

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. These applications impose unique QoS constraints, requiring low data rates but high network reliability in the face of unpredictable wireless environments. This paper presents two in-depth empirical studies on wireless channels in real homes, providing key design guidelines for meeting the QoS constraints of HAN applications. The *spectrum study* analyzes spectrum usage in the 2.4 GHz band where HANs 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 apartment. We find that the wireless conditions in these residential environments are 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 exists 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 in ten apartments. We find that there is not always a persistently reliable channel over 24 hours, and that link reliability does not exhibit cyclic behavior at daily or weekly timescales. Nevertheless, reliability can be maintained through infrequent channel hopping, suggesting dynamic channel hopping as a key tool for meeting the QoS requirements of HAN applications. Our empirical studies provide important guidelines and insights in designing HANs for 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.

Such HAN applications have increasingly adopted the IEEE 802.15.4 wireless personal area network standard [1] to provide wireless communication among sensors and actuators.

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 [2] have adopted 802.15.4 technology for use in residential automation applications. The IETF has promoted efforts to standardize IPv6 on top of 802.15.4 for integrating wireless sensors into the Internet [3], [4].

However, meeting the QoS requirements of these emerging applications often poses unique technical challenges, particularly in residential environments. A typical HAN application may feature a low data rate but require high network reliability. These QoS constraints are easily met by wired communication but are non-trivial when dealing with unreliable wireless channels. Notably, the low-power wireless sensor networks used in these applications are highly susceptible to *external* factors beyond the application's control, such as uncontrolled interference from other devices. In particular, 802.15.4 shares the unlicensed 2.4 GHz spectrum with Wi-Fi access points, Bluetooth peripherals, cordless phones, and numerous other devices prevalent in residential environments.

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 office 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

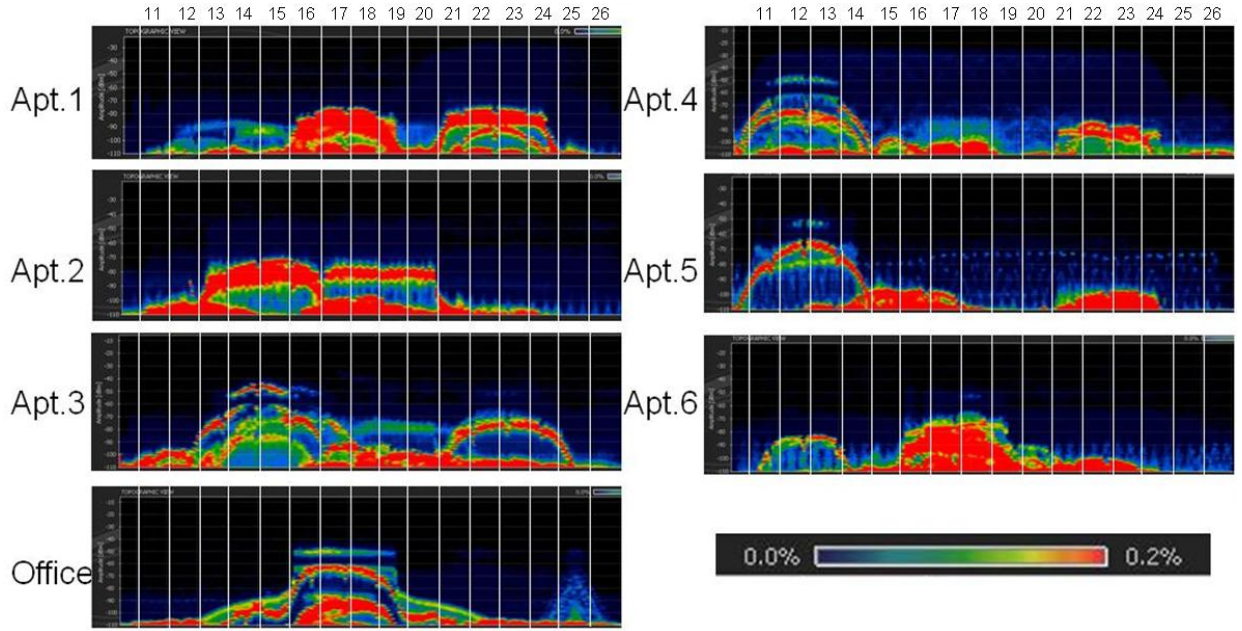


Figure 1. Histogram over 7 days' raw energy traces. X axis indicates 802.15.4 channels, Y axis indicates power, and color indicates how often a signal was detected at x GHz with an energy level of y dBm.

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 paper, we present a two-part empirical study which aims to characterize the real-world network performance of HANs, focusing specifically on devices based on the 802.15.4 standard. 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 how these challenging environments may directly affect applications' QoS, through an active probing study of wireless link reliability across all 16 channels in ten apartments. This second study focuses on packet reception ratio (PRR), which is both a direct indicator of link reliability and closely related to other important QoS metrics such as latency and energy consumption. From this active study, we make several more key observations which could greatly impact the QoS 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) Exploiting channel diversity by infrequent channel hopping at runtime

can effectively maintain long-term reliable communication. (4) Channel conditions are not cyclic. These findings indicate the importance of channel diversity in achieving reliable HAN deployments and provide design guidelines for meeting the QoS requirements of HAN applications.

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. Finally, we conclude in Section V by highlighting the implications of our findings on HAN design.

II. RELATED WORK

Several recent studies have aimed to characterize the impact of interference on wireless networks through controlled experiments [5]–[9]. [10]–[12] present theoretical analysis based on simulation study. Gummadi et al. [13] 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. [14] 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. [15] 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. For instance, our spectrum study in six apartments showed that — while Wi-Fi is a significant source of interference in

residential environments — non-Wi-Fi devices may also be non-negligible sources of interference. This result indicates that solutions specifically targeted at one type of co-existing wireless technology may not be generally applicable to all residential environments.

Bahl et al. [16] 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. Our own study focuses on the reliability of static, indoor wireless sensor networks designed for home environments, and on the unlicensed 2.4 GHz band used by IEEE 802.15.4 and shared by other wireless devices prevalent in residential 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, the diverse sources of interference (including the non-negligible impact of non-Wi-Fi devices), and the effectiveness of infrequent channel hopping in maintaining link reliability.

Papagiannaki et al. [18] performed an empirical study of home networks based on 802.11 technology. Our study considers devices based on the 802.15.4 standard, which operate at a much lower transmission power than 802.11 devices and hence are significantly more susceptible to interference. Our study therefore leads to a different set of observations that underscores the impact of spectrum usage on these low-power 802.15.4 networks.

Ortiz et al. [19] 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 to 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. [20] 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. [21] 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. Instead of body-area networks, our own study focuses on HANs designed for smart energy, which feature significantly different setups and wireless properties. Moreover, our study is performed under normal home activities, providing a realistic setting to evaluate HAN performance.

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. We note that this baseline is meant to illustrate how controlled testbed settings within an office environment 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 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 [22]. 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 during the residents’ normal daily activities, as well as in an office in Bryan Hall. 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 [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. (For brevity, we discuss this experiment in more detail in [23].) 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 trends and observations




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.

we make from these results should generally apply to other threshold values.

To assess the impact of ambient wireless signals on HANs, we 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. If such a channel exists, it could be used as a default, factory preset channel for HANs. For example, channel 26 is often assumed as a good default channel, because it does not overlap with the spectrum used by 802.11 in North America.


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 the 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 2 plots the occupancy rate of each channel in each location. 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. Notably, even channel 26 has a channel occupancy rate as high as 95.04%, contradicting the commonly-held assumption that channel 26 will be open. The uniformly high occupancy rate across channels 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 [24]. (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 in contrast has an occupancy rate of 100.0% in apartments 3 and 4; thus, there is no common



(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.

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 high-power interferer.

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

C. Does Spectrum Usage Change with Time?

We next explored whether the spectrum was stable in these residential settings. If spectrum is stable within a given




Figure 4. A comparison of the average channel occupancy rate between channels that overlap with Wi-Fi and channels that do not.

apartment, it would be possible for a technician to pick a single “best” channel for the HAN at deployment time and expect it to work well over a long time period.

To determine this, 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. Except for 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 a $\sigma = 14.9\%$ 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).

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

D. 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. If Wi-Fi is indeed the dominant interferer in residential settings, then HANs could leverage solutions which are specifically designed to avoid interference from Wi-Fi networks (e.g., [15]).

A visual 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

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 service set identifier (SSID) and 802.11 channel of all visible 802.11 access points (APs)¹. Based on this data, we are able to divide the 802.15.4 channels in each apartment into two groups: those that overlap with 802.11 APs detectable 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 4.

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. This observation can have important implications on the design of HANs, in that solutions specifically designed to deal with Wi-Fi interference may not be effective in all residential environments.

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.*

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.

Specifically, this study addresses the following questions. (1) Can a HAN find a single persistently reliable channel for wireless communication? (2) If no single channel can be used for reliable operation, can the network exploit channel

¹Although many APs may be configured not to broadcast their SSID, we have observed that the Wi-Spy software can still identify these “hidden” access points in practice.




Figure 5. Floor plan of an apartment used in the study.

diversity to achieve reliability? (3) Do channel conditions exhibit cyclic behavior over time?

A. Experimental Methodology

For this active study, we carried out a series of experiments in ten real-world apartments in different neighborhoods. (Due to the participating residents moving, only four of the apartments in this study are the same as those instrumented in the spectrum study.) Figure 5 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 [25] motes. Each mote is equipped with an IEEE 802.15.4 compliant Chipcon CC2420 radio [26]. 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 [27] using the CC2420 driver's default CSMA/CA MAC layer.

We measure the packet reception ratio (PRR), defined as the fraction of transmitted packets successfully received by the receiver. PRR is not only a direct indicator of link reliability, but also closely related to other important QoS metrics such as latency and energy consumption. To measure the 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 [14], 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”.

B. Is There a Persistently Good Channel?

We first analyzed our data from the perspective of finding a single, persistently good channel across all of the tested apartments. Again, if a common good channel exists across all apartments, then it could be used as a preset default channel for HANs. For this analysis, we grouped the data from all links in all apartments together and then subdivided it by channel. Figure 6 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 exhibit similar behavior and 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 an open 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 7 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




Figure 6. 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 7. 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.




Figure 8. The lowest PRR observed on each link's most reliable channel.

shows particularly high variance on channel 26, with a 25th-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 the links in our experiment are outside the “gray region” where small temporal changes in link quality can cause bursts of packet losses [28]. Moreover, each link had at least one channel with a high median PRR. Had the links been in the gray region, we would have expected lower median PRR or uniform variability across *all* channels.

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 8 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.

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 Channel Diversity Effective?

Our analysis above indicates that using a single channel is often not acceptable 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 9(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



(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 9. Retrospective channel-hopping analysis in different apartments.


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 average PRR). Figure 9(b) compares the proportion of windows which meet the 90% threshold under two retrospective strategies: an ideal channel-hopping strategy that maintains the PRR threshold with the minimum number of channel hops, and a strategy that fixes each link to its single “best” channel with the highest average PRR. (Note that both strategies make decisions based on the entire data trace retrospectively, and hence cannot be employed at run time; they are chosen here to analyze the *potential* benefit of channel hopping.) 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 notable 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 [29] and newer standards such as WirelessHART’s TSMP [30] and the forth-



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



(b) PMCC of PRRs during the same time in consecutive weeks.

Figure 10. The Pearson’s product correlation coefficient (PMCC) comparing the PRR at the same time on consecutive days or weeks.

coming IEEE 802.15.4e [31]. The results of our analysis confirm that this feature is indeed beneficial for maintaining link reliability in challenging residential environments.

Observation L3: *Channel hopping is effective in alleviating packet loss due to channel degradation. Infrequent channel hopping can effectively maintain reliable communication.*

D. Can Hopping be Scheduled Statically?

Because channel quality varies over time, we next explored whether it exhibits cyclic properties (e.g., due to recurrent human activities and schedules). 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 using same setup in one apartment over a period of 14 days. We then calculated the Pearson product-moment correlation coefficient (PMCC) [32], 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 10(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 10(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 L4: *Channel conditions are not cyclic, so channel-hopping decisions must be made dynamically.*

V. CONCLUSION

HANs based on wireless sensor network technology represent a promising communication platform for emerging home automation applications such as smart energy. These emerging applications often impose stringent QoS requirements in terms of network reliability, which are made challenging by the complex and highly variable wireless environments in typical residential environments. 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. The observations made in our study highlight the significant challenges that face HAN applications for achieving acceptable QoS in residential settings. Nevertheless, our observations also suggest that these 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 to a static channel assignment, whether made at the factory or at deployment time. (S1, L1, L2)
- 2) 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)
- 3) Reliable communication can be maintained through infrequent channel hopping. (L3)
- 4) Channel hopping cannot be performed based on a static, cyclic schedule. (L4) Instead, channel-hopping decisions should be made dynamically based on conditions observed at runtime. (S2, L2)

We believe that our findings and insights will provide general design guidelines and impact the development of HANs that are gaining increasing importance with the emergence of smart energy as the “killer app” for wireless sensor networks.

ACKNOWLEDGMENT

This work was supported by NSF under grants CNS-0448554 (CAREER) and CNS-1035773 (CPS), and by generous support from Broadcom Corporation and Emerson Climate Technologies.

REFERENCES

- [1] 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.
- [2] ZigBee Standards Organization, *ZigBee Smart Energy Profile Specification*, 2008.
- [3] “IPv6 over low power WPAN,” <http://datatracker.ietf.org/wg/6lowpan/charter/>.
- [4] “Routing over low power and lossy networks (ROLL),” <http://datatracker.ietf.org/wg/roll/charter/>.
- [5] 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.
- [6] 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.

- [7] ZigBee Alliance, “Zigbee and wireless radio frequency coexistence,” 2007.
- [8] I. Howitt and J. A. Gutierrez, “IEEE 802.15.4 low rate - wireless personal area network coexistence issues,” in *WCNC*, 2003.
- [9] 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.
- [10] S. Y. Shin, H. S. Park, S. Choi, and W. H. Kwon, “Packet error rate analysis of IEEE 802.15.4 under IEEE 802.11b interference,” in *WWIC*, 2005.
- [11] 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.
- [12] S. Pollin, M. Ergen, M. Timmers, A. Dejonghe, L. van der Perre, F. Cathoor, I. Moerman, and A. Bahai, “Distributed cognitive coexistence of 802.15.4 with 802.11,” in *Cognitive Radio Oriented Wireless Networks and Communications*, 2006.
- [13] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan, “Understanding and mitigating the impact of RF interference on 802.11 networks,” in *Sigcomm*, 2007.
- [14] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis, “An empirical study of low power wireless,” in *ACM Transactions on Sensor Networks*, 2010.
- [15] C.-J. M. Liang, B. Priyantha, J. Liu, and A. Terzis, “Surviving Wi-Fi interference in low power zigbee networks,” in *SenSys*, 2010.
- [16] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh, “White space networking with Wi-Fi like connectivity,” in *Sigcomm*, 2009.
- [17] D. Chen, S. Yin, Q. Zhang, M. Liu, and S. Li, “Mining spectrum usage data: a large-scale spectrum measurement study,” in *Mobicom*, 2009.
- [18] K. Papagiannaki, M. Yarvis, and W. S. Conner, “Experimental characterization of home wireless networks and design implications,” in *INFOCOM*, 2006.
- [19] J. Ortiz and D. Culler, “Multichannel reliability assessment in real world WSNs,” in *IPSN*, 2010.
- [20] 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.
- [21] R. C. Shah, L. Nachman, and C.-y. Wan, “On the performance of bluetooth and ieee 802.15.4 radios in a body area network,” in *BodyNets*, 2008.
- [22] Wi-Spy, <http://www.metageek.net/>.
- [23] 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>
- [24] [Online]. Available: <http://www.metageek.net/recordings>
- [25] J. Polastre, R. Szewczyk, and D. Culler, “Telos: Enabling ultra-low power wireless research,” in *IPSN*, 2005.
- [26] 2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver, Texas Instruments.
- [27] <http://www.tinyos.net/>.
- [28] J. Zhao and R. Govindan, “Understanding packet delivery performance in dense wireless sensor networks,” in *Sensys*, 2003.
- [29] Specification of the Bluetooth System, Version 4.0.
- [30] Technical Overview of Time Synchronized Mesh Protocol, White Paper, <http://www.dustnetworks.com>.
- [31] “IEEE 802.15.4e WPAN task group.” [Online]. Available: <http://www.ieee802.org/15/pub/TG4e.html>
- [32] S. M. Stigler, “Francis Galton’s account of the invention of correlation,” *Statistical Science*, vol. 4, no. 2, 1989.