

Washington University in St. Louis
Washington University Open Scholarship

All Theses and Dissertations (ETDs)

January 2011

Empirical Studies for Reliable Home Area Wireless Sensor Networks

Mo Sha

Washington University in St. Louis

Follow this and additional works at: <http://openscholarship.wustl.edu/etd>

Recommended Citation

Sha, Mo, "Empirical Studies for Reliable Home Area Wireless Sensor Networks" (2011). *All Theses and Dissertations (ETDs)*. 487.
<http://openscholarship.wustl.edu/etd/487>

This Thesis is brought to you for free and open access by Washington University Open Scholarship. It has been accepted for inclusion in All Theses and Dissertations (ETDs) by an authorized administrator of Washington University Open Scholarship. For more information, please contact digital@wumail.wustl.edu.

WASHINGTON UNIVERSITY IN ST. LOUIS
School of Engineering and Applied Science
Department of Computer Science and Engineering

Thesis Examination Committee:
Chenyang Lu, Chair
Yixin Chen
Christopher Gill

EMPIRICAL STUDIES FOR RELIABLE HOME AREA WIRELESS SENSOR
NETWORKS

by

Mo Sha

A thesis presented to the School of Engineering
of Washington University in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE

May 2011
Saint Louis, Missouri

copyright by

Mo Sha

2011

ABSTRACT OF THE THESIS

Empirical Studies for Reliable Home Area Wireless Sensor Networks

by

Mo Sha

Master of Science in Computer Science

Washington University in St. Louis, 2011

Research Advisor: Professor Chenyang Lu

Abstract: Home Area Networks (HANs) consisting of wireless sensors have emerged as the enabling technology for important applications such as smart energy and assisted living. A key challenge faced by HANs is maintaining reliable operation in real-world residential environments. In this thesis research, empirical studies on the spectrum usage in the 2.4 GHz band as well as 802.15.4 wireless channels are performed in diversified real residential environments. Based on the insights drawn from empirical studies, network design guideline and practical solution for Home Area Sensor Network are provided.

Acknowledgments

First of all, I would like to thank my advisor Dr. Chenyang Lu for his guidance, suggestions, continuous encouragement and constant support through out my study at Washington University in St. Louis, for his patience, kindness and professionalism while I chose my research direction in the early year, for helping me to learn how to be disciplined and persistent. Dr. Chenyang Lu is an excellent advisor in both research and life. Without his help, it would be impossible for me to achieve my work.

Thanks also go to members of our research group. Thanks to Greg Hackmann, Chengjie Wu, Sisu Xi, Yong Fu, Bo Li, Abu Sayeed Saifullah and Rahav Dor.

Last but not least, I would like give my thanks to my family: my dear mom, dad and my girlfriend for encouraging me, supporting me and being proud of me all the time. It is the foundation and ideals they instilled in me that made all of this come true.

Mo Sha

Washington University in Saint Louis
May 2011

Dedicated to my parents.

Contents

Abstract	ii
Acknowledgments	iii
List of Tables	vi
List of Figures	vii
1 Introduction	1
2 Related Work	4
3 Wireless Spectrum Study	6
3.1 Experimental Methodology	7
3.2 Is There a Common Idle Channel in Different Homes?	8
3.3 Does Spectrum Usage Change with Time?	9
3.4 Is Wi - Fi the Dominant Source of Spectrum Usage?	11
4 Multi-Channel Link Study	13
4.1 Experimental Methodology	13
4.2 Is There a Persistently Good Channel?	16
4.3 Is Retransmission Sufficient?	19
4.4 Is Channel Diversity Effective?	20
4.5 Can Hopping be Scheduled Statically?	22
4.6 How Should New Channels be Selected?	22
5 Conclusion	26
Appendix A Threshold Selection	28
References	31
Vita	36

List of Tables

3.1	The settings and dates where the spectrum data was collected.	6
4.1	The settings and dates where the link data was collected.	14

List of Figures

1.1	Histogram over 7 days' raw energy traces. X axis indicates 802.15.4 channels, Y axis indicates power, and color indicates how often a signal was detected at x GHz with an energy level of y dBm.	3
3.1	Channel occupancy rate. X axis designates channels, Y axis designates experimental settings, and color represents the proportion of readings above the occupancy threshold.	8
3.2	The standard deviation in channel occupancy rate at different timescales.	10
3.3	A comparison of the average channel occupancy rate between channels that overlap with Wi-Fi and channels that do not.	11
4.1	Floor plan of an apartment used in the study.	14
4.2	Box plot of the PRR for four channels in all ten apartments, calculated over 5-minute windows. Central mark in box indicates median; bottom and top of box represent the 25th percentile (q_1) and 75th percentile (q_2); crosses indicate outliers ($x > q_2 + 1.5 \cdot (q_2 - q_1)$ or $x < q_1 - 1.5 \cdot (q_2 - q_1)$); whiskers indicate range excluding outliers. Vertical lines delineate apartments.	15
4.3	Box plot of the PRR of five different links in the same apartment on four channels, calculated over 5-minute windows. Vertical lines delineate links.	15
4.4	The lowest PRR observed on each link's most reliable channel.	18
4.5	CDF of number of consecutive drops.	19
4.6	Retrospective channel-hopping analysis in different apartments.	21
4.7	The Pearson's product correlation coefficient (PMCC) comparing the PRR at the same time on consecutive days or weeks.	23
4.8	Correlation of channel reliability. The X and Y axes indicate channels; the color indicates the probability that channel x 's PRR $< 90\%$ when channel y 's PRR $< 90\%$	24
4.9	Correlation of channel reliability as a function of channel distance.	24
A.1	Relationship between RSS and PRR, as measured experimentally.	29
A.2	Relationship between SINR and PRR, as measured experimentally.	29

Chapter 1

Introduction

In recent years, there has been growing interest in providing fine-grained metering and control of home appliances in residential settings as an integral part of the smart grid. Wireless sensor networks offer a promising platform for home automation applications because they do not require a fixed wired infrastructure. Hence, home area networks (HANs) based on wireless sensor network technology can be used to easily and inexpensively retrofit existing apartments and households without the need to run dedicated cabling for communication and power. Similarly, assisted living applications such as vital sign monitoring and real-time fall detection leverage HANs to provide continuous health monitoring in the patient's home.

The lack of fixed infrastructure also poses key challenges which do not exist under traditional systems of wired sensors and actuators. Home-area applications require a degree of reliability which can easily be met by wired communication but are non-trivial when dealing with unreliable wireless channels. Residential settings present a particularly challenging environment for low-power wireless networks due to the varied and unpredictable nature of the wireless environment.

Figure 1.1 illustrates this challenge with raw spectrum usage traces collected from the 2.4 GHz spectrum in six apartments and an office building (described in more detail in Chapter 3). The office environment provides a relatively clean and predictable wireless environment, with only two major sources of noise: a campus-wide 802.11g network in the middle of the spectrum, and a 802.15.4 sensor network testbed at the upper end. In contrast, the residential settings present a much noisier and more varied environment; for example, apartments 4 and 5 show sporadic interference across the entire 2.4 GHz spectrum (represented by blue shapes spanning nearly the entire X

axis) which could complicate finding a persistently reliable communication channel. These results highlight a fundamental challenge of residential deployments: while the wireless devices in industrial and commercial settings are typically centrally managed, resulting in more predictable noise patterns, residential settings present numerous sources of environmental noise due to a lack of spectrum management. This challenge is compounded by the fact that wireless signals may traverse multiple neighboring residences, subjecting neighbors’ networks to interference beyond their control. For example, in just one apartment in our dataset, a deployed laptop was able to decode beacons from 28 distinct Wi-Fi access points.

In this thesis, a two-part empirical study is presented which aims to characterize the real-world performance of HANs. Specifically, we focus on the performance of devices based on the IEEE 802.15.4 wireless personal area network standard [50]. 802.15.4 radios are designed to operate at a low data rate and be inexpensively manufactured, making them a good fit for home automation applications where energy consumption and manufacturing costs are often at a premium. Industry standards such as ZigBee Smart Energy [38] and WirelessHART [21] have adopted 802.15.4 technology for use in residential and industrial automation applications. Motivated by these and other applications, the IETF has promoted efforts to integrate 802.15.4 networks into larger IPv6 networks, including the Internet [1, 3].

Our study is divided into two major parts. First, we carry out an analysis over spectrum analyzer traces collected in six apartments. This spectrum study of ambient wireless conditions in homes illustrates the challenge of finding a “clean” part of the shared 2.4 GHz spectrum in such settings. Our analysis demonstrates that the wireless environments in these apartments are much more crowded and more variable than an office setting. Moreover, while 802.11 WLANs contribute a significant fraction of the spectrum usage, we also identified signals across the 2.4 GHz band indicating non-negligible noise from non-802.11 devices. Second, we explore the concrete impact of these challenging environments on application performance, through an active probing study of wireless link reliability across all 16 channels in ten apartments. From this active study, we make several more key observations which could greatly impact the reliability of wireless sensor networks deployed in residential environments: (1) Link reliability varies significantly from channel to channel and over time. (2) In a

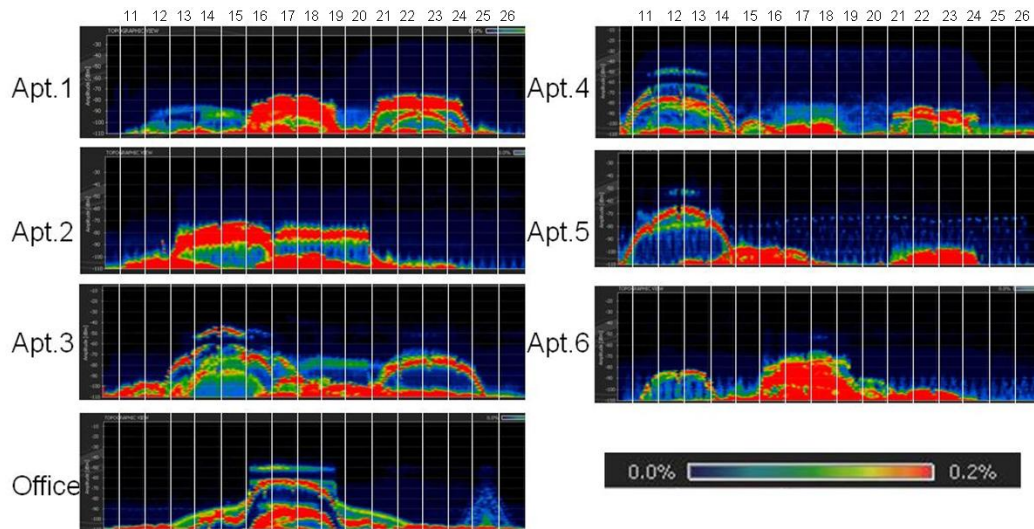


Figure 1.1: Histogram over 7 days’ raw energy traces. X axis indicates 802.15.4 channels, Y axis indicates power, and color indicates how often a signal was detected at x GHz with an energy level of y dBm.

typical apartment environment, there may not be a single channel which is persistently reliable for 24 hours. (3) Even the “best” channels suffer from bursty packet loss which cannot be overcome with retransmissions alone. (4) Exploiting channel diversity by switching channels a few times a day at runtime can effectively maintain long-term reliable communication. (5) Channel conditions are not cyclic. (6) Reliability is strongly correlated between adjacent channels. These findings indicate the importance of channel diversity in achieving reliable HAN deployments and provide guidelines for the design of reliable wireless sensor networks.

The rest of the thesis is organized as follows. Chapter 2 reviews related work. Chapter 3 discusses the findings of our passive spectral study. Chapter 4 then presents our active probing study and its performance results. Finally, we conclude in Chapter 5.

Chapter 2

Related Work

Several recent studies have aimed to characterize the impact of interference on wireless networks through controlled experiments [13, 28, 30, 47, 49]. [42, 48, 59] present theoretical analysis based on simulation study. Gummadi et al. [24] presents an empirical study on the impact of ZigBee and other interferers' impact on 802.11 links, proposing to alleviate interference with rapid channel-hopping in conjunction with 802.11b's existing support for Direct-Sequence Spread Spectrum (DSSS). Srinivasan et al. [52] examines the packet delivery behavior of two 802.15.4-based mote platforms, including the impact of interference from 802.11 and Bluetooth. Liang et al. [34] measures the impact of interference from 802.11 networks on 802.15.4 links, proposing the use of redundant headers and forward error correction to alleviate packet corruption. In contrast to these controlled studies, our own study examines the performance of HANs under normal residential activity. Moreover, our study considers ambient wireless conditions as a whole, rather than analyzing specific sources of interference.

Bahl et al. [15] presents a study of UHF white space networking, while Chen et al. [17] presents a large-scale spectrum measurement study followed by a 2-dimensional frequent pattern mining algorithm for channel prediction. These studies focus on supporting wide-area networks based on white space networking and the GSM band, respectively, while our own study focuses on the reliability of static, indoor wireless sensor networks designed for home environments. Accordingly, our study provides new insights into the reliability of HANs, including the high variability of residential wireless environments, the lack of persistently reliable wireless channels, and the ineffectiveness of retransmissions for maintaining link reliability. Papagiannaki et al. [40] performed an empirical study of home networks based on 802.11 technology.

Our study considers devices based on the 802.15.4 standard, which operates at a much lower transmission power than 802.11 and may exhibit very different behavior. Hence, we draw a different set of observations that underscores the impact of spectrum usage on these low-power 802.15.4 networks.

Ortiz et al. [39] evaluates the multi-channel behavior of 802.15.4 networks in a machine room, a computer room, and an office testbed. Ortiz's study finds path diversity to be an effective strategy to ensure reliability. Our own study finds that residential environments provide significantly different wireless conditions than an office, with the residential settings exhibiting more complex noise patterns and higher variability. This difference may be attributed homes being open environments with no centralized control on spectrum usage; many 2.4 GHz devices are used in homes, and the physical proximity of some residences means that strong interferers (such as 802.11 APs, Bluetooth devices, and cordless phones) may even affect the wireless conditions in other homes. Accordingly, our active study in Chapter 4 finds exploiting channel diversity to be an attractive strategy for ensuring reliability in residential environments. We note that channel and path diversity are orthogonal strategies; the two could be used together in particularly challenging wireless environments.

Hauer et al. [27] discusses a multi-channel measurement of Body Area Networks (BANs) and proposes a noise floor-triggered channel hopping scheme to detect and mitigate the effects of interference. Hauer's study features controlled indoor experiments along with outdoor experiments carried out during normal urban activity. Shah et al. [46] performed a controlled experiment to study the effect of the human body on BANs. Shah's study measures the effects of various activities (sitting, standing, and walking) and node placements (ear, chest, waist, knee, and ankle) on 802.15.4 radio performance. In contrast, our own study looks specifically at the multi-channel properties of indoor residential environments under normal home activities.

Chapter 3

Wireless Spectrum Study

In this Chapter, we present a study of the ambient wireless conditions in real-world residential environments. For this study, we collected 7 days' energy traces in the 2.4 GHz spectrum from six apartments in different neighborhoods. A detailed description of the experimental settings may be found in Table 3.1.

As a baseline for comparison, we also collected energy traces from an office in Bryan Hall at Washington University in St. Louis. We note that this baseline is meant to illustrate how controlled testbed settings may potentially be very different from real home environments; it is not meant to be a comprehensive study of office environments.

Specifically, this study addresses the following questions. (1) Is there a common area of the 2.4 GHz spectrum which is free in all apartments? (2) Does spectrum usage change with time? (3) Do residential settings have similar spectrum usage properties as office and industrial settings? (4) Is 802.11 the dominant interferer in residential environments?

Name	Begin Date	End Date
Apt. 1	2:00pm, Apr. 4, 2010	3:30pm, Apr. 19, 2010
Apt. 2	6:50pm, June 30, 2010	6:50pm, July 7, 2010
Apt. 3	9:05pm, May 12, 2010	11:29pm, May 20, 2010
Apt. 4	11:40am, June 6, 2010	12:40pm, June 13, 2010
Apt. 5	12:25pm, Apr. 20, 2010	10:50am, Apr. 28, 2010
Apt. 6	7:00pm, July 7, 2010	9:00pm, July 14, 2010
Office	1:15pm, July 16, 2010	1:20pm, July 23, 2010

Table 3.1: The settings and dates where the spectrum data was collected.

3.1 Experimental Methodology

We are primarily interested in the spectrum usage between 2.400 GHz and 2.495 GHz, which are the parts of the spectrum used by the 802.15.4 standard for wireless sensor networks. To analyze this part of the spectrum, we collected energy traces using a laptop equipped with a Wi-Spy 2.4x spectrum analyzer [12]. The Wi-Spy sweeps across the 2.4 GHz spectrum approximately once every 40 ms, returning a signal strength reading (in dBm) for each of 254 discrete frequencies. We continuously collected energy traces for 7 days in each apartment and Bryan Hall during the residents' normal daily activities. The resulting traces contained 15,120,000 readings for each of the 254 frequencies, resulting in a data set of approximately 2.5 GB per location. Figure 1.1 presents a histogram of the raw spectrum usage data in all seven datasets.

For the purposes of analysis, we apply a thresholding process like that employed in [17] to convert signal strength readings into binary values, with 0 denoting a channel being idle and 1 denoting a channel being busy. We found experimentally that a receive signal strength of -80 dBm is needed to create a high-quality link between a pair of Chipcon CC2420 radios; however, a noise level of -85 dBm or higher would be enough to induce packet drops on such a link. (We discuss this experiment in more detail in Appendix A.) Hence, throughout our analysis, we use -85 dBm as our threshold value to denote a busy channel. Using a constant threshold allows for a fair comparison across different apartments. While the specific numerical results of our analysis are dependent on the threshold, the trend and observations we make from these results should generally apply to other threshold values.

We also aggregate the data from the Wi-Spy's 254 channels into the 16 channels used by the 802.15.4 standard; i.e., an 802.15.4 channel is deemed busy if any of its corresponding Wi-Spy channels are busy.

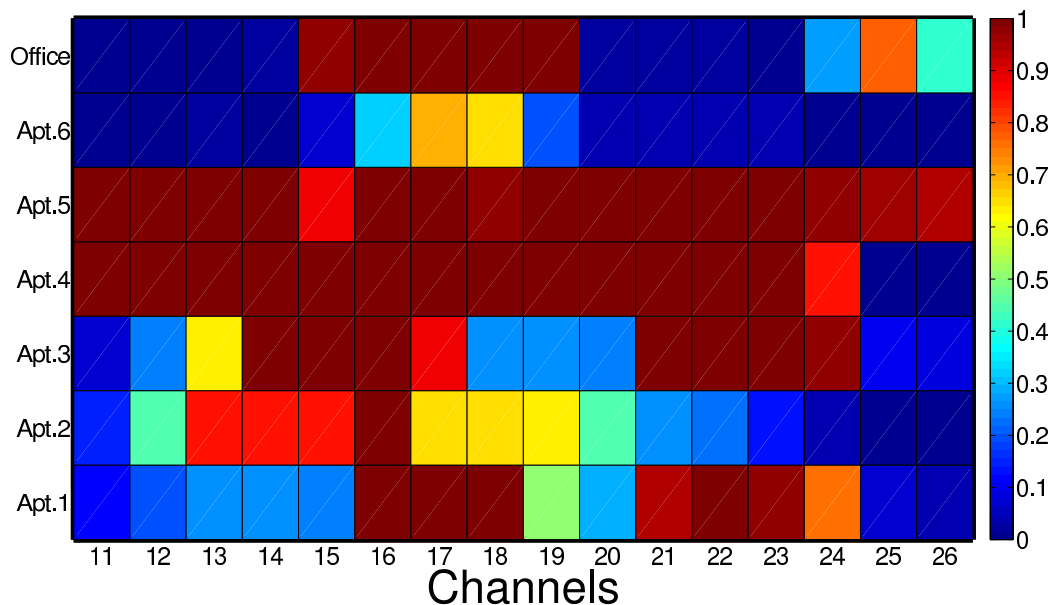


Figure 3.1: Channel occupancy rate. X axis designates channels, Y axis designates experimental settings, and color represents the proportion of readings above the occupancy threshold.

3.2 Is There a Common Idle Channel in Different Homes?

We first considered whether any 802.15.4 channel can be considered “clean” in all the tested residences. To determine this, we calculate the channel occupancy rate — i.e., the proportion of samples that exceeded the -85 dBm threshold — over all channels in the six apartments and office building. High occupancy rates correspond to a large proportion of samples where interference could have caused packet loss on an otherwise high-quality link.

Figure 3.1 plots the occupancy rate of each channel in each location. If we compare Figures 1.1 and 3.1, we can note various phenomenon that prevent finding a common idle channel. For example, apartment 5 has a channel occupancy rate above 95% for 15 of its 16 channels. This uniformly high occupancy rate is likely caused by a relatively high-power spread-spectrum signal across the whole 2.4 GHz spectrum, which appears in Figure 1.1 as a series of thin blue arches. Devices with such wireless footprints include Bluetooth transmitters, baby monitors, wireless speaker systems, and game

controllers [5]. (Unfortunately, by the very nature of residential environments lacking central management of wireless devices, there is no way to be certain about the sources of some of these phenomena.)

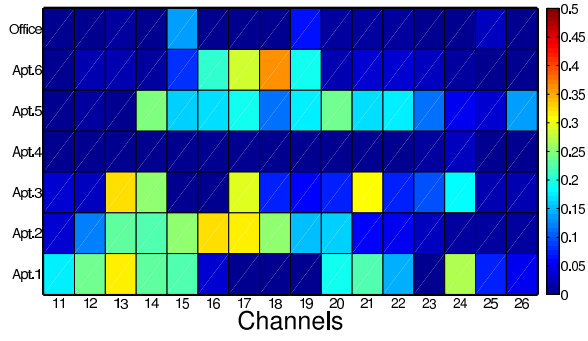
The only channel in apartment 5 with an occupancy rate below 95% is channel 15, which has an occupancy rate of 100.0% in apartments 3 and 4; thus there is no common good channel in these apartments. In the case of apartment 3, channel 15 is unusable due to it intersecting with the middle of multiple 802.11 APs, represented as superimposed arcs on the left side of apartment 3’s energy trace. For apartment 4, we see that only channels 25 and 26 have low occupancy rates; this phenomenon is likely caused by the tall blue shape across most of apartment 4’s energy trace, corresponding to some sporadic but very high-power and high-bandwidth interferer.

***Observation S1:** There may not exist a common idle channel across different homes, due to significant diversity in their spectrum usage patterns.*

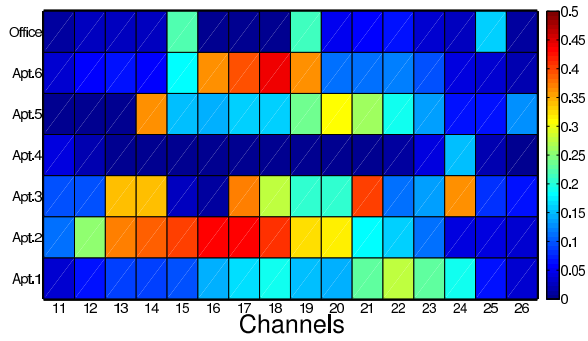
3.3 Does Spectrum Usage Change with Time?

We next explored whether the spectrum was stable in these residential settings. To do so, we calculated the standard deviation in occupancy (σ) for each apartment and each channel. Figure 3.2 plots the standard deviation from day-to-day, from hour-to-hour, and for every 5 minutes. We see that channel conditions in most apartments can be quite variable, regardless of the timescale used. Excluding apartment 4, σ ranges from 24.0%–36.2% for the worst channel at a daily timescale, from 27.4%–43.9% at an hourly timescale, and 36.4%–50.0% at a 5-minute timescale. Apartment 4 is stable across the spectrum on a day-to-day basis, with $\sigma \leq 2.5\%$ for all channels. However, even for this apartment, some variability emerges at shorter timescales, with channel 24 featuring $\sigma = 14.9\%$ on an hourly timescale and $\sigma = 36.0\%$ at a 5-minute timescale.

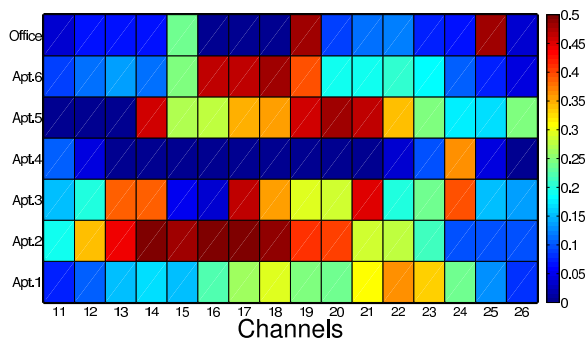
We also note that the office had much lower variability than all but apartment 4. For example, at a daily timescale, 10 of the 16 channels had $\sigma < 1.0\%$, and the most highly-variable channel had σ of only 13.7%. Indeed, even at a 5-minute timescale, only three channels reveal significant variability; these three channels are at the edge



(a) Daily standard deviation



(b) Hourly standard deviation



(c) 5-minute standard deviation

Figure 3.2: The standard deviation in channel occupancy rate at different timescales.

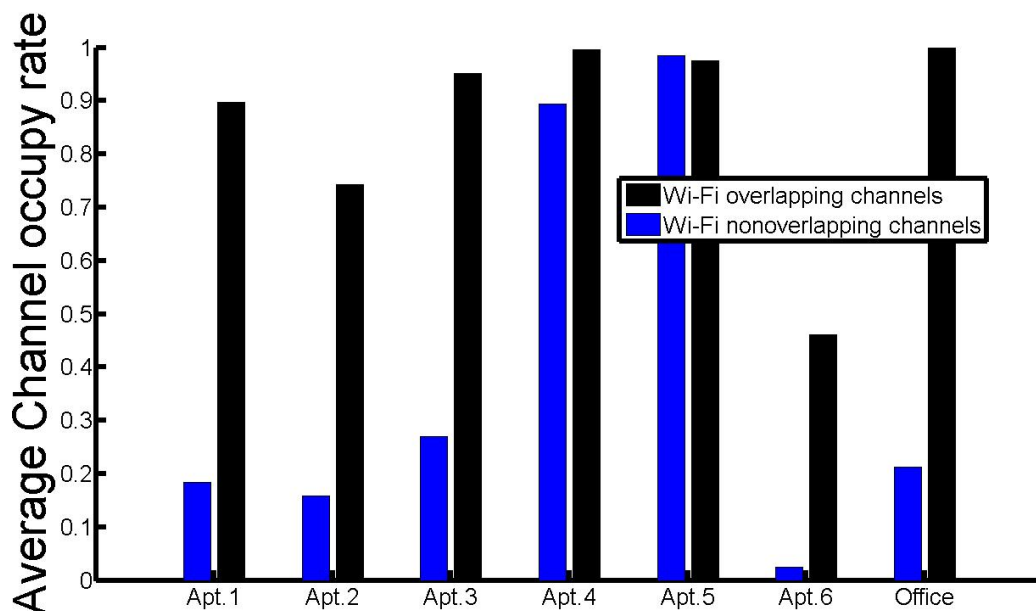


Figure 3.3: A comparison of the average channel occupancy rate between channels that overlap with Wi-Fi and channels that do not.

of the campus 802.11g network (15), at the center of the same network (19), and at the center of the building’s 802.15.4 testbed (25).

Observation S2: Spectrum occupancy in homes can exhibit significant variability over time, whether looking at timescales of days, hours, or minutes.

3.4 Is Wi - Fi the Dominant Source of Spectrum Usage?

Because of Wi-Fi’s ubiquity and relatively high transmission power, it is often treated as a dominant interferer. Thus, our final analysis of our passive spectrum data is to identify whether there are other significant sources of interference.

A simple inspection of Figures 1.1 and 3.1 suggests other important interferers besides Wi-Fi. Wi-Fi APs have a distinctive radiation pattern that manifests in Figure 1.1 as arcs the width of several 802.15.4 channels. For example, the energy traces for apartment 3 show two distinct arcs that are likely caused primarily by 802.11 APs

configured to two different channels. Referring to Figure 3.1, we see that these areas of the spectrum are indeed highly occupied. However, looking at the energy trace for apartment 5, we see evidence of Wi-Fi APs on only part of the spectrum; nevertheless, the channel occupancy rate is above 95% for nearly the entire spectrum. This phenomenon can be explained by the series of blue arcs across the 2.4 GHz spectrum, which indicate sporadic but high-powered spread-spectrum transmissions. (Again, by the nature of the environment, we cannot be certain about the source of this noise pattern.)

To quantify the relative impact of Wi-Fi, we leverage a feature of the Wi-Spy which logs the SSID and 802.11 channel of all visible 802.11 access points. Based on this data, we are able to assign each 802.15.4 channel in each apartment into two groups: those that overlap with 802.11 APs visible from the corresponding apartment, and those that do not. We then calculated the average channel occupancy rate for each of the two groups in each apartment, as shown in Figure 3.3.

In most of the apartments, there is a clear distinction between the overlapping and non-overlapping channels. For example, apartment 1 has an average occupancy rate of 89.7% for the overlapping channels compared to 18.3% for the non-overlapping ones. But strikingly, we find that the non-overlapping channels are not *always* significantly more idle than those which overlap with Wi-Fi APs. In apartments 4 and 5, the channel occupancy rates of the non-overlapping channels are similar to the overlapping ones; indeed, in apartment 5, the non-overlapping channels are slightly more occupied on average than the overlapping ones.

Observation S3: *While Wi-Fi is an important source of interference in residential environments, other interferers can also be non-negligible contributors to spectrum occupancy.*

Chapter 4

Multi-Channel Link Study

In this Chapter, we present a multi-channel link study in homes. The spectrum study presented in Chapter 3 focuses on characterizing the ambient wireless environment in homes. While link quality can be significantly influenced by interference from existing wireless signals, other factors such as signal attenuation and multi-path fading due to human activities can also impact the reliability of low-power wireless links. Our link study *directly* evaluates the multi-channel behavior of HANs by actively sending packets between motes equipped with 802.15.4 radios.

Specifically, this study addresses the following questions. (1) Can a HAN find a single persistently reliable channel for wireless communication? (2) If a good channel cannot be found, are packet retransmissions sufficient to deal with packet loss? (3) If no single channel can be used for reliable operation, can the network exploit channel diversity to achieve reliability? (4) Do channel conditions exhibit cyclic behavior over time? (5) Is reliability strongly correlated among different channels?

4.1 Experimental Methodology

For this active study, we carried out a series of experiments in ten real-world apartments in different neighborhoods, as listed in Table 4.1. (Due to the participating residents moving, only four of the apartments in this study are the same as in Chapter 3.) Figure 4.1 shows an example floor plan of one of the apartments used in the study; a similar topology was deployed in the other apartments. Each experiment was carried out continuously for 24 hours with the residents' normal daily activities.

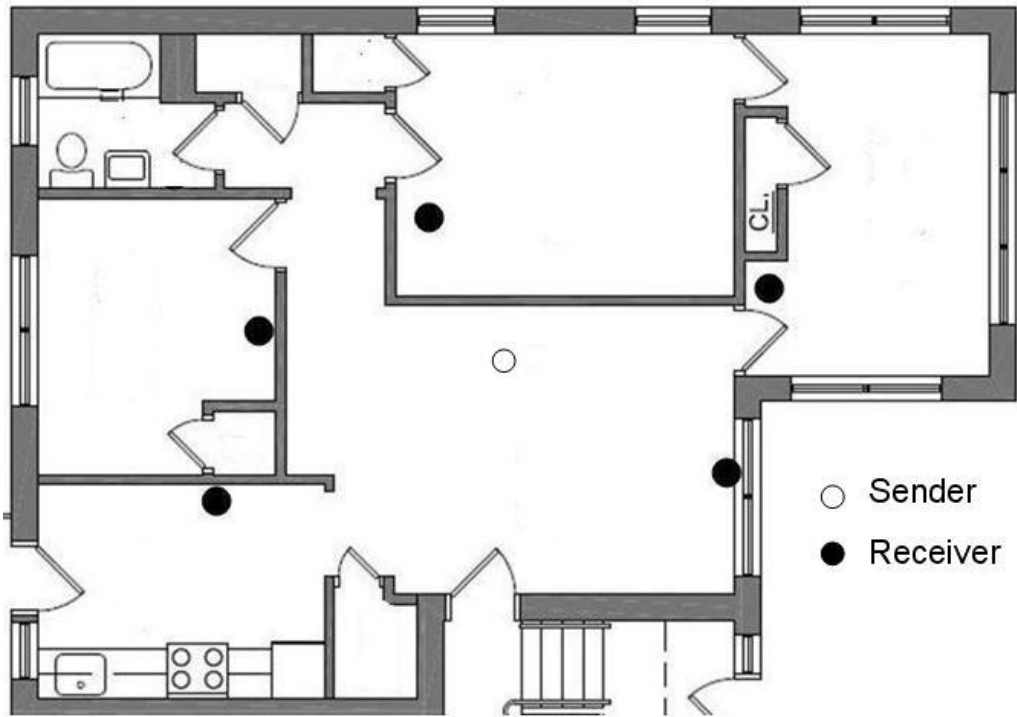


Figure 4.1: Floor plan of an apartment used in the study.

	Begin Date	End Date
Apt. 1	Sept. 30, 2009	Oct. 1, 2009
Apt. 2	Sept. 30, 2009	Oct. 1, 2009
Apt. 3	Oct. 3, 2009	Oct. 4, 2009
Apt. 4	Oct. 3, 2009	Oct. 4, 2009
Apt. 5	Sept. 30, 2009	Oct. 1, 2009
Apt. 6	Sept. 12, 2009	Sept. 13, 2009
Apt. 7	Oct. 3, 2009	Oct. 4, 2009
Apt. 8	Sept. 18, 2009	Sept. 19, 2009
Apt. 9	Oct. 6, 2009	Oct. 7, 2009
Apt. 10	Oct. 6, 2009	Oct. 7, 2009

Table 4.1: The settings and dates where the link data was collected.

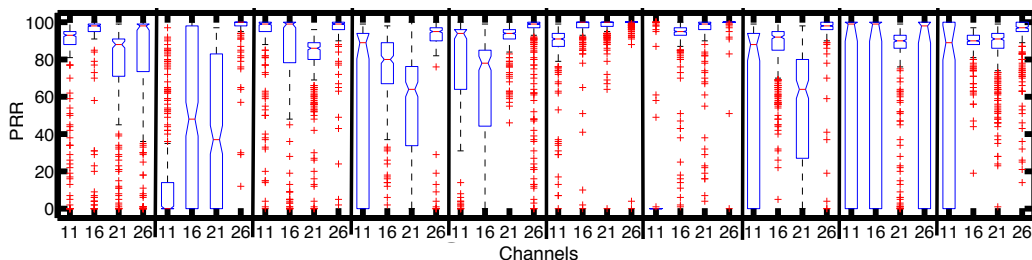


Figure 4.2: Box plot of the PRR for four channels in all ten apartments, calculated over 5-minute windows. Central mark in box indicates median; bottom and top of box represent the 25th percentile (q_1) and 75th percentile (q_2); crosses indicate outliers ($x > q_2 + 1.5 \cdot (q_2 - q_1)$ or $x < q_1 - 1.5 \cdot (q_2 - q_1)$); whiskers indicate range excluding outliers. Vertical lines delineate apartments.

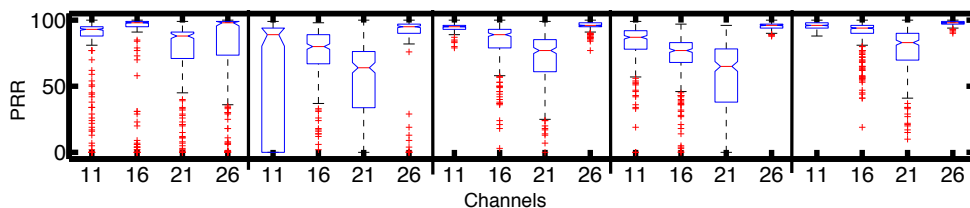


Figure 4.3: Box plot of the PRR of five different links in the same apartment on four channels, calculated over 5-minute windows. Vertical lines delineate links.

Our experiments were carried out using networks of Tmote Sky and TelosB [41] motes. Each mote is equipped with an IEEE 802.15.4 compliant Chipcon CC2420 radio [2]. IEEE 802.15.4 radios like the CC2420 can be programmed to operate on 16 channels (numbered 11 to 26) in 5 MHz steps. We leverage the CC2420’s Received Signal Strength (RSS) indicator in our experiments to measure the signal power of environmental noise. Our experiments are written on top of the TinyOS 2.1 operating system [6] using the CC2420 driver’s default CSMA/CA MAC layer.

To measure the Packet Reception Rate (PRR) of all channels at a fine granularity, we deployed a single transmitter node in each apartment which broadcast packets over each of the 16 channels. Specifically, the transmitter sent a batch of 100 consecutive packets to the broadcast address using a single wireless channel, then proceeded to the next channel in a round robin fashion. The process of sending 16 batches of 100 packets repeated every 5 minutes. The recipient nodes record the PRR over each batch of packets into their onboard flash memory. The use of a single sender and multiple recipients allowed us to test multiple links simultaneously while avoiding interference

between senders. (Inter-link interference is not a major concern in many HANs due to the low data rates that are typically employed; for example, 1 temperature reading every 5 minutes is sufficient for an HVAC system to control ambient temperature.)

It is worth noting that HAN applications such as smart energy require persistent, long-term reliability. Transient link failures are non-negligible — these failures represent periods where parts of a household may experience sporadic service or no service at all (e.g., changing the thermostat may have no effect until a wireless link is restored minutes or hours later). Hence, our study looks not just at the average PRR of each link but at its entire range of performance, including those outliers that indicate temporary failures.

In [52], links with a PRR below 10% were found to be poor-quality, and links with a PRR between 10% and 90% to be bursty. Accordingly, we use a PRR of 90% throughout this chapter as a threshold to designate links as “good” or “reliable”. Due to the numerous outliers found throughout Section 4.2, we expect our analysis to generally hold for any reasonable PRR threshold.

4.2 Is There a Persistently Good Channel?

Previous empirical studies [51, 52, 60] have looked at the issue of link variability in office testbeds. A potential cause of link variability in these environments is that some links exist in a “gray region” at the threshold of connectivity where small temporal changes in link quality can cause bursts of packet losses [60]. In contrast to previous studies, our study does not focus on links within gray regions, but rather on the impact of home environments on different wireless channels. As we discuss later, the links in our study have at least one channel with a high median PRR, and show different degrees of variability in different channels. This suggests that the links are likely outside the gray region, which would have caused lower median PRR or higher variation across all channels.

We first analyzed our data from the perspective of finding a single, persistently good channel across all of the tested apartments. Accordingly, for this analysis, we grouped the data from all links in all apartments together and then subdivided it by channel.

Figure 4.2 presents a box plot of the PRR in 4 channels in all the apartments, where the PRR has been calculated over 5-minute windows. (The remaining 12 channels are omitted for reasons of clarity.) From this figure, we see significant variations in PRR on the same channel when moving from apartment to apartment. For example, channel 11 achieves a median PRR $> 90\%$ in apartments 1, 3, and 9, albeit with many outliers; however, the same channel has a near-zero median PRR in apartment 2. Only channel 26 has a median PRR above the 90% threshold in all apartments.

We also see significant variations in PRR from channel to channel, even in the same apartment. Strikingly, these variations even affect channel 26, which is often considered a highly reliable channel since it is nominally outside the 802.11 spectrum in North America. Although channel 26 achieves uniformly high *median* PRR in all apartments, there are numerous points during the experiment where the PRR falls much lower. For example, apartment 9 has a 25th percentile PRR of 0.0% , indicating a substantial portion of the experiment where the channel experienced total link failure.

Further analysis showed that there is not likely to be a single good channel across multiple links in the *same* apartment. We regrouped the PRR data, this time looking at the performance of each link/channel pair individually. Figure 4.3 presents a box-plot of the PRR for all five links within one apartment; again, for reasons of clarity, we present the data from only 4 of the 16 channels. We observe that the median PRR on a given channel varies greatly across links, particularly for outlier points. Again, this variation even affects channel 26: all five links have at least one outlier below the 90% threshold, and four links have numerous outliers below the threshold. Link 1 shows particularly high variance on channel 26, with a 25-percentile PRR of only 73.5% in spite of a 98.0% median PRR. We also note that *all four* channels had numerous outliers below a PRR of 10% ; that is, any single channel selection would have led to at least one link experiencing near-total disconnection at some point during the day.

Interestingly, these large channel-to-channel variations suggest that “gray” links are not the dominant cause of variability in these apartments. If the variations were caused primarily by links being transitional, we would expect that the variability would be roughly the same across different channels. However, in our residential test

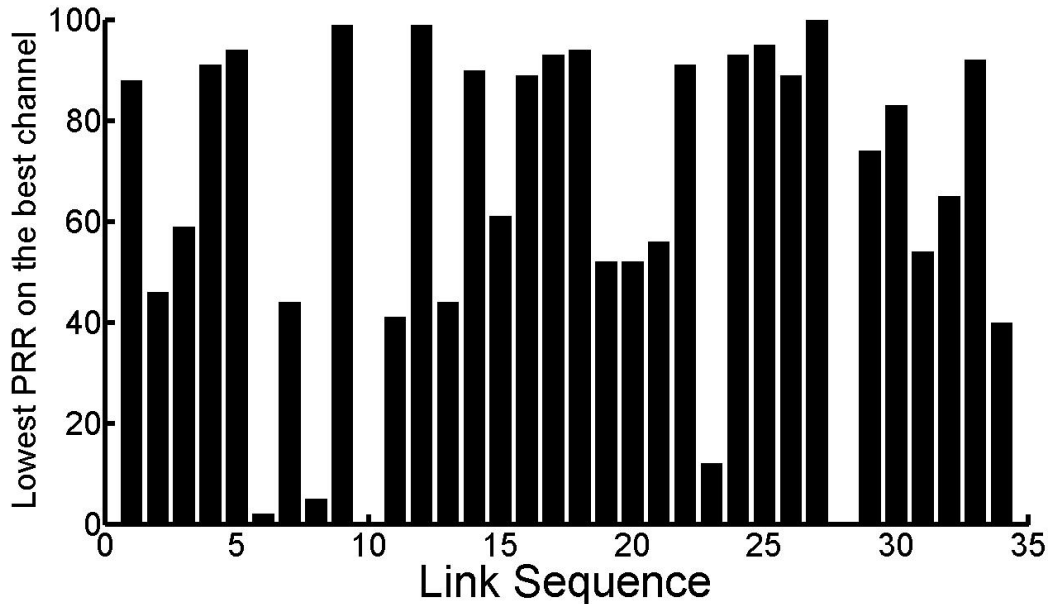


Figure 4.4: The lowest PRR observed on each link’s most reliable channel.

settings, we find that many links are significantly more variable on one channel than another.

Observation L1: Link reliability varies greatly from channel to channel.

Looking at the entire dataset across all apartments, we found that few links were able to achieve a consistently high PRR, even on their most reliable channels. Figure 4.4 plots the lowest PRR observed on each link’s most reliable channel: i.e., for the channel which achieves the highest average PRR over 24 hours, we plot the worst PRR out of all the 100-packet batches. Notably, only 12 of the 34 links in our dataset are able to persistently reach the 90% PRR threshold on even their best channel. Indeed, even lowering the threshold to 70%, more than half the links in our dataset would still have no persistently good channel. Again, large channel-to-channel fluctuations in link reliability suggest these temporal fluctuations are not caused by “gray” links.

Observation L2: Link reliability varies greatly over time, even within the same channel. Hence, even when selecting channels on a per-link basis, there is not always a single persistently reliable channel.

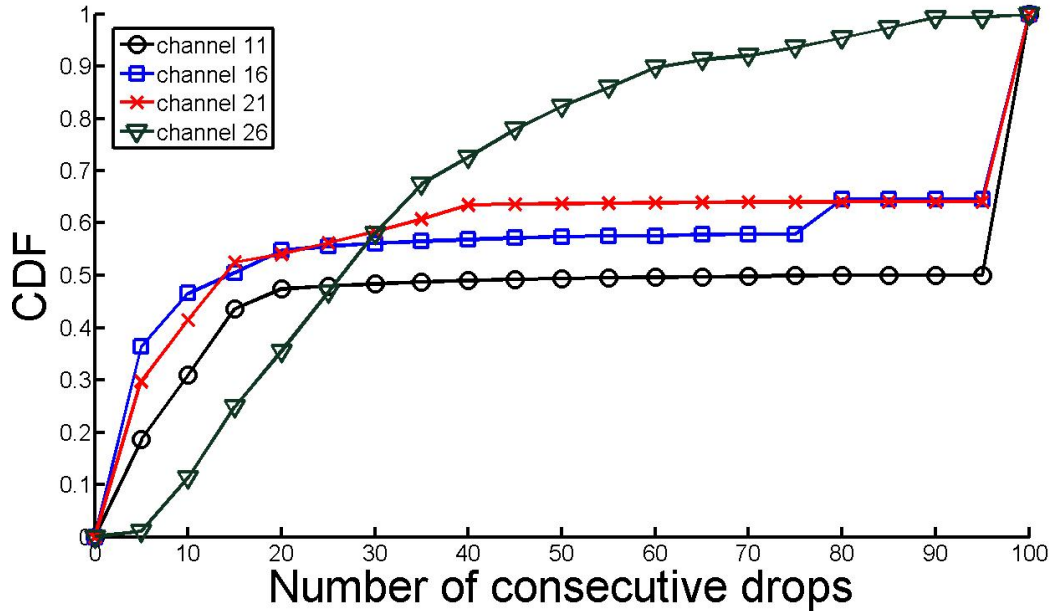


Figure 4.5: CDF of number of consecutive drops.

4.3 Is Retransmission Sufficient?

Because retransmissions are effective in alleviating transient link failures, we next analyze whether it would be effective in alleviating the link failures observed in our experimental traces. However, we found that retransmissions alone are insufficient in residential environments, due to the bursty nature of the packet losses.

Figure 4.5 illustrates this problem with the cumulative probability density (CDF) of consecutive packet drops for all links on four channels. Specifically, we measured consecutive packet losses *within each batch* of 100 packets; we did not include inter-batch losses due to the 5-minute gap between batches. Even on the best channel (channel 26), up to 85 consecutive packet drops were observed, and 10% of link failures lasted for more than 60 consecutive packets. On the remaining three channels, bursts of more than 95 consecutive packet drops were observed.

Observation L3: *Retransmissions alone are insufficient for HANs due to the burstiness of packet losses.*

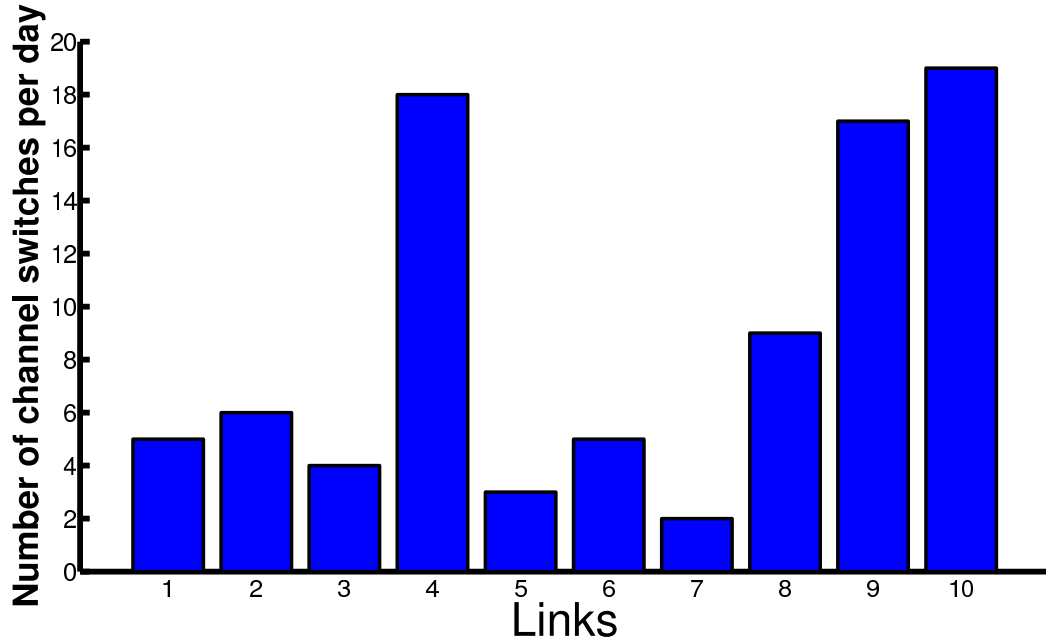
4.4 Is Channel Diversity Effective?

Our analysis above indicates that using a single channel is often not feasible when long-term reliability must be maintained. Thus, a natural question to ask is whether it is feasible to exploit channel diversity to achieve reliability in situations where single channel assignments are not practical.

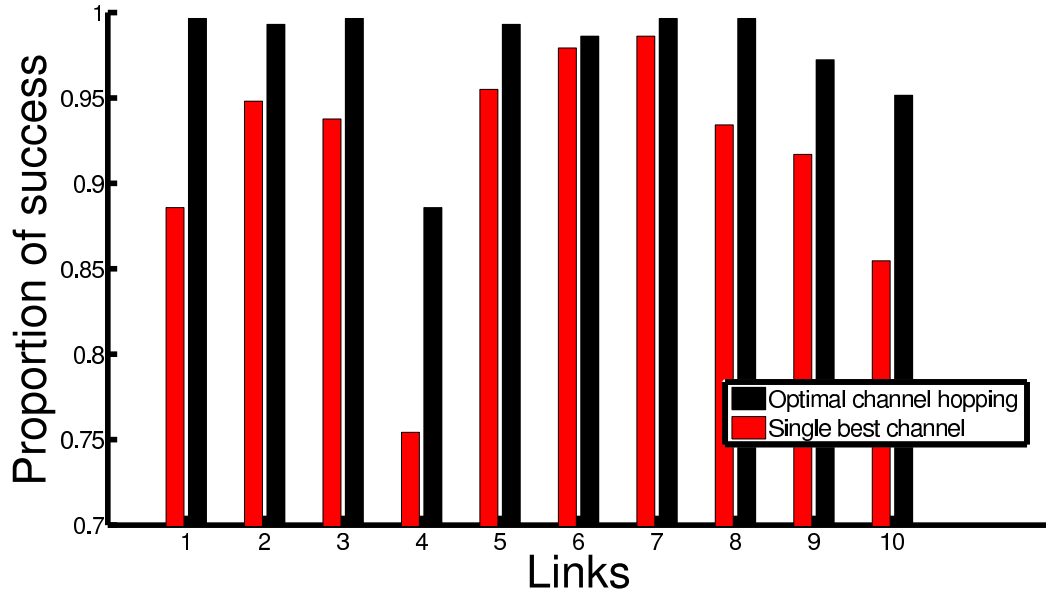
To understand the potential for channel hopping, we retrospectively processed our dataset to find the minimum number of channel hops needed to maintain a 90% PRR threshold. Figure 4.6(a) plots the number of channel hops required for 10 links in the dataset, one randomly selected from each apartment. We find that relatively few channel hops are needed to maintain link reliability; in no case is more than 20 hops required per day.

We note that there are periods where none of the 16 channels meet the PRR threshold, and hence no channel hopping occurs during these times. Nevertheless, channel-hopping can significantly reduce the number of link failures compared to picking the single “best” channel (i.e., that with the highest overall PRR). Figure 4.6(b) compares the proportion of windows which meet the 90% threshold under two retrospective strategies: the retrospective channel-hopping strategy, and a strategy that fixes each link to its single “best” channel. In some cases, the improvements achieved by channel hopping are modest. For example, links 6 and 7 only achieve a 0.7% and 1.0% higher success rate under channel hopping, largely because their success rates were already high without channel hopping. However, in most cases, we find at least moderate improvements in link success. For example, 6 out of the 10 links experience at least 5% fewer failures with channel hopping than with their single best channel; and links 1 (11.0%) and 4 (13.1%) have substantially higher success rates with channel hopping.

Channel hopping has been proposed in industry standards as a means for improving wireless link reliability, including established standards like Bluetooth’s AFH [10] and newer standards such as WirelessHART’s TSMP [11] and the forthcoming IEEE 802.15.4e [4]. The results of our analysis confirm that this feature is indeed beneficial for maintaining link reliability in challenging residential environments.



(a) Minimum number of channel hops required; one link randomly selected per apartment.



(b) The proportion of windows where the PRR threshold was met.

Figure 4.6: Retrospective channel-hopping analysis in different apartments.

Observation L4: *Channel hopping is moderately to greatly effective in alleviating packet loss due to channel degradation, depending on the link. Only a small number of channel hops per day are needed to effectively maintain reliable communication.*

4.5 Can Hopping be Scheduled Staticly?

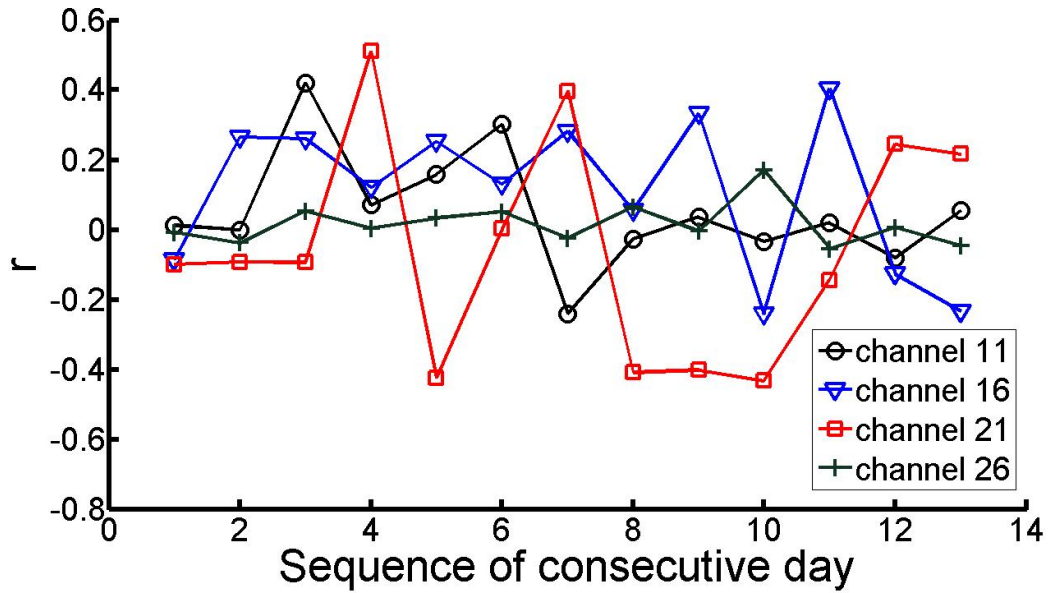
Because channel quality varies over time, we next explored whether it exhibits cyclic properties. If so, then channel-hopping could be implemented in a lightweight fashion by generating a static channel schedule for each environment. To perform this comparison, we carried out an extended experiment in one apartment over a period of 14 days. We then calculated the Pearson product-moment correlation coefficient (PMCC) [54], a common measure of dependence between two quantities, as r . Intuitively, r values near -1 or 1 indicate strong correlation, while values near 0 indicate independence.

Figure 4.7(a) plots r for PRRs calculated at the same times on subsequent days (e.g., 4 PM on Monday vs. 4 PM on Tuesday). Figure 4.7(b) compares the PRR during the same time in consecutive weeks (e.g., 4 PM on Monday vs. 4 PM on the next Monday). r is almost always smaller than 0.4 , regardless of the channel used; this indicates that there is no obvious correlation between consecutive days or consecutive weeks. Therefore, channel-hopping decisions must be made *dynamically* based on channel conditions observed at runtime.

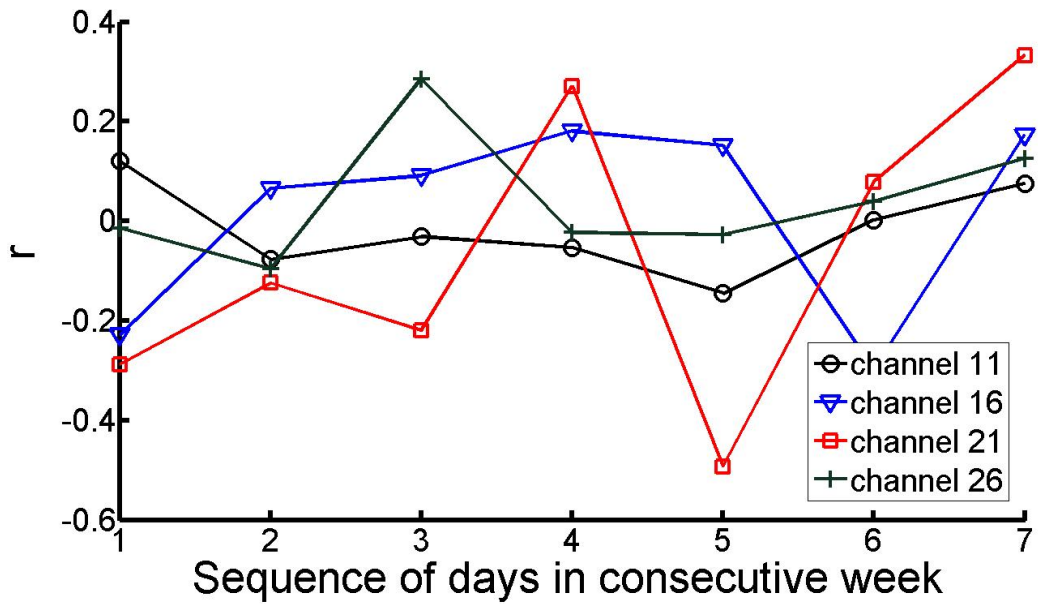
Observation L5: *Channel conditions are not cyclic, so channel-hopping decisions must be made dynamically.*

4.6 How Should New Channels be Selected?

Since channel-hopping must be performed dynamically, it is important to pick a good strategy for selecting new channels when the current channel has degraded beyond use. For the purposes of this analysis, we studied the effect of *channel distance* (the absolute difference between channel indices) on the *conditional probability* of channel



(a) PMCC of PRRs during the same time on consecutive days.



(b) PMCC of PRRs during the same time in consecutive weeks (e.g., $x = 1$ means consecutive Mondays).

Figure 4.7: The Pearson's product correlation coefficient (PMCC) comparing the PRR at the same time on consecutive days or weeks.

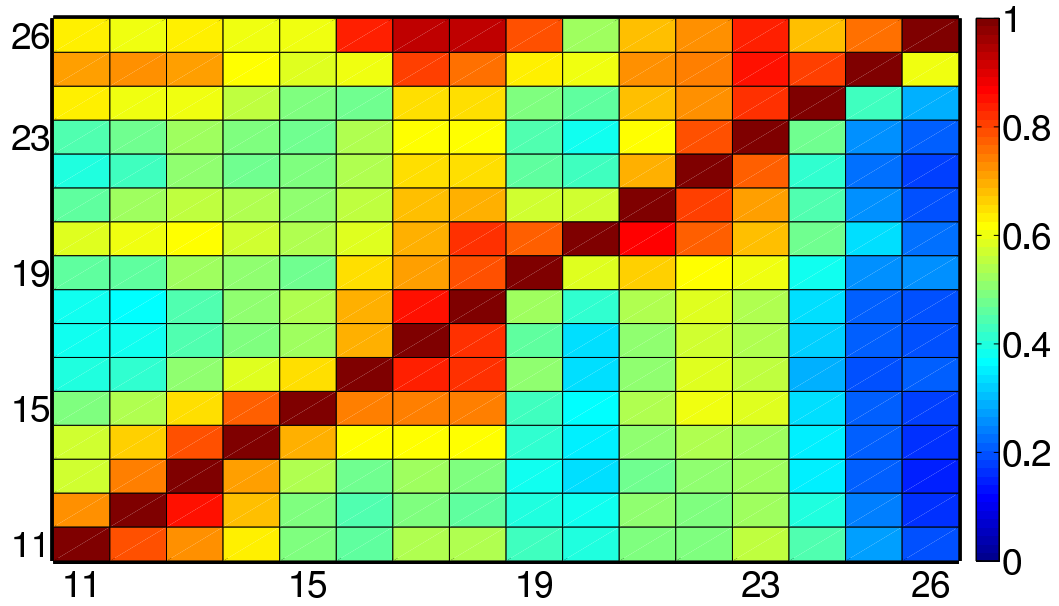


Figure 4.8: Correlation of channel reliability. The X and Y axes indicate channels; the color indicates the probability that channel x 's PRR $< 90\%$ when channel y 's PRR $< 90\%$.

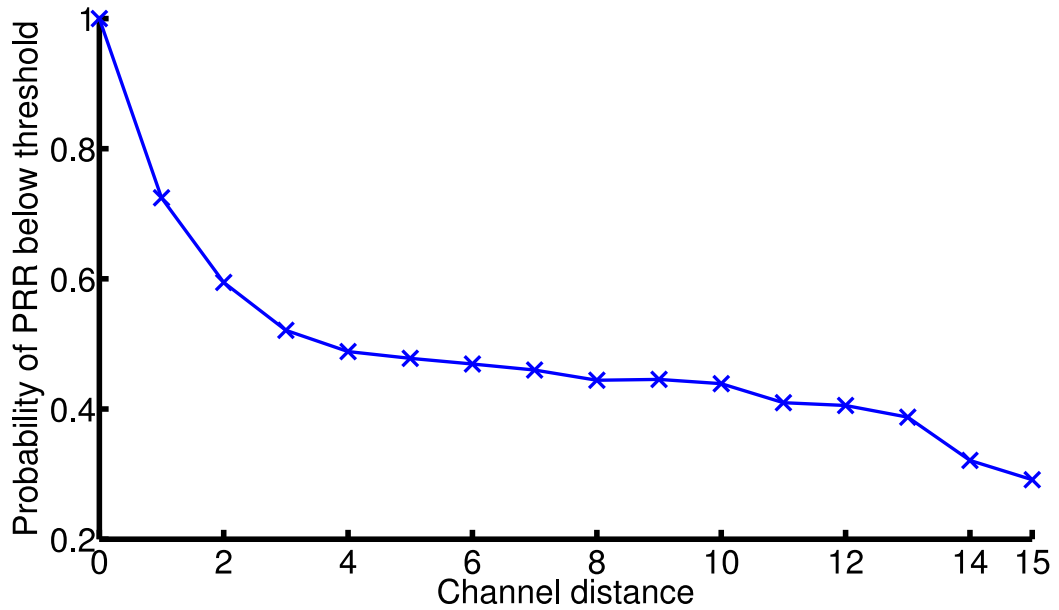


Figure 4.9: Correlation of channel reliability as a function of channel distance.

failure (the probability that channel x is below the PRR threshold when channel y is also below the threshold).

We observe that not all channels are equally good candidates for channel hopping: from Figure 4.8, we can see that performance is strongly correlated across adjacent channels. For instance, when channel 20 has poor PRR ($< 90\%$), there is a probability greater than 76.8% that channels 18, 19, 21, and 22 also suffer from poor PRR. In Figure 4.9, we plot the conditional probability of link failure as a function of channel distance. We observe that this probability can be as high as 70% between neighboring channels and 60% between every other channel, but drops off as channel distance increases.

Observation L6: *Reliability is strongly correlated across adjacent channels; channel-hopping should move as far away as possible from a failing channel.*

Chapter 5

Conclusion

HANs based on wireless sensor network technology represent a promising communication platform for emerging home automation applications. However, the complex and highly variable wireless environments in typical residential environments pose significant reliability challenges. This thesis presents an empirical study on the performance of HANs in real-life apartments, looking both at passive spectrum analysis traces and an active probing link study. The observations made in our study highlight the significant challenges facing home-area sensor networks in residential settings. Nevertheless, our observations also suggest that these reliability challenges may be tamed through the judicious use of channel diversity. Specifically, we may distill our findings into set of key design guidelines for developing reliable HANs:

1. Channel selection can have a profound impact on HAN reliability. Channel selection cannot be simply relegated a static channel assignment, whether made at the factory or at deployment time. (S1, L1, L2)
2. Channel selection should be made on a per-link, rather than per-network, basis. (L1)
3. Retransmissions alone cannot always compensate for a poor-quality channel. (L3)
4. Although Wi-Fi is a major source of channel usage, other wireless technologies may also contribute significantly to channel usage. Solutions which target a single interfering technology are not always sufficient in residential environments. (S3)

5. Reliable communication can be maintained through infrequent channel hopping. (L4)
6. Channel hopping cannot be performed based on a static, cyclic schedule. (L5)
Channel-hopping decisions should be made dynamically based on conditions observed at runtime, and should avoid channels adjacent to a deteriorating channel. (S2, L2, L6)

Appendix A

Threshold Selection

According to wireless communication theory, a packet can be successfully decoded if the signal-to-interference-plus-noise-ratio is above a certain threshold. To determine the threshold used to decide if a channel is busy or idle in our spectrum study, we study the impact of interference on packet reception empirically as follows. Let N_{dBm} be the total signal strength of the noise and interference measured at the receiver. Let RSS_{dBm} be the total signal strength associated with an incoming packet by the CC2420 radio, including the packet, noise, and interference. We can calculate the signal-to-interference-plus-noise-ratio ($SINR_{dB}$) as:

$$SINR_{dB} = 10 \log_{10} \frac{10^{RSS_{dBm}/10} - 10^{N_{dBm}/10}}{10^{N_{dBm}/10}} \quad (\text{A.1})$$

From Eq. (A.1), we get

$$10^{SINR_{dB}/10} = \frac{10^{RSS_{dBm}/10} - 10^{N_{dBm}/10}}{10^{N_{dBm}/10}} \quad (\text{A.2})$$

$$10^{N_{dBm}/10} = \frac{10^{RSS_{dBm}/10}}{10^{SINR_{dB}/10} + 1} \quad (\text{A.3})$$

Figure A.1 plots the correlation between receive signal strength and PRR as obtained experimentally between a pair of TelosB motes at varying distances and transmission powers. We see that $RSS_{dBm} = -80$ dBm places the link outside of the transitional “gray” region; similar results were observed in [52, 53]. Following the methodology in [45], we estimated the relationship between SINR and PRR experimentally using a pair of TelosB motes and a third interfering mote operating at varying distances

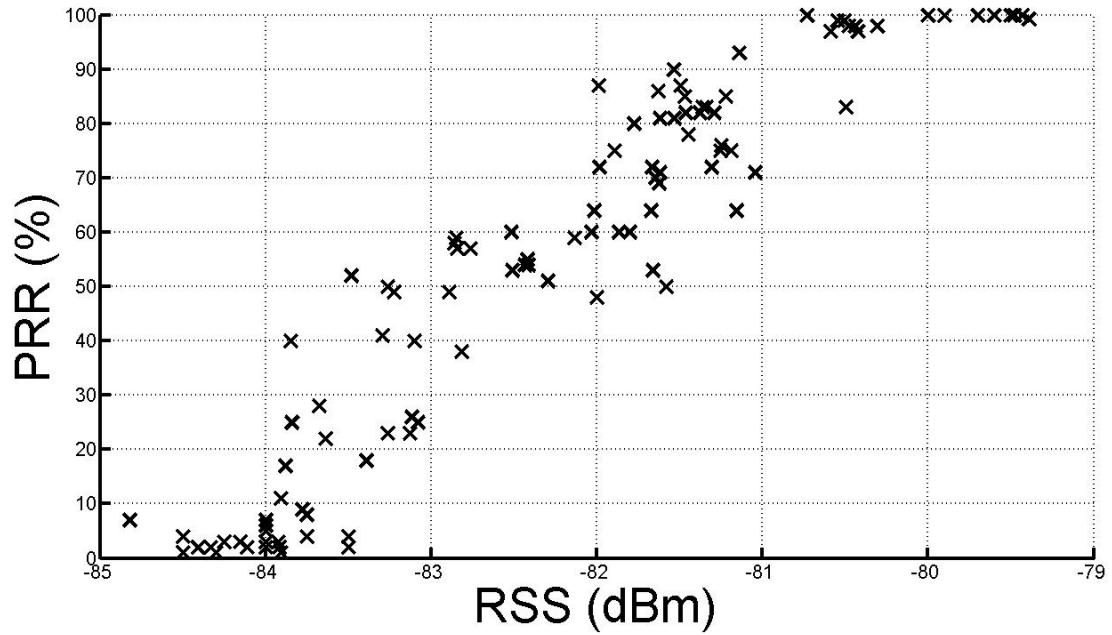


Figure A.1: Relationship between RSS and PRR, as measured experimentally.

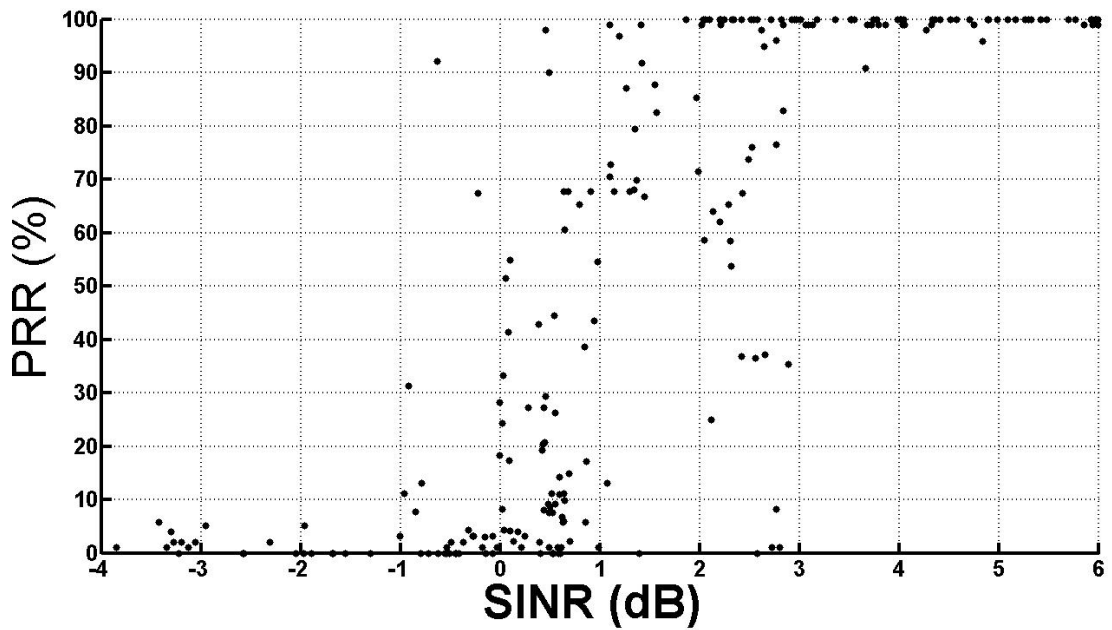


Figure A.2: Relationship between SINR and PRR, as measured experimentally.

and transmission powers. We plot this relationship in Figure A.2. A threshold of $SINR_{dB} = 4$ dB places the link outside of the transitional region; this result matches experiments performed in [45]. Therefore, we get

$$10^{N_{dBm}/10} = \frac{10^{-80dBm/10}}{10^{4dB/10} + 1} \quad (\text{A.4})$$

$$N_{dBm} = -85 \text{ dBm} \quad (\text{A.5})$$

Thus we choose -85 dBm as the threshold to distinguish a channel as busy or idle.

References

- [1] IPv6 over low power WPAN.
- [2] 2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver, Texas Instruments.
- [3] Routing over low power and lossy networks (roll). <http://datatracker.ietf.org/wg/roll/charter/>.
- [4] IEEE 802.15.4e WPAN task group. <http://www.ieee802.org/15/pub/TG4e.html>.
- [5] <http://www.metageek.net/recordings>.
- [6] <http://www.tinyos.net/>.
- [7] IEEE 802.11, <http://www.ieee802.org/11/>.
- [8] IEEE 802.15.1, <http://www.ieee802.org/15/pub/TG1.html>.
- [9] Low-Power Wireless Communication, <http://www.cs.berkeley.edu/~culler/AIIT/slides/W2-low-power.ppt>.
- [10] Specification of the Bluetooth System, Version 4.0.
- [11] Technical Overview of Time Synchronized Mesh Protocol, White Paper, <http://www.dustnetworks.com>.
- [12] Wi-Spy, <http://www.metageek.net/>.
- [13] ZigBee, White Paper, June 2007.
- [14] Muhammad Hamad Alizai, Olaf Landsiedel, Jo Agila Bitsch Link, Stefan Goetz, and Klaus Wehrle. Bursty traffic over bursty links. In *SenSys*, 2009.
- [15] Paramvir Bahl, Ranveer Chandra, Thomas Moscibroda, Rohan Murty, and Matt Welsh. White space networking with wi-fi like connectivity. In *Sigcomm*, 2009.
- [16] John Bicket, Daniel Aguayo, Sanjit Biswas, and Robert Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *Mobicom*, 2005.
- [17] Dawei Chen, Sixing Yin, Qian Zhang, Mingyan Liu, and Shufang Li. Mining spectrum usage data: a large-scale spectrum measurement study. In *Mobicom*, 2009.

- [18] Lance Doherty, William Lindsay, Jonathan Simon, and Kristofer S. J. Pister. Channel-specific wireless sensor network path analysis. In *ICCCN*, 2007.
- [19] Prabal Dutta and David Culler. Practical asynchronous neighbor discovery and rendezvous for mobile sensing applications. In *SenSys*, 2008.
- [20] Jakob Eriksson, Sharad Agarwal, Paramvir Bahl, and Jitendra Padhye. Feasibility study of mesh networks for all-wireless offices. In *MobiSys*, 2006.
- [21] HART Communication Foundation. WirelessHART technology, 2009. http://www.hartcomm.org/protocol/wihart/wireless_technology.html.
- [22] Paul Fuxjger, Danilo Valerio, and Fabio Ricciato. The myth of non-overlapping channels: interference measurements in ieee 802.11. In *Wireless on Demand Network Systems and Services, WONS*, 2007.
- [23] Omprakash Gnawali, Rodrigo Fonseca, Kyle Jamieson, David Moss, and Philip Levis. Collection tree protocol. In *SenSys*, 2009.
- [24] Ramakrishna Gummadi, David Wetherall, Ben Greenstein, and Srinivasan Seshan. Understanding and mitigating the impact of RF interference on 802.11 networks. In *Sigcomm*, 2007.
- [25] Shuo Guo, Yu Gu, Bo Jiang, and Tian He. Opportunistic flooding in low-duty-cycle wireless sensor networks with unreliable links. In *Mobicom*, 2009.
- [26] Gregory Hackmann, Octav Chipara, and Chenyang Lu. Robust topology control for indoor wireless sensor networks. In *SenSys*, 2008.
- [27] Jan-Hinrich Hauer, Vlado Handziski, and Adam Wolisz. Experimental study of the impact of WLAN interference on IEEE 802.15.4 body area networks. In *EWSN*, 2009.
- [28] Ivan Howitt and Jose A. Gutierrez. IEEE 802.15.4 low rate - wireless personal area network coexistence issues. In *WCNC*, 2003.
- [29] Ozlem Durmaz Incel, Stefan Dulman, Pierre Jansen, and Sape Mullender. Multi-channel interference measurements for wireless sensor networks. In *LCN*, 2006.
- [30] Xiangpeng Jing, Shanmuga S Anandaraman, Mesut Ali Ergin, Ivan Seskar, and Dipankar Raychaudhuri. Distributed coordination schemes for multi-radio coexistence in dense spectrum environments: An experimental study on the orbit testbed. In *DySPAN*, 2008.
- [31] Youngmin Kim, Hyojeong Shin, and Hojung Cha. Y-mac: An energy-efficient multi-channel mac protocol for dense wireless sensor networks. In *IPSN*, 2008.

- [32] Hieu Khac Le, Dan Henriksson, and Tarek Abdelzaher. A practical multi-channel media access control protocol for wireless sensor networks. In *IPSN*, 2008.
- [33] Philip Levis and David Gay. Tinyos programming. April 13, 2009.
- [34] C.-J. M. Liang, B. Priyantha, J. Liu, and A. Terzis. Surviving wi-fi interference in low power zigbee networks. In *SenSys*, 2010.
- [35] Ritesh Maheshwari, Shweta Jain, and Samir R. Das. A measurement study of interference modeling and scheduling in low-power wireless networks. In *SenSys*, 2008.
- [36] David Moss and Philip Levis. BoX-MACs: Exploiting physical and link layer boundaries in low-power networking. Technical Report SING-08-00, Stanford Information Networks Group, 2008.
- [37] Vishnu Navda, Aniruddha Bohra, Samrat Ganguly, and Dan Rubenstein. Using channel hopping to increase 802.11 resilience to jamming attacks. In *INFOCOM Minisymposium*, 2007.
- [38] ZigBee Standards Organization. *ZigBee Smart Energy Profile Specification*, 2008.
- [39] Jorge Ortiz and David Culler. Multichannel reliability assessment in real world WSNs. In *IPSN*, 2010.
- [40] K. Papagiannaki, M. Yarvis, and W. S. Conner. Experimental characterization of home wireless networks and design implications. In *INFOCOM*, 2006.
- [41] Joseph Polastre, Robert Szewczyk, and David Culler. Telos: Enabling ultra-low power wireless research. In *IPSN*, 2005.
- [42] S. Pollin, M. Ergen, M. Timmers, A. Dejonghe, L. van der Perre, F. Catthoor, I. Moerman, and A. Bahai. Distributed cognitive coexistence of 802.15.4 with 802.11. In *Cognitive Radio Oriented Wireless Networks and Communications*, 2006.
- [43] Lili Qiu, Yin Zhang, Feng Wang, Mi Kyung Han, and Ratul Mahajan. A general model of wireless interference. In *MobiCom*, 2007.
- [44] Mo Sha, Gregory Hackmann, and Chenyang Lu. Multi-channel reliability and spectrum usage in real homes: Empirical studies for home-area sensor networks. Technical Report WUCSE-2010-32, Washington University in St. Louis, 2010. <http://cse.wustl.edu/Research/Pages/technical-reports.aspx>.
- [45] Mo Sha, Guoliang Xing, Gang Zhou, Shucheng Liu, and Xiaorui Wang. C-MAC: Model-driven concurrent medium access control for wireless sensor networks. In *INFOCOM*, 2009.

- [46] Rahul C. Shah, Lama Nachman, and Chieh-yih Wan. On the performance of bluetooth and IEEE 802.15.4 radios in a body area network. In *BodyNets*, 2008.
- [47] Soo Young Shin, Hong Seong Park, Sunghyun Choi, and Wook Hyun Kwon. Packet error rate analysis of ZigBee under WLAN and Bluetooth interferences. In *IEEE trans. on wireless communications*, 2007.
- [48] Soo Young Shin, Hong Seong Parky, Sunghyun Choi, and Wook Hyun Kwon. Packet error rate analysis of IEEE 802.15.4 under IEEE 802.11b interference. In *WWIC*, 2005.
- [49] Axel Sikora and Voicu F. Groza. Coexistence of IEEE 802.15.4 with other systems in the 2.4 GHz ISM band. In *IMTC*, 2005.
- [50] IEEE Computer Society. *Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs)*, 2006.
- [51] Dongjin Son, B. Krishnamachari, and J. Heidemann. Experimental study of the effects of transmission power control and blacklisting in wireless sensor networks. In *SECON*, 2004.
- [52] Kannan Srinivasan, Prabal Dutta, Arsalan Tavakoli, and Philip Levis. An empirical study of low power wireless. In *ACM Transactions on Sensor Networks*, 2010.
- [53] Kannan Srinivasan and Philip Levis. RSSI is under appreciated. In *EmNets*, 2006.
- [54] Stephen M. Stigler. Francis Galton’s account of the invention of correlation. *Statistical Science*, 4(2), 1989.
- [55] Venkita Subramonian, Guoliang Xing, Christopher Gill, and Ron Cytron. The design and performance of special purpose middleware: A sensor networks case study. Technical Report WUCSE-2003-06, Washington University in St.Louis, 2003.
- [56] Yanjun Sun, Omer Gurewitz, Shu Du, Lei Tang, and David B. Johnson. Adb: An efficient multihop broadcast protocol based on asynchronous duty-cycling in wireless sensor networks. In *SenSys*, 2009.
- [57] L. Toscano, E. Lo Bello. Cross-channel interference in iee 802.15.4 networks. In *WFCS*, 2008.
- [58] Yafeng Wu, John A. Stankovic, Tian He, and Shan Lin. Realistic and efficient multi-channel communications in wireless sensor networks. In *Infocom*, 2008.

- [59] D. G. Yoon, S. Y. Shin, W. H. Kwon, and H. S. Park. Packet error rate analysis of IEEE 802.11b under IEEE 802.15.4 interference. In *VTC Spring*, 2006.
- [60] Jerry Zhao and Ramesh Govindan. Understanding packet delivery performance in dense wireless sensor networks. In *Sensys*, 2003.

Vita

Mo Sha

Date of Birth May 19, 1985

Place of Birth Beijing, China

Degrees B.Eng. Computer Science, July 2007
MPhil Computer Science, August 2009

Publications Mo Sha, Gregory Hackmann and Chenyang Lu (2011). ARCH: Practical Channel Hopping for Reliable Home-Area Sensor Networks. The 17th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS 2011), April 11 - 14, 2011, Chicago, IL, USA.

May 2011

Empirical Studies for HAN, Sha, M.S. 2011