Report Number: WUCSE-2010-32

2010

Multi-Channel Reliability and Spectrum Usage in Real Homes: Empirical Studies for Home-Area Sensor Networks

Mo Sha, Gregory Hackmann, and Chenyang Lu

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. This paper presents two in-depth empirical studies on the wireless channels in real homes. The spectrum study analyzes the spectrum usage in the 2.4 GHz band where wireless sensor networks based on the IEEE 802.15.4 standard must coexist with existing wireless devices. We characterize the ambient wireless environment in six apartments through passive spectrum analysis across the entire 2.4 GHz band... **Read complete abstract on page 2.**

Recommended Citation

Sha, Mo; Hackmann, Gregory; and Lu, Chenyang, "Multi-Channel Reliability and Spectrum Usage in Real Homes: Empirical Studies for Home-Area Sensor Networks" Report Number: WUCSE-2010-32 (2010). *All Computer Science and Engineering Research.*
https://openscholarship.wustl.edu/cse_research/46

Multi-Channel Reliability and Spectrum Usage in Real Homes: Empirical Studies for Home-Area Sensor Networks

Mo Sha, Gregory Hackmann, and Chenyang Lu

Complete 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. This paper presents two in-depth empirical studies on the wireless channels in real homes. The spectrum study analyzes the spectrum usage in the 2.4 GHz band where wireless sensor networks based on the IEEE 802.15.4 standard must coexist with existing wireless devices. We characterize the ambient wireless environment in six apartments through passive spectrum analysis across the entire 2.4 GHz band over seven days in each of the apartments. Notably, we find that the wireless conditions in these residential environments can be much more complex and varied than in a typical office environment. Moreover, while 802.11 signals play a significant role in spectrum usage, there also exist non-negligible noise from non-802.11 devices. The multi-channel link study measures the reliability of different 802.15.4 channels through active probing with motes. We discover that there is not always a persistently reliable channel over 24 hours; that reliability is strongly correlated across adjacent channels; and that link reliability does not exhibit cyclic behavior at daily or weekly timescales. Nevertheless, reliability can be maintained through a small number of channel hops per day, suggesting channel diversity as a key tool for designing robust HANs in residential environments. Our empirical studies provide important guidelines and insights for robust wireless sensor network design in residential environments.

2010-32

Multi-Channel Reliability and Spectrum Usage in Real Homes: Empirical Studies for Home-Area Sensor Networks

Authors: Mo Sha, Gregory Hackmann, 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. This paper presents two in-depth empirical studies on the wireless channels in real homes. The spectrum study analyzes the spectrum usage in the 2.4 GHz band where wireless sensor networks based on the IEEE 802.15.4 standard must coexist with existing wireless devices. We characterize the ambient wireless environment in six apartments through passive spectrum analysis across the entire 2.4 GHz band over seven days in each of the apartments. Notably, we find that the wireless conditions in these residential environments can be much more complex and varied than in a typical office environment. Moreover, while 802.11 signals play a significant role in spectrum usage, there also exist non-negligible noise from non-802.11 devices. The multi-channel link study measures the reliability of different 802.15.4 channels through active probing with motes. We discover that there is not always a persistently reliable channel over 24 hours; that reliability is strongly correlated across adjacent channels; and that link reliability does not exhibit cyclic behavior at daily or weekly timescales. Nevertheless, reliability can be maintained through a small number of channel hops per day, suggesting channel diversity as a key tool for

Type of Report: Other

Multi-Channel Reliability and Spectrum Usage in Real Homes: Empirical Studies for Home-Area Sensor Networks

Mo Sha, Gregory Hackmann, Chenyang Lu
Department of Computer Science and Engineering
Washington University in St. Louis, USA

***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. This paper presents two in-depth empirical studies on the wireless channels in real homes. The *spectrum study* analyzes spectrum usage in the 2.4 GHz band where wireless sensor networks based on the IEEE 802.15.4 standard must coexist with existing wireless devices. We characterize the ambient wireless environment in six apartments through passive spectrum analysis across the entire 2.4 GHz band over seven days in each of the apartments. Notably, we find that the wireless conditions in these residential environments can be much more complex and varied than in a typical office environment. Moreover, while 802.11 signals play a significant role in spectrum usage, there also exist non-negligible noise from non-802.11 devices. The *multi-channel link study* measures the reliability of different 802.15.4 channels through active probing with motes. We discover that there is not always a persistently reliable channel over 24 hours; that reliability is strongly correlated across adjacent channels; and that link reliability does not exhibit cyclic behavior at daily or weekly timescales. Nevertheless, reliability can be maintained through a small number of channel hops per day, suggesting channel diversity as a key tool for designing robust HANs in residential environments. Our empirical studies provide important guidelines and insights for robust wireless sensor network design in residential environments.**

I. 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 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 Section III). 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 paper presents a two-part empirical study of wireless channel properties in real-world residential settings. 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. We then explore the concrete impact of these challenging environments on application performance, through an active probing study of 802.15.4 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 typical apartment environment, there may not be a single channel which is persistently reliable for 24 hours. (3) Even the

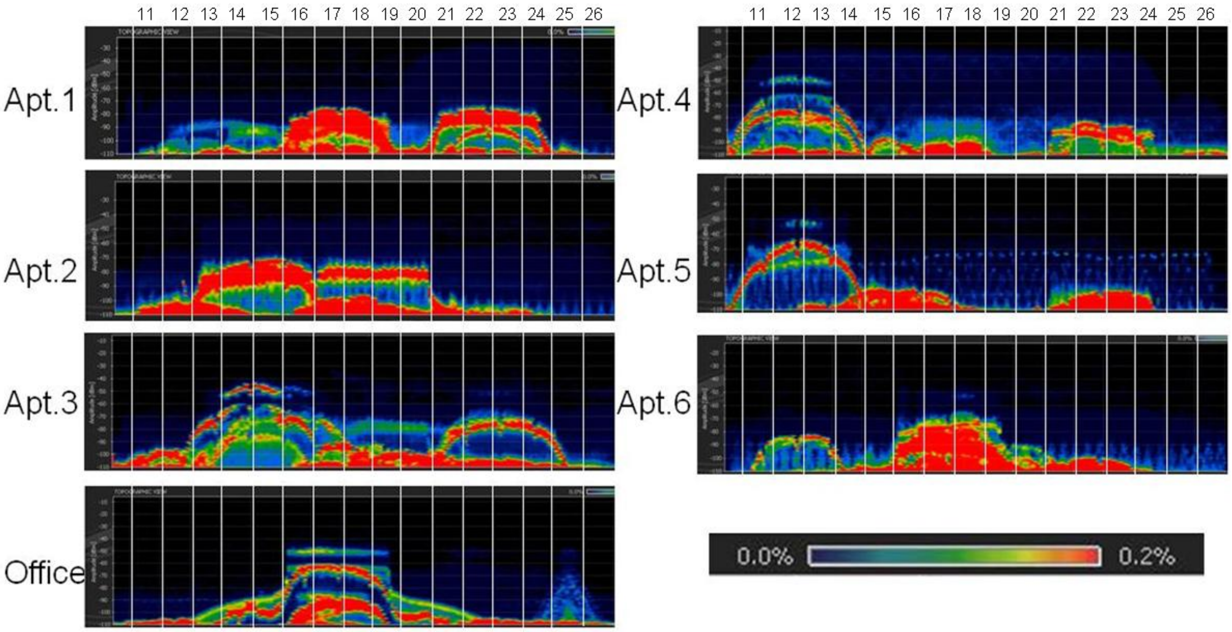

Figure 1. Histogram over 7 days' raw energy traces. X axis indicates 802.15.4 channels, Y axis indicates power, and color indicates how often a signal was detected at x GHz with an energy level of y dBm.

"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 paper is organized as follows. Section II reviews related work. Section III discusses the findings of our passive spectral study. Section IV then presents our active probing study and its performance results. Finally, we conclude in Section V.

II. RELATED WORK

Several recent studies have aimed to characterize the impact of interference on wireless networks through controlled experiments [1]–[5]. [6]–[8] present theoretical analysis based on simulation study. Gummadi et al. [9] 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. [10] examines the packet delivery behavior of two 802.15.4-based mote platforms, including the impact of interference from 802.11 and Bluetooth. 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. [11] presents a study of UHF white space networking, while Chen et al. [12] 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.

Ortiz et al. [13] 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 Section IV 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. [14] discusses a multi-channel measurement of Body Area Networks (BANs) and proposes a noise floor-

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 I
THE SETTINGS AND DATES WHERE THE SPECTRUM DATA WAS COLLECTED.

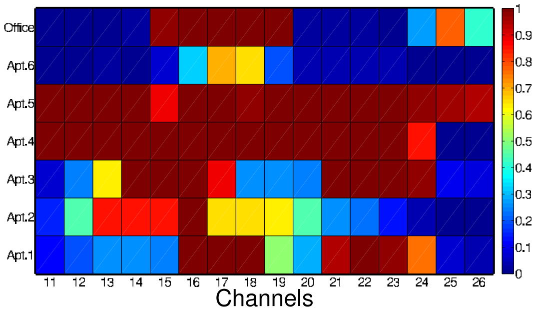

Figure 2. Channel occupancy rate. X axis designates channels, Y axis designates experimental settings, and color represents the proportion of readings above the occupancy threshold.

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. In contrast, our own study looks specifically at the multi-channel properties of indoor residential environments under normal home activities.

III. WIRELESS SPECTRUM STUDY

In this section, 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. As a baseline for comparison, we also collected energy traces from an office in Bryan Hall at Washington University in St. Louis. A detailed description of the experimental settings may be found in Table I.

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?

A. *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 [15]. 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 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 [12] to convert signal strength readings into binary values, with 0 denoting a channel being idle and 1 denoting a channel being busy. Environmental noise above the threshold (i.e., 1s) may introduce packet drops over otherwise good-quality links. Throughout our analysis, we use -85 dBm as our threshold value, which we identified experimentally between a pair of CC2420 radios. Details on this experiment may be found in Appendix A. 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.

B. *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. Figure 2 plots the channel occupancy rate in the six apartments and office building, i.e., the proportion of samples that exceeded the -85 dBm threshold.

If we compare Figures 1 and 2, 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 as a series of thin blue arches. Devices with such wireless footprints include Bluetooth transmitters, baby monitors, wireless speaker systems, and game controllers [17]. (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.*

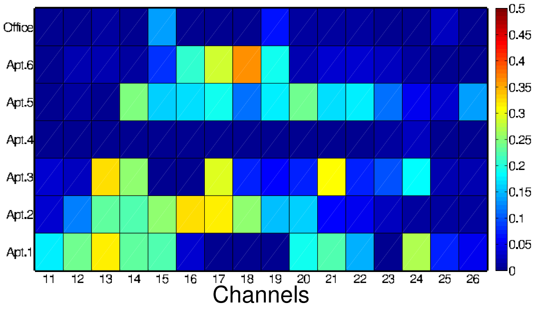

(a) Daily standard deviation

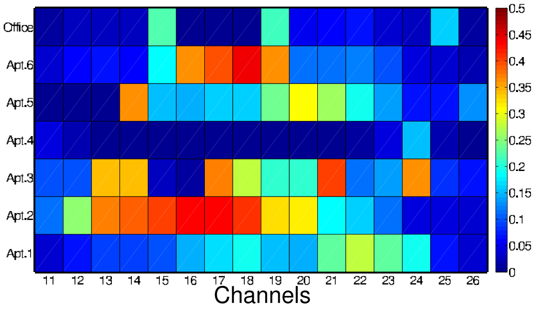

(b) Hourly standard deviation

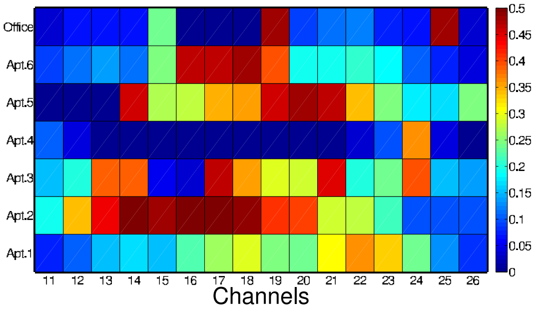

(c) 5-minute standard deviation

Figure 3. The standard deviation in channel occupancy rate at different timescales.

C. 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 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 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).

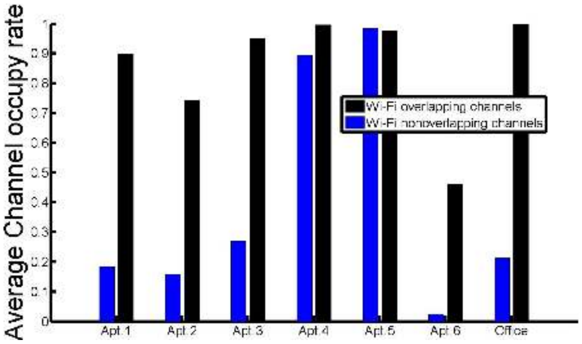

Figure 5. A comparison of the average channel occupancy rate between channels that overlap with Wi-Fi and channels that do not.

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

D. Is Channel Occupancy Temporally Correlated?

Although channel occupancy is highly variable even on a timescale of minutes, there may nevertheless be temporal correlations in channel usage on even shorter time scales (e.g., packet-to-packet). To determine if such a correlation exists, we computed the conditional channel usage function (CCUF) for each channel in each apartment. For $k > 0$, $CCUF(k)$ is the conditional probability that k consecutive busy readings are followed by another busy reading; for $k < 0$, $CCUF(k)$ is the conditional probability that $|k|$ consecutive idle readings are followed by another idle reading.

Figure 4 plots the CCUF for three apartments and four channels; results for other apartments and other channels are similar but omitted for space. We note that the CCUF curves are generally flat, with the probabilities almost always above 80% or at 0%. Intuitively, these flat shapes indicate the channel conditions are temporally independent: had there been a strong temporal dependence, we would expect to see the probability of a busy or idle reading affected by the number of preceding busy or idle readings [10].

Observation S3*: Channel usage is not strongly temporally correlated.*

E. 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 and 2 suggests other important interferers besides Wi-Fi. Wi-Fi APs have a distinctive radiation pattern that manifests in Figure 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

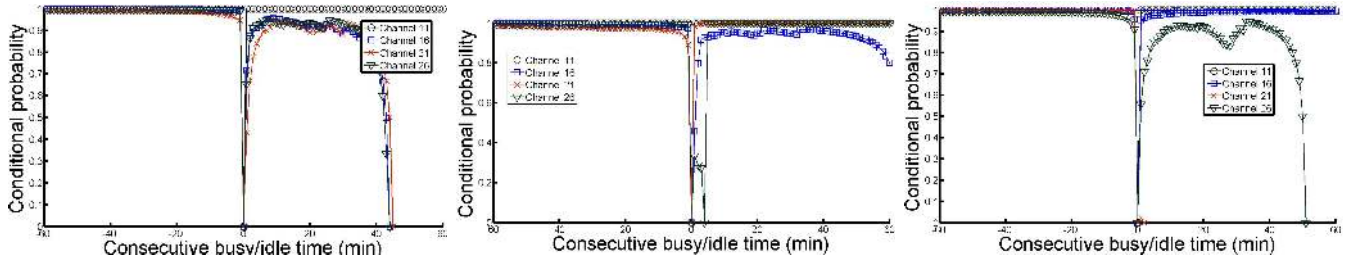

(a) Apt. 1 (b) Apt. 3 (c) Apt. 5

Figure 4. Conditional channel usage functions (CCUFs) in three different apartments. The X axis indicates consecutive busy or idle readings, where negative values represent consecutive busy readings and positive values represent consecutive idle readings. The Y axis provides the probability that the channel is currently busy/idle given x prior time slots which were all busy/idle.

caused primarily by 802.11 APs configured to two different channels. Referring to Figure 2, 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 5.

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 S4**: While Wi-Fi is an important source of interference in residential environments, other interferers can also be non-negligible contributors to spectrum occupancy.*

IV. MULTI-CHANNEL LINK STUDY

In this section, we present a multi-channel link study in homes. The spectrum study presented in Section III 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.

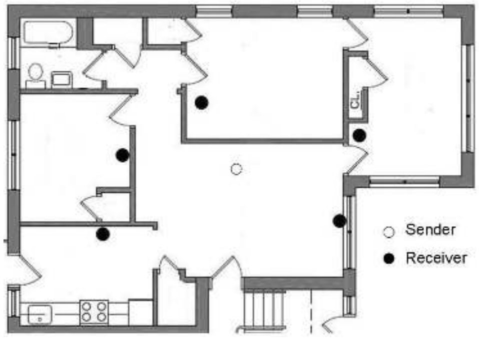

Figure 6. Floor plan of an apartment used in the study.

	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 II
THE SETTINGS AND DATES WHERE THE LINK DATA WAS COLLECTED.

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?

A. 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 II. (Due to the participating residents moving, only four of the apartments in this study are the same as in Section III.) Figure 6 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.

Our experiments were carried out using networks of Tmote Sky and TelosB [18] motes. Each mote is equipped with an IEEE 802.15.4 compliant Chipcon CC2420 radio [19]. 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 [20] 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 broadcasts 100 packets per channel to multiple recipient nodes, cycling over each of the 16 channels over 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 [10], 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 section as a threshold to designate links as "good" or "reliable". Due to the numerous outliers found throughout Section IV-B, we expect our analysis to generally hold for any reasonable PRR threshold.

B. Is There a Persistently Good Channel?

Previous empirical studies [10], [21], [22] 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 [21]. 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 7 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 8 presents a boxplot 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 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,

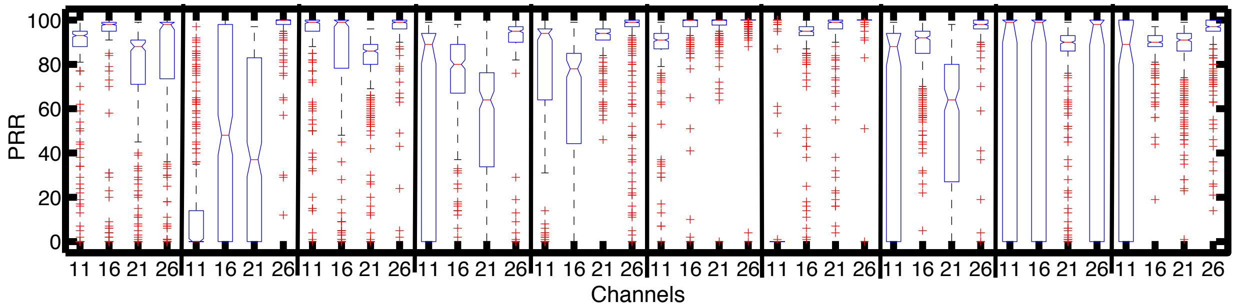

Figure 7. The PRR for four channels in all ten apartments, calculated over 5-minute windows. Vertical lines delineate apartments.

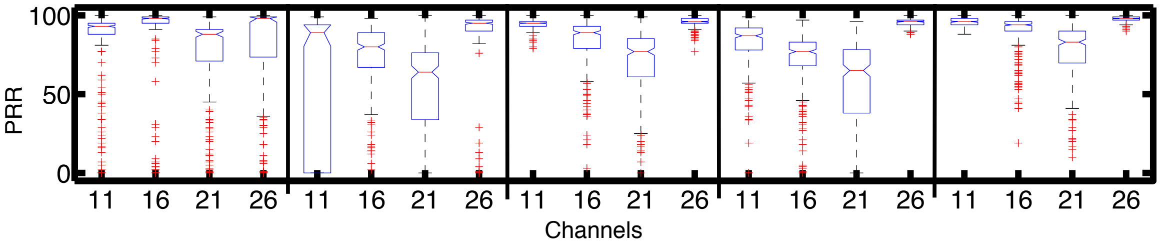

Figure 8. The PRR of five different links in the same apartment on four channels, calculated over 5-minute windows. Vertical lines delineate links.

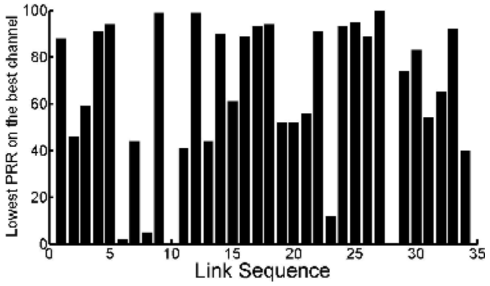

Figure 9. The lowest PRR observed on each link's most reliable channel.

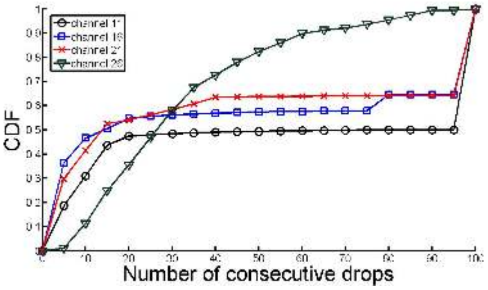

Figure 10. CDF of number of consecutive drops.

even on their most reliable channels. Figure 9 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.*

C. 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 10 illustrates this problem with the cumulative probability density (CDF) of consecutive packet drops for all links on four channels. 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.*

D. 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 with an optimal channel scheduling algorithm. For each link, the algorithm finds the

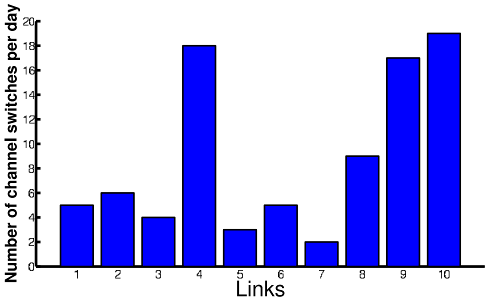

(a) Number of channel hops required under an optimal schedule; one link randomly selected per apartment.

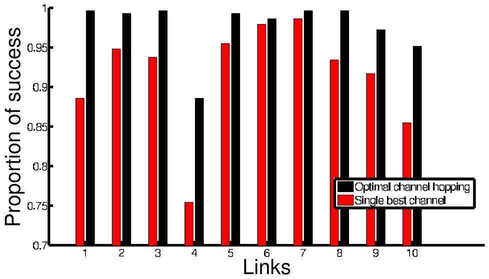

(b) The proportion of windows where the PRR threshold was met.

Figure 11. Optimal channel switching schedule in different apartments.

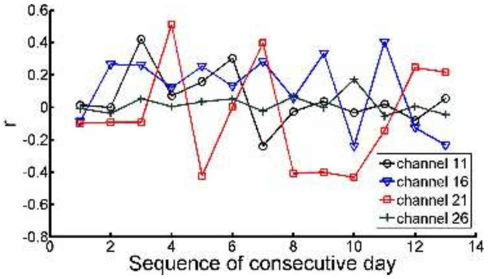

(a) PMCC of PRRs during the same time on consecutive days.

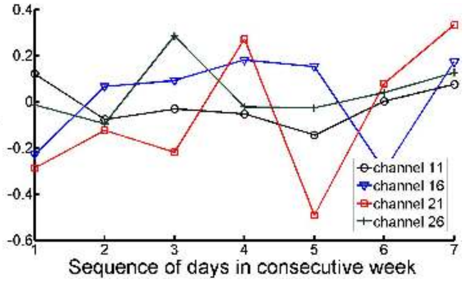

(b) PMCC of PRRs during the same time in consecutive weeks (e.g., $x = 1$ means consecutive Mondays).

Figure 12. The Pearson's product correlation coefficient (PMCC) comparing the PRR at the same time on consecutive days or weeks.

sequence of channels that exceeds the 90% PRR threshold as often as possible in our experimental data, while switching channels as few times as possible. (We note that this algorithm is for analysis only and cannot be implemented online, since it requires the whole data trace. However, the optimal channel hopping scheme indicates the potential benefits of channel hopping.) A formal definition of this algorithm and proof of its optimality are provided in Appendix B.

Figure 11(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 11(b) compares the proportion of windows which meet the 90% threshold under the optimal channel-hopping analysis and under the optimal single-channel analysis. 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 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 [23] and newer standards such as WirelessHART's TSMP [24] and the forthcoming IEEE 802.15.4e [25]. The results of our analysis confirm that this feature is indeed beneficial for maintaining link reliability in challenging residential environments.

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.*

E. Can Hopping be Scheduled Statically?

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) [26], 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 12(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 12(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

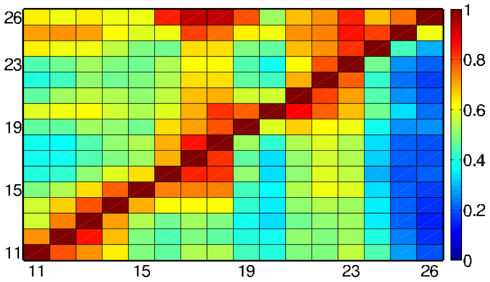

Figure 13. Correlation of channel reliability. The X and Y axes indicate channels; the color indicates the probability that channel x's PRR $< 90\%$ when channel y's PRR $< 90\%$.

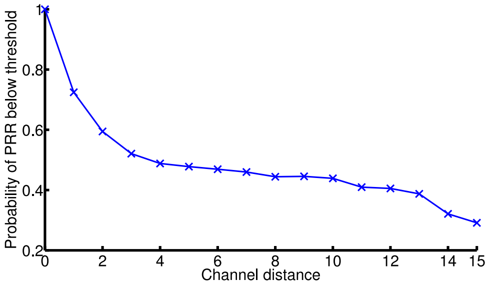

Figure 14. Correlation of channel reliability as a function of channel distance.

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.*

F. 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 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 13, 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 14, 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.*

V. 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 paper 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. Our study leads to several key insights for developing robust HANs based on WSNs. The results of our passive spectrum analysis indicate that residential settings can provide a much more complex and varied wireless environment than a representative office environment. Through our subsequent link study, we found that may not be a persistently reliable channel over 24 hours in real-world apartments; that retransmissions alone cannot compensate for poor channel conditions; and that link reliability in WSNs does not generally behave as a cyclic phenomenon. Despite these challenges, we also found that only a small number of judicious channel hops are required to maintain link reliability; we also observed a correlation in performance across nearby channels that could be used when selecting new channels. These findings suggest channel diversity as an effective option for achieving reliable HAN deployments. Based on the guidelines established in this study, we plan to develop robust HANs that leverage channel diversity in future work.

APPENDIX A
THRESHOLD SELECTION

According to communication theory, the probability that a transceiver r successfully receives an incoming packet ω, denoted by $p_r(\omega)$, is governed by the following physical model:

$$p_r(\omega) = Prob[\frac{\text{signal power of } \omega}{I_r + n_r} > \beta_r] \tag{1}$$

I_r is the interference experienced at r, which is equal to the power of other nodes' transmissions and other electromagnetic signals from the environment. n_r is a random variable that equals the power of ambient noise. β_r is a constant determined by modulation and transceiver sensitivity.

To study the impact of interference, we can apply the probabilistic model to construct a thresholding model. If N is the average ambient noise level measured at the receiver, we can then calculate the signal-to-interference-plus-noise-ratio ($SINR$) as:

$$SINR_{dB} = 10log_{10}\frac{10^{RSS_{dBm}/10} - 10^{N_{dBm}/10}}{10^{N_{dBm}/10}} \tag{2}$$

From Eq. (2), we get

$$10^{SINR_{dB}/10} = \frac{10^{RSS_{dBm}/10} - 10^{N_{dBm}/10}}{10^{N_{dBm}/10}} \tag{3}$$

$$10^{N_{dBm}/10} = \frac{10^{RSS_{dBm}/10}}{10^{SINR_{dB}/10} + 1} \tag{4}$$

Figure 15 plots the correlation between signal strength and PRR as obtained experimentally between a pair of TelosB motes. We see that $RSS_{dBm} = -80$ dBm places the link

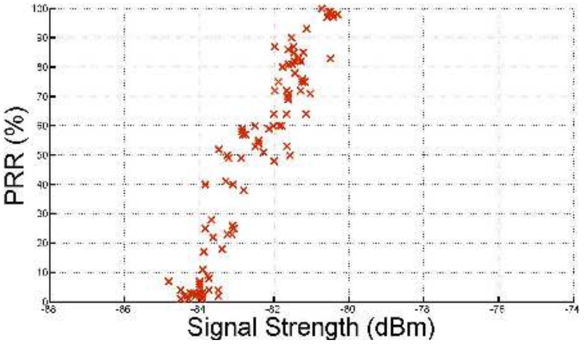

Figure 15. Relationship between signal strength and PRR, as measured experimentally.

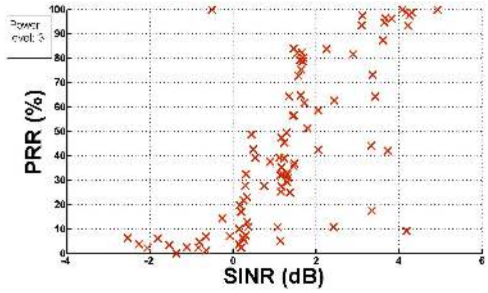

(a) Using low transmission power.

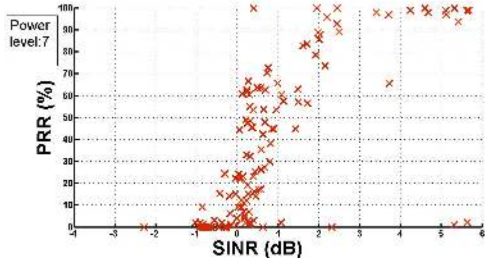

(b) Using high transmission power.

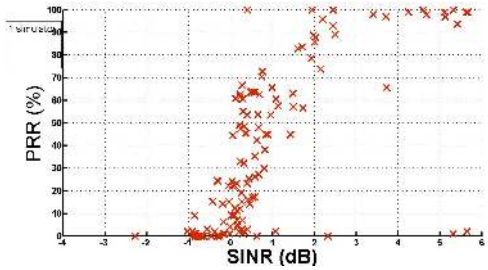

(c) Using 1 mote to simulate jammer.

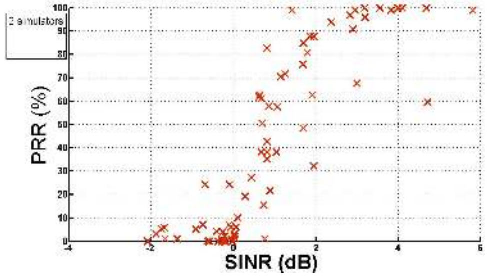

(d) Using 2 motes to simulate jammer.

Figure 16. Relationship between SINR and PRR under various experimental setups.

outside of the transitional "gray" region. Figure 16 likewise plots the correlation between SINR and PRR obtained experimentally using different transmission powers and external interferers. A threshold of $SINR_{dB} = 4$ dB likewise places the link outside of the transitional region in all experimental configurations. Therefore, we get

$$10^{N_{dBm}/10} = \frac{10^{-80dBm/10}}{10^{4dB/10} + 1} \tag{5}$$

$$N_{dBm} = -85 \text{ dBm} \tag{6}$$

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

APPENDIX B
OPTIMAL CHANNEL-HOPPING ALGORITHM

The optimal channel-hopping operates as follows. For each batch of 100 packets, we calculate a binary value for whether its PRR meets or exceeds the -85 dBm threshold discussed in Appendix A. These binary sequences are then input into the greedy algorithm shown in Algorithm 1. We initially check to see if any channel can meet the threshold at a given point in time; if not, there is no need to switch channels, since none will meet the threshold requirement. The algorithm then continuously searches for the longest sequence of consecutive *1*s (i.e., windows of uninterrupted reliability) until reaching the end of the dataset. A proof of the algorithm's optimality follows.

Algorithm 1 Optimal Algorithm

```
Input: $S = \{s_{m1}s_{m2}...s_{mn} | m \in [1,16]\}$ //binary sequences of 16 channels with length of $n$.
Output: $\Gamma$ //channel switching schedule.

if $\exists k \in [1,n], \forall s_{ik} = 0$ then
    Set $s_{1k} = s_{2k} = ... = s_{16k} = 1$; //if all channels cannot meet the threshold, no need to switch.
end if
Initialize $\Gamma = \emptyset, t = 1$;
repeat
    Find the longest sequence of consecutive 1's in $S$, which begins at $s_{ip}$ and ends at $s_{iq}$ where $p = t$ and $i \in [1,16]$;
    Set $t = q + 1, \Gamma = \Gamma \cup s_{ip}...s_{iq}$
until t>n
```

Proposition 1: Let S_i be the first sequence chosen by the greedy algorithm. There exists an optimal solution containing S_i.

Proof: Let Γ^* be any optimal solution. If $S_i \in \Gamma^*$, then the proof is finished. Otherwise, let S_j be the first solution in Γ^* with the minimal number of switches n. Construct a new solution Γ from Γ^* by discarding S_j and adding S_i. S_j and S_i begin at the same place but S_i has the longer consecutive sequence of *1*s; hence all bits equal to *1* in Γ^* will be the same in the new solution Γ (the rest of the solution did not change). Moreover, Γ contains just as many channel switches (n) as Γ^*, so Γ is still optimal.

After making the greedy choice S_i, we are left with a subproblem with a smaller length of sequences, and with no external constraints (since any solution is compatible with the choice of S_i).

Proposition 2: Let Γ^+ be an optimal solution to the subproblem, and let $\Gamma = \Gamma^+ \cup \{S_i\}$, where S_i is the greedy choice. Then Γ optimally solves the full problem.

Proof: Observe that $|\Gamma| = |\Gamma^+| + 1$. The usual contradiction argument follows.

REFERENCES

[1] X. Jing, S. S. Anandaraman, M. A. Ergin, I. Seskar, and D. Raychaudhuri, "Distributed coordination schemes for multi-radio co-existence in dense spectrum environments: An experimental study on the orbit testbed," in *DySPAN*, 2008.

[2] A. Sikora and V. F. Groza, "Coexistence of IEEE 802.15.4 with other systems in the 2.4 GHz ISM band," in *IMTC*, 2005.

[3] ZigBee, White Paper, June 2007.

[4] I. Howitt and J. A. Gutierrez, "IEEE 802.15.4 low rate - wireless personal area network coexistence issues," in *WCNC*, 2003.

[5] S. Y. Shin, H. S. Park, S. Choi, and W. H. Kwon, "Packet error rate analysis of ZigBee under WLAN and Bluetooth interferences," in *IEEE trans. on wireless communications*, 2007.

[6] S. Y. Shin, H. S. Parky, S. Choi, and W. H. Kwon, "Packet error rate analysis of IEEE 802.15.4 under IEEE 802.11b interference," in *WWIC*, 2005.

[7] 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.

[8] 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.

[9] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan, "Understanding and mitigating the impact of rf interference on 802.11 networks," in *Sigcomm*, 2007.

[10] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis, "An empirical study of low power wireless," in *ACM Transactions on Sensor Networks*, 2010.

[11] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh, "White space networking with wi-fi like connectivity," in *Sigcomm*, 2009.

[12] D. Chen, S. Yin, Q. Zhang, M. Liu, and S. Li, "Mining spectrum usage data: a large-scale spectrum measurement study," in *Mobicom*, 2009.

[13] J. Ortiz and D. Culler, "Multichannel reliability assessment in real world wsns," in *IPSN*, 2010.

[14] J.-H. Hauer, V. Handziski, and A. Wolisz, "Experimental study of the impact of WLAN interference on IEEE 802.15.4 body area networks," in *EWSN*, 2009.

[15] Wi-Spy, http://www.metageek.net/.

[16] M. Sha, G. Hackmann, and C. Lu, "Multi-channel reliability and spectrum usage in real homes: Empirical studies for home-area sensor networks," Washington University in St. Louis, Tech. Rep. WUCSE-2010-32, 2010. [Online]. Available: http://cse.wustl.edu/Research/Pages/technical-reports.aspx

[17] [Online]. Available: http://www.metageek.net/recordings

[18] J. Polastre, R. Szewczyk, and D. Culler, "Telos: Enabling ultra-low power wireless research," in *IPSN*, 2005.

[19] 2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver, Texas Instruments.

[20] http://www.tinyos.net/.

[21] J. Zhao and R. Govindan, "Understanding packet delivery performance in dense wireless sensor networks," in *Sensys*, 2003.

[22] D. Son, B. Krishnamachari, and J. Heidemann, "Experimental study of the effects of transmission power control and blacklisting in wireless sensor networks," in *IEEE SECON*, 2004.

[23] Specification of the Bluetooth System, Version 4.0.

[24] Technical Overview of Time Synchronized Mesh Protocol, White Paper, http://www.dustnetworks.com.

[25] "Ieee 802.15.4e wpan task group." [Online]. Available: http://www.ieee802.org/15/pub/TG4e.html

[26] S. M. Stigler, "Francis Galton's account of the invention of correlation," *Statistical Science*, vol. 4, no. 2, 1989.