflicting pair (if it does not exist), 2) Draw a uni-directional conflict edge from one node (i.e., link) to the other (direction depending on the conflict scenario), and 3) Annotate the conflict edge with a scenario used later to help estimate the impact of the conflict.

Before we further describe the subgraphs and conflict graph in more detail, it is important to note that both subgraph templates and the conflict graph represent relationships between pairs of wireless links. They do not represent pairs of radios, i.e., neither say that radio X interferes with radio Z. This is important to not oversimplify the environment, since interference depends on a pair of transmitters and receivers, not a single pair of radios.

To briefly illustrate the importance of this, we refer the reader to Figure 3.4. Despite the fact that T_1 and T_2 do not coordinate with each other and are within range, it would be over-simplified to say that T_1 and T_2 conflict with each other. In fact, it is not possible for T_1 to interfere with T_2's transmissions since T_2's only receiver (R_2) is out of range of T_1. Although transmissions on LE_2 from T_2 can interfere with T_1's transmissions on LE_1 (R_1 is within range of T_2), it would also be oversimplified to say that T_2 interferes with T_1 since T_2 does not impact T_1's link to R_3 (LE_3). In other words, T_2 does not *always* interfere with T_1.

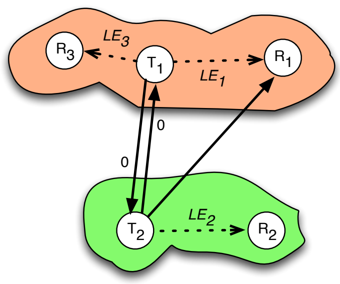

Figure 3.4: Link relationship and conflict example.

For these reasons, almost all aspects of our spectrum assignment system operate the level of wireless links, including our heterogeneous conflict graph and subgraph templates. In Figure 3.5, we show 5 example conflict templates that belong to our set H'. As shown, these templates describe relationships between two links: a base link, and an opposing (interfering) link. One common feature in all of these templates is the base receiver being within range of the opposing transmitter. Without this spatial link (SE_1), a potential conflict is not possible. Aside from this feature, each templates varies in the *exact* properties between the two links. For example, a template for both transmitters being within range but not coordinating (subgraph #1), both transmitters being out of each other's range (subgraph #4), and templates for various asymmetric situations (subgraphs #2, #3, #5). To assist the reader's understanding: subgraph #1 matches our conflict scenario given in Figure 3.4.

Although there are many possible templates needed to match all potential conflicts using subgraph isomorphism, the resulting behavior of many of these templates is the same. For example, properties represented by subgraphs #1 and #4 result in the same outcome: both transmitters being unable to coordinate with each other. Differing properties in subgraphs #3 and #5 both result in the base link coordinating with the opposing link, but not visa-versa. As briefly mentioned, knowing the resulting conflict behavior is important in our estimate of interference between the links. To our benefit, the resulting behavior of these templates can be classified in to 3 main categories and used to annotate our conflict graph (for weighting purposes later on):

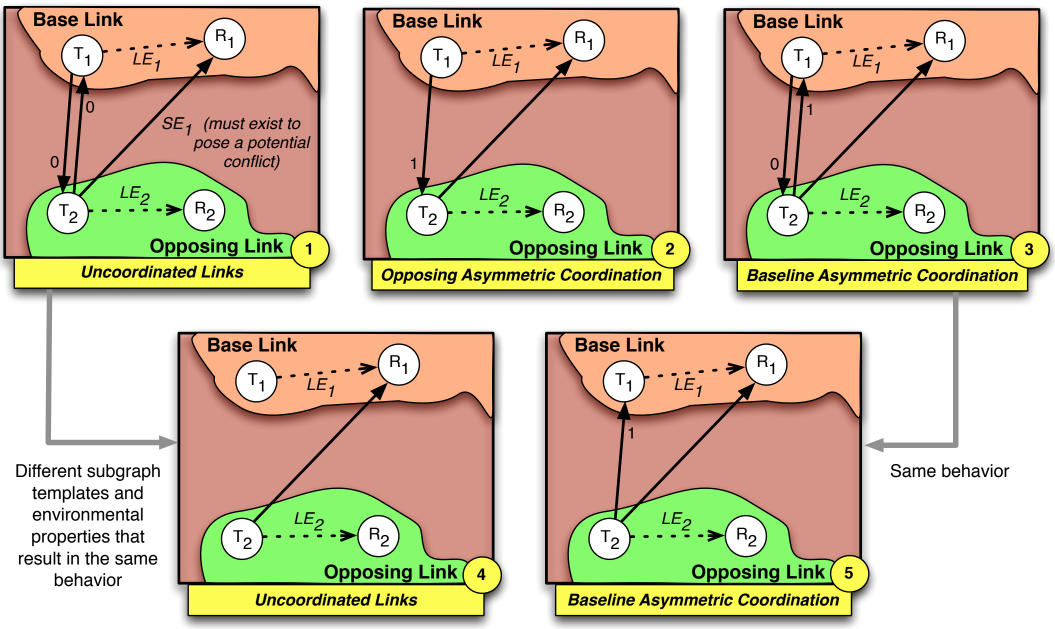

Figure 3.5: Examples subgraphs in our hypergraph-based model that indicate potential conflicts between links in the environment.

1. *Dual uncoordinated links* where both links do not coordinate with each other's transmissions, leading to the highest potential number of overlapping transmissions and interference (e.g., subgraph templates #1 and #4).

2. *Opposing asymmetric coordination* where the interfering link coordinates with the base link, but the base link does not coordinate with the interfering link. This means that the base link's transmitter can improperly begin transmitting during the other link's transmission, leading to a lost transmission at its receiver (e.g., subgraph template #2).

3. *Baseline asymmetric coordination* where the base link coordinates with the opposing link, but not visa-versa. This means that the base link will back off when it hears the opposing link active, but the opposing link can improperly begin transmitting amidst a transmission from the base link (e.g., subgraph templates #3 and # 5).

Therefore, we pre-classify each subgraph template in to one of these 3 scenarios, and annotate the edges in our conflict graph based on the subgraph that is matched. We use the annotations D, OA, and BA for each 3 conflict categories we have presented, respectively.

Finally, there are many subgraph matching algorithms to meet different graph constraints and computational requirements (e.g., distributed algorithms to run on clusters for significantly large graphs [106]). Given our relatively small graph sizes expected (e.g., with tens to hundreds of links), we use the naive algorithm presented by D. Epstein in [107] for simplicity.

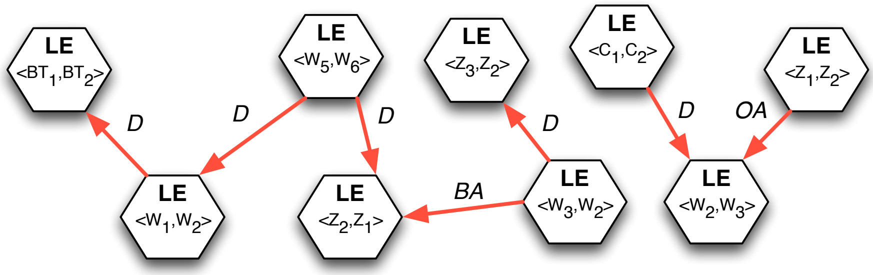

Figure 3.6: The resulting conflict graph with conflict type annotations for the environment shown in Figure 3.3. Generating using subgraph isomorphism and our conflict templates.

Using subgraph isomorphism, our subgraph templates, and our categories for annotations, we derive a more traditional conflict graph used as a basis for optimization from our larger hypergraph-based RF environmental model. In Figure 3.6, we show the resulting conflict graph for our hypergraph example (Figure 3.3). As expected, vertexes are links, and the edges denote a conflict from an opposing (interfering link) to a base link (with an appropriate annotation). For example, link edge $LE\{W_5, W_6\}$ conflicts with link edge $LE\{Z_2, Z_1\}$, and this conflict results in both transmitters being uncoordinated. We encourage the reader to go through additional conflicts in the graph on their to gain a better understanding of our templates and the matching process. However, in summary, this graph provides meaningful information about the conflicts in the environment, including their particular scenario to allow us to weight them as a basis for our optimization. Additionally, it is flexible to support other types of conflicts, and it is generic (i.e., no specifics of protocols or standards are used).

3.3.4 A Heterogeneous Predictive Channel Quality Metric

The last critical component needed by the assignment algorithm is the predictive channel quality metric. As we have already briefly discussed, the goal of this metric is to help internally evaluate potential frequency configurations by predicting the performance of radios, links, and networks sharing spectrum. Given a radio and its set of links in a heterogeneous environment, this metric should consider the impact of: 1) Other homogeneous radios and their links that will contend for airtime, and 2) Heterogeneous radios and their links that can degrade its performance due to potential active conflicts. Note that our conflict graph (like many others) does not consider the active frequencies of the links to classify it as a conflict. The conflicts in the graph are based on their spatial and coordination properties, and are only considered "active" if the two links are configured such that their channels overlap either completely or partially.

Given these two key considerations, we believe that the high-level structure for estimating the performance of a radio R_i and its links on an active frequency f is as follows. First, assume that R_i has an average airtime of A_i based on the demand of its links (i.e., those where it is a transmitter), and consider $C_i(f)$ to be the set of radios on frequency f that R_i coordinates with. Given consideration 1, radio R_i will receive an airtime that is at most the maximum of:

a) the residual airtime given the radios in $C_i(f)$ (i.e., 1 minus the sum of their airtime's), and b) its expected fair share with the radios in $C_i(f)$. This first part is a fairly simple estimation, made by many prior works (e.g., [15,21,55]). Then, given consideration 2, this airtime will be degraded by active conflicts when operating on frequency f. Denoting $\sigma_i(f) \in [0,1]$ to be the estimated fraction of airtime lost due to these conflicts, the total estimated performance (or "good airtime") of radio R_i on frequency f would be:

$$airtime_i(f) = \ max(1 - \sum_{c:\, R_c \in C_i(f)} A_c,\ \frac{1}{|C_i(f)|+1})\ *\ (1 - \sigma_i(f)) \tag{3.1}$$

With no sustained interference across R_i's links, $\sigma_i(f)$ will be 0 and estimated airtime will not be degraded. 60% sustained interference would degrade the usable airtime by this amount.

The Challenge: The key (and non-trivial) challenge in calculating Equation 3.1 is estimating $\sigma_i(f)$. As we discussed in our requirements section (Section 3.2), supporting general heterogeneity not only requires supporting diversity in each of the PHY, MAC, and application layers, but also the more complex interactions across all 3 layers. This is what makes estimating $\sigma_i(f)$ challenging: it is a result of these more complex interactions across all 3 layers between radios in the spectrum. In particular, the degree of interference depends on coordination on each of the radio's links with all other links in range (MAC layer), the traffic loads between each of these competing links (application layer), as well as various PHY layer properties such as the SINR on the links and their modulations which often provide different robustness properties based on their bitrate and error correction.

In the remainder of this subsection, we will describe how to estimate $\sigma_i(f)$ for a radio R_i, despite these many complex properties with the help of our hypergraph and conflict graph. It is impractical to believe that such an estimation could be exact, however, we believe that our estimation provides a reasonable weighting function for interference used by the assignment algorithm. This is an improvement over prior work where the estimation has either been Wi-F centric [15] or considered binary [56]. In line with our design principle, we show how to make this estimation generic where possible to support general heterogeneity and evolution. Additionally, the estimation is meant to be flexible: if exact information is known by the monitoring system, it can be applied rather than estimated.

The basis of our estimation in the interference component (i.e., $\sigma_i(f)$)

Consider the total airtime of radio R_i to be A_i, based on the airtime of its set of links K_i where R_i is considered the transmitter: $A_i = \sum_{j:\, L_j \in K_i} LinkAirtime_j$. If a radio R_i has a total airtime of 0.7, and one of its links $L_j \in K_i$ has an airtime of 0.3, then L_j accounts for 42.85% of radio R_i's total airtime (i.e., $LinkAirtime_j$ / A_i = 42.85%). If link L_j has a loss rate of $LinkLossRate_j = 0.5$ when radio R_i is operating on frequency f, then link L_j degrades 42.85% of radio R_i's total airtime by 0.5. Therefore, the fraction of total airtime lost due to sustained interference from link l on frequency f would be $0.4285 * 0.5 = 0.214$. If $LinkLossRate_j$ were 1, then the total fraction airtime lost due to this link would be its entire contribution, i.e., 42.85% of the total airtime instead of half of this. The total fraction of airtime lost due to sustained interference across all links can therefore be calculated as:

$$\sigma_i(f) = \sum_{j:\, L_j \in K_i} (LinkAirtime_j \;/\; A_i) * LinkLossRate_j \tag{3.2}$$

Decomposing *LinkLossRate$_j$*: we believe that $LinkLossRate_j$ (the challenging portion of estimating $\sigma_i(f)$) can be decomposed in to several fundamental factors across the MAC, PHY, and application layers (following our principle of design). At the MAC layer, we leverage our heterogeneous conflict graph to know what links potentially conflict with L_j due their inability to coordinate with L_j's transmitter, and their potential interference on L_j's receiver. We consider these links to belong to the set U_j (i.e., uncoordinated links). In our conflict graph, these nodes have a uni-directional conflict edge to link L_j. Next, we estimate the probability of loss from each link L_u on L_j, where $u\colon M_u \in U_j$. This probability is driven by remaining MAC, PHY, and application layer properties:

$$\forall u,\; ProbOfLoss_{ju} \;=\; ActiveConflict_{ju} \;*\; POverlap_{ju} \;*\; OLoss_{ju} \tag{3.3}$$

First and foremost, $ActiveConflict_{ju}$ is a binary indicator based on the operational frequencies of L_j and L_u, taking on a value of 1 if they overlap in the spectrum, and 0 otherwise (causing $ProbOfLoss_{ju}$ to be 0). Next, $POverlap_{ju}$ is the probability that transmissions on links L_j and L_u will overlap in time due to their uncoordinated behavior. This is a property of the application layer (e.g., the traffic on both links), and the conflict scenario annotated on the conflict edge from L_u to L_j (e.g., do both radios fail to back off, or does at least one back off?). Finally, $OLoss_{ju}$ is the probability that an overlapping transmission with L_u will cause a transmission failure for L_j. This is a PHY layer property, mainly driven by link L_j's SINR. These factors provide a probability of loss on link L_j from a link L_u where $u\colon M_u \in U_j$.

Finally, the probability of a successful transmission on L_j given N potentially conflicting links in $U_j = \{L_{u1}, L_{u2}, ..L_{uN}\}$ is the probability of no loss from L_{u1} (i.e., $1-ProbOfLoss_{ju1}$), times the probability of no loss from L_{u2}, ..., times the probability of no loss from L_{uN}. In other words, the probability of *no* overlapping transmission from all conflicting links that will cause a failure. Therefore, the probability of a loss on link L_j due to all of those potential conflicts is:

$$LinkLossRate_j = \; 1 - \prod_{u:\, M_u \in U_j} 1 - ProbOfLoss_{ju} \tag{3.4}$$

Estimating the probability of overlap and loss *(i.e., $POverlap_{ju}$ and $OLoss_{ju}$)*

The remaining pieces of information needed to compute $\sigma_i(f)$ are, for each link $L_j \in K_i$ and every potentially conflicting link $L_u \in U_j$, the probability of overlapping transmissions between links L_j and L_u, as well as the probability of loss due to their overlap. The monitoring system can provide either of these pieces of directly if available, and will likely result in a

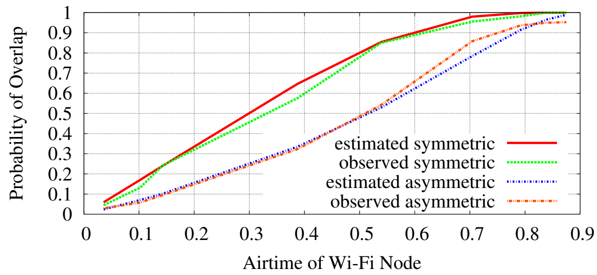

Figure 3.7: Comparing our probability of overlap estimate to actual observed values in dual uncoordinated ("symmetric") cases and asymmetric cases.

greater level of accuracy in the estimation of interference. However, if the information cannot be provided by the monitoring system, we provide a rather generic way to estimate these values, providing a reasonable weight to conflicts useful to the spectrum assignment process.

Estimating $POverlap_{ju}$: To estimate the probability that two uncoordinated links L_j and J_u will overlap, we leverage the observation that without coordination these links operate *entirely independently*. As a result, their events (i.e., packet transmissions) occur continuously and independent of one another. Therefore, heterogeneous radios can be modeled as independent Poisson processes by which, using knowledge of their average transmission lengths and airtimes, we can estimate their probability of overlap. This does *not* rely on specifics of the protocols, using only fundamentals features (airtimes and transmission lengths). Our derivation is similar to historical estimations of collision overlap in Ethernet and ALOHA networks without CSMA, however, a key difference is the potential for asymmetric coordination (i.e., one radio coordinating with the other, but not visa-versa).

First, consider λ_u to be the rate of transmissions from a conflicting link L_u based on its airtime. Then, consider V_{ju} to be the vulnerability window of transmissions from link L_j given a conflict with link L_u. The value of V_{ju} is based on the coordination behavior between L_j and L_u, as annotated in our conflict graph. If the annotation on the conflict edge from L_u to L_j is D, then both links are uncoordinated meaning L_u can transmit in the middle of L_j's transmission, and visa-versa. Therefore, the vulnerability window is $V_{ju} = T_j + T_u$ where T_j is the average transmission time on link L_j. If L_j coordinates with L_u but not visa-versa, the annotation on the conflict edge is BA and the vulnerability window will be $V_{ju} = T_j$, i.e., L_j is only vulnerable during its transmissions since L_u cannot sense them. Finally, the opposite scenario with an annotation of OA would lead to a vulnerability window of $V_{ju} = T_u$. Given our assumptions, the probability of overlap between L_j and L_u would therefore be:

$$POverlap_{ju} = 1 - e^{(-\lambda_u * V_{ju})} \tag{3.5}$$

To show the importance of modeling asymmetry and that this method can provide a reasonable estimate, we operate Wi-Fi and ZigBee links over the air, controlling their airtimes and coordination behavior. Using knowledge of the airtimes and average transmission lengths, we

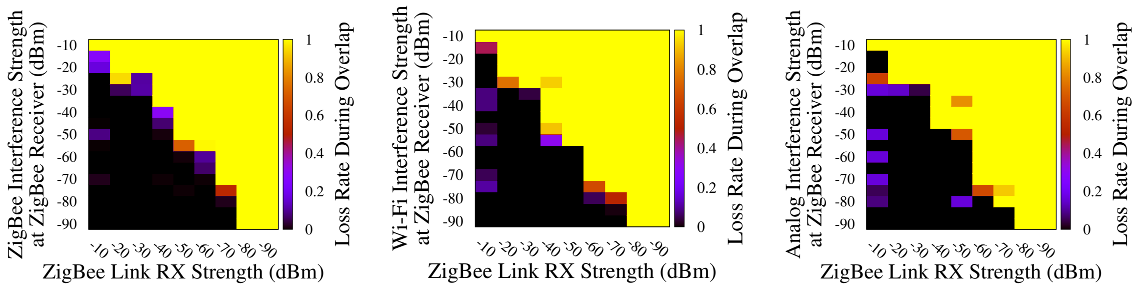

Figure 3.8: The probability of overlapping transmissions causing a loss between various technologies on a ZigBee link. This rate varies by the interference and reception strength.

estimate *POverlap* and compare it to an actual observed value. As shown in Figure 3.8, our estimations are close to the observed values. Additionally, in some cases, the asymmetric overlap rate is half that of the dual uncoordinated ("symmetric") rate. Therefore, if one ignores asymmetric behavior in the environment, the estimate can be significantly more inaccurate. Finally, note that all information used in this estimation is available in our hypergraph based model (through its meta-data and structure).

Estimating $OLoss(l, u)$: The last piece of information needed to estimate the interference between two conflicting links is the probability of loss during overlap. As mentioned, this is a property of PHY layer interactions between the links. Following our principles of design, we remain generic where possible, and include specifics where needed to maintain accuracy. In this case, we *must* consider specifics of the technologies used by the links since loss during overlap is a factor of many PHY layer properties such as the modulation and FEC used by the technologies which provide different levels of robustness.

We believe that the information needed, however, is not difficult to measure and can be pre-measured and provided by a lookup table. The key considerations to whether an overlap causes a loss between two links is: a) The technologies and their modulations in use, and b) The SINR at the receiver of the base link L_j. That is, the strength of its reception and the strength of the interference. To illustrate this, we refer the reader to Figure 3.8 which shows for various reception strengths on a ZigBee link and various interference strengths at the ZigBee link's receiver, what the probability of loss during overlap is with different interfering technologies. Clearly, the outcome depends on the reception and interference strength, showing that in general: the reception strength must be a few dB greater than the interference strength.

Since we have the reception and interference strengths between all links in our environment (provided by our smartphone-based monitor – Chapter 2), we can compute lookup tables of overlap loss rates given pairs of technologies and use the values provided by our monitor to index the tables. While this requires N^2 tables, many studies already pre-compute these tables to study interference in the research community (e.g., [7,15,16]). Computing these tables is not complex, yet they can ensure the greatest amount of realism. Therefore, we compute these tables between various technologies such as Wi-Fi, ZigBee, Bluetooth, and cordless phones for the purpose of our study.

Summary of our Metric: In summary, we have presented a reasonably generic metric for estimating the performance of a radio and its links on a specific frequency. This metric leverages a rather common calculation for estimating the fair share of airtime from coordinated links and radios, as well as our $\sigma_i(f)$ estimation to predict the sustained interference of a radio R_i based on conflicts with its links on frequency f. We will provide an evaluation of our metric's accuracy in our evaluation section (3.4).

3.3.5 Spectrum Assignment Algorithm

The last component in our heterogeneous spectrum management system is our assignment algorithm. We introduce a mixed integer program (MIP) that uses the hypergraph-based environmental model (Section 3.3.2) as input for the constraints in the system (e.g., the potential frequencies of each radio), and our conflict graph (Section 3.3.3) and metric (Section 3.3.4) as a basis to optimize, predict heterogeneous interference, and assign spectrum to improve performance. The optimization that we introduce follows our principles of design by continuing to describe fundamental properties.

Notation and modeling of our problem

Our representation of the hypergraph and its components are as follows. Let G be the hypergraph of our wireless environment with a total of I radios denoted R_i $(i = 1, ..., I)$. Each hyperedge H_e $(e = 1, ..., E)$ in graph G contains a set of wireless radios that it connects, such that $H_e = \{R_1, R_2, ...\}$. Next, our graph has S spatial edges SE_s $(s = 1, ..., S)$ where SE_{ji} denotes R_i being within range of R_j. If R_i coordinates with transmissions from R_j, then $R_j \in C_i$, i.e., the set of radios that it coordinates with.

Each communication link edge LE_l (l = 1, ..., L) represents the communication between two radios, where the link edge LE_{ij} denotes a communication link from R_i to R_j. Each communication link $LE_l \in G$ has an average desired airtime of $LinkAirtime_l$ and an average transmission length of T_l. The communication links where R_i is the transmitter belong to the set K_i, and the average airtime A_i of the radio R_i is based on the demand from its links: $A_i = \sum_{j:\, L_j \in K_i} LinkAirtime_j$.

The spectral bandwidth of R_i is denoted B_i, and the possible set of frequencies for each $R_i \in G$ belong to the set F_i. Let the variable f_{ix} represent whether the chosen frequency for R_i is at the index x in F_i (x = 1, ..., X). That is, $f_{ix} \leftarrow 1$ if $f_i = F_{ix}$, and $f_{ix} \leftarrow 0$ if $f_i \neq F_{ix}$. Clearly, we introduce the constraint $\sum_{x=1}^{X_i} f_x = 1$ since only one center frequency can be chosen.

To model whether radios R_i and R_j overlap in the spectrum, we first precompute $O_{i_x j_z} \in \{0,1\}$ which is, for all possible frequency and bandwidth combinations, whether the frequencies F_{ix} and F_{jz} overlap, given respective bandwidths of B_i and B_j. We use these precomputed values in combination with the variable indicators of what frequency is chosen for each radio (f_i and f_j), introducing a new binary variable o_{ij} that denotes whether R_i and R_j actively overlap in the spectrum.

Next, referencing the conflict graph derived from G via our subgraph analysis, each communication link $LE_l \in G$ has a set U_l of (at least partially) uncoordinated and conflicting links. These links are those that have a conflict edge to LE_l. Since these sets of links are pre-known to the optimization (derived during subgraph analysis), we use the annotations on each link to precompute the vulnerability window between all conflicting links. This window is denoted V_{ld} for a conflict from L_l to L_d.

Finally, we introduce the variable $airtime_i$ to be our estimated performance of R_i's configuration, its links, and its interactions with potentially active conflicts in the environment. This accounts for coordinated networks in range (i.e., sharing airtime) and a loss in airtime due to sustained interference σ_i using the basis of our predictive channel metric.

In our problem definition below, we do not include the derivation of σ_i since we have already included the equations and details in Section 3.3.4. Re-listing each of these equations in our constraints below would only be repetitive. However, note that in the derivation previously described, the binary indicator $ActiveConflict$ existed. In our optimization, this is replaced by our variable o which indicates whether two radios and their links overlap in the spectrum, i.e., inactivating the conflict if the links do not. Additionally, we use the signal strengths known on each link from the meta-data in the hypergraph, in addition to the protocol/standard used on each link, to appropriately look up $OLoss$ in the interference tables from the optimization (Section 3.3.4).

Objective: Our goal is to choose a frequency f_i for each radio $R_i \in G$ to maximize the performance of the radios and their networks in the environment. As we have discussed, this accounts for shared airtime and heterogeneous interference sustained due to the configuration of all other radios and their potential conflicts. In particular, we look to maximize the fraction of airtime received to the airtime desired: $airtime_i/A_i$ for each radio $R_i \in G$. By doing so, we do not greedily configure the environment to give more airtime to those that want more airtime, e.g., by maximizing $\sum airtime_i$ which would give preference towards configurations that resolve conflicts on links that have a greater desired airtime A_i. We ultimately look to be fair to each radio by maximizing the minimum fraction of desired airtime received. This will naturally resolve heterogeneous conflicts that will drive this fraction lower.

Problem *maxminairtime*(g):

$$\textbf{Maximize} \prod_{i:\ R_i \in g} \frac{airtime_i}{A_i}\ ,\quad \text{subject to}$$

$$\forall i,\ airtime_i = min(A_i,\ max(Residual_i, FairShare_i)) * (1 - \sigma_i) \tag{3.6}$$

$$\forall i,\ Residual_i = 1\ -\sum_{c:\ R_c \in C_i} A_c * o_{ic} \tag{3.7}$$

$$\forall i,\ FairShare_i = 1\ /\sum_{c:\ R_c \in C_i} o_{ic} \tag{3.8}$$

$$\forall e : H_e \in g,\ \ \forall R_i \in H_e,\ \ \forall R_j \in H_e,\ \ f_i == f_j \tag{3.9}$$

$$\forall i, \forall c,\ o_{ic} = \sum_{x=1}^{X_i} \sum_{z=1}^{X_j} O_{i_x c_z} \wedge f_{ix} \wedge f_{cz}\ \in \{0, 1\} \tag{3.10}$$

$$\forall i,\ \sum_{k=1}^{K_i} f_{ik} = 1 \tag{3.11}$$

$$\forall i, \forall r,\ o_{ir} \in \{0, 1\} \tag{3.12}$$

$$\forall i, \forall u,\ 0 \leq \sigma_i \leq 1 \tag{3.13}$$

$$\forall i,\ 0 \leq airtime_i \leq A_i \leq 1 \tag{3.14}$$

$$\forall i,\ 0 \leq Residual_i \leq 1 \tag{3.15}$$

$$\forall i,\ 0 \leq FairShare \leq 1 \tag{3.16}$$

Linear Modeling: To express in Equation 3.10 linearly, we replace $x_1 \wedge x_2 \wedge x_3$ with a corresponding binary variable y and introduce the following 5 constraints: (1) $y \leq x_1$, (2) $y \leq x_2$, (3) $y \leq x_3$, (4) $y \geq x_1 + x_2 + x_3 - 2$, and (5) $y \in \{0, 1\}$. To model a *min* function (e.g., Equation 3.6), we replace $min(x, y)$ with z and introduce the following 2 constraints: (1) $z \leq x$ and (2) $z \leq y$. Modeling the cross product found in *LinkLossRate* (Equation 3.4) in a linear way is done using a technique by Peterson [108]. The technique breaks down the computation of the cross product in to a series of products in which the result of a given product is replaced with a new variable, and this variable is used in the subsequent product. Lastly, in our objective function, we model $max(\ min(\frac{Airtime_i}{D_i})\ )$ as $max\ \eta$, where: $\forall i,\ \eta \leq \frac{Airtime_i}{D_i}$.

The program and its constraints: Above, we define a non-linear mixed integer program (MIP) formulation that takes as a parameter g, the hypergraph representation of the environment. Given the environment g, the MIP maximizes the minimum fraction of airtime received ($airtime_i$) to the desired airtime A_i of each radio R_i in the environment.

We briefly explain each of the constraints and relaxations within. (3.6) is the constraint on each radio R_i's estimated airtime. Note that we take the maximum of the residual and

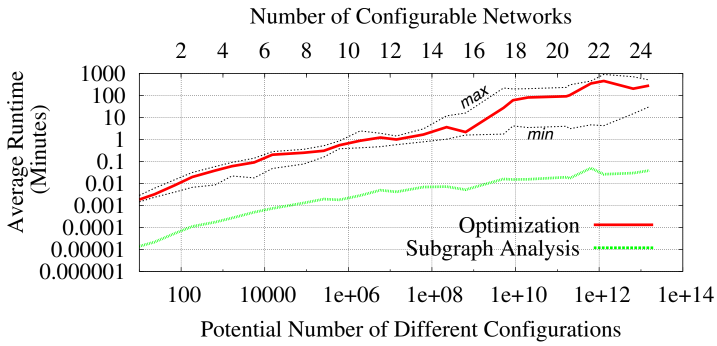

Figure 3.9: The average, minimum, and maximum total runtime of the optimization and subgraph analysis given an approximate number of configurable networks and a potential number of different configurations.

fairshare of airtime, but then take the minimum of A_i so that the radio never uses more than its demand. (3.7) calculates the residual airtime based on all other radios that R_i coordinates, belonging to the set C_i. (3.8) accounts for the fair share of airtime expected from these same coordinating radios. Since o_{ic} is binary, we can count the number of coordinating radios on the same frequency by summing o_{ic} across each radio $R_c \in C_i$. (3.9) introduces the constraint that radios connected by a hyperedge must have the same chosen frequency (i.e., they construct a network). (3.11) ensures that each radio can only have one center frequency, by ensuring that the sum across the indicators of which frequency index is active is 1. The remaining constraints (3.13 - 3.16) ensure that many of our variables take on values between 0 and 1. Additionally, (3.14) constrains $airtime_i$ to be less than or equal to its demand A_i.

Note that there is no additional or direct constraint to model whether a radio is "reconfigurable" or not (e.g., an external or neighboring radio or device). This is handled indirectly by having the frequency set F_i for a non-configurable radio R_i only contain its current frequency. In practice, the possible sets of frequencies for external radios are also considered to be unknown. This constrains the optimization to assign that radio to that particular frequency, since constraint 3.11 ensures that every radio must be assigned 1 frequency (and 1 only).

Performance of the optimization

An immediate concern in our optimization is its runtime, ultimately: *how long does it take to provide a set of spectrum assignments given various sizes of the environment?* (i.e., the number of networks it must configure). Given that there are several non-linear constraints in our optimization that make our mixed integer program non-linear, understanding its performance and limitations is important. To provide insight to the question of runtime, we run our algorithm against an increasing number of configurable wireless networks, static networks within range (i.e., unconfigurable, or neighboring, networks), as well as the total number of radios (our base unit in the environment and hypergraph).

As one might expect, the most important factor in the runtime of our algorithm is the potential number of different configurations. That is, the total number of different frequency combinations across the configurable networks. In Figure 3.9, we show the average, minimum, and maximum runtime in minutes given this potential number of configurations. This runtime is broken down as the total to perform subgraph analysis to generate the conflict graph, and the total time of the actual optimization that considers the conflict graph and hypergraph. On the top of figure, we also provide an approximate number of configurable networks that result in the potential number of configurations given on the x-axis. As shown, the algorithm will provide what it believes to be an optimal assignment in less than a minute with approximately 1 million different configurations, driven by 10 configurable networks. This is reasonable for the average home environment. 20 networks can be configured in an hour and a half, and once 20 networks is exceeded we reach a physical memory limit on our machine and the optimization time takes significantly longer: 16 hours with 4GB of memory. It was recently made possible to split the memory and computation of this algorithm across multiple machines (e.g., in the cloud) [109]. We have not experimented with this in detail as the framework was released after we completed our evaluation. However, results provided by the authors show reasonable reductions in solve time with parallelization of memory and computation [110].

In any environment, there are expected to be a set of networks that contribute to contention and interference, but are non-configurable (e.g., neighboring networks). These networks, however, still contribute to computation in our optimization, estimating their impact on the configurable networks. Therefore, we evaluate the additional runtime of our optimization given a total number of static networks varying from 0 to 60 in Figure 3.10. As shown, static networks contribute to the additional runtime of the optimization, but do not increase it significantly. Even with 40 static networks within range, only 2-3 minutes of runtime are added on average. It is not until we reach the physical memory limit of our machine due to the static networks that our runtime begins to increase significantly again.

Finally, we show the increase in runtime due to the number of radios in the environment. That is, how does the total number of radios impact runtime, given a number of configurable networks? To evaluate this, we fix the number of configurable networks to 5 and add radios to these configurable networks in round robin fashion, and monitor the total runtime. As shown in Figure 3.11, runtime will not increase significantly even after adding 70 radios to the configurable networks in our optimization. Again, it is not until we reach our physical memory limit that the runtime begins to increase significantly.

Discussion and limitations of our optimization

Given the runtime of our algorithm, we believe that it is suitable for home environments and smaller complexes, potentially providing the capability to more efficiently organize chaotic environments such as apartment buildings. It may be possible to partition larger enterprise or campus environments and run our optimization on these partitions to optimize

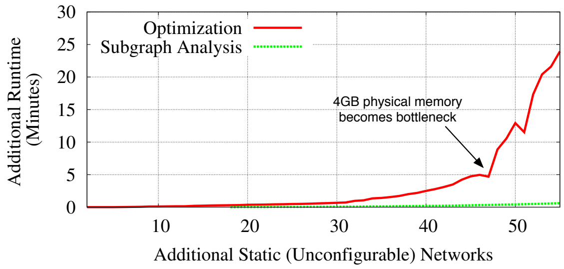

Figure 3.10: Additional runtime given a number of static (unconfigurable) networks (e.g., neighboring networks).

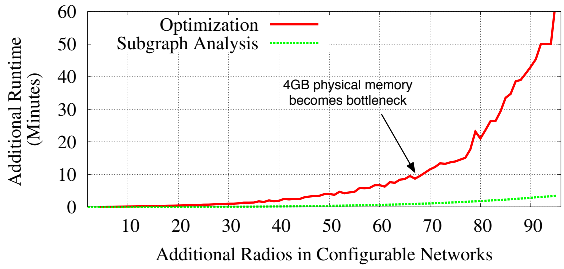

Figure 3.11: Additional runtime given more radios participating in a fixed number of configurable networks.

them. However, this is outside the scope of our work. Additionally, the recent work to parallelize the memory and computation of our open source solver, *SCIP* [111], may show that it can solve larger environments leveraging the cloud.

By using a MIP representation, we believe that our optimization is flexible to additional constraints. There are aspects of the spectrum that we have simplified, however. For example, we assume that the signal characteristics (e.g., propagation) are the same on all frequencies. We do not model a drop in signal quality when re-assigning to a higher frequency (e.g., from 2.4GHz to 5GHz), or an improvement in quality when re-assigning to a lower frequency (e.g., from 2.4GHz to the white space TV bands). However, we believe that our MIP-based approach provides enough flexibility to add these constraints in the future. We discuss other limitations and future work in more detail in Section 3.5.

3.4 Evaluation

In this section, we present an evaluation of our heterogeneous spectrum assignment system. Our goal is to characterize the system's ability to find reasonable and efficient spectrum assignments in heterogeneous environments, improving the performance and fairness of networks within it. In particular, these assignments should avoid heterogeneous conflicts when possible (e.g., through frequency isolation), and otherwise intelligently place the networks and radios in the spectrum to at least reduce the number of conflicts and their impact on performance.

Evaluation Methodology: To perform this evaluation, we introduce several controlled heterogeneous environments and scenarios that vary in their topology, demand, and degree of interference. We classify these environments in to 4 main categories:

1. *Targeted* scenarios where we introduce various types of conflicts between heterogeneous networks that have been reported as common in prior works (e.g., in [7,15,16]). These

targeted scenarios focus our evaluation of the assignment algorithm on its ability to reconfigure a single network to avoid the conflict scenario we introduce. This allows us to provide the reader with a deeper understanding of how our predictive channel metric estimates various types of interference, and how our single algorithm provides assignments that avoid various diverse types of conflicts.

2. *Weakly constrained* scenarios include multiple heterogeneous networks with the potential for conflicts, however, there is sufficient spectrum to isolate incompatible networks (i.e., the spectrum is weakly constrained). These scenarios illustrate the spectrum management algorithm's capabilities in relatively isolated environments. This category of environments begins to highlight abilities of the entire optimization at assigning multiple networks.

3. *Moderately constrained* scenarios include the configuration of multiple heterogeneous networks with the potential for conflicts present if not placed intelligently, due to a moderately constrained spectrum in terms of demand.

4. *Severely constrained* scenarios are over provisioned scenarios where frequency isolation of heterogeneous technologies is not always possible to avoid all conflicts. In some cases, networks must be placed intelligently to reduce interference, reduce contention, and provide a level of fairness across the networks (i.e., no network should receive significantly less airtime, or more interference).

To evaluate configurations across these scenarios, we create the environments in a heterogeneous testbed where the networks are configured to generate a specific demand, and they are placed to generate various types of conflicts we introduce and study. Our smartphone-based heterogeneous monitor (Chapter 2) collects the necessary information from these environments, and provide it to our system which first builds the hypergraph (Section 3.3.2) and then performs subgraph analysis to generate the conflict graph (Section 3.3.3). The hypergraph and conflict graph are then used as inputs to our optimization (Section 3.3.5), which provides the suggested spectrum assignments for each radio. We reconfigure the radios to match the assignments and show the resulting throughput, airtime, and loss rates of the networks and devices. Where applicable, we compare the resulting configuration with configurations generated by other algorithms, such as where networks enter the spectrum in a first-come-first-serve manner, as well as against a basic largest first algorithm to assign the spectrum.

Targeted scenarios and evaluation

Starvation and high loss due to *analog transmitters*: We begin with two scenarios that include analog transmitters, i.e., devices that continuously transmit, causing starvation in CSMA-networks due to either continual back-off from them, or high amounts of loss (observed in many studies[7,8,112]). To evaluate our system's ability to account for these conflicts, we first introduce an analog cordless phone within spatial range of a Wi-Fi transmitter. In

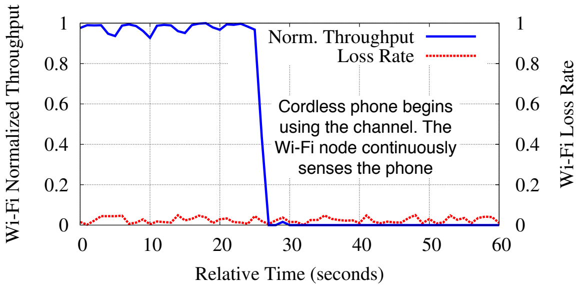

Figure 3.12: Cordless phone causing throughput drop on Wi-Fi network due to continual back-off.

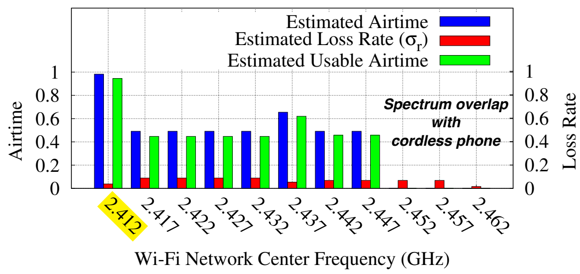

Figure 3.13: Estimated airtime, loss rate, and usable airtime for the Wi-Fi network with the phone in the upper part of the spectrum.

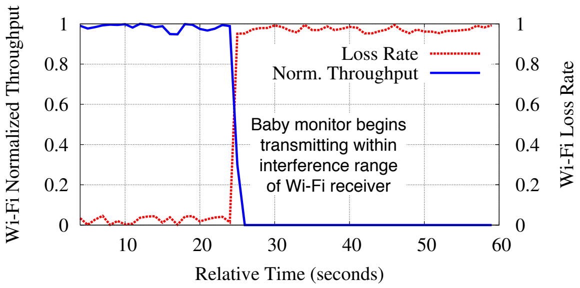

Figure 3.14: Baby monitor causing throughput drop on Wi-Fi network due to increased interference on the receiver.

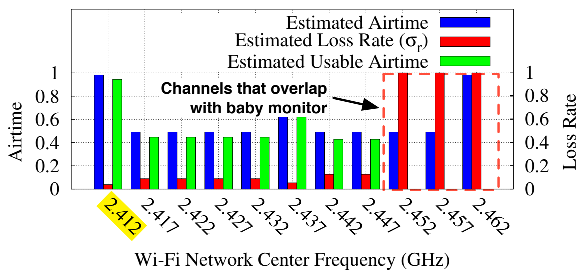

Figure 3.15: Estimated airtime, loss rate, and usable airtime for the Wi-Fi network with the baby monitor.

Figure 3.12, we plot the normalized throughput and loss rate of the Wi-Fi link before and after we turn the cordless phone on at a relative time of 25 seconds. As shown, the throughput of the Wi-Fi link drops significantly, but the loss rate remains low and stable.

In Figure 3.13, we break down the channel estimations made by our optimization when deciding how to configure the Wi-Fi network that contains the link. The *"estimated usable airtime"* is the value of our calculated channel quality metric (i.e., Equation 3.1). However, we also include its two main components: the estimated available airtime due to contention (graphed as "Estimated Airtime"), and the estimated loss rate due to heterogeneous conflicts – our σ value in the main equation (graphed as "Estimated Loss Rate"). As shown, our metric predicts the minimal usable airtime as starvation in its estimated airtime given contention, *not* loss. Given other active networks in the spectrum, the optimization suggests a center frequency of 2.412 GHz. We will soon present results that validate the estimations across all of the potential channels.

Next, we replace the cordless phone with a baby monitor in the spectrum, another analog device which has been the focus of studies. We move it away from the Wi-Fi link's transmitter,

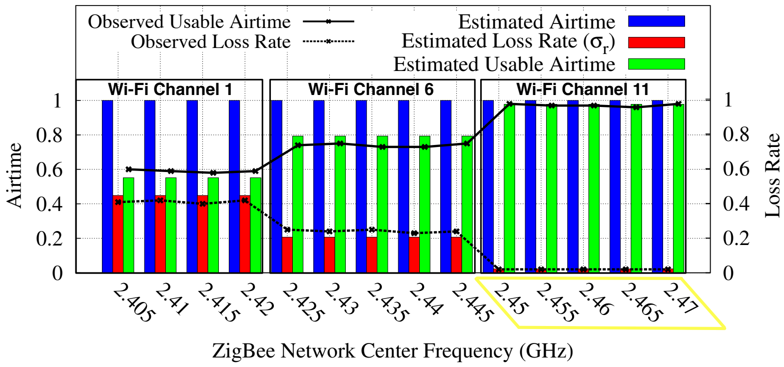

Figure 3.16: ZigBee channel estimations that properly reflect loss and usable airtime given fully hidden and asymmetric scenarios with Wi-Fi networks.

and in to interference range of one of its receivers. In Figure 3.14, we show the result of turning the baby monitor on at a relative time of 25 seconds. Again, the throughput drops significantly, however, loss rate increases in this conflict scenario. As shown in Figure 3.15, our optimization accurately estimates the result of this conflict scenario, shown in the set of higher frequencies. Estimated usable airtime is again predicted to be minimal, however, this is now the result of a higher loss rate, *not* a lower estimated airtime due to contention.

Inability to coordinate vs. asymmetric coordination: Prior studies have reported many asymmetric coordination scenarios between heterogeneous wireless networks, i.e., where one network backs off to another network, but not visa versa. This is commonly reported between heterogeneous devices that have different transmission powers, such as ZigBee and Wi-Fi [16,34,39]. Ultimately, ZigBee networks have low desired airtime, i.e., they do not transmit often and will likely still be able to meet their desired airtime requirement after retransmissions. However, it is still desirable to configure them to operate on channels with the least interference, since retransmissions affect their power consumption which is important given ZigBee nodes are often battery powered. Here, we evaluate our system's ability to detect and distinguish these scenarios, accurately reflecting the loss rate across all channels.

First, we place a Wi-Fi network on Wi-Fi channel 1 that conflicts with a ZigBee network in our environment, such that neither backs off to each either and the Wi-Fi network creates loss on the ZigBee network (i.e., they are destructively hidden). We place another Wi-Fi network on channel 6 that has the same demand as the first Wi-Fi network, however it is spatially placed at a distance reported in prior works that causes the ZigBee network to back-off to the Wi-Fi network, but not visa-versa [16] (i.e., it is a destructive asymmetric scenario). Finally, we setup a third Wi-Fi network on channel 11 with the same demand, but far enough apart such that the SINR on the ZigBee link should create no loss (i.e., hidden, and non-destructive).

We use our monitor to collect the information from the environment as described. We construct the hypergraph, and then run the optimization. Additionally, we configure the

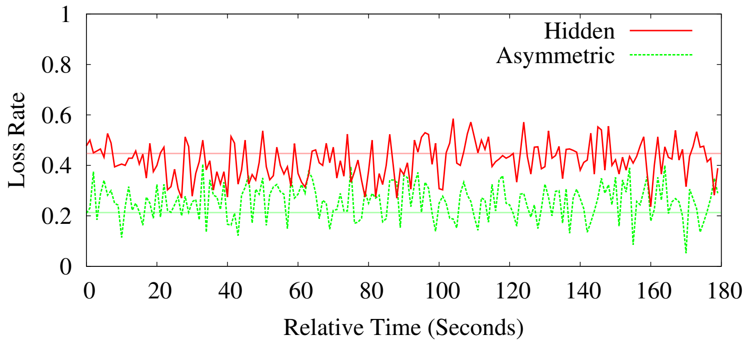

Figure 3.17: The estimated loss rate given a fully hidden and conflicting Wi-Fi network on a ZigBee network, as well as an asymmetric scenario between a Wi-Fi network and ZigBee network.

ZigBee network to operate on each of its potential center frequencies, and record the observed usable airtime and loss rates. The resulting channel estimations made by our optimization, as well as the true observed values, are illustrated in Figure 3.16. As shown, the optimization properly estimates the different loss rates from the Wi-Fi networks, accounting for whether the conflict scenario was completely hidden, asymmetric, or hidden but not strong enough to cause conflict. The optimization will assign the network on a channel that overlaps with Wi-Fi channel 11.

Of course, the dynamic nature of the traffic between the networks will cause the loss to fluctuate. However, we provide a reasonable estimate and importantly reflect the difference in loss rates given the scenario. To illustrate this, we refer the reader to Figure 3.17 where we plot the estimated loss rates of the hidden and asymmetric scenarios as horizontal flat lines, and the actual observed loss rate over 3 minutes. As shown, the loss fluctuates given dynamics of the traffic, and we slightly underestimate the loss of the asymmetric scenario, and slightly overestimate the loss of the hidden scenario. However, they are reasonable estimates and importantly reflect the different conflict scenario.

Accounting for multiple interferers: Next, we setup a targeted scenario to configure a ZigBee network where there are multiple Wi-Fi interferers across the spectrum, and we move the cordless phone and baby monitor to be active interferers in the ZigBee network's upper channels centered at 2.475 GHz and 2.48 Ghz, respectively. In particular, we setup high sources of Wi-Fi interference on channel 1, moderate sources of interference on channel 11, and low sources of interference on channel 1. Our predictive channel metric and optimization are designed to account for the multiple sources of interference on each of these channels that, together, increase the loss rate. In Figure 3.18, we show the estimated, and observed, airtimes and loss rates. As shown, we again provide reasonable approximations of airtimes and loss rates across all of the channels. The estimates are not exact, overestimating and understanding in several instances, however, we believe they are reasonable to perform proper spectrum management.

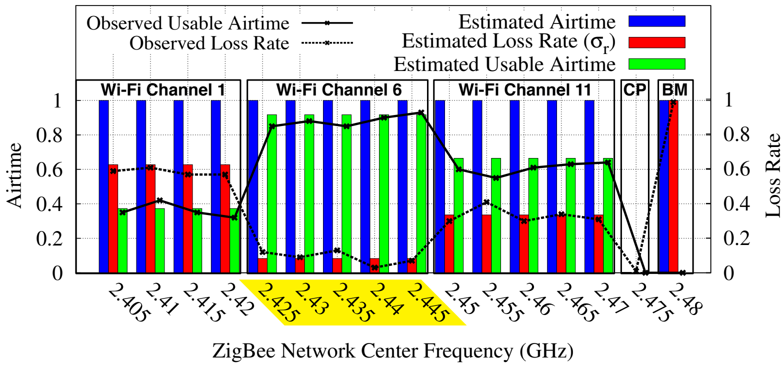

Figure 3.18: Estimated and observed airtimes and loss rates for a ZigBee network in our optimization, shown to estimated reasonably assign a channel with low interference.

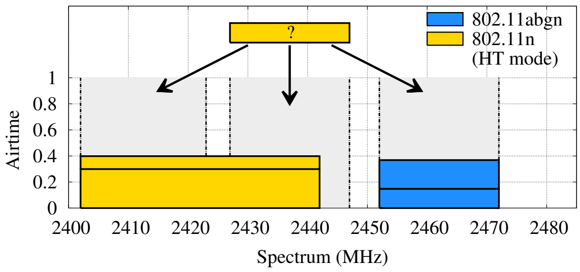

Figure 3.19: Digital networks must be carefully assigned, aligned properly, and not grouped with heterogeneous networks.

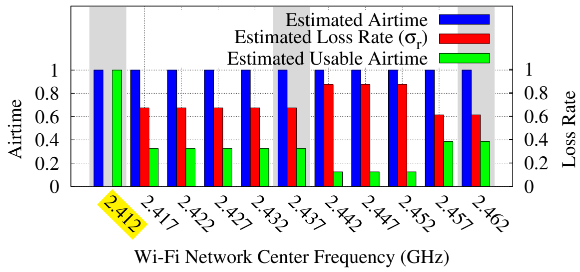

Figure 3.20: Our assignment algorithm will ensure digitally coordinating networks are aligned to properly coordinate.

Avoiding interference given analog vs. digital coordination: As we described earlier, coordination can happen through analog power sensing, or digitally. Our optimization and predictive channel quality metric account for the coordination type when assigning spectrum. We construct the network setup pictured in Figure 3.19 with two static (unconfigurable) 802.11n 40 MHz networks in HT mode (i.e., they digitally coordinate with a non-legacy preamble) with a primary channel of 1, and two static 802.11 networks in legacy mode at channel 11. Then, we introduce a configurable 20 MHz 802.11n HT mode network constrained to the 2.4 GHz band that is within spatial and interference range of the static networks.

Focusing on the non-overlapping channels 1, 6, and 11, our optimization must carefully assign the 20 MHz 802.11n network by properly estimating the impact of it operating with heterogeneous networks and not properly aligning with networks that it digitally coordinates with. We darken non-overlapping channels 1, 6, and 11 in Figures 3.19 and Figures 3.20. As shown, if the network is placed in to channel 11, it will receive interference from the legacy 802.11 networks. If placed on channel 6, it will receive interference from the 40 MHz

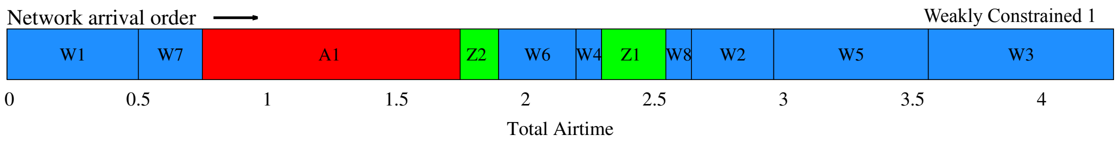

Figure 3.21: Networks in weakly constrained scenario 1, ordered by arrival (for FCFS purposes).

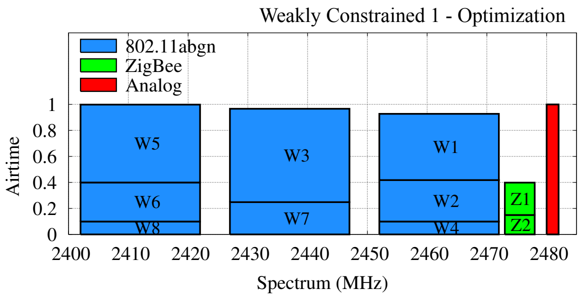

Figure 3.22: Our optimization isolates heterogeneous technologies given enough spectrum, and fits homogeneous networks.

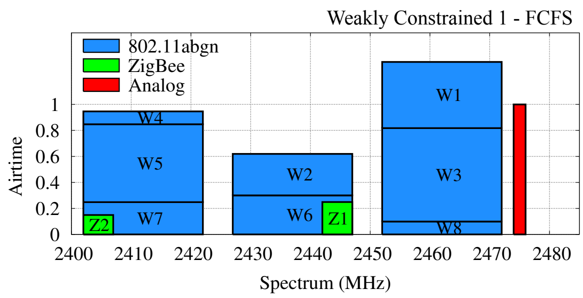

Figure 3.23: Example configuration when networks greedily choose channel assignments in a first-come, first served, manner.

802.11n networks since it will partially overlap with the 40 MHz network's secondary channel (where it cannot digitally coordinate). If properly aligned with the digitally coordinating 40 MHz networks on channel 1, our optimization accurately estimates no loss coordination (with enough airtime) to receive in full, its desired airtime (i.e., a value of 1). The higher loss rates in-between channels 6 and 11 are due to the configurable 802.11 network conflicting with the 40 MHz networks and the legacy networks.

Weakly constrained scenarios and evaluation

Next we shift our focus to environments that involve multiple configurable networks, instead of focusing on configuring a single network to avoid common conflicts reported in prior works. In these weakly constrained scenarios, we focus on environments where there is sufficient spectrum to isolate heterogeneous technologies to avoid conflict. Importantly, we compare the organization that our optimization provides against "first-come, first served" organizations that reflect today's environments. That is, networks greedily choose what channel is best for them when they are introduced in the environment, and they are unlikely to reconfigure their assignment (even if a different channel will provide better performance at a later time). When deploying in a FCFS manner, we have the networks greedily choose the channel with the least usage. Additionally, the FCFS assignment is done agnostic to heterogeneous conflicts.

Weakly Constrained 1: We start with a configuration between mixed 802.11 networks

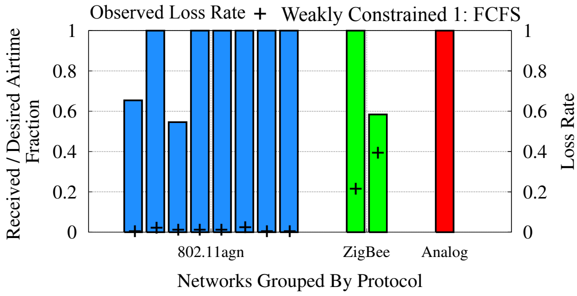

Figure 3.24: FCFS observed values of loss rates and received airtimes for each network, normalized by the networks desired airtime.

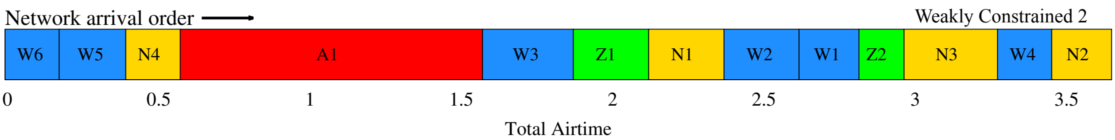

Figure 3.25: Networks in the weakly constrained scenario 2, ordered by arrival (for FCFS purposes).

(that all support the legacy preamble), two ZigBee networks, and an analog cordless phone that have the respective airtimes shown in Figure 3.21. In Figure 3.22, we show the resulting spectrum organization provided by our optimization: heterogeneous networks are isolated, and the 802.11 networks are packed in a way that all networks are able to meet their desired airtime. Alternatively, we show the resulting configuration given a FCFS configuration in Figure 3.23 (arriving in the order shown in Figure 3.21). As shown, a lack of sensing between the heterogeneous networks can lead to their sharing of a channel. Additionally, the FCFS arrival leads to the inefficient assignments of the 802.11 networks: not all networks receive their desired airtime due to contention given the assignments. We illustrate the resulting observed airtime fractions and loss rates in Figure 3.24. 802.11 networks 1 and 3 receive only a fraction of their desired airtime due to contention, and the ZigBee networks receive high loss rates due to their operation in channels shared by 802.11 networks.

Weakly Constrained 2: Next, we reconfigure some of the legacy supporting 802.11 networks to be high-throughput 802.11n networks i.e., they use the newer digital preamble that makes them incompatible with legacy networks. The resulting set of networks and their airtimes are shown in Figure 3.25. Leveraging the resulting conflict graph from our management system that accurately reflects conflicts between the digitally coordinating networks and legacy networks, our optimization continues to isolate the conflicting networks as shown in Figure 3.26. The result of a FCFS assignment of channel is shown in Figure 3.27. Again, the inability of the heterogeneous networks to sense each other leads to a mixed spectrum of conflicting networks, that leads to the observed high loss rates and reduced airtimes shown

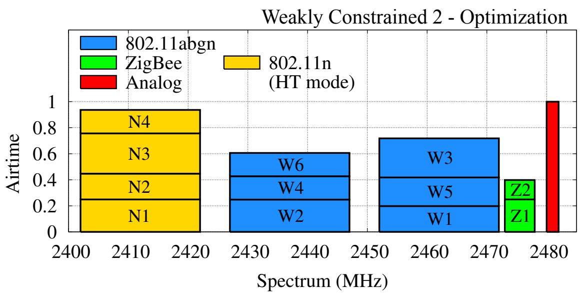

Figure 3.26: Heterogeneous technologies are again separated with enough spectrum available (e.g., 802.11n HT mode vs. legacy).

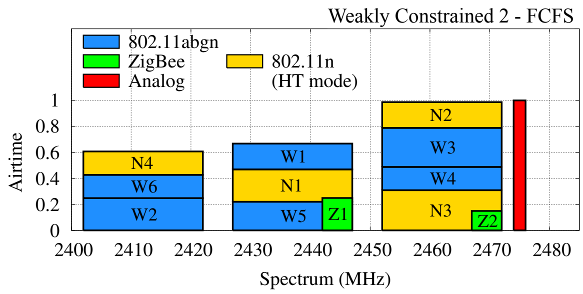

Figure 3.27: Resulting assignments given a FCFS configuration of the weakly constrained scenario 2.

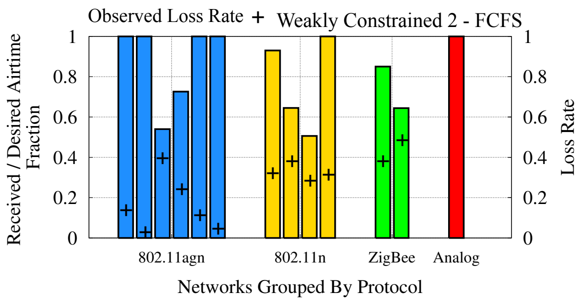

Figure 3.28: FCFS observed values of loss rates and received airtimes for each network, normalized by the networks desired airtime.

in Figure 3.28. The legacy supporting 802.11 networks receive slightly lower loss rates due to their use of basic spectrum sensing that allows them to back-off to the high-throughput networks. However, the high-throughput networks receive higher loss rates due to their strict digital-only sensing. The ZigBee networks also receive high loss rates from both sets of 802.11 networks. Although one of these ZigBee networks is still able to meet its desired airtime, losses lead to retransmissions, which can affect the ZigBee network's performance in terms of power consumption (a concern for the often battery powered ZigBee radios).

Weakly Constrained 3: In our final weakly constrained scenario, we introduce 40 MHz high-throughput Wi-Fi networks, and configure some of these networks to be able to use the 5 GHz band. These networks, their airtimes, and their order of arrival are shown in Figure 3.29. We present the resulting configuration given our spectrum assignment system and optimization in Figure 3.30: heterogeneous networks are again isolated, as well as digitally coordinating networks being properly aligned i.e., the 20 MHz high-throughput 802.11 networks are aligned with the primary channel of the 40 MHz networks. Our optimization avoids placing the 20 MHz high-throughput networks in the secondary channel of the 40 MHz

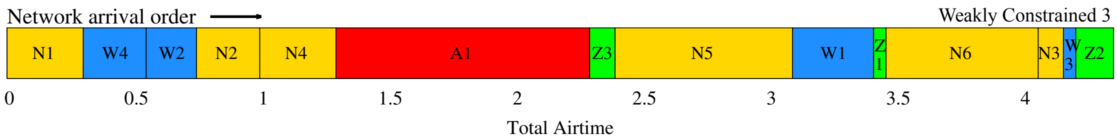

Figure 3.29: Networks in the weakly constrained scenario 3, ordered by arrival (for FCFS purposes).

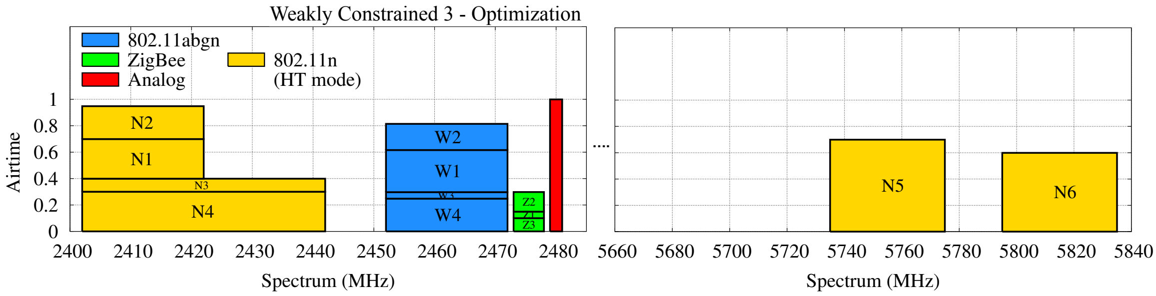

Figure 3.30: Network placement given our optimization. The scenario includes 40 MHz networks and the use of high-throughput 802.11 networks in the 5 GHz spectrum band.

networks, which would lead to their inability to coordinate. Due to potential contention, the networks that support the 5 GHz frequencies are pushed in to the 5 GHz spectrum band as shown. ZigBee networks are isolated, in addition to the analog cordless phone.

The FCFS configuration is shown in Figure 3.31, with its observed airtimes and loss rates in Figure 3.32. Again, the FCFS configuration suffers from low airtime fractions, and high loss rates. In particular, the 40 MHz 802.11n networks suffer greatly due to interference from legacy networks, as well as interference from the 802.11n HT network that resides in the secondary channels. Although the first ZigBee network receives all of its desired airtime due to retransmissions, its loss rate is over 50%, meaning that all transmissions will likely take 2 tries to be successful, which can significantly increase the power consumption of the radio on the network. The resulting configuration shows many inefficiencies that could be avoided by using our spectrum assignment system and algorithm.

Moderately constrained scenarios and evaluation

In the upcoming moderately constrained scenarios we will present, we increase the spectrum usage from the networks such that our optimization must intelligently place networks to try and avoid conflicts. The demand prevents our algorithm from simply isolating the networks as we have shown in the weakly constrained scenarios. Due to the inefficiencies in FCFS configurations shown extensively in our results, we focus solely on our optimization's results in these scenarios. We will revisit FCFS configurations (as well as largest-first insertion) in the severely constrained scenarios.

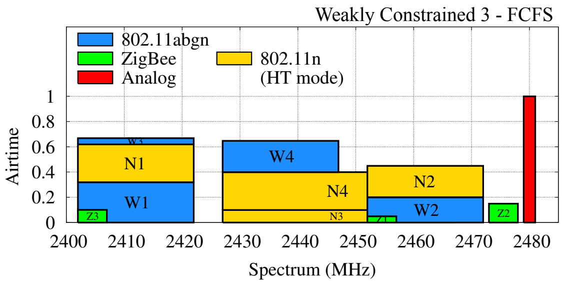

Figure 3.31: The resulting FCFS configuration of the weakly constrained scenario 3, showing overlap from heterogeneous networks.

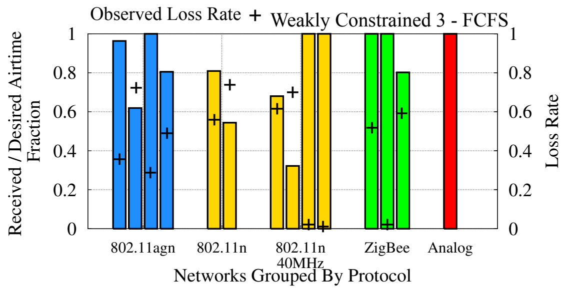

Figure 3.32: Resulting performance and loss rates of networks given the FCFS for weakly constrained scenario 3.

<u>Moderately Constrained 1:</u> In the first moderately constrained scenario, we reconfigure the first weakly constrained scenario such that the analog cordless phone operates in the ZigBee network's higher potential channels, forcing the ZigBee networks to operate within the 802.11 channels to avoid complete starvation from overlapping with the cordless phone. Our optimization should reconfigure the networks to avoid conflicts by considering SINR, providing an interference free configuration even though complete frequency isolation is not possible.

We show the resulting configuration provided by our spectrum assignment system in Figure 3.33 and ask the reader to compare the channel assignments of the 802.11 networks with those in Figure 3.22. With the ZigBee networks forced in to the 802.11 channels, our optimization chooses a different placement for the 802.11 networks to avoid potential conflicts. We draw conflict arrows between the 802.11 networks and the ZigBee networks to highlight which 802.11 networks interfere with the two ZigBee networks. As shown, our assignment avoids placing the ZigBee networks where they would conflict with 802.11 networks. To do so, our algorithm in particular reassigned several networks. Wi-Fi networks $W5$, $W7$, and $W8$ were reconfigured to share a channel, creating a channel where $Z1$ would receive no interference from conflicts. Likewise, $W3$ and $W2$ were moved together to create another interference-free channel for $Z2$.

Additionally, we note that our optimization still considered contention between the Wi-Fi networks and placed them in a way that still allowed them to meet their desired airtimes, given the shown configuration. Due to brevity, we do not show the resulting loss rates. However, the ZigBee networks received their desired airtimes with no loss, and the Wi-Fi networks also received their desired airtime.

<u>Moderately Constrained 2:</u> Next, we take the third weakly constrained scenario and add two additional legacy 802.11 networks to it, pushing the demand higher which does not allow complete isolation, requiring an intelligent placement to avoid conflict. If one places these two networks in 802.11 channel 11 to operate with the other legacy networks shown in

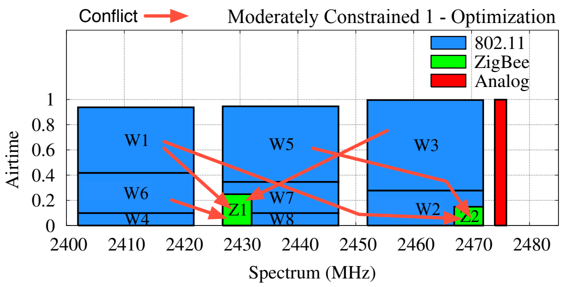

Figure 3.33: With the ZigBee networks forced to operate in 802.11 channels, our optimization still assigns them to avoid conflict.

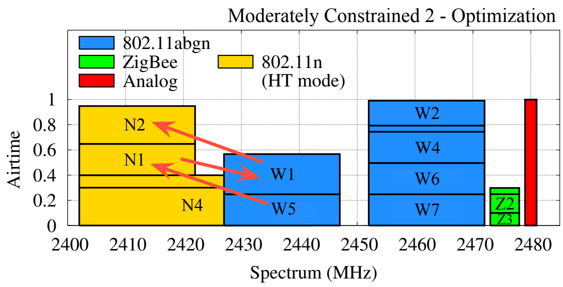

Figure 3.34: With additional demand from legacy W-Fi networks, these networks are re-assigned to intelligently avoid conflict also.

Figure 3.30, all legacy networks will receive a reduced airtime due to the constrained spectrum and contention. Instead of taking this choice of isolation to avoid conflict, leading to reduced performance, our optimization determines that 802.11 networks $W1$ and $W5$ (where $W5$ was newly introduced), do not conflict with the 40 MHz networks allowing them to operate in their secondary channel freely. We show the resulting configuration in Figure 3.34 and draw the conflicts to show the intelligent placement which avoids conflict. Again, due to brevity we do not plot the resulting loss rates, however, networks $W1$ and $W5$ received no observed loss rates, as well as the two 40 MHz networks not receiving loss from these networks. This shows that under more constrained spectrum, our optimization can still find efficient organizations by intelligently placing networks to avoid conflict.

Severely constrained scenarios and evaluation

Finally, we introduce severely constrained scenarios where spectrum demand prevents configurations that allow isolation, or intelligent placement to avoid all potential conflicts. Instead, the networks must be intelligently placed to avoid as many conflicts as possible, and networks should be placed to receive the least possible loss. We compare the performance of our optimization at achieving these goals, and compare the configurations with resulting FCFS configurations, in addition to two new points of comparison: 1) A largest-first insertion of networks in to the spectrum where the networks are able to leverage our channel quality metric that predicts loss due to heterogeneous networks, and 2) Our optimization that a Jain's fairness-based objective function to try and ensure that no network receives significantly unfair airtime. The prior comparison allows us to compare to intelligently placing networks without the runtime of the full optimization, and the latter comparison allows us to compare to an objective of fairness across the networks.

Severely Constrained 1: We begin with the configuration of networks shown in Figure [**?**], which shows a high total airtime demand and heterogeneous networks that can lead to conflict. We show the resulting airtimes of the 4 different assignment methods in Figure 3.36.

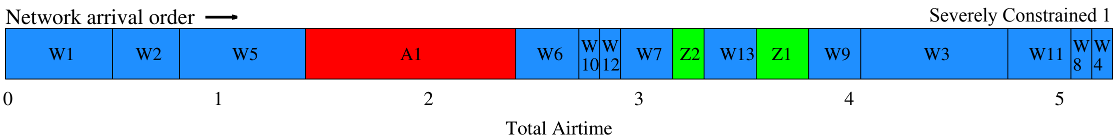

Figure 3.35: Networks in the severely constrained scenario 1, ordered by arrival (for FCFS purposes).

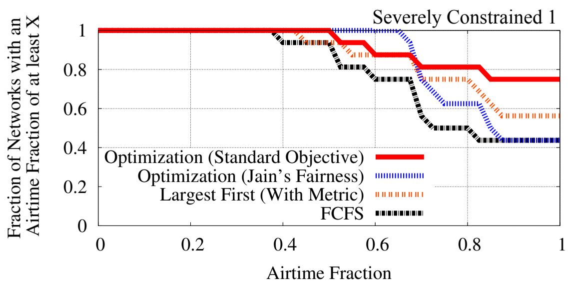

Figure 3.36: Resulting airtime fractions given the four different assignment methods under the severely constrained scenario 1.

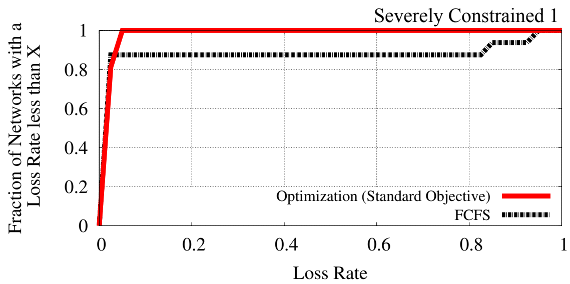

Figure 3.37: Loss rates can be exceptionally high with a FCFS assigned environment, compared to our optimization.

The figure shows, for each method, the fraction of networks that receive at least a certain amount of airtime. First, our optimization with the standard objective finds a configuration that has the most networks meet their desired airtime: 75% of networks. Comparing to the FCFS configuration where only 43% of networks meet their desired airtime. Even with our predictive channel quality metric, a largest-first method of insertion and assignment cannot meet the performance of our optimization (with either objective). Note that the organization provided by our optimization with the standard objective results in a few networks that receive relatively lower (and un-fair) airtime fractions. Using the Jain's fairness objective with our optimization results in a configuration where no network receives as low of airtimes, but to achieve this several better performing networks receive lesser airtime fractions: only 58% of networks receive their desired demand with the Jain's fairness objective, compared to the 75% with the standard objective.

In Figure 3.37, we show the resulting loss rates for the networks, focusing on the optimization with the standard objective and the FCFS method of assignment. As shown, neither method is able to avoid complete loss across all of the networks, however our optimization ensures that no network received a loss rate greater than 8%. This is compared to the FCFS method of assignment where some networks receive significant loss rates between 80-90% (which has been observed by prior works in heterogeneous environments [7,8,16]). These results show the benefit of our spectrum assignment system, as well as its flexibility to use different objectives if fairness is desired over better overall performance for all networks.

Severely Constrained 2: We conclude our evaluation results with the severely constrained

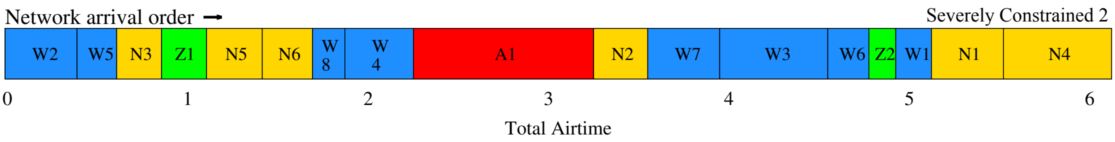

Figure 3.38: Networks in the severely constrained scenario 2, ordered by arrival (for FCFS purposes).

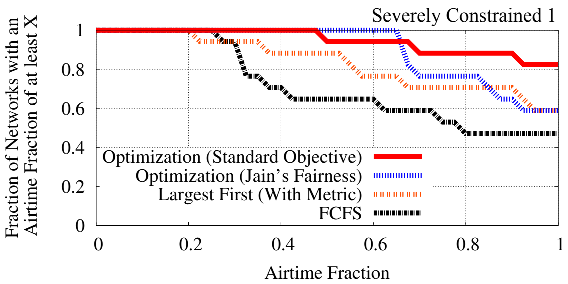

Figure 3.39: Resulting airtime fractions given the four different assignment methods under the severely constrained scenario 2.

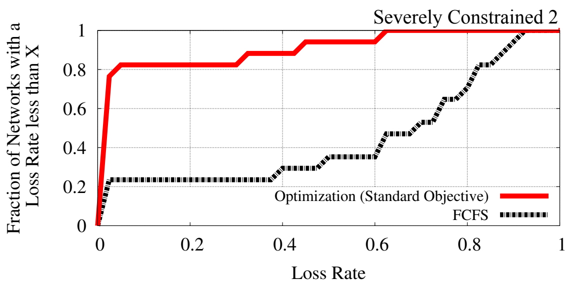

Figure 3.40: Our optimization continues to avoid higher loss rates experienced in FCFS assignment.

scenario of networks depicted in Figure 3.38. There is a high degree of heterogeneity and overall desired airtime. Intelligent placement will be needed to reduce loss rates given conflicts (since isolation will not be possible). We show the resulting airtime fractions of the networks given our four insertion methods in Figure 3.39. Again, our optimization with the standard objective finds a configuration where the most networks are able to meet their desired airtime, where 82% of networks meet their demand. A FCFS configuration has less than half of the networks meeting their demand, and 60% of networks meet their demand using our optimization with the Jain's fairness objective and the largest first method of insertion. The Jain's fairness objective continues to provide the best performance for the least-performing network, but again pulls down the performance of other networks to do so.

In Figure 3.40, we compare the loss rates observed using our optimization with the standard objective as compared to the FCFS insertion. With severely constrained networks given their demand and the available spectrum, as well as the high degree of heterogeneity, FCFS is shown to result in high loss rates for many of the networks. Only 20% of networks receive a loss rate lower than 40%, with some networks receiving loss rates greater than 80% again. This is compared to our optimization where high loss rates are still observed (e.g., greater than 50%), but 80% of networks still receive minimal loss rates (e.g., less than 10%). This shows the strength of our optimization to provide configurations that reduce loss and improve overall performance, despite being severely constrained.

Evaluation Summary

In this section, we presented many diverse points of evaluation across various types of heterogeneous conflicts commonly reported in today's environments and prior studies (e.g., [7, 8,16]). Unlike these prior works, however, we explored the ability of spectrum management to avoid conflicts rather than coexistence techniques. Our evaluation showed that a properly designed spectrum management system for heterogeneous networks can avoid these various types of conflicts with a single algorithm (i.e., solution). Our targeted scenarios showed this ability, focusing on various commonly reported conflicts, and our weak, moderate, and severely constrained multi-network evaluation points showed our ability to configure multiple networks to avoid various potential conflicts in the spectrum. The resulting configurations showed isolation when possible, and configurations that avoid high loss rates when isolation is not possible due to high demand. Comparing to FCFS configurations, we were able to significantly reduce loss rates and increase performance of networks with the potential for heterogeneous conflicts.

3.5 Chapter Summary

Following the thesis of this dissertation, we explored the potential of spectrum management as a long-term solution to reducing interference between the many (and quickly evolving) unlicensed heterogeneous wireless technologies. Again, this strategy is contrary to the common approach seen today of developing coexistence techniques to alleviate interference between these heterogeneous networks (an approach that requires N^2 solutions). The system and work we presented in this chapter show the potential of spectrum management: through careful design, our system and algorithm (a single solution) can address various types of conflicts between various heterogeneous technologies, reducing interference and improving efficiency.

To achieve this goal, we introduced novel contributions in 3 key areas: 1) A hypergraph-based RF environmental model that can represent an environment with heterogeneous technologies, their interactions, and their constraints, 2) A predictive channel quality metric that can estimate the performance of a network, accounting for contention from coordinating networks and interference from heterogeneous technologies, and 3) A mixed integer program (MIP) based optimization that accounts for the various constraints across heterogeneous technologies, and can efficiently (re)organize the spectrum to decrease contention and avoid cross-technology interference.

Through evaluation with real heterogeneous radios in a wireless testbed, we illustrated that our spectrum assignment system (comprised of these 3 main components) can account for and avoid various types of conflicts and interference that we have traditionally relied upon coexistence techniques to alleviate. Importantly, this illustrates that a single solution (spectrum management) can address and alleviate the general problem of interference between heterogeneous networks.

Limitations and Future Research: There are several limitations and simplifications in our system that need to be addressed to support growing trends in the spectrum, and future research that can be done to improve its efficiency and understand its limitations. However, our system and optimization is designed to be flexible enough to support these additional trends with simple changes. We outline a few major areas that need to be addressed in the future to overcome, improving the potential of our spectrum assignment system.

- *Fixed spectral bandwidth:* Currently, the spectral bandwidth of a radio B (in our optimization) and operating on any of its center frequencies is considered to be fixed. For many standards and protocols, the bandwidth is fixed as we model it. However, to deal with the increased demand for bandwidth and throughput, Wi-Fi and its 802.11n/ac protocols have introduced the option of wider bands such as 40 MHz, 80 MHz, and even 160 MHz. Clearly, one can set B to its desired bandwidth, however, it is possible to extend our optimization framework to have B be a set of potential bandwidths. Then, the algorithm can choose B to optimize its throughput while avoiding contention and interference. Prior work has shown large values of B are not always desired and can provide lower throughput due to increased contention and interference [113].

- *Disseminating frequency assignments:* Although we collect the necessary information to perform spectrum management, and assign spectrum through the system we introduced in this chapter, we do not introduce a way to disseminate these assignments and reconfigure the end devices. In enterprise environments, there is expected to be a network administrator with a mechanism and system that enables them to reconfigure the frequencies of the devices in their domain. However, in chaotic environments and, particularly in the home, there is no unifying mechanism or standard that allows a central spectrum management system to reconfigure a diverse and heterogeneous set of networks. This is important future work to deploy our system.

- *Thorough sensitivity analysis:* We evaluate our system using accurate information collected by our smartphone-based monitor. However, our system is expected to work with various types of monitors as long as they can produce the hypergraph-based representation of the environment we outlined in our system design. As future work, it is important to understand the impact of the accuracy in the information provided to our spectrum assignment system. In particular, the accuracy of the meta-data provided such as the link reception strengths and interference strengths, and the loss estimate tables that provide the expected loss rate given an SINR value. The study should evaluate how sensitive these values are, and others (e.g., traffic loads, expected coordination, and spatial range), to understand how accurate the information needs to be to provide effect and efficient spectrum assignments.

- *TDMA Networks:* In our work, we focused primarily on CSMA networks. However, TDMA networks also create and suffer from interference. While we do not consider such networks in our optimization and assignment, it is possible to use our predictive metric to choose what channels a TDMA and frequency hopping network should use or avoid. For example, by using the metric across all of its channels and then having it adaptively avoid channels where expected performance does not meet some threshold.

Chapter 4

Enabling Primary Coexistence in the White Spaces

Up until this point in our dissertation, we have motivated, studied, and built better heterogeneous monitoring and spectrum management techniques to alleviate the general problem of heterogeneous interference, particularly between unlicensed devices. Spectrum management proved to still be effective and efficient since, in many cases, it is possible to reconfigure network frequencies in a way that they may still overlap with other heterogeneous networks, but receive small enough amounts of interference to not create loss. As we have argued, spectrum management is a good long-term solution to general heterogeneity since it can handle changes in technology without a need to change the solution.

In this chapter, we provide a contrasting study of heterogeneous interference in an increasingly popular scenario where the assumptions, rules, and regulations are extremely different than what we have addressed so far in this dissertation. This scenario is in the "white space" i.e., licensed spectrum with unlicensed access in its idle parts (described in Chapter **??**). In these bands, one must now consider the existence of *spectrum primaries* (i.e., licensed and heterogeneous users). Unlike the unlicensed devices that we have addressed thus far, primaries have strict rules that protect them against interference (regulated by the FCC [13,18,20]).

Although we have shown spectrum management to be effective and efficient at reducing heterogeneous interference between unlicensed devices, here, we will show that spectrum management can be extremely inefficient at protecting licensed users. As a result, rules and regulations put in place to guarantee interference-free operation through spectrum management end up threatening the very goal of the white spaces: *additional spectrum.*

Instead, we argue that coexistence techniques are better suited to protect licensed users in the white spaces. This is for several key reasons and differing assumptions than the problem of general heterogeneity we have dealth with so far:

1. While the coexistence-based approach is N^2 regardless of the spectrum band, N is small in the white spaces. First, the number of spectrum primaries in any licensed band is small (e.g., only 2 in the TV white spaces), and second: the *unlicensed* devices in these bands have certain restrictions that make their contribution to N small. In the TV white spaces, unlicensed devices are restricted to be wideband. This eliminates various technologies and will likely lead to many OFDM-based unlicensed devices in the spectrum. These devices can use similar (or the same) coexistence techniques.

2. Unlike unlicensed devices and their standards which constantly change (e.g., 802.11), licensed devices do not evolve quickly (if at all). Licensed devices are granted operation based on a specific technology. For example, wireless microphones (spectrum primaries in the TV white spaces) must use narrowband FM signals. Digital microphones are not licensed to operate in this part of the spectrum. Without much or any evolution in the primaries, coexistence techniques to address them will likely be long-lived.

3. Coexistence techniques can be more spectrum efficient by allowing the licensed and unlicensed devices to operate on the same channel. In particular, this difference in efficiency shows when there are strict rules against interference. Like isolation in the frequency domain, however, coexistence techniques can still provide interference-free guarantees if designed properly.

For these reasons, in this chapter, we explore a coexistence-based approach to solving heterogeneity with licensed devices in the white spaces. In particular, our goal is to recover a significant amount of spectrum availability that is lost in the bands due to current spectrum management-based protection mechanisms. This chapter takes an important step towards ensuring the white spaces meet their main goal of providing additional spectrum, and demonstrates that this is possible through a practical and deployable coexistence system that can still guarantee interference-free operation around the spectrum primaries. Note that spectrum management can still be used in the white spaces to mitigate interference between the unlicensed devices. This makes our work up until this chapter still applicable in this band.

Chapter outline: Section 4.1 introduces the problem of spectrum availability in the white spaces and heterogeneity's threat to it. We present a background on the spectrum primaries in Section 4.2, as well as related work in avoiding interference with spectrum primaries. In Section 4.4, we present the first in-depth analysis of how wireless data transmissions impact the primaries, and use our findings to design an interference-free coexistence system that can reclaim the spectrum surrounding the primaries in Section 4.5. We present a full prototype of our coexistence system in Section 4.6, and show through evaluation that it is effective at avoiding interference, and efficient at reclaiming significant amounts of white space (*over 100 MHz in dense scenarios*). We conclude the chapter in Section 2.7 with the potential impact of our work, lessons learned, and future research.

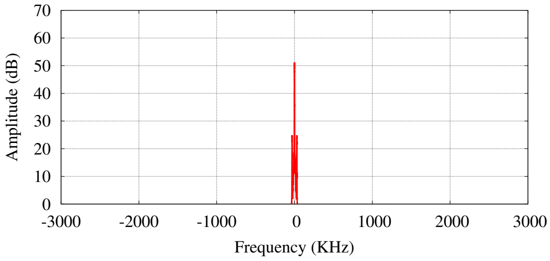

Figure 4.1: A single microphone at the center of a 6 MHz white space channel.

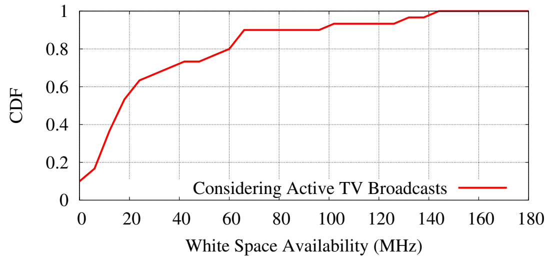

Figure 4.2: Current white space availability in the largest 30 US cities (by population).

4.1 The TV White Spaces and Inefficient Management

The proliferation of wireless devices has led to an impending spectrum crisis [114]. To provide more spectrum, the FCC and spectrum regulators worldwide are exploring techniques to reuse unoccupied TV channels (white spaces) for data communication [17]. The FCC finalized its rules on September 23, 2010 [20], while the UK, Brazil, Finland, Singapore and other countries are working on white space rules as well. While such efforts have been significant, the rules in place to protect primary users threaten the very goal the white spaces are trying to achieve: additional spectrum availability.

This is particularly true regarding the ruling's handling of wireless microphones (mics), which along with TV broadcasts are primary users of this spectrum. In the First Order from December 2008, the FCC ruled conservatively towards mic protection by enforcing a white space device to vacate an entire TV channel (6MHz) in the presence of even a single mic (at most 500KHz). Referring the reader to Figure 4.1, we show a single microphone at the center of a 6 MHz white space channel. This figure illustrates just how inefficient vacation is, especially if there is only a single mic in the channel. In the Second Order, the FCC made the ruling even more conservative by stating that two channels will now be exclusively reserved for mics. This is *in addition* to allowing licensed mics to operate in *any* channel under database (or sensing) and channel vacation protection.

Particularly in populated areas, these rules significantly reduce spectrum availability for white space devices. To illustrate this, we refer the reader to Figure 4.2. This figure shows the CDF of white space availability (given active TV broadcasts) in the top 30 US cities by population. In 12 of these 30 cities, there are only 2 or fewer unoccupied TV channels, and in 21 of these cities, there are no more than 5. Dedicating two TV channels for wireless microphones effectively eliminates white spaces in a large fraction of major US cities. Moreover, as microphones will still be able to operate as primary users in any of the other TV channels, the amount of white spaces in some locations of the remaining cities will effectively be reduced to zero. It is clear that such conservative measures run counter to the goal of white space networking, and are likely to be a major impediment to its widespread adoption.

Problem Scope and Challenge: In this chapter, we show that such conservative and inefficient spectrum management is unnecessary. Instead, full microphone protection can be achieved through coexistence, while still enabling white space devices (WSDs) to reclaim large fractions of the spectrum. It is well-known that in different parts of the spectrum that wideband devices which often OFDM can use subcarrier suppression [42,45,115] to eliminate interference with narrow-band devices. In the white spaces, such an approach would potentially allow a WSD to coexist with a mic in the same TV channel, and use the remaining 95% of its spectrum [115].

Unfortunately, such solutions cannot easily be deployed in the white spaces for several reasons. First, existing techniques that adaptively determine the degree of required suppression require interfering with the narrowband devices (i.e., SWIFT [42]). This is unacceptable as it would cause harmful audible interference with the mic. Second, the amount of suppression required to protect the mic's transmission depends on the mic's received signal power, as well as the white space device's interference, at the mic receiver. Both values are unknown at the WSD, change over time, and cannot be estimated using channel reciprocity techniques [116] given mic systems are one-way (receiver never transmits).

Summary of Approach and Contributions: To overcome the challenges, we present SEISMIC (Spectrum Efficient Interference-free System for MICs), a system that allows white space devices to coexist with mics and "recover" a close to optimal fraction of the spectrum in the TV channel, while fully protecting the mic in all circumstances. Our system design is based on an in-depth characterization of the impact of white space transmissions and RF interference on mic audio recordings. To allow cooperation between the mic system and WSDs, SEISMIC uses a simple device called a MicProtector to measure the interference at the mic receiver and a low-complexity signaling protocol to notify white space devices of *impending disruption*. Using this explicit signaling feedback, secondary users can suppress the proper frequency to avoid disrupting the mic's audio quality.

The SEISMIC approach combines several attractive properties. First and foremost, it is safe since the explicit feedback avoids harmful interference at the mic receiver, independent of the placement of mics and WSDs. Secondly, is purely reactive, restricting secondary white space communication only when needed and to the degree necessary. That is, SEISMIC optimizes spectrum usage by minimizing the number of subcarriers that are suppressed based on actual measured WSD signal levels. Finally, in spite of the addition of the new MicProtector device, SEISMIC is a very practical solution: It does not require replacing legacy mic equipment or advanced registration of mics before events, and cost is likely to be small (no low-threshold sensing!).

We find that SEISMIC allows white space devices to converge to within 25KHz of optimal suppression 72% of the time and within 75KHz 93% of the time, with zero-interference to microphones. We show that this allows a WSD to get up to 95% of the bandwidth of a channel that is completely lost with today's restrictive FCC regulations, and even up to 85% of the channel in many (10+) mic scenarios. While these results are specific to mics in TV white spaces, we note that the SEISMIC design is more general and can be easily adapted

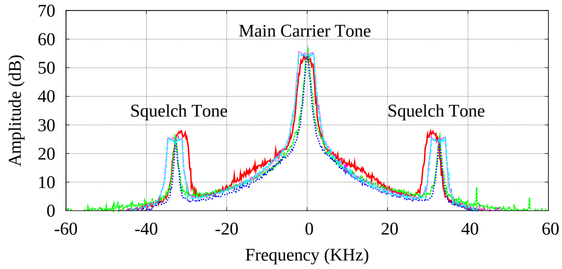

Figure 4.3: Spectrum signature of 6 idle microphones.

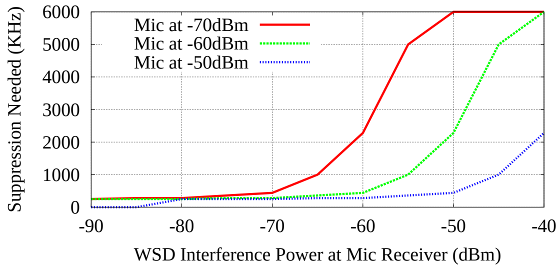

Figure 4.4: The suppression needed by the WSD varies.

Microphone	Form Factor	Idle Width	Squelch/Peak
Audio-Technica ATW-T210	Handheld	65 KHz	30 dB
Electro-Voice BPU-2	Beltpack	62 KHz	25 dB
Sennheiser E935	Handheld	69 KHz	30 dB
Sennheiser EW100	Handheld	66 KHz	30 dB
Sennheiser SK2000XP	Beltpack	69 KHz	30 dB
Shure UR-2	Handheld	65 KHz	30 dB

Table 4.1: Microphones used throughout our study.

for efficient coexistence with other primary users. For example, SEISMIC can enable WSDs to transmit at greater than 40 mW on channels adjacent to those occupied by TVs.

4.2 Background and Related Work

In this section, we provide the necessary background on wireless microphone signals for the reader to understand this chapter, followed by related work in detecting and avoiding wireless microphones.

Microphone Signal Background: A wireless mic system consists of a mic transmitter and a mic receiver. The mic transmitter converts audio into RF using frequency modulation (FM). The mic receiver decodes the FM signal to retrieve the transmitted audio signal. Figure 4.3 shows the RF spectrum of the signal when six mics are idle. The signal consists of a center signal that carries data and two side tones, called *squelch tones*. The mic receiver decodes the mic signal only when the squelch tones are successfully received. This helps protect against garbled sound when there is interference, and prevents risk of audio amplifiers and speakers when the mic signal is low [117]. Table 4.1 summarizes the properties of the 6 mics used in our work.

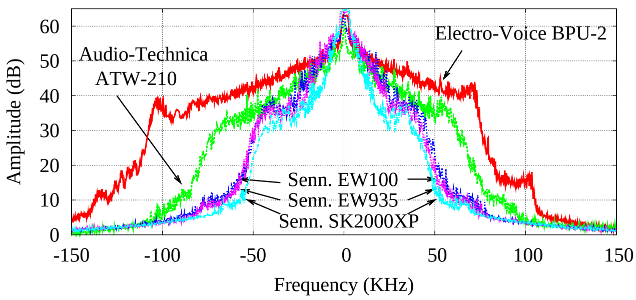

Figure 4.5: The maximum spectrum use of 5 of the microphones.

Although every mic has a similar RF signature when idle, Figure 4.5 shows the *maximum* spectrum values for the mics after 10 minutes of continuous audio recording. As shown, different mics hop differently across the spectrum given the same audio.

Licensed mics operate under the FCC Part 74 rules in the US, and similar rules worldwide. This rule restricts the operation of wireless mics to at most a 200 KHz wide signal, and a max transmit power of 250mW, although most mics use a max of 10mW. Unlicensed mics are allowed to operate under the FCC Part 15 rules at a maximum transmit power of 50 mW. Worldwide, mics are allowed to operate on any unoccupied TV channel. In the US, the FCC recently modified the rules and reserved two TV channels exclusively for wireless mics [20]. These two channels vary by region, and cannot be used by WSDs. When two channels are not enough, organizers can reserve additional channels 30 days in advance.

Previously proposed solutions to protect mics:

1. Sense for wireless microphones: This was amongst the first proposed and most popular initial solutions proposed to avoid interference to mics [21,49,118,119,120,121,122]. The FCC's initial ruling set this threshold to -114 dBm over 200KHz, and the Second Order reduced it to -107 dBm. OFCOM is considering a sensing threshold of -126 dBm. Such a sensing-based approach has several drawbacks. First, spectrum sensing is an expensive operation in terms of cost an energy [20], and was the primary reason for removing this requirement in the Second Order. Second, the sensing threshold is extremely conservative, both because the signal propagation environment is unknown and because it is meant to protect the mic receiver, whose location is unknown by sensing the mic. Third, sensing for mics at below the noise level (as mandated by existing regulations) is prone to false positives [123]. Therefore, WSDs might end up vacating an entire TV channel even when there is no mic present in the vicinity. Finally, and most importantly, this approach is inefficient, since WSDs need to vacate an entire TV channel.

2. Microphone Beaconer: To address the second and third concern above, mic companies, such as Motorola and Shure, have proposed the use of a separate beaconer device to reduce the sensing threshold [50,124,125]. The beaconer uses the first 500 KHz of every TV channel to signal the presence of a mic using a 250 mW signal. White space devices vacate every TV channel that has this beacon. This approach still suffers from the other two drawbacks from above – sensing needs extra hardware and entire channel needs to be vacated. Another shortcoming of this approach is that WSDs may vacate the channel even when their transmission does not interfere with the mic. We elaborate on this in Section 4.4.

3. Two-Reserved & On-Demand Reservation: This is the approach taken by the FCC in the Second Order. Reserving two TV channels for wireless mics (even when there are no mics in the vicinity), as specified in the FCC's Second Order [20], significantly reduces the amount of white spaces in urban areas. In the top 30 urban areas, our analysis showed that 12 (40%) had only two unoccupied TV channels, 60% had three or less, and 70% had 5 or less[1], so two channels represents a significant fraction of the white space spectrum. This is a serious limitation since success in urban areas, where more WSD users are expected, is seen as a likely driver for white space device use in rural areas [20]. pace availability is further reduced since event organizers can also reserve any TV channel for mic operation.

Despite this conservative approach that seemingly favors mic users, the audio community is concerned about the ruling as well [126]. The ruling leads to increased cost since most mic users will have to replace their existing equipment. This is because the 2 reserved channels will be geo-dependent as the first 2 free channels in the 180MHz of spectrum, however mic systems usually have a limited 40MHz front-end which is likely to not be in the range of the reserved channels. Furthermore, users who pay a large sum of money for their mic placement [21] will have to redo the mic placement. Mic operators, including those who handle big events such as the Super Bowl, are unhappy about having to reserve TV channels 30 days in advance [126]. The RF environment changes frequently, and they adjust frequencies until the last minute.

4.3 Designing Spectrum-Efficient Coexistence for the White Spaces

The goal of SEISMIC is to maximize the amount of spectrum available for white space communication, while ensuring no interference to the primary users. The WSD can suppress a portion of frequency around a narrow-band primary user's transmission to avoid interference with the primary user [42,45,115,127]. This would suggest that if the secondary can learn about the exact transmission frequency and bandwidth of the primary user, the primary user could vacate a sufficiently large "guard-band," and use the remainder of the channel without interfering. Here, a database or some beaconing-device could inform of such information.

Unfortunately, such non-adaptive, open loop solutions are very inefficient. The amount of spectrum to suppress depends on the SINR at the mic receiver, i.e. the ratio between the received signal strength from the primary transmitter and the collective interference power

[1] Geo-location database: http://whitespaces.msresearch.us/

generated by the WSDs. Since, the mic receiver's SINR is not known at the WSD, it needs to make the most conservative, worst-case assumptions. This results in suppressing too much spectrum. To show the inefficiency of such a static, open-loop solution we measure the amount of suppression needed with various SINR values when using a high-quality WSD prototype from Adaptrum and real mic system. Since the WSD does neither know the interference power it creates on the mic nor the received signal strength at the mic receiver, it has no idea where in Figure 4.4 it operates. Thus, it has to make worst-case assumptions to avoid interference: it must suppress the entire channel—which is exactly what the FCC requires.

4.3.1 Towards an Adaptive Solution

Several adaptive solutions have been proposed to determine the proper amount of suppression [42,43,128,129]. They can be broken into two groups: (i) those that adapt based on how the the primary user reacts to interference [42,128], and (ii) those that use channel reciprocity to estimate the interference the secondary user generates, while assuming the worst case for the (unknown) signal component (mic or TV).

Unfortunately, none of these solutions are suitable for coexistence in the white spaces. First, we cannot allow any disruption of primary users (both mics and TV broadcasts). Second, even if we were allowed to interfere temporarily, mics and TV broadcasts are passive transmissions; they do not back off. Third, channel reciprocity cannot be done with primaries in the white spaces since all white space primaries are one-way systems in which the receiver never transmits. Hence, the WSD cannot estimate the interference component of SINR at the primary receiver. Even if you could solve that problem,[1] the signal component of the SINR is still unknown. The WSD would have to be conservative.

Thus, the challenge is how to devise a system in which the WSD can adapt their behavior based on the SINR-values (RF interference and signal strength) at the primary receiver; and to convey this information to the WSD in a *disruption-free* manner from the passive device. This should be done without adding significant complexity to the WSD or primary.

4.3.2 SEISMIC Design

In the design of SEISMIC, we try to approximate an ideal solution in which WSDs have explicit feedback about the SINR-values (RF interference and signal strength) at the primary receiver, allowing them to suppress the minimal number of subcarriers while avoiding any disruption of the mic system. Adaptive systems based on closed loop feedback must include three logical components: measurement, analysis, and adaptation. Since our system is distributed, we also need a signaling component to exchange information.

SEISMIC implements these components as follows:

[1] For example by deploying some "beaconer" device co-located with the mic receiver [50,124,125].

1. **Measurement:** We introduce a low-complexity enhancement to the mic receiver called a *MicProtector*, that monitors RF interference and mic signal strength. Being co-located with the primary receiver, it accounts for all relevant factors when monitoring its SINR. (Section 4.5.2)

2. **Analysis:** With SINR, SEISMIC accordingly determines how the system can be optimized, i.e., how to minimize the number of suppressed subcarriers while avoiding audible interference. (Sections 4.4 and 4.5.2)

3. **Adaptation:** The WSD follows a protocol in which it adjusts both the number of suppressed subcarriers and the transmit power of its transmission. (Section 4.5.3)

4. **Signaling:** When needed, the MicProtector sends feedback to the WSD using a novel *signaling* mechanism (*strobing*) to warn of impending disruption. (Section 4.5.4)

While we depict the MicProtector as a standalone device, future mic systems can build such functionality in to the receiver. The option of using a standalone device is attractive since it allows deployment without replacing all mic systems. This is similar to the use of converter boxes to cope with the DTV transition. Mic manufacturers have been willing to adopt an additional device to signal the presence of the mic [50,125]. In private communication, mic operators at large events have been more willing to add this device than replace their existing mic systems because of the high cost of replacing equipment and of replanning frequencies.

In the next section we present a detailed measurement study on the the impact of RF interference on audio quality; this leads to the *analysis* component. The other components and their integration, are discussed in Section 4.5.

4.4 The Impact of Secondary Interference

In this section, we provide the first in-depth analysis of how wireless data transmissions impact wireless mics. Such an understanding is critical towards *analysis* and proper *adaptation.* Using a controlled environment (§4.4.1), we introduce variable RF interference by independently controlling the *power*, *duration*, and *frequency.* We measure the amount of audible interference using the *Perceptual Evaluation of Speech Quality* (PESQ) metric [130] (§4.4.1).

4.4.1 Experimental Setup

Our experimental setup is illustrated in Figure 4.6. We use a PC to play sound samples. We place a wireless mic close to the PC speakers, mimicking a person speaking into the mic. To ensure that any audio disruption is caused only by WSD interference and not from other sources, we placed the PC speakers and the mic in an anechoic chamber. The mic receiver

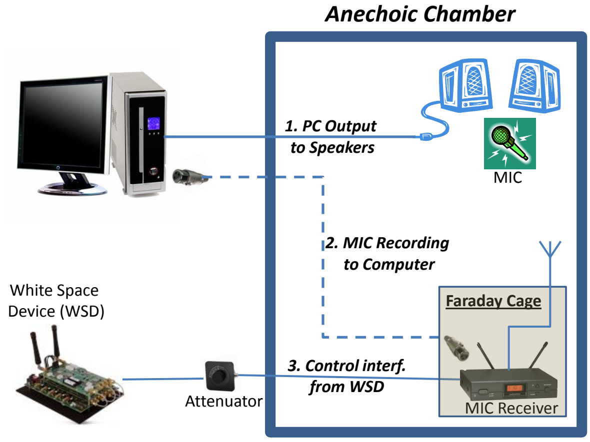

Figure 4.6: Anechoic chamber setup for audible interference tests.

is connected to the PC using an XLR cable, where we save and process the resulting mic recording.

To study the impact of WSD interference on mic recordings, we used WSDs to transmit in a conducted setup to the mic receiver. The WSD was wired directly to the mic receiver to introduce the interference in a controlled manner. We also connected a spectrum analyzer to measure the RF spectrum and channel power. Since mic receivers have exposed antenna elements, we isolated them by placing the receiver in a Faraday cage. To control the power of the transmitted and received signals we placed two RF attenuators: between the WSD and its antenna and between the mic receiver and its antenna. We ran tests with two WSDs: Adaptrum, and a USRP2 with a TV TX/RX (WBX) front-end.

We used PESQ to quantify the impact of the interference introduced by the WSD on the mic recording. PESQ is a signal processing algorithm that provides an estimate for the Mean Opinion Score (MOS), a widely used measure of subjective sound quality, i.e. how humans perceive the quality of sound. PESQ outputs a number from one to five to mimic the MOS results. In our work we used the wideband version of PESQ algorithm called WB-PESQ, standardized in 2005 by ITU-T Recommendation P.862.2. For ease of explanation, we present *normalized* PESQ values. A score of 1 is perfect quality, and score of 0 represents heavy disruption.

4.4.2 Interference in Power

In the first set of experiments, we have the WSD generate interference continuously (worst case in *time*) and adjust both the power of the mic signal (P_m) and power of the white space interference (P_n) at the mic receiver. For reasons that will become apparent in the results, we focus on the amplitude of the white space interference in relation to two separate power components of the mic signal: mic peak power (P_m) and the power of the squelch squelch

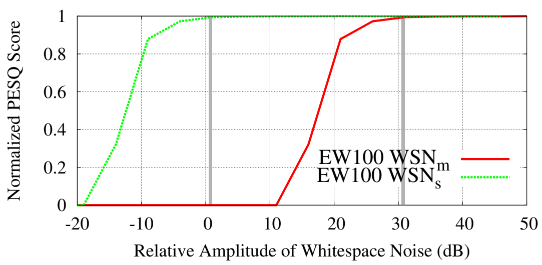

Figure 4.7: Audible interference, varying WSN. Areas where score reaches 1 are highlighted.

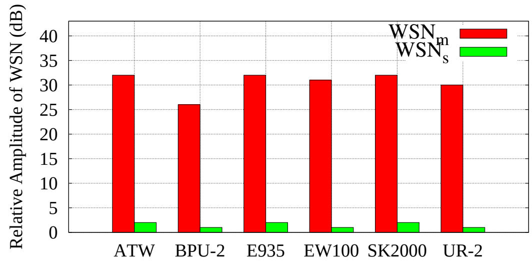

Figure 4.8: Relative amplitude of the WSN for each mic when a perfect score of 1 is reached.

tones (P_s) (§4.2). Therefore, the amplitude of the white space interference in relation to either component, *measured at the mic receiver*, is computed as: $WSN_m = P_m - P_n$ and $WSN_s = P_s - P_n$, where the WSN is in dB. So, if P_m is -30dBm, P_s is -60dBm and P_n is -40dBm, then WSN_m is 10dB and WSN_s is -20dB.

Noise Generation and Measurement: We use the variable attenuator to control the RF interference generated by the WSD. The interference level is set to 100mW (20dBm), and for each test we attenuate it in 5dB steps until the normalized PESQ value begins to approach 1. We then decrease the step size to 1dB for accuracy. We measure the power values (P_n, P_m, P_s) at the mic receiver using the spectrum analyzer. Since it is attached to the RF input ports, we can measure the power as close to the RF chain as possible, accounting for factors such as attenuation in the cables.

Power Results: Figure 4.7 shows the normalized PESQ score as a function of the white space interference level on the Sennheiser EW100 mic, marking the points at which the PESQ value becomes perfect with vertical grey lines. Our results show that if the WSN amplitude is greater than the mic signal peak (i.e., $P_n > P_m$), the interference is severe enough to cause the mic receiver to stop transferring audio (due to the squelch tones) as is seen from a normalized PESQ score of 0. However, once the peak is approximately 10dB above the white space noise (i.e., $WSN_m \geq 10dB$), the noise becomes less severe and the voice in the audio track becomes noticeable. Most surprisingly, *once the mic squelch tone power was 1dB above the white space noise (i.e., $WSN_s \geq 1dB$), the normalized PESQ score achieved a perfect value of 1.* This is *despite the fact that 19dB of RF interference still present in the operating band of the mic.*

We repeated the same experiment for the other mics. Figure 4.8 shows the same result holds: as soon as the white space interference level is a few dB below the squelch tones, we get a perfect PESQ score. The result even holds for the BPU-2 mic, which has squelch tones that are separated by 25dB from the mic signal peak (30dB for the other mics).

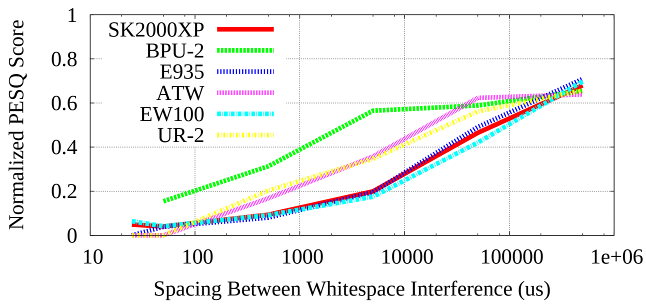

Figure 4.9: Resulting interference with variable spacing of transmissions in time.

To verify that this result is independent of the power of the mic, we repeated the experiments with the attenuator between the mic receiver and its antenna set to three different levels: 0dB, 20dB, and 40dB. The results, shown in Fig. 4.11, confirm this independence. For all mics and all three mic signal levels, the PESQ is perfect as long as the WSD interference is 1-2dB below the squelch tone signal levels.

Observing Capture: This result should not be a surprise given the well studied phenomenon of *FM capture*, which allows for zero reduction in audio quality, despite the possible presence of significant noise. For FM demodulation, frequency shift is measured by tracking the strongest frequency component in a limited band. As long as the main carrier power exceeds the noise by an amount which allows for clean tracking of the frequency shifts, then the FM receiver will "capture" the signal with zero noise [131]. Such behavior was acknowledged by the FCC in the First Order [18] (Paragraph 38): "FM receivers exhibit a 'capture effect' in which they respond to only the strongest signal received on a frequency and reject any weaker interfering signals."

The capture effect is independent of the squelch tones, but the squelch tones are convenient in determining the allowable level of interference. While 1dB of separation may seem small, the actual power difference is significant (allowing capture) due to the decibel being a logarithmic unit.

4.4.3 Interference in Time

To evaluate the impact of the packet durations, we configure the USRP2 to mimic 802.11-like interference with respect to symbol timings and OFDM subcarriers. With the USRP2's master clock of 100MHz decimated by 8 and an FFT of size 64, we achieve an OFDM symbol time of: $(8/100MHz * 64) = 5.12\mu s$. To ensure the USRP2 can ramp up its transmitter, we use 3 successful symbols in length as our minimum to achieve a minimum interference duration of 15.36μs, comparable to four 802.11a/g/n symbols which are 16μs in length, i.e. a very minimal "frame." We ran experiments with all six mics and changed the timing of the interference by controlling the inter-frame gap using a sub-microsecond scheduler [37].

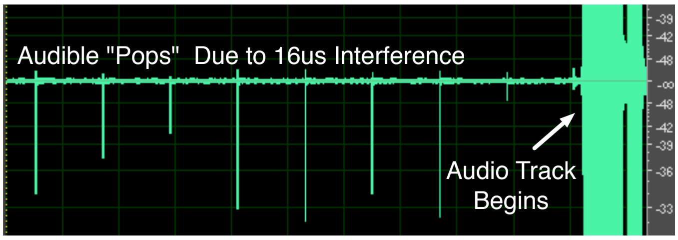

Figure 4.10: Audible pops occur even with extremely short amounts of interference (e.g., only 16 μs long)

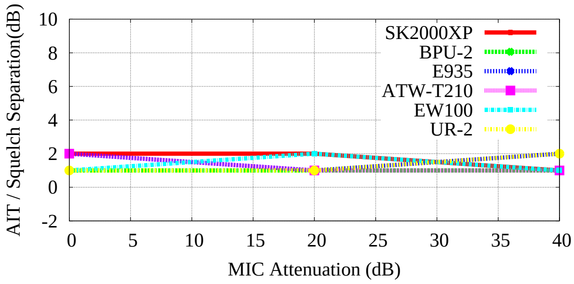

Figure 4.11: Audible interference threshold (AIT) remains stable when varying mic signal atten.

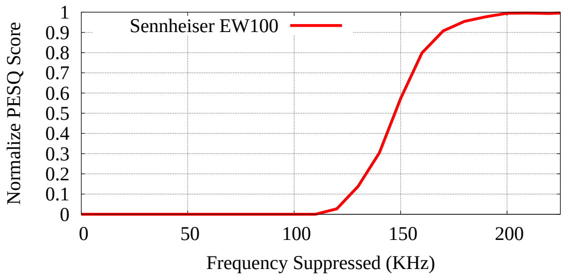

Figure 4.12: Senn. EW100 requires 200KHz of non-interfering noise for no audible interference.

Shown in Figure 4.9, the small inter-frame spacings (10s of μs, similar to 802.11 IFS and backoff) cause PESQ scores near zero, while even for spacings as high as 500ms, the normalized PESQ only reaches 0.7. The cause for this is shown in Figure 4.9. Even spaced 500ms apart, the 16 μs of interference create audible "pops" in the audio recording during each insertion of interference. Clearly, it seems like it is not possible to send extremely short packets over the transmission of the microphone (at a higher power), for example, to bootstrap the device in a channel and find microphones in range.

4.4.4 Interference in the Frequency Domain

We now evaluate how much spacing in the frequency domain is needed for a mic system to have zero audible interference. Interference is constant in time and the interference power is set to 10dB above the squelch tones. Initially, the WSD interferes across the entire TV channel (i.e., 6MHz). Then, we incrementally suppress frequency at the center of the mic's band outwards in 5KHz steps. To get accurate measurements, we ensure that the power falloff is steep. We attenuate the mic signal and interference from the WSD so the power reaches the noise floor within 2KHz (Figure 4.13).

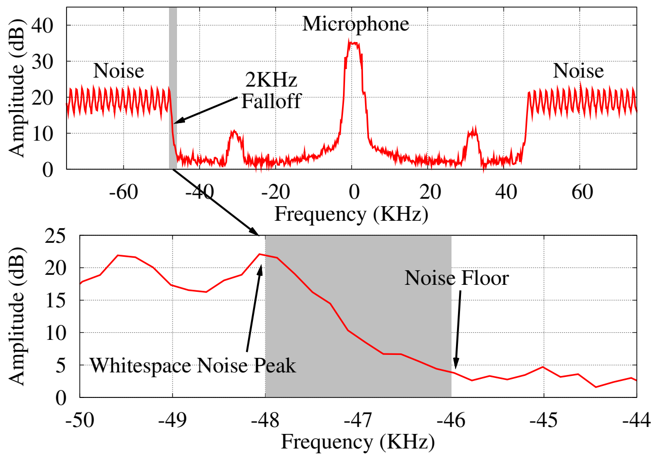

Figure 4.13: Illustrating the falloff to reach the noise floor in 2KHz.

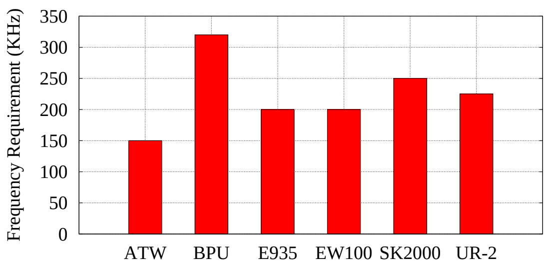

Figure 4.14: The frequency tolerance thresholds for all of the mics for no audible interference.

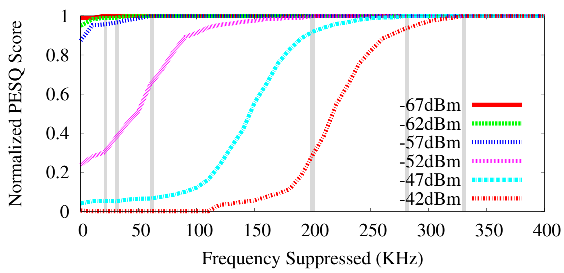

Figure 4.15: Varying WSD interference power at mic receiver, suppression needed varies greatly.

Frequency Domain Results: Figure 4.12 shows the impact of the amount of suppressed spectrum on the normalized PESQ score of the Sennheiser EW100. We see that the EW100's audio quality is severely affected (PESQ=0) when there is less than approximately 110KHz of interference-free spectrum at the center of the mic signal. Beyond this point, audio quality begins improving and once there is 200KHz of free spectrum, there is zero audible interference on the mic. We perform this same experiment for all mics. The results, Figure 4.14, show that the minimal amount of interference-free spectrum ranges from 150KHz (ATW-T210) to 325KHz (BPU). Moreover, we note that models from the same manufacturer can require different amounts of interference-free spectrum (e.g., Sennheiser's E935, EW100, and SK2000XP).

Varying WSD's power and power leakage: The proximity of the WSD affects its power, and thus power leakage past the suppressed subcarriers. To evaluate this, we repeated

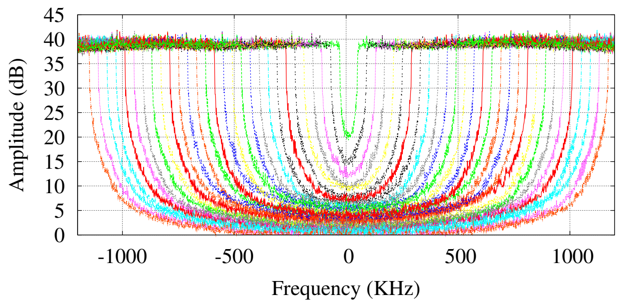

Figure 4.16: The spectrum profile of the Adaptrum transmitter, suppressing frequency at 80KHz steps.

the frequency suppression test on the Sennheiser EW100, but vary the noise power at the receiver from -42dBm to -67dBm in 5dB steps. At each step, we sweep the amount of frequency suppressed from 0KHz to 400KHz in 5KHz steps and compute the normalized PESQ score.

We present the results in Figure 4.15, highlighting the point at which we achieve zero audible interference for the various noise powers. The results show that, as expected, the amount of frequency suppressed at the transmitter needed to achieve zero audible interference will be different depending on the noise power at the receiver. At the strongest power level, -42dBm, we need to suppress a little over 330KHz at the transmitter. At the lowest power, -67dBm, we only need to suppress 20KHz. Measured, although not shown, at -77dBm (in two more power steps), 0KHz needs to be suppressed for zero audible interference because the interference power is already more than 1 dB below the squelch tones (WSN_s).

To provide further visualization of this, we suppress frequency on the Adaptrum transmitter at 80KHz steps and record the resulting power in the frequency domain without attenuating the transmitter (Fig. 4.16). The result is that, depending on how much is suppressed, the power within the band can still be significant. For example, suppressing 250KHz leaves 9dBm of additional noise power even at the center frequency. As the white space transmitter is attenuated (to emulate a WSD farther away), the form of the signatures will stay the same, however the noise floor will shift up. So if the white space transmitter was attenuated by 10dB, the center frequency at 250KHz of suppression would now experience 0dBm of additional noise (whereas it would have experienced 9dBm of additional noise with no attenuation). This shows that the number of subcarriers that need to be suppressed to achieve N-amount of interference-free frequency at the mic receiver will be dependent on the noise power at the receiver.

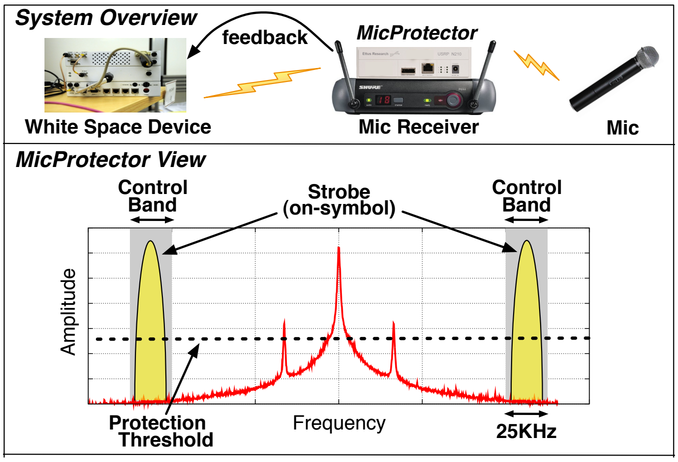

Figure 4.17: Overview of SEISMIC and MicProtector.

4.5 SEISMIC: Towards Ideal Coexistence with Microphones

We explained in Section 4.3 how spectrum-efficient coexistence between mics and WSDs must be a feedback-driven, closed loop design that allows the WSD device to adapt based on the SINR properties at the mic receiver. In this section, we first revisit the SEISMIC design and then elaborate on the three key SEISMIC components in Sections 4.5.2–4.5.4.

4.5.1 System Overview

As discussed in Section 4.3, any spectrum efficient solution to the microphone coexistence problem in white spaces requires either additional hardware or changes to legacy systems. In our case, we use a simple device called a *MicProtector* which resides near the mic receiver (e.g., on top of the receiver in Figure 4.17), near an array of mic receivers common in productions (§4.5.7), or built in to future mic receivers. The MicProtector is responsible for both the *measurement* and *analysis* components in the closed loop control, in addition to providing feedback on the analysis to the WSD. To do so, the MicProtector monitors the interference power and mic signal (i.e., SINR), and employs a *Protection Threshold* to notify a WSD of *impending* disruption to the mic's audio. Based on the study presented in the prevoius section, the protection threshold is set below the mic's squelch tones.

To notify of impending interference, a low complexity pulse-based signaling mechanism (§4.5.4) is used to communicate with the WSD. We call this *strobing*, and it requires only carrier sense-like functionality. The strobes are transmitted in *control bands* surrounding the mic (see Figure 4.17), which we also use for measuring SINR. Since the strobes are raised in both control bands, the WSD can determine the mic's operational band (i.e., frequency and bandwidth).

To ensure that a WSD never exceeds the *Protection Threshold* and causes an audio disruption, the WSD and and MicProtector engage in a protocol. Whenever a WSD starts transmitting on a new frequency band, it does so at minimum power, and then increases this power gradually. As the WSD ramps up its power, *if* the WSD is in disruption-range of a mic, the interference level at the mic receiver will slowly approach the Protection Threshold at which point the WSD will be notified of *impending disruption*. With each impending disruption notification, the WSD suppresses additional frequency. In doing so the system approaches the "ideal state" of suppressing a minimal number of subcarriers.

4.5.2 Detecting Impending Interference

The MicProtector must accurately and quickly measure SINR to notify the WSD of impending disruption before it occurs. The technical challenge of doing so is that the mic signal is constant and the FM nature shifts power in the band, making interference estimation directly in the band difficult. Given our goal of enabling coexistence between wideband WSDs and mics, interference will be wideband. Coexistence between narrowband WSD and narrowband primaries (mics) is a completely separate challenge, which our work is not looking to address. We are assuming wideband WSD that have OFDM capabilities to perform subcarrier suppression. Narrowband devices are unlikely to use OFDM, and are not WSDs that could follow our protocol. We would suggest that such narrowband follow the channel vacation rule.

Under this assumption, we can accurately detect the level of interference independent of the mic's signal by measuring the power directly outside of the band. Since the operational band is small (∼200KHz), estimation directly outside the band in our system is expected to be accurate: prior work [132] has shown frequency selective fading can be severe (30dB) across 20MHz frequency ranges but remains modest (<1dB) for the smaller 200KHz frequency range.

To perform this measurement, we introduce *control bands* at the MicProtector which are 25KHz bands on both sides of the mic's operational band (see Fig. 4.17). Using these bands, the MicProtector can accurately measure the interference power generated from WSDs in range. Given that noise is additive, measuring the interference power of multiple WSDs is handled through the measurement in the control bands. Noise will be cumulative in the SINR measurement.

The MicProtector must monitor the squelch tone power, as shown in §4.4.2, audible disruption is caused when the interference level reaches the squelch tones. To do so, it measures the power in the frequency area of the squelch tones, which are approximately at a ±32KHz offset from the center of the mic's band and subtracts the interference power.

Finally, the MicProtector must be able to warn a WSD of *impending interference*, i.e., there is a *Protection Threshold* below the squelch tones. If the mic signal were stable and there was no delay in WSD adaptation, this threshold could be placed 1dB below the squelch tones. However, in the time it takes a WSD to adapt, the mic signal could drop due to changes in the environment or mobility; or the WSD's signal may increase. Therefore, the protection threshold needs to be more conservative to protect against fluctuation. In §4.6,

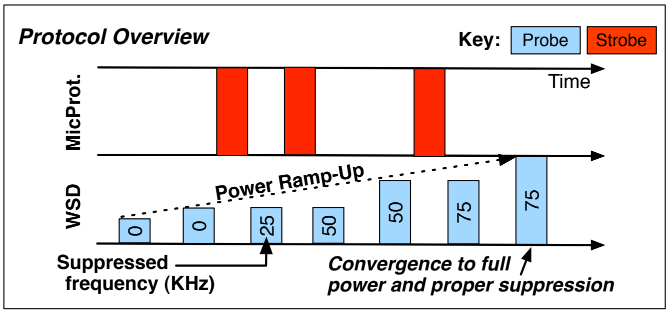

Figure 4.18: Overview of SEISMIC adaptation protocol.

we show that using a conservative threshold of 10dB below the squelch tones achieves all these goals. However, we also show (§4.6.1) that even if we wanted to select an even more conservative threshold (e.g., 20dB below the squelch tones), significant spectrum gains can still be achieved.

4.5.3 Adaptation Protocol

The goal and challenge of the adaptation protocol is to reuse the surrounding frequency around a mic's transmission without ever creating an audible disruption. Such a task is non-trivial. When first entering a channel, if a WSD were to transmit at full power without knowing mic placement or what SINR values it could create, it could easily exceed a mic's protection threshold and create an audible disruption.

To overcome this, SEISMIC exploits the FM capture effect in mic systems where RF interference below the squelch tones is disruption-free. From this, we design *underlay probe packets* to the mic system, which reside under the mic signal. Such packets implicitly ask the mic system: "is this frequency usage at this power level acceptable?"

To converge without causing a disruption when first entering a channel, the WSD begins at minimal power (P) and transmits a probe packet.[1] After a probe transmission, the WSD waits ΔT for an impending interference notification. Without notification, the WSD increases its transmission power by ΔP and transmits another probe packet. The ΔT time between each step is dependent on the time it takes to reliably detect impending interference notifications. In our SDR-based implementation (§4.6), we require ΔT to be $320\mu s$. However, this time could be significantly reduced in a hardware implementation (10s of μs). For ΔP, we find 2dB to be a reasonable increment, ensuring interference is increased slowly without significantly increasing convergence time. Through evaluation, ΔP=2dB achieves 16ms average convergence time.

[1] If the signal strength at the mic receiver is very weak, the initial lowest power level could create audible disruption at the mic receiver. We address this scenario in Section 4.5.5.

Algorithm 1 Adaptation Algorithm at WSD:

S: Spectrum used by WSD, initially the entire desired spectrum.

P: Transmit power used by WSD, initially at minimum level.

ΔT: Ramp up time interval .

ΔS: Amount of additional spectrum suppressed in each iteration.

ΔP: Power increment in each iteration.

Ramp-up:

```
1: while P below desired power level and S ≠ {} do
2:    wait for time ΔT
3:    transmit underlay signal on spectrum S using power P.
4:    if strobe M(F_Mic) received then
5:       Suppress an additional ΔS of spectrum around F_Mic.
6:    else
7:       Increase P by ΔP;
8:    end if
9: end while
```

If a notification of impending interference is received (i.e., interference power reached protection threshold), the WSD *must* suppress ΔS frequency, or back down its power. Ultimately, ΔS will be dependent on the parameters of the WSD. Using subcarrier suppression for a discontiguous waveform, ΔS can be no smaller than the width of a subcarrier (i.e., suppressing in smaller steps is not possible). We use a ΔS of 25KHz in our USRP2 WSD implementation (§4.6), which also matches our Adaptrum industry WSD subcarrier size. Note that the larger ΔS is, the more likely the WSD will suppress un-needed frequency. The smaller ΔS is, the WSD will achieve a closer-to-optimal amount of suppression. This process continues until convergence, illustrated in Figure 4.18.

Several comments are in order:

- By design, if the initial minimal power level does not cause disruption at the mic, the protocol is guaranteed to ensure no mic disruptions. We discuss in Section 4.5.5 how we can guarantee disruption-freedom in *all* cases.
- The protocol converges to an optimal state, or a close approximation, i.e., full power with minimal suppression.
- The protocol works even in the presence of multiple mics and multiple MicProtectors. Whenever a strobe signal $M(F_{Mic})$ is received, the WSD blocks off additional spectrum

around the mic centered at $M(F_{Mic})$. I.e., there can be multiple “holes” in the spectrum used by the WSD.

- As we show in Section 4.6, the ramp up time interval can be implemented to be short; so convergence is fast.

There are two more details to the protocol. First, when a new mic enters the channel, it may be within disruption range of a WSD. To prevent disruption, when the MicProtector is initialized with the mic system, it sends out a special strobe pattern (§4.5.4) which acts as a reset. Detecting the reset forces *all* WSDs in to the probing and ramp-up phase.

When can a WSD can reclaim suppressed spectrum as mics leave the channel? Given that the frequency with which mics enter and leave a channel is typically prolonged (e.g., a concert or a lecture), the process does not need to happen often or quickly. To reclaim spectrum, the WSD can simply re-initialize its transmission power, un-suppress frequency, and then restart the adaptation protocol using more spectrum. Notice that with this method of reclaiming spectrum, SEISMIC is inherently robust and conservative: WSDs react to impending interference in the most conservative way (by suppressing more spectrum), but can reclaim spectrum only by resetting their power level to the minimal level and restarting the entire protocol anew.

4.5.4 Strobing: Notifying Impending Disruption

The previous sections have shown that SEISMIC relies on a signaling technique from the MicProtector to notify of impending disruptions. The signal must be simple, robust, and spectrum efficient; yet able to convey the necessary information (mic's operational band and center frequency). Specifically, requiring the support of a complex protocol (e.g., 802.11) limits WSD and mic system design. The signaling should also happen in-band to remain efficient and avoid the WSD needing to tune to another frequency.

To meet these goals, we introduce a technique we refer to as *strobing.* It adds minimal complexity on both sides. It only requires basic power generation at the MicProtector, and simple carrier sense-like power detection at WSD. Furthermore, with thoughtful placement of the strobe signals in the control bands, the strobes can convey the necessary information for the WSD to adapt in a spectrum efficient manner.

Stobes resemble On/Off-Keying (OOK) and Morse codes, in which the power of a tone is quickly raised and lowered (i.e., a strobe light) in a pre-determined pattern to convey a signal. Patterns are generated using alternations of on- and off-symbols, where an on-symbol is the presence of a tone and the off-symbol is the absence of the tone. On- and off-symbol lengths are fixed in time, and unique patterns are generated by alternating the power (or presence) of the tone, for example: [1,0,1,0,1,0,...] or [1,1,0,0,1,1,0,0,...]. This is effectively changing the rate of the strobing, as a factor of the fixed symbol length. We provide a simple time-domain example at the top of Figure 4.19 with hard symbol transitions.

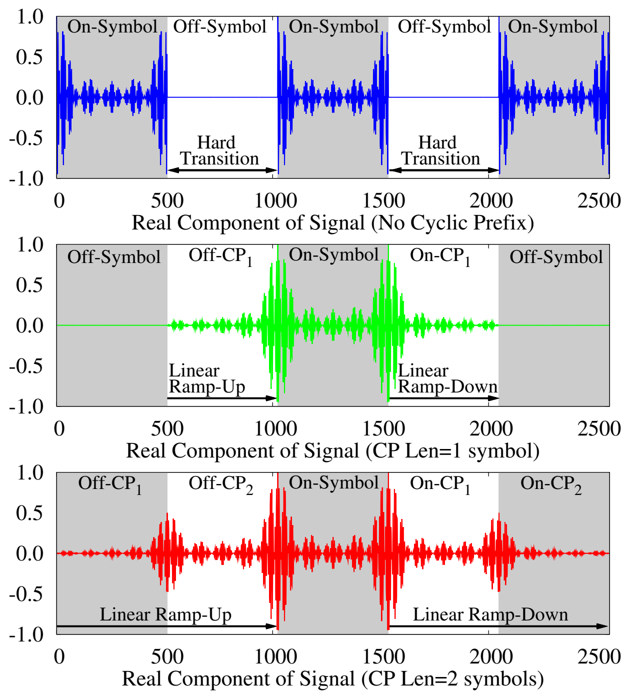

Figure 4.19: Strobes with varying cyclic prefixes in the time-domain.

Simply, the presence of a strobe can signal of impending disruption from the MicProtector to a WSD. By strobing and simply changing the strobe rate, we create unique signals which do not mimic WSD behavior. Notice that by generating a strobe in the two control bands, both the center frequency and bandwidth of the mic are conveyed. The middle point between the strobes is the center frequency, and the distance between them is the bandwidth (see Figure 4.17).

Strobe Patterns: With a single MicProtector in a channel, generating two identical strobes in the control bands can properly convey location and width of the mic's band. However, with more than one MicProtector strobing, where one band starts and another ends is difficult to determine (e.g., which strobe starts or ends a band?). In a planned environment where mics are spaced more than 500KHz apart, bands may be more easily distinguished. However, in unplanned environments where mics may be placed closer in frequency, their bands will be indistinguishable. To eliminate this problem, we use two different strobe patterns (i.e., rates) in the control bands: a start- and end-of-band pattern.

Safely Generating Strobes: The immediate concern of strobing is to ensure that the strobes sent by the MicProtector never interfere with the mic signal. We must ensure that the tones are generated in a way that no power is leaked in to the mic's band. We find

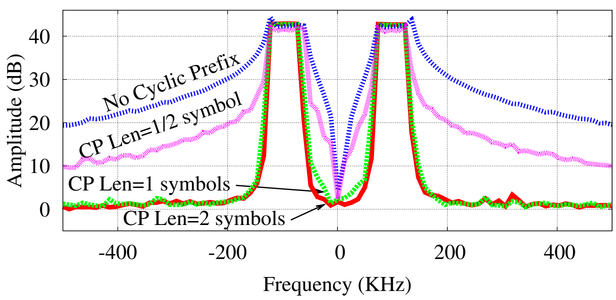

Figure 4.20: Strobes with varying CP in the frequency-domain.

that a cyclic prefix is *critical* to ensuring this. From extensive evaluation using a nanosecond level sample capture, we find that hard symbol transitions create leakage in to the mic band and surrounding spectrum. Hard transitions are shown in the time domain at the top of Figure 4.19, and the resulting interference in the frequency domain is shown with the corresponding line in Figure 4.20. Clearly, the result shows interference in the mic band.

We find that from generating what we refer to as a linear power ramping cyclic prefix (LPR CP), we can eliminate this leakage in to the mic band. An LPR Off-CP is used to gradually scale the power up from an off-symbol to an on-symbol, and an LPR On-CP is used to gradually scale the power back down from an on-symbol to an off-symbol. The linear power scaling is done by scaling the complex samples in software (i.e., on the DSP or in the software of a software-defined radio) to avoid complications of fine-grained power control in hardware. The LPR CP is illustrated in the time domain in Figure 4.19, and the resulting frequency usage in Figure 4.20. As shown, using an LPR CP of size 2 removes this critical leakage, protecting from interference even at high TX power.

Strobe Detection: Strobe detection is similar to carrier sense-like functionality and pattern detection. After a probe, the WSD monitors the state of the channel broken down into 25KHz bins. The matching is done in an absolute manner, marking each bin as a 1 or 0 and over time matching a strobe pattern. This is very parallelizable in hardware.

4.5.5 Low-Power Mic Signals

If the mic signal is low and the squelch tones are barely above the noise floor, even a single (or multiple) new WSD's transmitting probe packets at minimal power could create disruption at the mic. To avoid this, we introduce a final set of unique strobes which are proactively generated when the mic signal is low. The signal when the protection threshold is below the noise floor at the MicProtector, since this threshold represents the point at which interference above it threatens disruption. When a WSD detects a low-power strobe signal it vacates the channel. The low-power signal ends when the threshold goes above the noise floor.

4.5.6 Multiple White Space Devices

So far, we have described the operation of SEISMIC in a scenario with a single interfering white space device. Clearly, in order for SEISMIC to be practical, it must also be robust in the presence of multiple WSDs. One may wonder whether the system is still sufficiently protective if many WSDs simultaneously ramp up their power.

Fortunately, *SEISMIC can handle any number of WSDs* and still guarantees full protection to the mic. Consider the following inductive argument. At some time T, there are n WSDs transmitting in proximity of the mic. Let the cumulative interference level I_T created by all these WSDs at the mic receiver be below the protection threshold. In the worst-case, all n WSDs simultaneously ramp up their transmission power once before the MicProtector is able to send a strobe signal. In this case, it is guaranteed that the new cumulative interference level I_{T+1} is still below the mic's squelch tones.

This is true because the adaptation in SEISMIC uses *multiplicative increases* in transmission power levels. By dB's relative definition, any additive increase in dB corresponds to a multiplicative increase of power in mW. For example, an additive increase by 3dB corresponds to (roughly) 2x power in mW, and additive increase by 2dB is approximately 1.6x. Consequently, the interference power (in mW) of each of the n WSDs is increased at most by a multiplicative factor of ΔP ($\sim$1.6x), and hence, the cumulative interference of all WSDs is $I_{T+1} \leq \Delta P \cdot I_T$. Thus, assuming the adaptation protocol is correct for a single WSD, it is also correct for n WSDs. In fact, observe that the more simultaneously transmitting WSDs, the smaller the variance and hence the more robust the protocol. The above computation only takes into account the effect of ramping up the transmission powers, but the exact same reduction from the n-WSD case to the 1-WSD case can also be made for the effects of mobility and/or fading.

4.5.7 Multiple Microphones

SEISMIC also works in the presence of multiple mics (and thus multiple MicProtectors). In this case, the potential danger is that the strobing signals of one MicProtector could interfere with another mic in close proximity (thus causing disruption) because MicProtectors themselves do not actually follow the SEISMIC adaptation protocol before transmitting their strobing signals. Fortunately, it turns out that in practice, this is not an issue. Due to intermodulation interference[1], proper frequency coordination in a location with multiple mics is essential, and coordination software for wireless mics ensure third-order and fifth-order harmonics are eliminated. A consequence of this coordination is that nearby mics will always be placed such that their frequencies are at least 500KHz apart, which leaves more than sufficient space for SEISMIC's control bands.

[1]Intermodulation is a type of interference in which a receiver picks up two dissimilar frequencies that interact within the receiver's electronics to produce sum and difference frequencies, including harmonics of these frequencies, which results in a whilsting noise.

4.5.8 Partial Deployment

An immediate concern of SEISMIC is deployment: the protocol relies on feedback from the MicProtector to suppress frequency. If a mic receiver does not have a MicProtector, a WSD will continue ramping up without suppression since it does not receive a notification of impending interference. An unlikely non-partial deployment solution to this problem would be to require *all* licensed mic systems to include the MicProtector, or risk WSD interference.

Instead, it is possible to partially deploy SEISMIC to enable protection over all mics (with or without the MicProtector), while also allowing coexistence with mic systems that have a MicProtector. Ultimately, the more mic systems that adopt the MicProtector over time, the more efficient the white space spectrum will become. To do so, licensed mics register in the database as being SEISMIC enabled or not. When entering a channel, the WSD consults the database to learn of *possible* mics within range. If all mic receivers are SEISMIC enabled, the WSD participates in the SEISMIC protocol. Otherwise, it must vacate the channel. A single non-SEISMIC mic system in the presence of many SEISMIC enabled systems reduces efficiency, however the average number of mics in range of a WSD is likely to be low.

4.6 Prototype and Performance Evaluation

We implement a full prototype of the MicProtector and SEISMIC on the USRP2. We build a custom software stack for the key components: (i) measurement & analysis , (ii) strobe generation and detection, and (iii) the client with power ramping and suppression. We use the WBX daughterboard which is a full transceiver in the frequency range of 50MHz to 2.2GHz. We conduct all experiments over the air on TV channel 21, using an approved experimental license. Our SEISMIC parameters: ΔT=320μs, ΔS=25KHz, and ΔP=2dB.

We evaluate SEISMIC in several dimensions in this section. We begin with an evaluation of the protection threshold's robustness. Then, live over-the-air experiments with a white space device running the SEISMIC protocol, and a mic system equipped with our MicProtector prototype. From this, we show the system's efficiency and robustness to avoid disruptions in challenging scenarios. We conclude with a simulation using real mic placement data from 3 major events.

4.6.1 Impact of the Protection Threshold

The *protection threshold* at the MicProtector is set to allow a WSD to ramp up to proper suppression and operate without ever exceeding the power of the mic's squelch tones. Without a buffer between this threshold and the squelch tones, variations in the mic's signal power or the WSD's interference level due to mobility, fading, etc., could cause the interference level to go above the squelch tones. There is an inherent trade-off when choosing this protection threshold. The lower the threshold, the more conservative the protection. The higher the protection threshold, the better the spectrum efficiency. To illustrate this, we perform a simple over-the-air experiment in which we vary the protection threshold and the interference

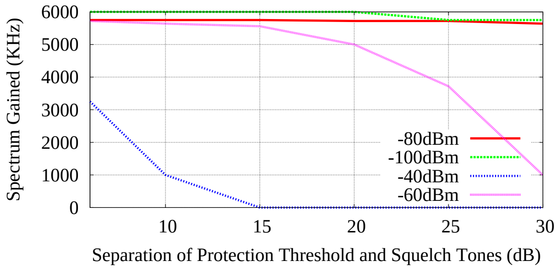

Figure 4.21: Impact of the protection threshold on spectrum efficiency.

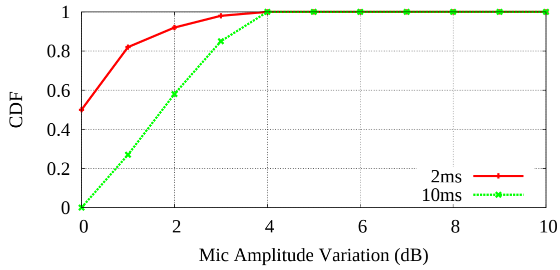

Figure 4.22: Variation in microphone signal with mobility.

from the WSD. As we see in Figure 4.21, for most cases the WSD can reuse most of the spectrum. When the WSD interference is high (e.g., -40 and -60 dBm) WSDs get reasonable spectrum only when the separation between the protection threshold and squelch tones is low.

Evaluating an Appropriate Protection Threshold: In determining the appropriate threshold, one has to consider the signal variation over the max transmission time of a WSD. The WSD cannot adapt during this time to ensure the interference does not exceed the squelch tones. To provide some insight, we performed a live mic experiment. Over a 60 second period, we walked to and from the mic receiver and swung the mic in fast movements to trigger quick signal variations. We calculated the maximum variation over 2ms and 10ms periods (i.e., max WSD TX times), and present a CDF in Figure 4.22. This shows that over both periods the maximum variation we find is 4dB. The WSD could also be ramping up its power and probing, at a 2dB step. Despite the probes being much shorter in time ($\sim$10s of μs), we still account for this and now consider the minimum 4dB+2dB=6dB. We add 4dB to account for other variations in the WSD power, and find 10dB to be sufficient.

4.6.2 Spectrum Efficiency Scenarios

Given a 10 dB protection threshold, we evaluate the spectrum efficiency achieved by SEISMIC under different scenarios. We cover the range of scenarios by varying the two components of SINR at the mic receiver: the mic's signal at the receiver, and the WSD's interference at the mic receiver. For the former, we vary the received power of the mic, and the latter is the power before frequency suppression. The resulting spectrum that can be used by the WSD is shown in Figure 4.23. When the mic's squelch tone power is high, no suppression is needed and the entire 6 MHz of spectrum can be used. In most cases, only 250 KHz of spectrum needs to be suppressed. On increasing the WSD IBS power, it begins to overpower the mic signal, and therefore lesser spectrum is available for the WSD in order to protect the mic. Finally, the sharp cliff at the left occurs because SEISMIC is protecting the mic in low

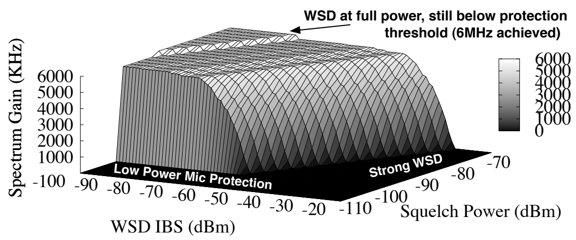

Figure 4.23: SEISMIC's spectrum gain at different interference powers and microphone signal powers.

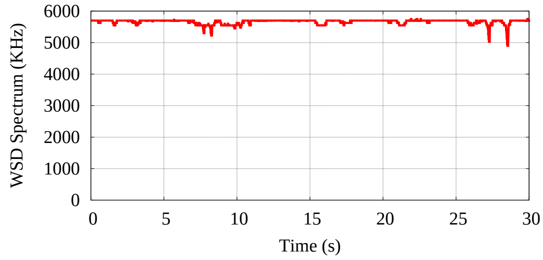

Figure 4.24: Sample SEISMIC spectrum under moderate interference and mobile mic (10-30 ft).

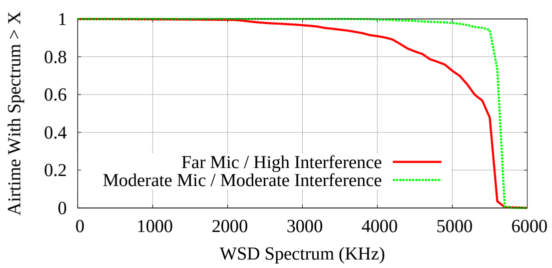

Figure 4.25: Fraction of airtime with WSD having a given amount of spectrum with SEISMIC.

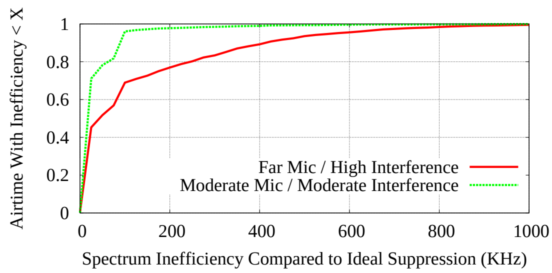

Figure 4.26: Spectrum inefficiency compared to the ideal suppression as a fraction of airtime.

power (i.e., when protection threshold drops below the noise floor, causing WSDs to vacate - §4.5.3).

4.6.3 Live Experimentation with SEISMIC's Effectiveness

To evaluation SEISMIC's effectiveness at avoiding interference with a microphone, we evaluate SEISMIC in a live setup. We use the MicProtector prototype paired with a Sennheiser mic system and our WSD running the SEISMIC protocol. We evaluate under two experimental setups: (1) moderate WSD interference (-70dBm) and mobile mic operation between distances of 10-30 feet, and (2) under a more challenging scenario with high WSD interference (-50dBm) and mobile mic operation between distances of 50-70 feet. Under mobility, we walk with the mic, lower the mic to hip level, raise it and speak in to it, turn our bodies, etc. A protection threshold of 10dB under the squelch tones is used, which we motivate in §4.6.1. Experiments are conducted for 5 minute time periods.

Results (Effectiveness): The moderate WSD interference scenario where the mic is operated within 10-30 feet of the receiver is common in concerts (audio equipment is on/behind stage), and in lecture halls (mic receiver is near podium). Figure 4.24 illustrates the amount of spectrum that can be used by a WSD over a 30 second period. As shown, the available spectrum is both high and stable, despite the mobility of the mic. This is because the WSD suppression is adequate and the mic signal is strong. We also plot a CDF of the available spectrum in Figure 4.25 (Moderate). The gain is significant. The WSD used >5.5MHz of spectrum for nearly 93% of its airtime.

The second scenario is more challenging. As we see in the time series of Figure 4.27, the mic's squelch tones can at times be very low, e.g. at 12 seconds. Based on our protection threshold of 10 dB, and the noise floor of our USRP based MicProtector at -98 dBm, our system notifies of a low-power mic at -88 dBm (shown as a dotted line in the Figure).[1] In these situations, such as at 12 and 24 seconds, the WSD vacates the entire TV channel. At other times, the WSD ramps up the power and uses the available spectrum. As we see in Figure 4.25, the WSD is able to use 5MHz of spectrum 75% of the time, and 4 MHz of spectrum 90% of the time. Even in this challenging scenario with heavy fluctuations of mic signal power, the protection threshold and SEISMIC protocol ensured zero audible microphone interference.

To highlight the efficiency of determining the proper number of subcarriers to suppress, we used information at the MicProtector to compute the optimal amount of frequency suppression and compared it to the spectrum the WSD actually used. Figure 4.26 shows that in the moderate scenario, the WSD is able to converge to within 25KHz of optimal suppression 72% of the time and within 100KHz 97% of the time. In the more challenging scenario (Figure 4.26), the WSD is able to converge to within 400KHz 89% of the time.

Finally, we evaluated the time it takes the WSD to reclaim spectrum after a mic leaves the channel. We allow the WSD to converge with the Sennheiser mic in the channel using SEISMIC, and then turn the microphone off. We repeated this 25 times and found an average time of 272ms. We believe that this time is sufficient to re-probe for the white space spectrum and to reclaim it.

4.6.4 SEISMIC's Efficiency with Many Mics

We now study the benefit of SEISMIC in heavy mic usage scenarios, which may typically be found in cities or on campuses. To evaluate these benefits, we obtained mic registration data for 3 major events: the 2008 NBA All Star Game (191 mics), the 2010 BCS Championship Bowl (108 mics), and the 2010 Worldwide Partner Conference (77 mics). The channel placement of these mics is shown in Figure 4.31.

Efficiency Setup: To quantify SEISMIC's spectrum efficiency, we develop a simulation environment from the event data in which a MicProtector exists for every mic. Now we ask,

[1]We note that in production systems, the noise floor over 400 KHz will be much lower, and can operate at lower squelch tones.

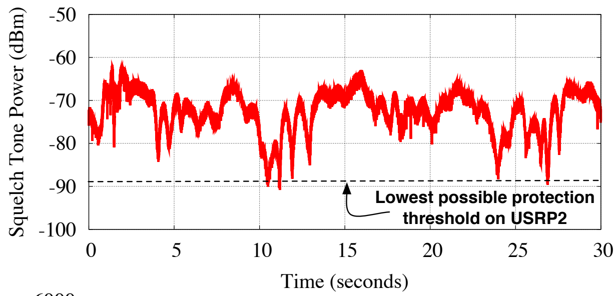

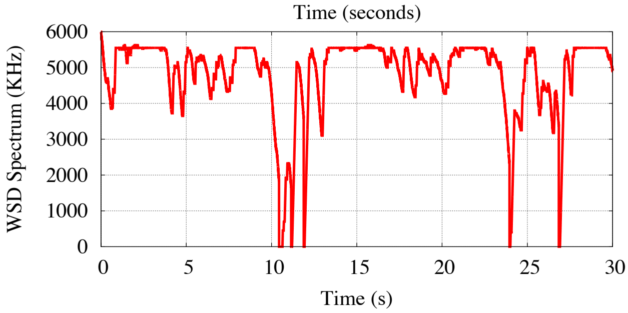

Figure 4.27: Robustness to avoid disruption, adapting the channel.

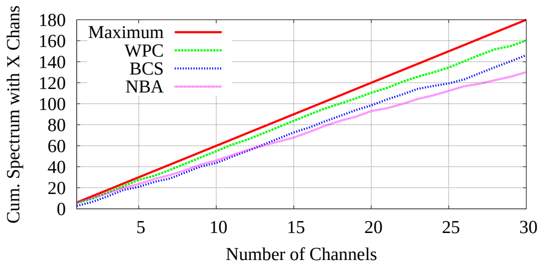

Figure 4.28: Average SEISMIC client spectrum availability at each event.

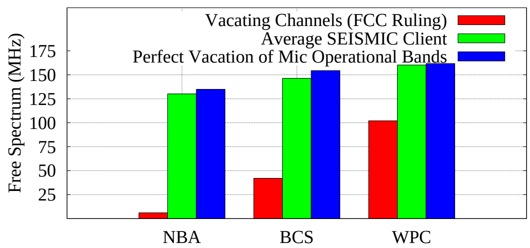

Figure 4.29: Comparing the available spectrum with FCC ruling and perfect suppression.

if X WSDs are in this environment, each with various noise powers at the mic receivers, what is the average amount of usable white space spectrum Y? Two important characteristics are computed: (i) received squelch tone power at a mic's receiver, and (ii) WSD interference at each mic receiver. We weight mic signal strengths towards better-to-average (yet still have low signal mics), and generate the WSD powers uniform randomly between -110dBm and -20dBm. We account for WSD suppression to reduce interference on one mic, can contribute to reduced interference on another mic.

Efficiency Results: Given that we know mic placement but not the active TV broadcasts, we evaluate SEISMIC by varying the number of channels available in each event. These results are presented in Figure 4.28, such that if only 10 channels were available in the area, SEISMIC achieves 44MHz compared to a max of 60MHz. The results show a significant gain

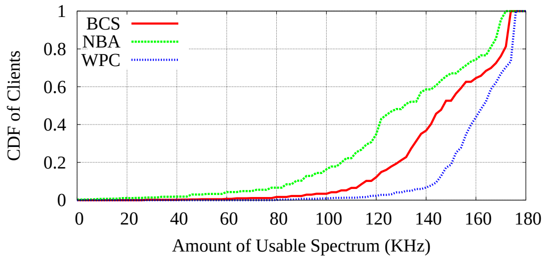

Figure 4.30: CDF of available spectrum at each event using SEISMIC.

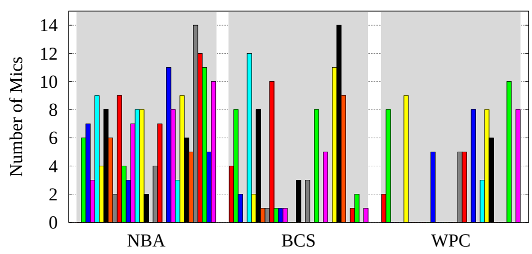

Figure 4.31: The number of mics per TV channel at each event.

in spectrum and promise for ensuring white space in highly dense areas where there are only 2 or 3 channels available.

Further exploring the possibilities of SEISMIC, we assume no TV broadcasts are active and present the resulting spectrum gains based on the mic placement data in comparison with vacating channels and "perfect" vacation of a mic's operational band. As shown, SEISMIC can provide up to 21x the amount of spectrum over channel vacation (e.g., NBA event) and come close to "perfect" vacation. This is due to some WSDs not needing to suppress any subcarriers since their interference remains underneath the MicProtector's protection threshold. To provide further insight in to this, in Figure 4.30, we show the available spectrum across the 1000 SEISMIC clients we simulate in each event. At the most dense event (NBA game), 50% of all clients have at least 130KHz. Only 5% of clients have less than 50MHz of spectrum. This highlights SEISMIC can significantly increase spectrum availability for white space networking.

4.7 Chapter Summary

In this chapter, we explored the efficiency and effectiveness of spectrum management and coexistence techniques in the TV white spaces to provide interference-free guarantees for the spectrum's primaries. As shown through our work, a WSD device needs to vacate the channel in the presence of a microphone given its limited information about the microphone's signal quality (at its receiver), and its interference on the microphone. While we show that this is effective at avoiding interference, it is highly spectrum inefficient: sacrificing up to 95% of the spectrum in a channel.

Motivated by this severe inefficiency that threatens the additional spectrum provided by the white space, we explored the possibility of more spectrum-efficient coexistence techniques. We conducted the first in-depth analysis of WSD transmissions on wireless microphones, where we made several key observations that led to SEISMIC: a spectrum efficient interference-free system for mics. SEISMIC allows a WSD to enter a channel without any knowledge of microphones in it, and by following a novel coexistence protocol we introduced:

ramp up its transmission power, suppressing spectrum usage around the microphone to share a channel and avoid interference. To know how much to suppress, the WSD receives feedback about its level of interference from the mic system using the MicProtector that we introduce.

Through evaluation, we showed that SEISMIC is effective at preventing interference, and efficient at reclaiming spectrum. Even with a mobile microphone whose signal fluctuates quickly,SEISMIC is able to quickly adjust its frequency suppression to avoid interference. In terms of efficiency, SEISMIC can achieve up to 21x the amount of spectrum when compared to spectrum management which must enforce channel vacation. Using microphone data from major events (e.g., the college football championship), we also show that SEISMIC could potentially provide an additional 130 MHz of spectrum on average.

Finally, we note that we have made a video of SEISMIC operating in the same channel with wireless microphones to demonstrate its effectiveness.[1]

Potential Impact: SEISMIC can enable a significantly more spectrum-efficient use of the available white space spectrum in the TV bands, while coexisting in a disruption-free manner with mics. In particular, no channels would need to be reserved for mics. For this reason, we believe that the FCC should amend its ruling to allow WSDs to operate on the same TV channel as long as its power is below the squelch tones of the mic at the mic receiver; and that the white space protocols (e.g. IEEE 802.11af and IEEE 802.22), should be modified to ensure the power limits. Such changes are not unattainable. The FCC has shown its willingness to make changes to the ruling through its removal of the sensing requirement in the Second Order. To accomplish this, we have demonstrated SEISMIC to the FCC, including Chairman Genachowski, various mic operators who plan events such as the Super Bowl and mic manufacturers. In this context, it is encouraging to note that mic manufacturers such as Shure show great interest in a solution such as SEISMIC. Through this effort, we hope to enable more spectrum efficient white space networking.

Limitations: Moving forward with SEISMIC, there are several challenges and limitations of our work.

- *Spectrum policy changes*: As we just discussed, our proposed coexistence protocol requires changes to current regulations to, at the most basic level, allow WSDs to share a channel with wireless microphones. Without this basic change, SEISMIC cannot operate in the spectrum. We are hopeful, however, that this change will eventually be made towards better spectrum efficiency.

- *OFDM-based WSDs*: To coexist with the wireless microphones in the same channel, we based our solution on subcarrier suppression: a technique that assumes an OFDM PHY-layer on the WSDs. As a result, our solution is limited to devices with an OFDM PHY. Although a limitation, the push for "Super Wi-Fi" in these bands will likely

[1] KNOWS. http://research.microsoft.com/en-us/projects/KNOWS/

lead to many OFDM devices. Additionally, the regulation that states WSDs must be wideband. This will also contribute to the dominance of OFDM in the band.

- *End-point modifications*: Like all coexistence techniques, modifications are needed to at least one of the endpoints. In our work, we require modifications to both endpoints: the WSD and the microphone system. Although modifications are required, we do not believe they are overly complex. Additionally, we show that the MicProtector functionality can reside on a standalone/attachable device to the mic system. This means that, while the WSD still needs direct modifications, we can provide SEISMIC support at the microphone system through a simple external device attachment. Such an attachment is good for the support of legacy systems, also.

- *Mix with non-SEISMIC microphone systems:* SEISMIC will create interference on wireless microphones operating on the same channel without a MicProtector. SEISMIC will ramp up its power and, without a MicProtector on the mic system, the WSD will never receive a notification to suppress its frequency. As a result, if there is a channel that consists of 6 microphones, 5 of which are SEISMIC enabled, and one that is not, a SEISMIC enabled WSD cannot join the channel. If it does, it will suppress frequency around the 5 SEISMIC enabled microphones properly, but it will end up interfering with the non-SEISMIC microphone system. Therefore, it cannot join the channel if there is a single non-SEISMIC device.

Future Research: In this chapter, we presented a study of spectrum management and coexistence techniques to provide spectrum efficient coexistence for spectrum primaries in the TV bands. While we believe this is a significant step for white space technology, there are likely to be many other white space bands in the future where the spectrum primaries vary. This is likely to require different solutions than the one we proposed, dependent on the application and signal type of each particular primary. As additional white space bands become available, new interference studies will be needed to understand the impact of data transmissions from WSDs, by which new coexistence techniques can be developed.

Chapter 5

Conclusions and Future Work

Throughout this dissertation, we addressed key challenges in interference driven by the inability of heterogeneous wireless technologies to effectively share the spectrum. We addressed these challenges between heterogeneous and unlicensed devices, as well as secondary devices with (licensed) spectrum primaries in the white spaces. Mitigating this type of heterogeneous interference is critical to ensuring efficient utilization of the spectrum as we find more applications for wireless technology, and cater the radios and their protocols to the application. This chapter concludes the dissertation with a summary of our approach to these challenges, as well as the contributions and a discussion on remaining open problems.

5.1 Contributions

In the context we have described above, this dissertation presented the following thesis: *better monitoring and spectrum management may provide a "single" and more long-term solution to interference between heterogeneous technologies in unlicensed spectrum, whereas coexistence protocols may be more suitable in providing spectrum-efficient interference avoidance with spectrum primaries.*

Given this statement, we explored three key aspects in reducing cross-technology interference between unlicensed and licensed devices in two case studies:

1. The design of a monitoring system to detect heterogeneous wireless networks and map their interference within an environment.
2. A spectrum assignment algorithm that is generic and easily evolvable to manage dense environments with many unlicensed heterogeneous technologies.

3. An interference-free coexistence protocol between unlicensed devices and wireless microphones in the white spaces that significantly improves spectrum efficiency over currently deployed spectrum management mechanisms.

Below, we highlight the contributions made in each of these areas.

5.1.1 Better Monitoring for Unlicensed Heterogeneous Devices

Key Insight and Motivation: One of the biggest challenges in beginning to address interference between heterogeneous wireless technologies and networks (e.g., by applying spectrum management) is detecting them in an environment and understanding the interference between them. In particular, this has become critical to the home environment where the density of heterogeneous wireless networks is significantly high with a complete lack expertise, tools, and information about the networks and their interference. Current monitoring systems for heterogeneous technologies are overly complex and costly for the home environment and its average user, typically requiring an understanding of signal level characteristics and a costly deployment of multiple sensors throughout the environment. Driven by these concerns, our goal in Chapter 2 was to develop an accurate and usable monitoring system for unlicensed and heterogeneous technologies in the home, capable of detecting and mapping interference between them.

Core Contributions: To achieve these goals, we developed a novel wireless monitoring system based on the *smartphone* that is capable of deriving where signals from heterogeneous devices go in the home and what they interfere with. We presented a 3 phase design (training, monitoring, and diagnostics) that balanced user involvement with system complexity. Our design creates device abstractions with user recognizable identifiers through the use of cross-layer information to help bring aspects of the RF environment up to a level the user can understand. By continuing to collect information in the monitoring phase as the user walks near their devices (e.g., with the phone in their pocket), we are able to keep an up-to-date map of the RF environment with little to no user involvement. Our 10-home user study and heterogeneous testbed showed the system to be both usable and accurate, and to demonstrate the usefulness of the information collected we used force-directed graphing to create a device-level map of the user's home by which diagnostics were overlaid.

5.1.2 Spectrum Management for Unlicensed Heterogeneous Networks

Key Insight and Motivation: As we motivated through the thesis of this dissertation, a spectrum management based approach to solving the general problem of interference between heterogeneous unlicensed technologies has more desirable and long-term properties than an N^2 coexistence mechanism based approach. Our smartphone-based monitoring system was the first critical step towards proper spectrum management by collecting the necessary information, however, current spectrum assignment models are predominantly homogeneous and/or Wi-Fi centric (i.e., they only focus on reconfiguring a Wi-Fi network to avoid interference from heterogeneous technologies). A more comprehensive algorithm that supports the many unlicensed heterogeneous technologies is needed, in addition to one that can support the evolution of these protocols and future technologies without significant (or any) changes to the general system and algorithm.

Core Contributions: In Chapter 3, we presented a system for assigning an environment of unlicensed and heterogeneous wireless technologies such that they will be isolated (if possible) or placed in ways that they will receive better performance. To support the many technologies that unlicensed wireless devices support, the system is based on the principles that the system and its components must describe fundamental properties of heterogeneous technologies and the environment (e.g.,"Do the two devices coordinate?") – not specifics of technologies, and the system must remain generic where possible to support the evolution of the protocols and spectrum over time (e.g., the introduction of new protocols or spectrum bands).

Following these principles, we introduced several components in our spectrum assignment system to achieve our goals and support heterogeneous networks:

1. A hypergraph-based model of the RF environment that represents the heterogeneous radios and networks within it, as well as their constraints (e.g., possible frequencies).
2. Subgraph searching of the hypergraph-based model of the environment with templates to detect conflicts and generate a more traditional conflict graph.
3. A generic predictive channel quality metric for heterogeneous networks.
4. A spectrum assignment algorithm based on a mixed integer program that uses the hypergraph for its constraints, and conflict graph with predictive channel quality metric to efficiently assign spectrum to avoid conflicts and improve performance.

5.1.3 Spectrum-Efficient Coexistence with Spectrum Primaries

In Chapter 4, we shifted our focus from addressing interference between unlicensed devices to interference avoidance between these devices and spectrum primaries in the white spaces. As the demand of wireless applications continues to strain spectrum availability, we will likely continue to open up additional spectrum for unlicensed devices through "white space" spectrum access in licensed bands. This means that these many unlicensed devices will have access to additional spectrum in licensed bands with the exception of avoiding interference with the primary users in those bands.

Measurement results that we presented in the beginning of Chapter 4 showed that current spectrum management-based solutions to interference avoidance with spectrum primaries are extremely spectrum inefficient. In particular, the spectrum management-based techniques and regulations to protect wireless microphones in the TV white spaces can leave up to 95% of a channel and the spectrum idle. Following our thesis, we argued that coexistence techniques are better suited for providing spectrum efficient solutions to interference avoidance with spectrum primaries. However, current coexistence techniques between unlicensed devices would not be sufficient in the white spaces due to their inability to guarantee complete interference avoidance (e.g., during bootstrapping). A new coexistence technique would be needed to meet more strict requirements with spectrum primaries.

Core Contributions: To develop an effective and efficient coexistence mechanism between unlicensed devices and spectrum primaries in the white space, we began Chapter 4 with the first in-depth study of how data transmissions from unlicensed devices impact wireless microphones. The study was conducted by controlling the data transmission's interference power, frequency use, and duration. The study highlighted interference of any duration whose power stayed below the microphone's squelch tones would avoid all audio disruption and that, given the current information and closed-loop system between microphones and unlicensed devices, channel vacation is the only way to ensure interference avoidance. This motivated our closed-loop SEISMIC system design that introduced a MicProtector that monitors interference at microphone systems and provides feedback to nearby white space devices (WSDs) that suppress frequency around microphones they will potentially interfere with (as notified by in-band "strobes" from the MicProtector). By ramping up the transmission power of a WSD when entering a new channel and continuing to adapt frequency usage with strobe notifications from MicProtectors, we can guarantee interference-free and spectrum-efficient coexistence. Live evaluation with real microphones showed the system to be effective at avoiding interference and efficient at reclaiming spectrum. A simulation with real microphone placement data further showed the potential of the system's ability to reclaim spectrum.

5.2 Implications and Outlook

Based on the results of this dissertation, our work suggests that better monitoring and spectrum management can provide long term solutions to interference between the many evolving heterogeneous and unlicensed devices, whereas coexistence techniques can provide spectrum-efficient solutions to interference between primary and secondary devices in the white spaces. This is to the contrary of the majority of past work that has developed coexistence techniques between unlicensed devices [8,16,39,41,43,44,48], and spectrum management between primary and secondary devices [21,22,49,50,51,52,53].

Clearly, there are many implications and challenges moving forward given our suggestion to refocus efforts. Most importantly, to improve the outlook of this approach, additional efforts are needed that follow this general approach in other spectrum bands and environments to provide further confidence. In particular, we highlighted the potential of the approach through two case studies with focuses in two specific environments: one between unlicensed devices in the home, and one between primary and secondary devices in the TV white spaces and urban areas.

With a limited set of environments and assumptions between the sets of technologies, we have only begun addressing a small part of the larger problem space that we introduced in Chapter 1. For example, we did not address the challenges of monitoring significantly larger environments such as urban or enterprise environments, and we did not address how to develop a scalable algorithm or optimization to assign the spectrum in these larger environments. Additionally, there are other spectrum bands with unlicensed access than 2.4 and 5 GHz that we focused on, which may bring new challenges towards monitoring or spectrum management. Focusing on the coexistence of wireless microphones and wideband secondary users in the white spaces meant that we could reasonably assume an FM-based primary user, and a OFDM-based secondary user. It is likely that other future white space spectrum bands will have different primaries with different signal types, and potentially wideband and narrowband secondary users.

To improve upon the various narrow aspects of this dissertation, we first need to study monitoring and spectrum management between unlicensed and heterogeneous devices in other environments (e.g., other than the home), in other spectrum bands (e.g., outside of 2.4 and 5 GHz), and as new unlicensed devices and technologies arise. With the quick evolution of technologies and unlicensed devices: 1) Does our spectrum management algorithm support them? and 2) If not, what significant changes are needed, or fundamental properties need to be described to properly support them? Time is the true test to our system and its fundamental design principles to support heterogeneous unlicensed devices and their evolution. Deploying assignments from spectrum management systems (including ours) is an on-going challenge, particularly in chaotic environments, to make spectrum management between unlicensed and heterogeneous devices practical.

Second, we need to study the potential of coexistence techniques in other (future) white space spectrum bands (e.g., other than the TV white spaces), with other potential primary and secondary devices (e.g., other than the wireless microphone primaries we study). While we have shown significant spectrum efficiency benefits of coexistence in the TV white spaces, the wireless microphone primaries we consider are particularly narrow (a few hundred kilohertz). Additionally, the coexistence technique we employ builds upon this assumption: wide secondary devices with narrow primary devices. Other spectrum bands without this assumption will likely require another coexistence technique.

Finally, our work and any future work on coexistence between primary and secondary devices hinges on regulations allowing for this potential interaction between the primaries and secondaries. Currently, regulations in the TV white spaces have been focused on frequency isolating primaries and secondaries entirely. As a result, current regulations in the white spaces prevent coexistence techniques from being deployed. This was a primary reason for motivating spectrum management between unlicensed devices (i.e., coexistence techniques are rarely deployed). A major difference in this case, however, is the greater need for spectrum efficiency. Hopefully further studies and techniques show the importance of the greater efficiency provided by coexistence techniques in the white spaces, pressuring the support for such protocols through regulations (i.e., removing restrictions preventing coexistence techniques).

5.3 Future Work

There are several directions for future research given various narrow aspects of our work, as well as future work that arises from the implications and outlook we have described in the previous section. These directions are meant to be long term venues for research, whereas we have provided more short term issues that need to be addressed in monitoring for heterogeneous networks, spectrum management, and coexistence with spectrum primaries at the end of each of their respective chapters.

The larger problem space: We addressed immediate problems of interference between heterogeneous networks that are small in scope compared to the larger problem space, and sometimes narrow compared to the possibilities of future conflicts between heterogeneous networks. For example, by focusing our monitoring and spectrum management work to the small geographic area of the home where all devices are owned and known to a family, as well as our narrow focus on a single spectrum primary: wireless microphones. Clearly, as we have discussed, the problem of interference between heterogeneous networks and devices is significantly larger. In particular, additional work is needed in larger geographic areas (as we will further discuss below), including environments outside of the home, and in future white space spectrum bands (also discussed below). For example, while we have suggested further exploring spectrum management to address general heterogeneity between networks, we have only shown its potential in the home. For spectrum management to be a better long-term solution than coexistence techniques, further studies are needed to show its capabilities or shortcomings in other environments. Additionally, the wideband nature of secondary devices and narrowband nature of primary users in the white spaces showed significant benefits towards coexistence. Other bands where both primary and secondary devices are wideband may show different results, and will likely need a different approach and technique towards coexistence.

Comprehensive heterogeneous monitoring in other environments: In Chapter 2, we presented a monitoring system for heterogeneous devices leveraging the smartphone to overcome the lack of expertise and equipment needed to monitor the home. Clearly, the design of this work was narrowly focused on the home environment, which led to the use of the smartphone at the base of our monitoring system. The smartphone may not be the best monitor for all environments (e.g., the enterprise), and when one considered geographic areas other than the home: it cannot be assumed that a device like the smartphone would come within close proximity of all potential devices to derive where signals go and what they interfere with. While WifiNet introduced a monitoring system for heterogeneous networks in enterprise environments, it was WiFi-centric and it only derived what heterogeneous signals reached the access points in the enterprise (and their interference strengths). Additionally, better monitoring of heterogeneous technologies in urban environments may help decrease their chaotic nature, which our spectrum assignment system can help re-assign.

Spectrum management in other spectrum bands: The work on spectrum management that we presented in Chapter 3 focused primarily in the 2.4 GHz band, with basic evaluation including the 5 GHz band. Our spectrum management system and algorithm are, again, designed based on fundamentals to support other spectrum bands. However, we provide no study of our algorithm in other bands with unlicensed devices with an understanding of its limitations. For example, in the 900 MHz band or as high as 60 GHz. These bands may not be considered as dense, but different protocols exist in these bands that we have not provided an in-depth analysis with (e.g., DECT). Studies in these bands can further provide confidence to our approach, or highlight limitations.

Scalability of spectrum management in larger environments: Our work in spectrum management focused on the small geographic area of the home that allowed us to explore a global optimization of frequency assignments. In larger geographic areas (e.g., a large enterprise, or urban area), a global optimization may be difficult to scale considering the many potential (and growing number) of frequencies that radios support. Overcoming this scalability problem is a major barrier to addressing larger environments. It may be possible to address this scalability problem with a distributed game theoretic approach, however. With this approach, networks could distributedly use information provided about the environment to compare and determine (locally) what outcome would arise from the radio using the information to behave selfishly and find the best channel for itself, and using the information to contribute to the common good: choosing a channel that would benefit itself and neighboring radios or networks equally.

Deploying configurations from spectrum assignment: A large, essential, and critical piece of functionality missing from spectrum assignment systems today (particularly in the home and chaotic environments), is a mechanism to deploy suggested assignments from spectrum managers and their algorithms. This is a significant challenge given the need for the spectrum management system to directly communicate assignments across heterogeneous technologies that have various communication layers (e.g., at the PHY). Even if the many unlicensed devices and this manager shared a communication layer, there is no unified layer or protocol to parse incoming spectrum reassignment requests. Providing a way to authenticate requests and only deploy certain assignments is also a major challenge to this approach in chaotic environments. A design that allows heterogeneous devices to accept data (global) information from monitoring systems may be needed, or a remote and authorized cloud-based service that takes this information and reassigns spectrum based on it.

Frequency-aware spectrum assignment and adaptation: The algorithm and predictive metric for spectrum assignment that we presented in Chapter 3 do not support and distinguish different properties of the spectrum bands. For example, when calculating the SINR of a link and its expected performance on a different channel, it is assumes the same propagation characteristics as the channel it is currently operating on. This is a reasonable assumption today given the concentration of devices in 2.4 and 5 GHz, however, as more spectrum bands begin to be supported we must be able to more carefully distinguish properties across these bands and consider them in assignment. This is likely to require a completely new study, additional research to develop a frequency-aware predictive metric, and an algorithm that is cognizant of these differences. This work is based on the observation of continuing to support new frequency bands that this dissertation does not focus on.

Considering a fusion of coexistence and spectrum management: In this dissertation, we have considered coexistence and spectrum management in isolation. That is, using coexistence solely to avoid interference with spectrum primaries, and spectrum management to reduce interference in the unlicensed bands. However, a fusion between these techniques may be beneficial. In particular, one might perform spectrum management in the unlicensed bands and the management algorithm may determine that particular resulting conflicts are best solved by enabling an otherwise disabled coexistence technique (e.g., disabled to prevent inefficiencies when not in conflict). Additionally, the algorithm may determine a more efficient solution is possible if certain coexistence techniques are used by particular networks. In white space spectrum bands, one must also consider spectrum management amongst the secondary devices that are also operating with coexistence techniques to avoid interference with spectrum primaries. Our work did not evaluate the impact of spectrum management across devices that may also be behaving differently due to certain coexistence mechanisms being enabled to collectively avoid interference with primaries.

Coexistence in other licensed (white space) bands: The TV white spaces have shown the potential of providing additional spectrum through dynamic spectrum access with the presence of primary users. To continue to meet the growing demand, it is likely that other spectrum bands will be opened for white space use. Future studies will be needed on how to provide spectrum-efficient coexistence with the primary users in these bands, which are likely to have different interference properties than the microphone systems we have studied. Developing coexistence protocols with these primaries may be necessary to ensure efficient use of the spectrum. In particular, studies of coexistence techniques that deal with wider band primaries are needed. The focus of our work in Chapter 4, in addition to other works [8,42], focus on narrowband primaries. Additionally, the wireless microphone primary in the white spaces was analog. Dealing with digital primaries in other bands will require new interference studies and can lead to new challenges to ensure interference-free and spectrum-efficient coexistence. Each of these studies should compare to a spectrum management-based approach in these new bands to understand the benefit of our general approach in focusing on spectrum management as a more spectrum-efficient solution for avoiding interference between primaries and secondaries in the white spaces.

Safe regulations that allow for coexistence techniques: Aside from technical contributions as future work, there is interesting future work in spectrum policy and regulations that could ensure interference-free coexistence techniques in the white spaces. In particular, regulations and rules that allow for various types of techniques to be deployed, rather than specifying a single protocol that the primaries and secondaries must follow to achieve spectrum-efficient coexistence. In other words, the regulation should not require SEISMIC, but rather allow protocols like SEISMIC to exist in the spectrum. How one regulates what protocols are allowed or ensures that they do not interfere with each other is a significant challenge. In general, a study of what policies are needed to allow coexistence in white space bands is needed.

Lessons for future white space bands: Finally, given that the TV white spaces are the first spectrum band of their type, it is important to take away key lessons from our work that can guide policy, organization, and future research on white space spectrum bands. In other words, if a future band was being considered for white space technology, what does our work and studies suggest? First, spectrum-efficient rules and regulations are needed to avoid significant amounts of inactive spectrum to protect primary users. Information about where spectrum primaries, where their signals propagate, and whether secondary devices will impact these signals has also been shown to be critical. As we saw through the evolution of the TV white space regulations, sensing was too difficult in practice, leading to the need for more global and centralized data about primary users and their signals to make the deployments more practical. Therefore, the type of information needed and how it is provided to perform such dynamic spectrum access should be carefully considered in future white space bands. Either through a database, or a more localized and distributed feedback mechanism such as our strobes in SEISMIC to convey information. As of now, it has been assumed that primaries are not modified in any way to allow white space spectrum access. However, small changes (e.g., allowing them to notify of impending interference) can allow better and potentially more effective interference avoidance. Future bands should consider whether primary users can be (simply) modified in any way to improve spectrum efficiency or interference avoidance. Finally, as discussed above, regulations should consider the possibility of coexistence and not enforce frequency isolation if spectrum-efficient coexistence is possible, and it can guarantee interference avoidance.

References

[1] CTIA - The International Association for the Wireless Telecommunications Industry. **The CTIA Files Ex Parte to FCC to Request More Spectrum**. `http://ctia.org/media/press/body.cfm/prid/1866`. 1, 5, 13

[2] ZigBee Alliance and Wireless Standard. `http://www.zigbee.com/`. 1

[3] IEEE 802.11 set of standards. **Telecommunications and information exchange between systems in local and metropolitan area networks.** `http://www.ieee802.org/11/`. 1

[4] The Bluetooth short range wireless connectivity standard. `http://bluetooth.com`. 1

[5] The Z-Wave standard. `http://www.z-wave.com/modules/ZwaveStart/`. 1

[6] The DECT standard. `http://www.dect.org/`. 1

[7] Shravan Rayanchu, Ashish Patro, and Suman Banerjee. **Airshark: detecting non-WiFi RF devices using commodity WiFi hardware**. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, IMC '11, pages 137–154, New York, NY, USA, 2011. ACM. 2, 4, 9, 18, 21, 22, 32, 45, 49, 73, 78, 79, 90, 92

[8] Shyamnath Gollakota, Fadel Adib, Dina Katabi, and Srinivasan Seshan. **Clearing the RF smog: making 802.11n robust to cross-technology interference**. In *Proceedings of the ACM SIGCOMM 2011 conference*, SIGCOMM '11, pages 170–181, New York, NY, USA, 2011. ACM. 2, 4, 9, 11, 14, 15, 21, 45, 49, 57, 79, 90, 92, 130, 134

[9] **Notes on Industrial, Scientific and Medical (ISM) frequency band. OF-COM.** `http://www.ofcom.org.uk/static/archive/ra/topics/research/rrac/members/rrac01_19.doc`, 2002. 4

[10] Daniel Halperin, Srikanth Kandula, Jitendra Padhye, Paramvir Bahl, and David Wetherall. **Augmenting data center networks with multi-gigabit wireless links**. In *Proceedings of the ACM SIGCOMM 2011 conference*, SIGCOMM '11, pages 38–49, New York, NY, USA, 2011. ACM. 4

[11] Xia Zhou, Zengbin Zhang, Yibo Zhu, Yubo Li, Saipriya Kumar, Amin Vahdat, Ben Y. Zhao, and Haitao Zheng. **Mirror mirror on the ceiling: flexible wireless links for data centers**. In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, SIGCOMM '12, pages 443–454, New York, NY, USA, 2012. ACM. 4

[12] **Wireless HD. http://www.wirelesshd.org/**. 4

[13] Federal Communications Commission (FCC). **Title 47, Part 15 of FCC Federal Regulations in Telecommunication.** `http://www.gpo.gov/fdsys/pkg/CFR-2010-title47-vol1/pdf/CFR-2010-title47-vol1.pdf`, 2013. 4, 95

[14] Kenneth Carter. **Unlicensed to kill: a brief history of the Part 15 rules**. Info, Vol. 11, No. 5, pp. 8-18, 2009. 4

[15] Shravan Rayanchu, Ashish Patro, and Suman Banerjee. **Catching whales and minnows using WiFiNet: deconstructing non-WiFi interference using WiFi hardware**. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, NSDI'12, pages 5–5, Berkeley, CA, USA, 2012. USENIX Association. 4, 14, 16, 17, 18, 19, 22, 23, 45, 47, 49, 50, 54, 57, 60, 69, 73, 78

[16] Chieh-Jan Mike Liang, Nissanka Bodhi Priyantha, Jie Liu, and Andreas Terzis. **Surviving wi-fi interference in low power ZigBee networks**. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, SenSys '10, pages 309–322, New York, NY, USA, 2010. ACM. 4, 9, 11, 12, 14, 15, 18, 21, 45, 47, 49, 57, 73, 78, 81, 90, 92, 130

[17] **FCC press release, FCC Adopts Rules for Unlicensed Use of Television White Spaces**. `http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-260A1.pdf`, November 2008. 5, 13, 14, 97

[18] **Unlicensed Operation in the TV Broadcast Bands**. `http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-04-113A1.pdf`, May 2004. 5, 6, 14, 95, 106

[19] **FCC Adopts Rule For Unlicensed Use Of Television White Spaces, http://www.fcc.gov/**. 5, 6

[20] **FCC press release, FCC Frees Up Vacant TV Airwaves For "Super Wi-Fi" Technologies and Other Technologies**. `http://www.fcc.gov/Daily_Releases/Daily_Business/2010/db0923/DOC-301650A1.pdf`, September 2010. 5, 6, 95, 97, 100, 101

[21] Paramvir Bahl, Ranveer Chandra, Thomas Moscibroda, Rohan Murty, and Matt Welsh. **White space networking with wi-fi like connectivity**. In *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, SIGCOMM '09, pages 27–38, New York, NY, USA, 2009. ACM. 5, 6, 15, 54, 56, 69, 100, 101, 130

[22] Rohan Murty, Ranveer Chandra, Thomas Moscibroda, and Paramvir (Victor) Bahl. **SenseLess: A Database-Driven White Spaces Network**. **11**, pages 189–203, Piscataway, NJ, USA, February 2012. IEEE Educational Activities Department. 6, 14, 15, 130

[23] Ranveer Chandra, Thomas Moscibroda, Paramvir Bahl, Rohan Murty, George Nychis, and Xiaohui Wang. **A campus-wide testbed over the TV white spaces**. *SIGMOBILE Mob. Comput. Commun. Rev.*, **15**(3):2–9, November 2011. 6

[24] **FCC approves Google's white space wireless database.** `http://www.engadget.com/2013/06/30/fcc-approves-google-white-space-wireless-database/`. 6

[25] **FCC approves latecomer Microsoft as white-space database provider.** `http://www.fiercebroadbandwireless.com/story/fcc-approves-latecomer-microsoft-white-space-database-provider/2011-08-03`. 6

[26] **'White Space' Super Wi-Fi Coming Soon.** `http://www.forbes.com/sites/elizabethwoyke/2011/01/25/white-space-wireless-access-deployment-nears/`, January 2011. 6

[27] **Microsoft-led consortium to trial super WiFi network across the pond.** `http://www.engadget.com/2011/06/27/microsoft-led-consortium-to-trial-super-wifi-network-across-the/`. 6

[28] **'Super Wi-Fi' Has 100km Range, Is Coming to Save the Day.** `http://www.techhive.com/article/236829/super_wi_fi_has_100km_range_iscoming_to_save_the_day.html`, July 2011. 6

[29] Byung jae Kwak, Nah oak Song, Leonard E. Miller, and Senior Member. **Performance analysis of exponential backoff**. *IEEE/ACM Transactions on Networking*, **13**:343–355, 2005. 7

[30] Patrick Verkaik, Yuvraj Agarwal, Rajesh Gupta, and Alex C. Snoeren. **Softspeak: making VoIP play well in existing 802.11 deployments**. In *Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, NSDI'09, pages 409–422, Berkeley, CA, USA, 2009. USENIX Association. 8

[31] Hardin G. **'Tragedy of the Unmanaged Commons'**. In *Trends in Ecology & Evolution*, 1994. 9

[32] George Nychis, Ranveer Chandra, Thomas Moscibroda, Ivan Tashev, and Peter Steenkiste. **Reclaiming the white spaces: spectrum efficient coexistence with primary users**. In *Proceedings of the Seventh COnference on emerging Networking EXperiments and Technologies*, CoNEXT '11, pages 1:1–1:12, New York, NY, USA, 2011. ACM. 9

[33] I. Howitt. **Bluetooth performance in the presence of 802.11b WLAN**. *Vehicular Technology, IEEE Transactions on*, **51**(6):1640 – 1651, November 2002. 9, 10, 12, 14, 21, 49

[34] M. Petrova, Lili Wu, P. Mahonen, and J. Riihijarvi. **Interference Measurements on Performance Degradation between Colocated IEEE 802.11g/n and IEEE 802.15.4 Networks**. In *Networking, 2007. ICN '07. Sixth International Conference on*, pages 93 –93, 2007. 11, 49, 81

[35] Sandra Fiehe, Janne Riihijärvi, and Petri Mähönen. **Experimental study on performance of IEEE 802.11n and impact of interferers on the 2.4 GHz ISM band**. In *Proceedings of the 6th International Wireless Communications and Mobile Computing Conference*, IWCMC '10, pages 47–51, New York, NY, USA, 2010. ACM. 11, 21

[36] **Report of the Unlicensed Devices and Experimental Licenses Working Group. FCC Spectrum Policy Task Force, final report.** `http://transition.fcc.gov/sptf/files/E&UWGFinalReport.pdf`, 2002. 13

[37] George Nychis, Thibaud Hottelier, Zhuocheng Yang, Srinivasan Seshan, and Peter Steenkiste. **Enabling MAC protocol implementations on software-defined radios**. In *Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, NSDI'09, pages 91–105, Berkeley, CA, USA, 2009. USENIX Association. 13, 106

[38] Soo Young Shin, Hong Seong Park, Sunghyun Choi, and Wook-Hyun Kwon. **Packet Error Rate Analysis of ZigBee Under WLAN and Bluetooth Interferences**. *Wireless Communications, IEEE Transactions on*, **6**(8):2825–2830, 2007. 14, 49

[39] Xinyu Zhang and Kang G. Shin. **Enabling coexistence of heterogeneous wireless systems: case for ZigBee and WiFi**. In *Proceedings of the Twelfth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, MobiHoc '11, pages 6:1–6:11, New York, NY, USA, 2011. ACM. 14, 15, 81, 130

[40] N. Golmie, N. Chevrollier, and O. Rebala. **Bluetooth and WLAN Coexistence: Challenges and Solutions**. *IEEE Wireless Communications Magazine*, **10**:22–29, 2003. 14, 49

[41] Jeffrey Wojtiuk. **Bluetooth and WiFi integration: Solving co-existence challenges**. *Defense electronics magazine*, 2004. 14, 15, 130

[42] Hariharan Rahul, Nate Kushman, Dina Katabi, Charles Sodini, and Farinaz Edalat. **Learning to share: narrowband-friendly wideband networks**. In *Proceedings of the ACM SIGCOMM 2008 conference on Data communication*, SIGCOMM '08, pages 147–158, New York, NY, USA, 2008. ACM. 14, 16, 47, 98, 101, 102, 134

[43] S.M. Mishra, R.W. Brodersen, S.T. Brink, and R. Mahadevappa. **Detect and avoid: an ultra-wideband/WiMAX coexistence mechanism**. *Communications Magazine, IEEE*, **45**(6):68 –75, 2007. 14, 15, 102, 130

[44] A Kumar. **Enhancing coexistence, quality of service, and energy performance in dynamic spectrum access networks**. In *ACM MobiCom*, 2004. 14, 15, 130

[45] Lei Yang, Wei Hou, Lili Cao, Ben Y. Zhao, and Haitao Zheng. **Supporting demanding wireless applications with frequency-agile radios**. In *Proceedings of the 7th USENIX conference on Networked systems design and implementation*, NSDI'10, pages 5–5, Berkeley, CA, USA, 2010. USENIX Association. 14, 98, 101

[46] A. Attar and A.H. Aghvami. **A Framework for Unified Spectrum Management (USM) in Heterogeneous Wireless Networks**. *Communications Magazine, IEEE*, **45**(9):44–51, 2007. 14

[47] S. Buljore, V. Merat, H. Harada, S. Filin, P. Houze, K. Tsagkaris, V. Ivanov, K. Nolte, T. Farnham, and O. Holland. **IEEE P1900.4 System Overview on Architecture and Enablers for Optimised Radio and Spectrum Resource Usage**. In *New Frontiers in Dynamic Spectrum Access Networks, 2008. DySPAN 2008. 3rd IEEE Symposium on*, pages 1 –8, 2008. 14

[48] IEEE 802.19. **Workgroup on standards for coexistence between wireless standards of unlicensed devices**. `http://www.ieee802.org/19/`. 15, 130

[49] D. Borth, R. Ekl, B. Oberlies, and S. Overby. **Considerations for Successful Cognitive Radio Systems in US TV White Space**. In *New Frontiers in Dynamic Spectrum Access Networks, 2008. DySPAN 2008. 3rd IEEE Symposium on*, pages 1–5, 2008. 15, 100, 130

[50] G.J. Buchwald, S.L. Kuffner, L.M. Ecklund, M. Brown, and E.H. Callaway. **The Design and Operation of the IEEE 802.22.1 Disabling Beacon for the Protection of TV Whitespace Incumbents**. In *New Frontiers in Dynamic Spectrum Access Networks, 2008. DySPAN 2008. 3rd IEEE Symposium on*, pages 1–6, 2008. 15, 101, 102, 103, 130

[51] David Gurney, Greg Buchwald, Larry Ecklund, Steve Kuffner, and John Grosspietsch. **Geo-location Database Techniques for Incumbent Protection in the TV White Space**. In *DySpan*, 2008. 15, 130

[52] Yuan Yuan, Paramvir Bahl, Ranveer Chandra, Philip A. Chou, Ian Ferrell, Thomas Moscibroda, Srihari Narlanka, and Yunnan Wu. **KNOWS: Kognitiv Networking Over White Spaces**. In *DySPAN*, 2007. 15, 130

[53] Hyoil Kim and Kang G. Shin. **Fast Discovery of Spectrum Opportunities in Cognitive Radio Networks**. In *DySPAN*, 2008. 15, 130

[54] Paramvir Bahl, Ranveer Chandra, Jitendra Padhye, Lenin Ravindranath, Manpreet Singh, Alec Wolman, and Brian Zill. **Enhancing the security of corporate Wi-Fi networks using DAIR**. In *Proceedings of the 4th international conference on Mobile systems, applications and services*, MobiSys '06, pages 1–14, New York, NY, USA, 2006. ACM. 16, 18, 19, 22, 49

[55] Yu-Chung Cheng, John Bellardo, Péter Benkö, Alex C. Snoeren, Geoffrey M. Voelker, and Stefan Savage. **Jigsaw: solving the puzzle of enterprise 802.11 analysis**. In *Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '06, pages 39–50, New York, NY, USA, 2006. ACM. 16, 19, 22, 49, 56, 60, 69

[56] Chunyi Peng, Haitao Zheng, and Ben Y. Zhao. **Utilization and fairness in spectrum assignment for opportunistic spectrum access**. *Mob. Netw. Appl.*, **11**(4):555–576, August 2006. 17, 54, 55, 56, 57, 60, 69

[57] E. Rozner, Y. Mehta, A. Akella, and Lili Qiu. **Traffic-Aware Channel Assignment in Enterprise Wireless LANs**. In *Network Protocols, 2007. ICNP 2007. IEEE International Conference on*, pages 133–143, 2007. 17, 55, 56, 57

[58] Arunesh Mishra, Eric Rozner, Suman Banerjee, and William Arbaugh. **Exploiting partially overlapping channels in wireless networks: turning a peril into an advantage**. In *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, IMC '05, pages 29–29, Berkeley, CA, USA, 2005. USENIX Association. 17, 54, 56

[59] Kaushik Lakshminarayanan, Samir Sapra, Srinivasan Seshan, and Peter Steenkiste. **RFDump: an architecture for monitoring the wireless ether**. In *Proceedings of the 5th international conference on Emerging networking experiments and technologies*, CoNEXT '09, pages 253–264, New York, NY, USA, 2009. ACM. 22, 48

[60] Steven Siying Hong and Sachin Rajsekhar Katti. **DOF: a local wireless information plane**. In *Proceedings of the ACM SIGCOMM 2011 conference*, SIGCOMM '11, pages 230–241, New York, NY, USA, 2011. ACM. 22, 48

[61] TI News Release. **TI demonstrates world's first Android development platform bringing ZigBee to Smartphones and Tablets**. `http://www.ti.com/ww/in/news_detail/2011/news_detail_20110304.html`. 24, 36

[62] Souvik Sen, Romit Roy Choudhury, and Srihari Nelakuditi. **SpinLoc: spin once to know your location**. In *Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications*, HotMobile '12, pages 12:1–12:6, New York, NY, USA, 2012. ACM. 30

[63] **IEEE Registration Authority OUI Public Listing**. `http://standards.ieee.org/develop/regauth/oui/public.html`. 32

[64] Zhenyun Zhuang, Kyu-Han Kim, and Jatinder Pal Singh. **Improving energy efficiency of location sensing on smartphones**. In *Proceedings of the 8th international conference on Mobile systems, applications, and services*, MobiSys '10, pages 315–330, New York, NY, USA, 2010. ACM. 35

[65] Frank Dabek, Russ Cox, Frans Kaashoek, and Robert Morris. **Vivaldi: a decentralized network coordinate system**. In *Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '04, pages 15–26, New York, NY, USA, 2004. ACM. 41

[66] **20 Myths of Wi-Fi Interference.** `http://www.cisco.com/en/US/prod/collateral/wireless/ps9391/ps9393/ps9394/prod_white_paper0900aecd807395a9_ns736_Networking_Solutions_White_Paper.html`. 48

[67] **MetaGeek — Home of Wi-Spy and inSSIDer**. `http://www.metageek.net/`. 48

[68] **AirMagnet AirMedic and Spectrum XT**. `http://www.airmagnet.net/products`. 48

[69] **Bandspeed AirMaestro spectrum analysis solution**. `http://www.bandspeed.com`. 48

[70] **Cisco Spectrum Expert**. `http://www.cisco.com/en/US/products/ps9393/index.html`. 48

[71] **Blind signal detection and identification over the 2.4 GHz ISM band for cognitive radio.** *Masters thesis, University of Southern Florida*, 2009. 48

[72] Frank Brickle. **Automatic Signal Classification for Software Defined Radios**. November/December 2003 issue of QEX, 2003. 48

[73] **DttSP: the basic modulation, demodulation, signal conditioning, and synchronization processes required to operate a high performance transceiver using DSP as the detection and synthesis stages**. `http://dttsp.sourceforge.net/`. 48

[74] Gladstone Marballie. **Symbol Timing and Coarse Classification of Phase Modulated Signals on a Standalone SDR Platform**. *Ph.D. thesis, Virginia Polytechnic Institute and State University*, 2010. 48

[75] Paramvir Bahl, Ranveer Chandra, Jitendra Padhye, Lenin Ravindranath, Manpreet Singh, Alec Wolman, and Brian Zill. **DAIR: A framework for managing enterprise wireless networks using desktop infrastructure**. In *HotNets*, 2005. 49

[76] J.D. Freeman and S. Simi. **Remote monitoring of indoor environment using mobile robot based wireless sensor network**. In *Computer Science Education (ICCSE), 2011 6th International Conference on*, pages 1080–1084, 2011. 49

[77] YuanYuan Li and L.E. Parker. **Detecting and monitoring time-related abnormal events using a wireless sensor network and mobile robot**. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pages 3292–3298, 2008. 49

[78] Kaushik Lakshminarayanan, Srinivasan Seshan, and Peter Steenkiste. **Understanding 802.11 performance in heterogeneous environments**. In *Proceedings of the 2nd ACM SIGCOMM workshop on Home networks*, HomeNets '11, pages 43–48, New York, NY, USA, 2011. ACM. 49

[79] A. Baid, S. Mathur, I. Seskar, S. Paul, A. Das, and D. Raychaudhuri. **Spectrum MRI: Towards diagnosis of multi-radio interference in the unlicensed band**. In *Wireless Communications and Networking Conference (WCNC), 2011 IEEE*, pages 534–539, 2011. 49, 50

[80] Mythili Vutukuru, Kyle Jamieson, and Hari Balakrishnan. **Harnessing exposed terminals in wireless networks**. In *USENIX NSDI '07*, 2008. 53, 56

[81] Arunesh Mishra, Suman Banerjee, and William Arbaugh. **Weighted coloring based channel assignment for WLANs**. *SIGMOBILE Mob. Comput. Commun. Rev.*, **9**(3):19–31, July 2005. 53, 54, 55, 56, 58

[82] S. Ramanathan. **A unified framework and algorithm for channel assignment in wireless networks**. *Wirel. Netw.*, **5**(2):81–94, March 1999. 54

[83] Maaly Hassan and Andrew Chickadel. **A Review of Interference Reduction in Wireless Networks Using Graph Coloring Methods**. March 2011. 55, 56, 57

[84] Rohan Murty, Jitendra Padhye, Alec Wolman, and Matt Welsh. **Dyson: an architecture for extensible wireless LANs**. In *Proceedings of the 2010 USENIX conference on USENIX annual technical conference*, USENIXATC'10, pages 15–15, Berkeley, CA, USA, 2010. USENIX Association. 55, 57, 59

[85] Rohan Murty, Jitendra Padhye, Ranveer Chandra, Alec Wolman, and Brian Zill. **Designing high performance enterprise Wi-Fi networks**. In *Proceedings of the*

5th USENIX Symposium on Networked Systems Design and Implementation, NSDI'08, pages 73–88, Berkeley, CA, USA, 2008. USENIX Association. 55, 57, 59

[86] Yong Ding, Yi Huang, Guokai Zeng, and Li Xiao. **Channel assignment with partially overlapping channels in wireless mesh networks**. In *Proceedings of the 4th Annual International Conference on Wireless Internet*, WICON '08, pages 38:1–38:9, ICST, Brussels, Belgium, Belgium, 2008. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering). 55, 56, 57

[87] S. Ramanathan. **A unified framework and algorithm for channel assignment in wireless networks**. *Wirel. Netw.*, **5**(2):81–94, March 1999. 55

[88] Kamal Jain, Jitendra Padhye, Venkata N. Padmanabhan, and Lili Qiu. **Impact of interference on multi-hop wireless network performance**. In *Proceedings of the 9th annual international conference on Mobile computing and networking*, MobiCom '03, pages 66–80, New York, NY, USA, 2003. ACM. 56

[89] Hongkun Li, Yu Cheng, Chi Zhou, and P. Wan. **Multi-dimensional Conflict Graph Based Computing for Optimal Capacity in MR-MC Wireless Networks**. In *Distributed Computing Systems (ICDCS), 2010 IEEE 30th International Conference on*, pages 774–783, 2010. 56

[90] Xi Liu, Anmol Sheth, Michael Kaminsky, Konstantina Papagiannaki, Srinivasan Seshan, and Peter Steenkiste. **DIRC: increasing indoor wireless capacity using directional antennas**. In *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, SIGCOMM '09, pages 171–182, New York, NY, USA, 2009. ACM. 56

[91] Arunesh Mishra, Vivek Shrivastava, Suman Banerjee, and William Arbaugh. **Partially-overlapped Channels not considered harmful**. In *ACM Sigmetrics*, 2006. 56

[92] Balabhaskar Balasundaram and Sergiy Butenko. **Graph Domination, Coloring and Cliques in Telecommunications**. 56

[93] Vivek Shrivastava et al. **CENTAUR: realizing the full potential of centralized wlans through a hybrid data path**. In *ACM MobiCom*, pages 297–308, 2009. 56

[94] H. Hernandez and C. Blum. **Distributed graph coloring in wireless ad hoc networks: A light-weight algorithm based on Japanese tree frogs' calling behaviour**. In *Wireless and Mobile Networking Conference (WMNC), 2011 4th Joint IFIP*, pages 1–7, 2011. 56

[95] Der-Rong Din and Shian-Shyong Tseng. **A Simulated Annealing Algorithm for Extended Cell Assignment Problem in a Wireless ATM Network**. In *Applications of Evolutionary Computing Lecture Notes in Computer Science*, 2001. 56

[96] M. Cesana, D. Maniezzo, P. Bergamo, and M. Gerla. **Interference aware (IA) MAC: an enhancement to IEEE802.11b DCF**. In *Vehicular Technology Conference, 2005*, 2005. 56

[97] Sooyeol Im and Hyuckjae Lee. **Dynamic spectrum allocation based on binary integer programming under interference graph**. In *Personal Indoor and Mobile Radio Communications (PIMRC), 2012 IEEE 23rd International Symposium on*, pages 226–231, 2012. 56, 57

[98] Sang-Seon Byun and I. Balasingham. **Approximations of multiobjective optimization for dynamic spectrum allocation in wireless sensor networks**. In *Consumer Electronics (ICCE), 2010 Digest of Technical Papers International Conference on*, pages 427–428, 2010. 56

[99] R. Akl and A. Arepally. **Dynamic Channel Assignment in IEEE 802.11 Networks**. In *Portable Information Devices, 2007. PORTABLE07. IEEE International Conference on*, pages 1–5, 2007. 57

[100] Lise Getoor and Christopher P. Diehl. **Link mining: a survey**. *SIGKDD Explor. Newsl.*, **7**(2):3–12, December 2005. 65

[101] Matthias Brocheler, Andrea Pugliese, and V. S. Subrahmanian. **COSI: Cloud Oriented Subgraph Identification in Massive Social Networks.** In *ASONAM*, pages 248–255. IEEE Computer Society, 2010. 65

[102] Jure Leskovec, Ajit Singh, and Jon Kleinberg. **Patterns of influence in a recommendation network**. In *Proceedings of the 10th Pacific-Asia conference on Advances in Knowledge Discovery and Data Mining*, PAKDD'06, pages 380–389, Berlin, Heidelberg, 2006. Springer-Verlag. 65

[103] Zhiping Zeng, Jianyong Wang, Lizhu Zhou, and George Karypis. **Out-of-core coherent closed quasi-clique mining from large dense graph databases**. *ACM Trans. Database Syst.*, **32**(2), June 2007. 65

[104] Xifeng Yan, X. Jasmine Zhou, and Jiawei Han. **Mining closed relational graphs with connectivity constraints**. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, KDD '05, pages 324–333, New York, NY, USA, 2005. ACM. 65

[105] Caleb C. Noble and Diane J. Cook. **Graph-based anomaly detection**. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '03, pages 631–636, New York, NY, USA, 2003. ACM. 65

[106] Zhao Zhao, Guanying Wang, A.R. Butt, M. Khan, V.S.A. Kumar, and M.V. Marathe. **SAHAD: Subgraph Analysis in Massive Networks Using Hadoop**. In *Parallel Distributed Processing Symposium (IPDPS), 2012 IEEE 26th International*, pages 390–401, 2012. 67

[107] David Eppstein. **Subgraph isomorphism in planar graphs and related problems**. In *Proceedings of the sixth annual ACM-SIAM symposium on Discrete algorithms*, SODA '95, pages 632–640, Philadelphia, PA, USA, 1995. Society for Industrial and Applied Mathematics. 67

[108] C. C. Peterson. **A note on transforming the product of variables to linear form in linear programs**. 1971. 75

[109] **Ubiquity Generator Framework**. `http://ug.zib.de/`. 77

[110] Yuji Shinano, Stefan Heinz, Stefan Vigerske, and Michael Winkler. **FiberSCIP - A shared memory parallelization of SCIP**. September 2013. 77

[111] **SCIP: Solving Constraint Integer Programs**. `http://scip.zib.de/`. 78

[112] Ramakrishna Gummadi, David Wetherall, Ben Greenstein, and Srinivasan Seshan. **Understanding and mitigating the impact of RF interference on 802.11 networks**. In *Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '07, pages 385–396, New York, NY, USA, 2007. ACM. 79

[113] Lara Deek, Eduard Garcia-Villegas, Elizabeth Belding, Sung-Ju Lee, and Kevin Almeroth. **The impact of channel bonding on 802.11n network management**. In *Proceedings of the Seventh COnference on emerging Networking EXperiments and Technologies*, CoNEXT '11, pages 11:1–11:12, New York, NY, USA, 2011. ACM. 93

[114] **Connecting America: The Nation Broadband Plan (FCC)**. http://download.broadband.gov/plan/national-broadband-plan.pdf. 97

[115] J.D. Poston and W.D. Horne. **Discontiguous OFDM considerations for dynamic spectrum access in idle TV channels**. In *New Frontiers in Dynamic Spectrum Access Networks, 2005. DySPAN 2005. 2005 First IEEE International Symposium on*, pages 607–610, 2005. 98, 101

[116] Glenn Judd, Xiaohui Wang, and Peter Steenkiste. **Efficient channel-aware rate adaptation in dynamic environments**. In *Proceedings of the 6th international conference on Mobile systems, applications, and services*, MobiSys '08, pages 118–131, New York, NY, USA, 2008. ACM. 98

[117] **Audio-Technica Tutorial on FM Mic Squelching**. http://www.audio-technica.com/cms/site/557bfecb9ab0cfee/index.html/. 99

[118] Srihari Narlanka, Ranveer Chandra, Paramvir Bahl, and Ian Ferrell. **A Hardware Platform for Utilizing the TV Bands with a Wi-Fi Radio**. In *IEEE LANMAN*, June 2007. 100

[119] **Demonstration of a Prototype Dynamic Spectrum Access System, Philips Research**. In *DySPAN demo session*, 2008. 100

[120] Shaoyi Xu, Yanlei Shang, and Haiming Wang. **Application of SVD to sense wireless microphone signals in a wideband cognitive radio network**. In *Signal Processing and Communication Systems, 2008. ICSPCS 2008. 2nd International Conference on*, pages 1–7, 2008. 100

[121] Harpreet S. Dhillon, Jeong-O Jeong, Dinesh Datla, Michael Benonis, R. Michael Buehrer, and Jeffrey H. Reed. **A sub-space method to detect multiple wireless microphone signals in TV band white space**. *Analog Integr. Circuits Signal Process.*, **69**(2-3):297–306, December 2011. 100

[122] Dan Zhang, Lijun Dong, and Narayan B. Mandayam. **Sensing Wireless Microphone with ESPRIT from Noise and Adjacent Channel Interference.** In *GLOBECOM*, pages 1–5. IEEE, 2010. 100

[123] **FCC press release, Evaluation of the Performance of Prototype TV-Band White Space Devices**. `http://hraunfoss.fcc.gov/edocs_public/attachmatch/DA-08-2243A3.pdf`, November 2008. 100

[124] Wu Yu-chun, Wang Haiguang, and P. Zhang. **Protection of Wireless Microphones in IEEE 802.22 Cognitive Radio Network**. In *Communications Workshops, 2009. ICC Workshops 2009. IEEE International Conference on*, pages 1–5, 2009. 101, 102

[125] **Mitigating the Effects of Unlicensed Devices on Wireless Microphones, Ahren Hartman and Edgar Reihl SHURE Incorporated, SHURE Incorporated**. 2009. 101, 102, 103

[126] **Private communication with J. Stoffo and J. Eskew of Professional Wireless Systems, http://www.professionalwireless.com/**. 101

[127] **Ultra-wideband Technology for Short-Range, High-Rate Wireless Communications, http://www.ieee.or.com/Archive/uwb.pdf**. 101

[128] Rong-Rong Chen and Xin Liu. **Coexisting with CSMA-Based Reactive Primary Users**. In *New Frontiers in Dynamic Spectrum Access Networks, 2008. DySPAN 2008. 3rd IEEE Symposium on*, pages 1 –7, 2010. 102

[129] S.M. Mishra, R. Tandra, and A. Sahai. **Coexistence with Primary Users of Different Scales**. In *New Frontiers in Dynamic Spectrum Access Networks, 2007. DySPAN 2007. 2nd IEEE International Symposium on*, pages 158 –167, 2007. 102

[130] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra. **Perceptual Evaluation of Speech Quality (PESQ)-A New Method for Speech Quality Assessment of Telephone Networks and Codecs**. In *IEEE Conf. on Acoustics, Speech, and Signal Processing*, 2001. 103

[131] K Leentvaar and J Flint. **The Capture Effect in FM Receivers**. In *IEEE Transactions on Communications*, 1976. 106

[132] Daniel Halperin, Wenjun Hu, Anmol Sheth, and David Wetherall. **Predictable 802.11 packet delivery from wireless channel measurements**. **40**, pages 159–170, New York, NY, USA, August 2010. ACM. 111