

TIIM: Technology-Independent Interference Mitigation for Low-power Wireless Networks

Anwar Hithnawi, Hossein Shafagh
Department of Computer Science
ETH Zurich, Switzerland
{hithnawi, shafagh}@inf.ethz.ch

Simon Duquennoy
SICS Swedish ICT AB
Kista, Sweden
simonduq@sics.se

ABSTRACT

The rise of heterogeneity in wireless technologies operating in the unlicensed bands has been shown to adversely affect the performance of low-power wireless networks. Cross-Technology Interference (CTI) is highly uncertain and raises the need for agile methods that assess the channel conditions and apply actions maximizing communication success. In this paper, we present TIIM, a lightweight *Technology-Independent Interference Mitigation* solution that detects, quantifies, and reacts to CTI in realtime. TIIM employs a lightweight machine learning classifier to (i) decide whether communication is viable over the interfered link, (ii) characterize the ambient conditions and apply the best coexistence mitigation strategy. We present an in-depth experimental characterization of the effect of CTI on 802.15.4 links, which motivated and influenced the design of TIIM. Our evaluation shows that TIIM, while exposed to extensive and heterogeneous interference, can achieve a total PRR improvement of 30% with an additional transmission overhead of 5.6%.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless communication*

Keywords

Interference Mitigation, 802.15.4, Cross-Technology Interference, Machine Learning

1. INTRODUCTION

The ubiquitous and tetherless access to information enabled by the wireless medium, and recent advances in wireless communication, have led to a rapid surge in wireless data traffic congesting the unlicensed bands. This traffic is generated from heterogeneous radios that follow different protocols and communication primitives. A few examples include WiFi, Bluetooth, IEEE 802.15.4, 2.4 GHz cordless phones, surveillance cameras, game controllers, and 2.4 GHz

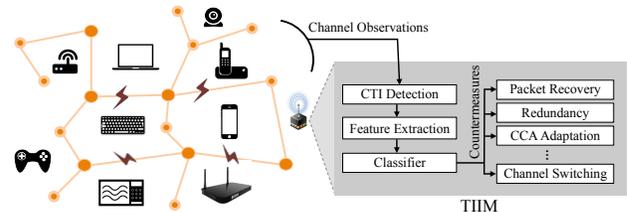


Figure 1: Low-power communication links suffer from crowded ISM bands. TIIM dynamically applies interference mitigation measures specific to channel conditions.

RFID. This rises a unique set of communication challenges, notably co-existence, cross-technology interference, and fairness amidst high uncertainty and scarcity of interference-free channels (see Figure 2).

RF interference is a classical and fundamental communication problem, which has traditionally been managed in the unlicensed bands through spectrum fragmentation, where different protocols would largely aim at communicating over non-overlapping segments of the spectrum, and employing carrier sense to avoid interfering with active communication. However, these solutions were designed under the assumption of low spatial density of radios and presence of few radio technologies that respect these customary rules. These assumptions do not hold anymore. As a consequence, CTI is emerging as a major problem in the unlicensed bands [3, 4, 10, 28]. This has motivated researchers to work on novel solutions to address the CTI problem by considering both link layer (power-control, rate-control, scheduling, packet and error recovery [9, 15, 20]) and physical layer solutions (interference cancelation [7, 28], spectrum adaptation [18, 32]). The most widely adopted solution is to avoid interfered frequencies by employing spectrum sensing to identify interference-free channels.

The lack of interference-free channels led researchers to work on sophisticated spectrum adaptation schemes and develop novel classification approaches that provide information about the interference source [1, 17, 29, 31]. It has been shown that when the interference source is known, specialized mitigation approaches can improve the network performance. Almost every radio communication technology has hidden repeating patterns that form a signature for that particular technology. Researchers explored these properties to build interference classification tools that can report on the root source of the interference problem. This approach yields interesting results but is bound to

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
IPSN '15, April 14 - 16, 2015, Seattle, WA, USA.
Copyright 2015 ACM 978-1-4503-3475-4/15/04 ...\$15.00.
<http://dx.doi.org/10.1145/2737095.2737104>

a fixed set of interfering technologies that are known at design time.

Looking at the design space of spectrum co-existence solutions, and based on the observations we made while empirically studying¹ the impact and the interaction patterns of CTI on low-power wireless networks (i.e., 802.15.4), we suggest the consideration of the following aspects when addressing the CTI problem: (i) PHY signals not only encode bits, but also contain rich information about the ambiance, which is particularly enlightening in case of interference. (ii) The mere presence of interference is not always harmful, metrics such as energy detection can falsely trigger communication to back off and introduce unnecessary deferrals. Thus, it is important to consider measures that can better quantify the impact of CTI. (iii) Given the scarcity of the frequencies allocated to wireless networks, it is desirable to allow concurrent transmissions that potentially can be correctly recovered. (iv) There is no one-size-fits-all solution. The high degree of diversity in radio technologies results in different implications on the wireless link that need to be addressed with different strategies. (v) The impact of the same source of interference can quickly change due to mobility (e.g., interferer moves away) or due to change in the configuration (e.g., WiFi bit-rate, or application traffic pattern). Thus, frequent adaptation is required.

In this paper, we present TIIM, as illustrated in Figure 1, an adaptive interference mitigation system, that selects interference mitigation strategies directly based on measured medium properties, skipping the interference classification step. Hence, TIIM is independent of the interference technology it is combating. To this end, we train the system to detect interference patterns and map these to a link-layer interference mitigation strategy that works best for this particular pattern, regardless of the interference technology.

Contributions. This paper makes the following contributions:

- Characterization of Cross-Technology Interference. To motivate the need for and show the feasibility of our approach, we perform a detailed measurement study consisting of more than 2.5 million transmissions. We characterize the interaction patterns between a prevalent set of RF interferers and 802.15.4 communication links in an anechoic chamber and in office environment.
- Design and evaluation of TIIM, a system that detects, quantifies and reacts to CTI in realtime. TIIM’s design consists of the following steps: (i) exploring the feature space of Cross-Technology Interference, (ii) constructing a lightweight decision tree classifier that learns the conditions where particular countermeasures perform best and uses this knowledge to select countermeasures for unseen channel instances at runtime. In our evaluation, TIIM archives an average accuracy of 92.9% in inferring the correct countermeasure for the current channel, helping to increase packet reception rate with a controlled overhead.

In the rest of the paper, we elaborate on TIIM’s key intuition, empirically characterise and analyse CTI patterns on 802.15.4 links, and present TIIM’s design, followed by performance evaluation. We conclude the paper with

¹We make the collected traces for the empirical CTI study and the modified radio drivers available at http://www.inf.ethz.ch/~hanwar/CTI_Study_Traces/

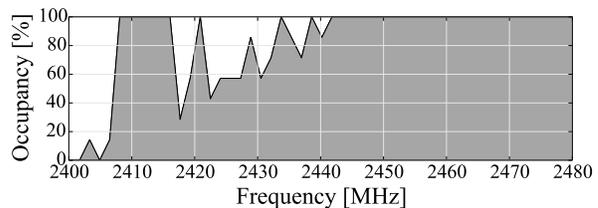


Figure 2: Averaged channel occupancy in the 2.4 GHz band over one week (26.-31. August 2013). Data from Microsoft Spectrum Observatory in an enterprise building in Brussels, Belgium.

a discussion of the limitations and opportunities of designing CTI-aware and adaptive link-layer protocols.

2. BACKGROUND

In this section, we cover some background on wireless interference, unlicensed bands, and IEEE 802.15.4.

2.1 Cross-Technology Interference in a Nutshell

The broadcast nature of the wireless medium makes it inherently vulnerable to interference from spatially close concurrent transmissions that overlap in time and frequency. This can consequently reduce or even prevent completely the ability of receivers to decode information from signals. Wireless communication can be subject to disturbance by interference from intra-technology, cross-technology, or noise sources. Wireless technologies strive to avoid interference and typically apply a set of mechanisms to achieve fairness and reduce interference within the same technology (e.g., reserve the medium, allocate channels, and probe for idleness). However, most protocols are not designed with coexistence in mind. This is mainly because of the infeasibility of interference coordination due to the absence of communication means between these diverse technologies (i.e., speak different PHY protocols). Consequently, CTI is emerging as a major problem in the unlicensed bands [3, 4, 10, 28].

The unlicensed bands are small segments of the radio spectrum that were reserved internationally for the use of RF energy for *Industrial, Scientific, and Medical* (ISM) purposes and have been widely utilized for unlicensed short-range wireless radios. In the following, we spot the light on certain communication properties that are adopted by many radios, and make it particularly challenging for low-power technologies such as 802.15.4 to coexist: (i) wide-band: many devices transmit in frequency bands significantly wider than 802.15.4. For example, to cope with the high demand of throughput over WiFi, the recent amendments of 802.11 allow the configuration of 40 MHz-wide channels in the 2.4 GHz band. As another example, microwave ovens affect almost 50% of the 2.4 GHz band. (ii) high-power: today’s high-power interferers in the unlicensed bands pose a serious threat to 802.15.4 networks, as they can cause 802.15.4 links to experience complete loss of connectivity. Although the FCC lightly regulates this aspect by setting an upper limit of 30 dBm for transmit-power in the unlicensed bands, energy leaks from microwave ovens can reach up to 60 dBm. This is significantly higher than typical output power of 802.15.4 radios which is 0 dBm.

We briefly recall how signals are transmitted and received over the wireless channel. The following assumes

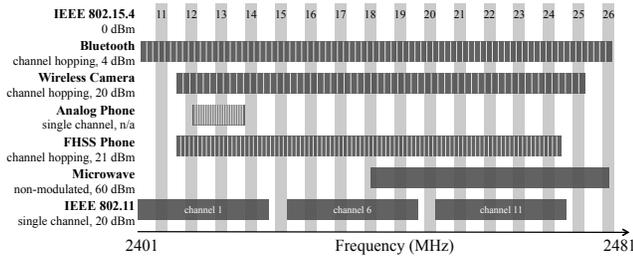


Figure 3: RF channels of IEEE 802.15.4 and the selected set of prevalent RF interferers in the 2.4 GHz ISM band studied in this paper.

Minimum-Shift Keying (MSK) signal. Note that we omitted unnecessary details to simplify this communication primer. Radios convert binary data into modulated signals. These signals are generally represented as a discrete and complex function:

$$s[n] = A_s e^{i\theta_s[n]}, \quad (1)$$

where A_s is the amplitude of the transmitted sample n , $\theta_s[n]$ is its phase. Note since MSK embeds all the information in the phase, A_s is constant for all samples, hence the signal carries constant energy.

After the signal traverses the channel, the receiver receives:

$$y[n] = Hs[n], \quad (2)$$

where H is a complex number that approximates the effect of the wireless channel (attenuation and phase) from the transmitter’s antenna to the receiver’s antenna [16]. In the presence of an unknown interferer, i.e., the desired signal interfered with unknown signal, the signal at the receiver is represented as follows:

$$y[n] = Hs[n] + H'i[n], \quad (3)$$

where $i[n]$ is the interfering signal and H' is the approximation of the channel from the interferer’s transmitter to the receiver. When these two signals interfere, their energies add up:

$$E[|y[n]|^2] = E[|Hs[n] + H'i[n]|^2] \quad (4)$$

This insight on additive energy of interfering signals highlights the ambient information the signal carries along and that can assist in detecting interference and localizing interfered symbols within interfered packets.

2.2 IEEE 802.15.4

IEEE 802.15.4 PHY. For devices operating in the 2.4 GHz band, the IEEE 802.15.4 standard [5] defines the *Offset Quadrature Phase-Shift Keying* (O-QPSK) modulation scheme with a half pulse shaping. The transmitter’s radio transforms binary data into modulated analog signals by adapting spreading and modulation. The data is first grouped into 4-bit symbols, which are mapped to one of 16 *Pseudo-random Noise* (PN) sequences that are 32-bit long. Each bit in a PN sequence is then modulated to the carrier signal using O-QPSK, which is equivalent to MSK. For demodulation, the receiver’s radio converts each half-sine pulse signal into a chip. The radio performs soft decisions at the chip level [2] providing PN sequences. The de-spreading is

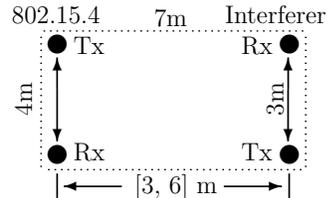


Figure 4: Controlled experiments setup for CTI characterization in an anechoic chamber.

performed by mapping the PN sequence to the symbol with the highest correlation.

802.15.4 transmission occurs in one of the 27 non-overlapping allocated channels. Out of these, 16 (from 11 to 26) are allocated in the 2.4 GHz band, each with 2 MHz bandwidth and 5 MHz channel spacing (see Figure 3). The remaining 11 channels are allocated in sub-GHz bands.

IEEE 802.15.4 MAC. IEEE 802.15.4 has several MAC-layer protocols, defined both in the original standard and its 2012 amendment 802.15.4e. In its simplest form, 802.15.4 employs contention-based CSMA/CA communication. Before a node starts transmission, it waits for a random back-off period to assure that the medium is idle. For this, it relies on *Clear Channel Assessment* (CCA). If CCA declares the channel to be free, the transmission is carried out, otherwise it defers the transmission for a random backoff time. For data verification, the receiver computes a 16-bit CRC check over the payload of a received packet. It discards packets that do not pass the check and accordingly withholds the ACK transmission.

More sophisticated 802.15.4e MAC layers such as TSCH, CSL, or RIT also employ acknowledged transmissions, exponential backoff, and (optionally for TSCH) CCA.

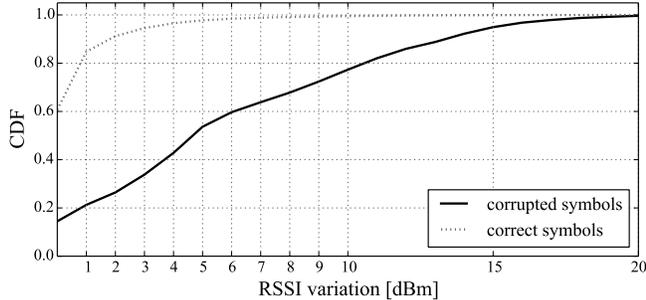
3. TIIM OVERVIEW

We now present a high-level overview of our interference mitigation system TIIM. The intuition underlying TIIM is that each interference mitigation approach works well under specific channel assumptions of error patterns, such as error rate, signal to interference ratio, or occupancy level. Each interference instance, independent of the technology, leaves a particular signature in the channel that shapes the channel properties in a unique way. The goal of TIIM is to automatically select at runtime the most effective mitigation strategy for the current interfered channel and get the best out of the interfered channel in the crowded spectrum.

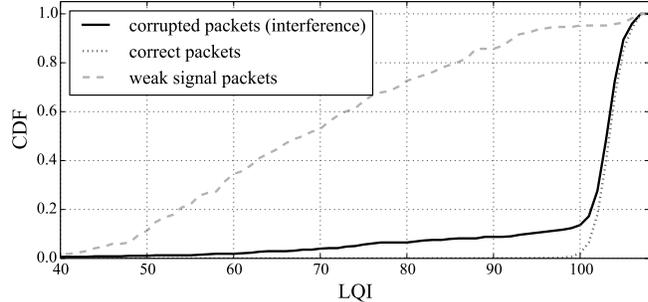
The core component of TIIM is a lightweight decision tree classifier that is trained to learn under which channel conditions (i.e., signatures) a particular mitigation scheme empirically works the best. TIIM uses the decision tree online to predict the best channel mitigation strategy for yet unseen instances of an interfered channel. TIIM’s design consists of the following steps:

CTI Characterization. The first step in the design of TIIM is to characterize *Cross-Technology Interference* (CTI). We run experiments where we expose a 802.15.4 communication link to various types of interferers, both in an anechoic room and in office environment. We collect channel properties and communication statistics at high frequency. Section 4 presents the results of this characterization in details.

Learning Phase. We simulate every considered mitigation



(a) CDF of RSSI variations for correct and corrupted symbols due to interference.



(b) CDF of per-packet Link Quality Indicator (LQI).

Figure 5: Observations from our traces on interference detection. (a) shows that interfered packets often experience high RSSI variations. (b) shows the clear distinction between LQI of corrupted packets due to interference and those corrupted due to weak signal.

strategy against the traces collected in the characterization step, and compute both their gain and cost. In this phase of supervised learning, the decision tree classifier learns for each channel feature which particular mitigation strategy scores highest.

Runtime. At runtime, nodes monitor their current channel condition mostly through signal strength sampling at high frequency during packet reception. Whenever interference is detected, they feed the decision tree with channel statistics as input and obtain a decision about the mitigation strategy to employ.

Our work is inspired by the core idea behind interference classification approaches, such as SoNIC [17] and Airshark [31], which use measurement samples drawn from commodity hardware to detect the type of interference source. While these approaches can provide useful information on how to potentially mitigate CTI, they can not combat CTI autonomously. They require either user intervention or querying a central entity that maintains the mapping of an interference source to the corresponding countermeasure. It is yet not clear how such approaches can be utilized in an automated way. TIIM departs from the above in that it skips the interference classification step. Instead, it infers the best mitigation strategy from channel properties directly, independent of the technology causing the interference.

4. CHARACTERIZING CROSS-TECHNOLOGY INTERFERENCE

In this section, we characterize how arbitrary interfering signals interact with 802.15.4 communication and focus on identifying distinct features of interfered channels and packets in practical systems that could assist in: (i) detecting and quantifying interference, (ii) pinpointing the viability of opportunistic transmission in interfered channels, (iii) selecting a countermeasure that works best for the current underlying interference patterns.

4.1 Controlled Experiments Setup

We run our experiments in an anechoic chamber, in order to have full control on the sources of errors, type of channel distortions, and to isolate the impact of surrounding interference sources. We consider a simple network setup, as depicted in Figure 4, which consists of one transmitter and one receiver for both 802.15.4 and the considered interfering

technology, i.e., a pair of 802.15.4 nodes and a pair of interferer nodes. We base our sender and receiver applications on Contiki OS [8] and directly interface them to the node’s radio driver. We consider different traffic patterns and different configurations of packet length and transmission power. **Interfering Technologies.** We focus on a set of interferer technologies that are prevalent in today’s environments. Our considered set consists of low/high power, narrow/wide band, analog/digital, channel hopping/fixed frequency, and CSMA/non-CSMA interferers. This represents common underlying properties adopted by most radio technologies. Figure 3 summarizes the features of the considered RF technologies in our study. In the following, we briefly highlight some of their properties.

- *IEEE 802.11.* We create WiFi interference using a Netgear WNR3500L router and a laptop that supports IEEE 802.11 b/g/n in the 2.4 GHz ISM band. We use the network tool `iperf` [19] to generate saturated TCP traffic and non-saturated UDP traffic that resemble file download and VoIP, respectively.

- *Bluetooth.* To evaluate the interference generated by Bluetooth on 802.15.4, we use two HTC Desire phones transferring a large file. Bluetooth uses the adaptive frequency hopping technique across 79 MHz of bandwidth in the 2.4 GHz ISM band. The hopping occurs at a rate of 1600 hops/s, hence it occupies a 1 MHz channel for 625 μ s.

- *Digital cordless phone (FHSS).* We experiment with the Uniden DCT6485-3HS cordless handset system. The phone base and handset communicate using frequency hopping over 90 channels of 800 kHz width in the range [2407.5 - 2472] MHz.

- *Analog cordless phone.* We experiment with the Vtech GZ2456 cordless handset system. The phone base transmits in the 900 MHz band and receives in the 2.4 GHz band. The phone handset accordingly transmits and receives using the reverse order of frequency ranges.

- *Wireless Camera.* We use the Philips SCD 603 digital video baby monitor. It comprises a 2.4 GHz wireless camera and a wireless video receiver. The wireless camera uses frequency hopping over 61 channels, where each channel has a width of 1.125 MHz.

- *Microwave oven.* We use a residential microwave oven, the Clatronic MWG 758. We heat a cup of water in the microwave to emulate an interference typical to that emitted by these appliances.

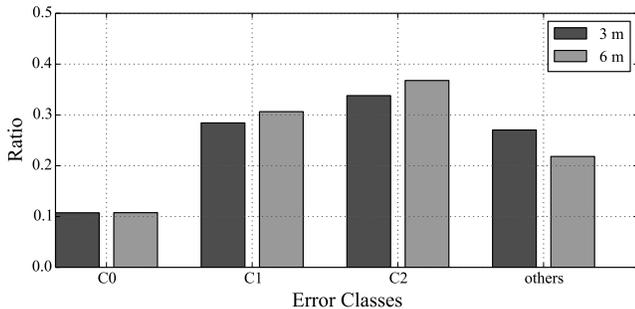


Figure 6: Error classes from CCA-enabled traces, averaged over all sources of interference at distance 3 m, and respectively at 6 m. C0: corrupted symbols larger than 1/3 of the payload’s length, C1: few (1-12) corrupted symbols (suitable for FEC), C2: single error burst not longer than 1/3 of the payload length (suitable for packet merging).

4.2 Interference Detection

We investigate how 802.15.4 radios can detect whether a received packet has been subject to interference or not. The 802.15.4 transmitted signal encodes information in phase rather than amplitude, therefore its amplitude (i.e., energy) is constant within the coherence time. When two signals interfere, their energies add up and this results into variations in the energy level of the received signal. Confined with the PHY and link layers of the OSI stack, the standard design of off-the-shelf radios treats the PHY layer as a black box that provides decoded bits (i.e., MAC layer PDU) and a limited PHY information and deprives the access to signal level information.

For detecting interference, we explore two different possibilities: (i) Capturing energy variations during packet reception by sampling the radio’s RSSI register. We modified the CC2420 driver in Contiki [8] to capture RSSI values at a rate of one sample per symbol (i.e., one reading each 16 μ s). The sampling is performed from the *Start of Frame Delimiter* (SFD) to the last symbol of the packet. Figure 5(a) depicts the CDF of the RSSI variations of interfered erroneous packets and correct packets as observed in our traces. We see a clear correlation between RSSI variation and packet corruption. For instance, 90% of the non-interfered packets have a variation under 2 dBm, while more than 70% of the interfered packets experienced variations higher than 2 dBm. In case the variance of RSSI is greater than a threshold (default to 2 dBm), our system recognizes the received packet as an interfered packet.

(ii) The second alternative is to exploit how per-packet channel metrics are computed in off-the-shelf radios. The Link Quality Indicator (LQI) is confined to average the readings of few symbols (in case of CC2420 [2], over the first 8 symbols following the SFD field). Consequently, it cannot capture the spike in the received power due to interference, as long as the interference spike does not fall within these few symbols. As a result, per packet LQI readings are not reflecting the impact of interference. In fact, LQI provides indications of a good and stable channel in most of the interfered packets. Similar observations have been reported for 802.11 in [30]. Figure 5(b) highlights this insight with respect to LQI considering weak signals and interfered signals. It shows that 85% of corrupted packets due to weak

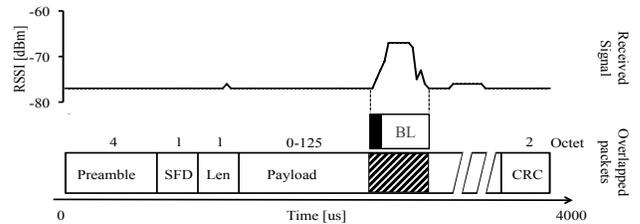


Figure 7: An example of the correlation of signal variations with the symbol errors within a received frame. TIIM makes use of this knowledge to request the retransmission of only the co-located corrupted symbols.

signal have an LQI of about 90 or less, whereas only 10% of packets suffering interference have an LQI of 85 or less. To detect interference, the system can monitor the LQI of received corrupted packets. Frequent erroneous packets with good link metric (LQI > 90) could be used to detect interfered packets.

Interference detection is a key component for addressing interference. Interference mitigation schemes come with an overhead that should be avoided in the absence of interference. We exploit the LQI approach of interference detection to trigger the recovery phase in TIIM, and the RSSI variations to classify interfered packets in the recovery phase.

4.3 Exploring the Feature Space

In the following, we explore possible features of communication links that can potentially mirror occupancy and error patterns in interfered channels.

Persistency and Occupancy: The primary spectrum usage modalities that exist in the unlicensed bands are comprised of: (i) Persistent. Technologies adopting this form operate in dedicated frequency bands and generate static energy, thus monopolizing the medium completely. Legacy analog devices adopt this modality. This form of interference causes a complete loss of connectivity for the interfered low-power nodes [10], primarily because the continual energy emission prevents the carrier sense from declaring the channel to be free.

(ii) Non-persistent. Technologies adopting this form operate either in dedicated frequency bands and generate traffic with time varying load, or exploit frequency diversity and hop across the spectrum. They exhibit a time-variant ON and OFF pattern of energy emission due to underlying communication patterns, such as frequency hopping, continual inter-frame spacing (e.g., SIFS, DIFS), back-off slots, and varying load, or periodic ON and OFF cycles of noise radiation, as for the microwave oven. This translates to exchanged packets being either correctly received (i.e., the shorter the transmission time, the higher the chances) or partially overlap with the interfering signal leading to packet loss or corruption. It is clear that even at this level, realizing whether the interference is persistent or non-persistent should be followed by adopting different mitigation schemes.

Properties of Corrupted Packets: Features that can represent an estimate of the error bit-rate and error patterns in packets can serve as an important meta-information for error recovery mechanisms. We define a nominal feature that can take four values representing different classes of error patterns: C₀: error rate > 33%, C₁: few corrupted

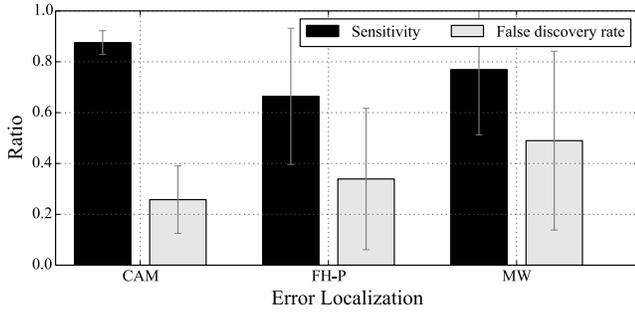


Figure 8: Error localization for CCA-disabled traces from anechoic chamber. TIIM uses fine-grained RSSI sampling to localize the area of corrupted symbols in the payload. Sensitivity is a metric indicating the ratio of correctly detected corrupted symbols, whereas the false discovery ratio shows the ratio of correct symbols among the detected symbols. Different technologies experience varying bit-error localization performance.

symbols, between 1 and 12 corrupted symbols, C_2 : one error burst², not larger than 33% of the payload, and *others*. Figure 6 depicts the ratio of each of these error classes aggregated over the traces from the controlled experiment.

To compute this feature, we can either rely on retransmissions to identify the corrupted parts of a packet or, as illustrated in Figure 7, by analyzing energy surge bursts in the received sampled RSSI as these surges correlate to corrupted parts of a packet in case of missing or corrupted retransmissions.

Interference Estimation: The mere presence of interference is not always harmful. Hence, finding metrics that can better quantify the actual impact of interference can largely influence the way we address interference and potentially increase spectral efficiency. We define a metric that considers the reception status of packets during an observation window. Assuming a total number n of packets transmitted during the observation window, n_i is the number of interfered packets, n_s number of corrupted packets due to other channel impairments (e.g. weak signal), and n_l number of lost packets. The estimated CTI impact is: $\text{estimated_interference} = (n_i + n_l)/n$. As we cannot clarify the source for lost packets, i.e., packets that the receiver failed to detect their preamble, we took a conservative approach and accounted them as impacted by interference.

4.4 Countermeasures

In the following, we briefly cover a set of link-layer mitigation schemes that we consider in the design of TIIM. These mitigation schemes have been proposed and evaluated in the literature in the context of increasing the resilience of 802.15.4 against interference. TIIM is not bounded to this set of countermeasures, and can be trained and extended with further countermeasures. In this prototype, TIIM is trained to select one of the following mechanisms or a combination of them:

- **Reed-Solomon Forward Error Correction (FEC):** The Reed-Solomon (RS) code is a block-based error correcting code that is particularly effective at correcting burst errors.

²We define an error burst as co-located corrupted symbols with max. up to 5 correct symbols in between.

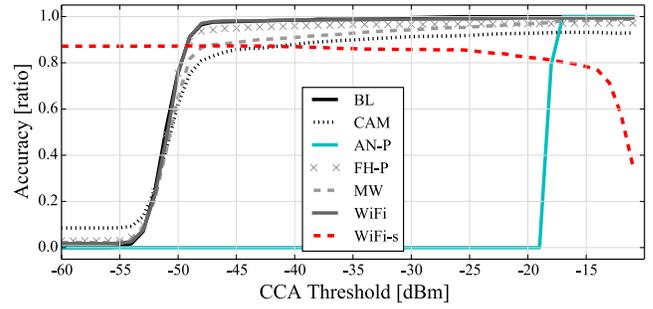


Figure 9: Accuracy of different CCA thresholds for anechoic room traces with interferers at distance 6 m. Accuracy is defined as $\frac{TP+TN}{\#\text{transmissions}}$. True Positives (TP) are the cases where the channel was free and the transmission successful. True Negatives (TN) indicate cases where the channel was busy and the transmission was lost or corrupted. A fixed CCA threshold does not serve well under all channel conditions.

RS code divides a message m into n blocks of defined size and adds extra redundant parity of t blocks to the message. RS code can correct up to $t/2$ and detect up to t block errors. The overhead cost of RS code is constant, both correct and corrupted packets bear the redundancy overhead. RS code works well for error patterns that fall under the recovery capacity of the parity check. In TIIM, we use 12 Bytes of parity.

- **RSSI-based Packet Recovery (PM):** In the presence of interference, a sender often has to retransmit packets several times until the receiver decodes a correct copy. Partial packet recovery [11, 20, 21] and packet merging [9, 27] aim at reducing the amount of redundantly in received data in this process. As depicted in Figure 6, about 40% of the corrupted packets we witnessed in the controlled study traces, are of class C_2 (it varies for different interferers). This class of errors stands for a single error burst in the corrupted packet, where the rest of the packet is error-free and would potentially gain from packet recovery mechanisms.

In TIIM, the receiver selectively requests the segments of a packet where symbols are likely to be corrupted [20]. The segments are identified based on hints from the PHY layer, namely surges in RSSI readings. Figure 7 illustrates an example of RSSI variations (from our controlled experiment) during the reception of an 802.15.4 packet. It shows how the surge in RSSI corresponds to the error location due to interference by Bluetooth. Figure 8 depicts the accuracy we achieved in localizing corrupted symbols by utilizing sampled RSSI in our traces.

- **Adaptive CCA Thresholding (no-CCA):** 802.15.4 networks generally use carrier sensing before transmission, in order to reduce collisions. We evaluate the efficiency of carrier sense in 802.15.4 radios under CTI, from our saturated CCA disabled experiment traces. We investigate the relation between the sampled energy at the sender before transmission and the actual success or failure of packet transmissions.

Figure 9 summarizes the results of carrier sense efficiency experiments. The carrier sense works well in many scenarios, but can possibly result in false positives (i.e., channel free but transmission either corrupted or lost) and false negatives (i.e., channel busy but transmission successful). The following scenarios are worth looking at: (i) *Frequency Hopping*

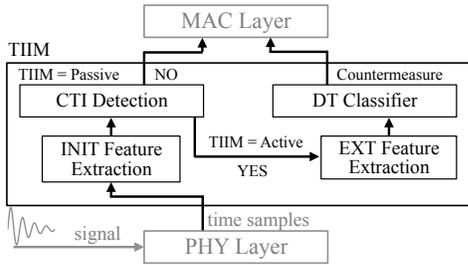


Figure 10: TIIM’s design integrated into the OSI model. TIIM remains passive while observing channel conditions, i.e., initial (INIT) features. Upon detection of interference it turns active, collects further channel metrics, i.e., extended (EXT) features, for a given time window, and inquires the *Decision Tree* (DT) classifier for a countermeasure.

(FH) interferers: In the presence of frequency hopping interferers, carrier sense is not effective. FH interferers do not react to 802.15.4 transmissions. Accordingly the chances a packet encounters corruption given channel is sensed free or occupied is the same. (ii) Analog interferers: In the case of analog phone at distance 6 m, although the energy level in the channel is high, PRR is barely affected from interference. Hence, carrier sense causes unnecessary deferrals. In Figure 9, we observe that for the analog interferer (AN-P) which transmits with high power, any threshold below -18 dBm results in 100% false negatives. Thus, an adaptive CCA scheme that can assess and prevent harmful concurrent transmissions, while allowing safe concurrent transmissions, could largely enhance spectral efficiency. TIIM follows rather a radical strategy and recommends *no-CCA* when it detects that using CCA is causing high false negatives.

- *Channel Switching or No Action*: TIIM can infer that the interference in the channel is not harmful thus no action is required or that the interference in the channel is severe thus communication over this channel is not viable even with assistance of link-layer interference mitigation mechanisms. In such situations, it gives the recommendation of channel switching. A more informed decision about the selection of a suitable channel out the spectrum segment affected by the active interferer is being addressed in our future work.

5. TIIM ARCHITECTURE

So far we have concentrated on describing TIIM at a high-level and discussing design decisions and some essential empirical observations on the 802.15.4 interfered channels. We now provide an overview of TIIM’s components and its operation modes before detailing the classification algorithm and discussing TIIM’s integration into the system. All aspects of TIIM have been carefully chosen and designed with runtime and memory efficiency in mind. We focus to leverage the following four primary goals:

- (i) Improving *spectral efficiency* and *packet reception ratio* in the presence of CTI. (ii) *Compatibility*: we design TIIM such that it can be implemented as software modifications on top of commodity hardware. (iii) *IEEE 802.15.4 PHY compatibility*: compliance to the existing standards which allows seamless integration into existing systems. (iv) Supporting *heterogeneous CTI patterns*, oblivious to interference source type, distance, or configuration.

5.1 Modes of Operation

TIIM operates in two modes: passive and active. Having two modes allows to avoid imposing additional computational overhead to interference free communication. The system runs mainly in the passive mode which monitors a number of channel metrics, namely, the *Link Quality Indication* (LQI) value of corrupted frames at the receiver side, and packet losses and CSMA deferrals at the sender side. TIIM detects harmful interference using a simple threshold mechanism (τ_{active}) (see Section 4.2). Upon detecting harmful interference, TIIM switches to the active mode. An overview of TIIM is shown in Figure 10.

TIIM’s active mode operates on a window of communication events W_{active} . In our experiments, we use a window length of 5 seconds. We found this to be a good tradeoff between time to react to interference and confidence of selection. Note that the time window length can be adapted to the level of activity in the channel and its optimization is out of scope of this work.

During active mode, TIIM continuously processes time samples collected during packet reception and communication statistics for the time of the observation window, and computes a set of features as described in Section 4.3. Then, the node feeds the decision tree classifier with the features and triggers the mitigation strategy inferred by the classifier. Finally, TIIM switches to the passive mode to monitor the performance of the activated mitigation strategy.

5.2 Inferring Countermeasures

To infer the best (set of) countermeasure(s) for a given input feature set (see Section 4.3 for the feature space discussion), we use a supervised learning approach and construct a decision tree classifier. At runtime, the trained classifier assigns unseen instances of interfered channel, i.e., observation window W_{active} , to one of the 6 classes. The classes represent the set of mitigation strategies we consider in this prototype of TIIM (see Section 4.4): (i) *no-CCA*: disables the carrier sense. (ii) *FEC*: applies forward error correction with fixed block of 12 Bytes of redundancy. (iii) *nC-FEC*, applies forward error correction and disables the carrier sense simultaneously. (iv) *PM*: applies RSSI-based packet merging. (v) *nC-PM*: applies RSSI-based packet merging and disables the carrier sense simultaneously. (vi) *no-action*: takes no action as the potential gain of countermeasures is not significant or interference is not harmful.

5.2.1 Feature Selection

In Section 4.3, we empirically explored the feature space of the CTI classification problem and highlighted the set of features that best describe the problem. We considered packet specific features, spectrum specific features, and communication link features.

The initial set of features consisted of 25 features derived from domain specific knowledge. Training the classifier using all obtainable features is not a good practice as this can result in overfitting.

To reduce the feature set, we use the following selection techniques: First, we evaluate the level of intercorrelation among features and exclude redundant features. We identify a subset of 5 features that are uncorrelated among each other, yet correlated in predicting the same class. Then, to increase the accuracy of the initial subset, we apply exhaustive search to evaluate all possible remaining subsets in

Feature	Description	Purpose
(1) Occupancy level (2) Duty cycle	number of busy samples / total number of samples {T,F}, patterns of x consecutive idle samples	Temporal behaviour of interferers
(3) Energy span during packet reception (4) Energy level during packet reception (5) RSSI regularity during packet reception	range(RSSI _{normalized}) median(range(RSSI _{normalized})) weighted-average(mode(RSSI _{normalized}))	Capture RSSI temporal characteristics
(6) Packet corruption rate (7) Packet loss rate	ratio(P _{erc}) ratio(P _l)	Detect and quantify harmful interference
(8) Packet length (9) Error rate (10) Error burstiness	average(packet_length) ratio(err_typ == C1) ratio(err_typ == C2)	Capture error characteristics
(11) Energy perception per packet (12) Energy perception level per packet (13) Backoffs	range(CCA) median(CCA) ratio(CCA_deferral)	Detect unnecessary deferrals

Table 1: Features utilized by TIIM. These features are calculated over an observation window W_{active} . To remain environment-agnostic, each packets’s fine-grained RSSI series are normalized. Features 1-2 can either be induced from RSSI series of multiple packets, or as stand-alone sampling of the medium for a short period of time. Features 3-5, and 12-13 are collected during packet receptions, and aggregated in a representative way. Features 6-8, and 13 are communication statistics. Features 9-10 are based on per packet RSSI computations.

combination with these 5 fixed selected features. We benchmark all features that contribute to decision trees with high accuracy and select those with the highest rank. Table 1 summarizes the final set of features we use to train our classifier and which are used by TIIM during runtime.

5.2.2 Data Labeling and Ground Truth

We divide our dataset into training and test sets. Each sample in the dataset (x_i, y_i) represents the feature vector x_i that describes a window of communication events in an instance of interfered link, and y_i the corresponding label which indicates the best mitigation strategy for this instance. Using the training set, the supervised learning algorithm aims at constructing a good model that learns the complex pattern in the training set to predict class y' for an unseen feature vector x' .

In Section 5.2.1, we elaborated on how to construct and calculate the feature vectors $\{x_1, \dots, x_N\}$ (see Table 1 for complete list of features). Now we discuss how we automatically label the samples in our dataset. To label our dataset for each window of observations, we simulate the outcome of each of the considered mitigation strategies. This yields for each instance a corresponding gain (i.e., PRR increase), and cost (i.e., communication overhead). The labeling algorithm quantifies the benefit of each countermeasure and selects an optimal countermeasure $A \in \{no-CCA, FEC, nC-FEC, PM, nC-PM, no-action\}$ that achieves the highest gain-cost balance, as defined by application requirements.

The application expresses its requirement through a gain $g()$ maximization and cost $c()$ minimization equation: $f(A) = g(A) - c(A) \times \alpha$, where the best countermeasure is the one with the highest $f(A)$. The configuration parameter α defines the weight of the cost. Note that $g(A)$ and $c(A)$ are normalized, and both are in the range $[0, 1]$. With $\alpha = 0$ the cost is not at all considered, which results in the highest possible total gain. The higher α , the more the algorithm emphasizes on minimizing the total cost. In our experiments, we consider $\alpha = 0.5$ which is a good tradeoff between total cost and total gain.

5.2.3 Decision Tree Classifier

TIIM’s classifier is based on decision tree (DT) [26] classification model. The model can be seen as sequential binary

decisions, corresponding to traversing binary trees. Decision trees consist of inner nodes (i.e., split nodes) and leaf nodes (i.e., classes). The split functions can be seen as questions that add incrementally to the certainty of the correct class.

Each node evaluates a computationally inexpensive function on one of the input features and as a result forwards the currently evaluated data recursively down in the tree until the data reaches one of the leaf nodes.

Albeit decision tree classifiers are not necessarily the best classifiers in terms of accuracy, they are relatively efficient in terms of computational and memory overheads; with careful optimization they can run on severely constrained devices. We use the C5.0 algorithm [6]. Our tree consists of 200 leaves in case trained by anechoic chamber traces and 300 in case of office traces.

6. EVALUATION

In this section, we present the experimental evaluation of TIIM. We begin in Section 6.1 by describing the experimental setup and briefly describe the trace collection methodology. Then, in Section 6.2, we elaborate on the accuracy of the decision tree classifier in inferring the correct countermeasures. In Section 6.3, we perform trace-driven simulation to demonstrate the prospective gain of applying TIIM compared with the gain of applying fixed mitigation strategies, followed by evaluating the system gain online.

6.1 Experimental Settings

Up to this point, the analysis has been carried out in an anechoic chamber, an environment that is shielded from external radio interference and diminishes multipath propagation effects. This allows us to recognize patterns and identify the impact of each of the considered interfering technologies. In the following, we address the evaluation of TIIM in a typical environment that incorporates the impact of external uncontrolled interferers, other channel impairments (e.g., multipath effect), and multiple sources of interference.

We performed all our experiments in ETH Zurich’s computer science building. Figure 11 shows the layout of the experimental setup. There are two stationary sensor nodes located in an office room with a line-of-sight link of 4.5 m.

We consider two types of experiments: (i) Single active interferer: in this run, we use each of the considered interferers

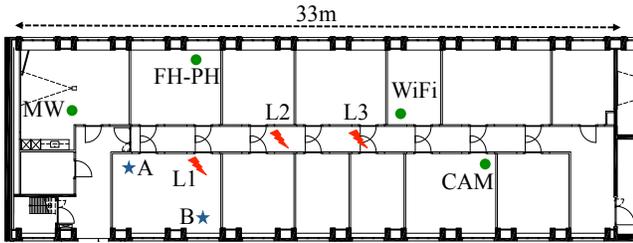


Figure 11: Layout of the office experiment setup. A and B are TelosB nodes located in an office with a line-of-sight link of 4.5 m. A is the sender and B the receiver. The interferers are located at locations L1 (3.5 m), L2 (6 m), and L3 (10 m), where L1 is within line-of-sight to A and B and the other two locations within non-line-of-sight. Circles indicate the location of our multiple interferer scenario.

to generate interference individually. The interferers are located in this run first at location L1 (3.5 m), then L2 (6 m), and L3 (10 m). Both locations L2 and L3 are in non-line-of-sight to the sensor nodes, while L1 is within the line-of-sight. (ii) Multiple active interferers: in this run, we consider interference generated from multiple sources running simultaneously. The positions of the interferers are highlighted as circles. During the experiments, the nodes were exposed to interference from various uncontrolled sources existing in the building. To mention some, the university’s WiFi network which is present on 802.11 channels 1, 6, and 11, Bluetooth mice and keyboards, and a small 802.15.4 heating control system deployed in the same floor.

Methodology: Our focus is to capture channel statistics over the 802.15.4 link between node A and node B. Node A sends short packets (20 byte) and long packets (100 byte)³ at 100 ms intervals to node B.

We first perform experiments without controlled interference, then with a single interferer activated at a time, and finally with activating multiple interferers. For the single interferer run, we consider all the interferers mentioned in Section 4.1. For the multiple interferers scenario, we consider a subset of these interferers, as highlighted in Figure 11. In all office environment experiments, the sender uses its maximum transmission power (0 dBm).

We instruct the sender to disable carrier sense and log the CCA value at the time of transmission. This allows us to perform trace-driven simulation for countermeasures involving carrier sense enabled and disabled. We instruct the receiver’s radio to pass packets with failed CRCs rather than discarding them to enable us processing erroneous packets. Moreover, the modified radio driver samples RSSI at a rate of 62.5 kHz during packet reception along logging other relevant PHY and link-layer metrics. Overall we capture 64 hours of extensive CTI experiments. We collect fine-grained channel and communication measurements to allow systematic evaluation and comparison of TIIM and the considered countermeasures under the exact adverse link dynamics.

Rather than detailing on recovery results per technology, we focus on discussing the adaptability of TIIM under dynamic and various interference implications with the goal

³Note that, the experiments carried out in the anechoic chamber considered more configurations of different packet lengths, transmission powers, and traffic patterns.

Predicted as	(a)	(b)	(c)	(d)	(e)	(f)
(a) no-action	95.0	3.9	0.0	0.2	0.5	0.4
(b) no-CCA	0.0	97.2	0.0	1.7	0.0	1.0
(c) FEC	11.4	4.5	68.2	11.4	2.3	2.3
(d) nC-FEC	0.1	16.8	0.2	80.8	0.0	2.2
(e) PM	31.1	12.2	1.1	0.0	41.1	14.4
(f) nC-PM	0.5	36.6	0.2	5.2	0.1	57.5

Table 2: Confusion matrix of the decision tree on the traces collected in office environment.

of maximizing the overall performance gain and minimizing the overall overhead cost.

6.2 TIIM Accuracy

We first discuss the performance of TIIM’s core component, the decision tree classifier, in inferring the correct countermeasure.

We use half of the data points in the office environment for training the classifier. The data points include various types of interference instances, as described above. The other half of the data points are used to evaluate the prediction accuracy. Table 2 shows the confusion matrix for DT classification. The decision tree achieves a mean classification accuracy of 92.9%.

One reason that TIIM’s accuracy is higher than traditional interference classification approaches [17] is that TIIM does not need to differentiate between radios causing similar channel signatures. TIIM aims at detecting channel signatures that can benefit from a certain countermeasure.

While TIIM achieves high accuracy in inferring the correct countermeasure for most of the classes, it performs poorly for *Packet Merging* (PM). This is mainly due to the low occurrence of incidents that were labeled as PM in our dataset, e.g., only 0.18% are labeled as PM in our office environment traces. Since PM is under-represented in our dataset, the DT could not learn well its characteristics.

The low number of PM instances in our traces is mainly due to some hardware-based inaccuracies affecting the localizing of corrupted symbols within a packet. Consequently, we were limited from achieving the full potential of PM.

In general, achieving high accuracy in localizing the positions of errors by solely relying on off-the-shelf radios is hard. For instance, we encountered in some cases of RSSI sampling non-consistent delays that are reflected in a slight drift between the RSSI surge position and the actual location of the error burst which affected the accuracy of our scheme of error localization. We believe that designing radios that allow better interfacing between PHY and upper communication layers can yield better performance for RSSI-based packet recovery schemes.

The consequence of prediction inaccuracies in our system is not necessarily high. TIIM has the chance to re-adjust its suggested countermeasure after the time window of observing the channel. Hence, the worst case scenario is that for the next window time, TIIM introduces a cost that does not yield any gain.

To verify whether the accuracy achieved by the classifier is not tied to the training environment, we train the classifier on our dataset from the anechoic chamber and evaluate it on the dataset from the office environment. The mean classification accuracy of 94.1% is even slightly higher than the accuracy achieved when the DT was trained and evaluated on data points from the same environment.

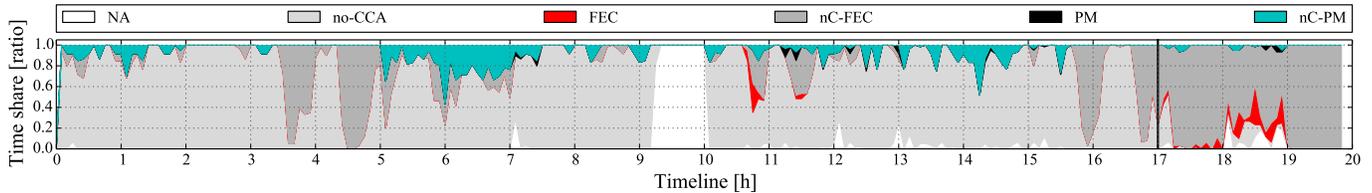


Figure 12: Offline performance of TIIM (resolution 5 min).

Countermeasure	PRR	Cost ratio
Base	0.566	0.000
no-CCA	0.863	0.106
FEC	0.572	0.198
nC-FEC	0.882	0.291
PM	0.567	0.005
nC-PM	0.723	0.037
TIIM	0.873	0.056

Table 3: Performance of TIIM as compared to static mitigation assignment. The dynamic countermeasure selection of TIIM allows it to reach an interesting trade-off between PRR and cost.

6.3 Results

We discuss the evaluation results of TIIM from two aspects. First, we discuss the potential benefits of leveraging the detection capability of TIIM, as opposed to employing a fixed countermeasure through trace-driven evaluation over 64 hours of CTI extensive runs. Second, we present the overall system performance gain achieved while running TIIM. **Evaluation Metrics:** We employ the following metrics to evaluate the performance of TIIM.

- *Packet Reception Ratio (PRR)*: This is the ratio of successfully received packets over the total number of transmitted packets during a specific time period.
- *Gain*: compares the achieved PRR to the baseline PRR (default 802.15.4 PRR under interference).
- *Cost*: the ratio of transmission overhead introduced by the countermeasure to the base transmission. For instance, for FEC a fixed transmission overhead of 12 Byte per packet is considered. For no-CCA, the transmission of positive deferrals is considered as cost. Positive deferrals are those transmissions that were lost or corrupted, but would have been deferred with CSMA. PM requires retransmission of the localized corrupted symbols. Hence, we consider the overhead cost of the new frame.
- *Adaptability*: to detail the ability of TIIM to adapt quickly to unanticipated changes in the interfered channel, we illustrate the degree of the system adaptability by showing its dynamic behavior in a timeline plot.

6.3.1 Adaptive Interference Mitigation

We now evaluate the overall prospective performance gain for an adaptive interference mitigation system as compared first to the performance of standard 802.15.4 and then to the gain of applying a fixed interference countermeasure. We perform a trace-based simulation using the 64 hours of CTI traces. We run each of the countermeasures and calculate its corresponding cost and gain.

As depicted in Table 3, the PRR of our traces under interference lies around 56%. Applying a static countermeasure could potentially be the best solution in case the channel conditions remain static and best for that countermeasure. However, due to changes of interference patterns in

our traces, the static countermeasures can cause high cost overheads. For instance, FEC in combination with noCCA causes about 30% additional transmission. TIIM achieves almost the highest PRR gain with a cost overhead of 5.6%.

The adaptability of TIIM enables it to perform best in a dynamic channel. Figure 12 shows the timeline of selected 20 hours of our traces and how TIIM adapts the selection of countermeasures according to its assessment of the channel.

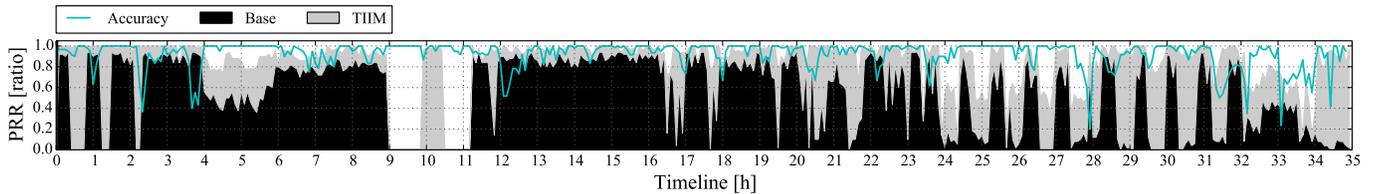
In the following, we spot some interesting observations. Starting from hour 0 to hour 3, we notice that no-CCA is dominated with transitions to nC-PM and nC-FEC. From hour 3.5 to hour 5, nC-FEC becomes the dominating countermeasure. We notice the first use of PM at hour 7. Particularly, the no action between hour 9 and 10 is interesting: it happens while the high-power analog phone is active at distance 3 m, causing severe interference. In this period, communication is not viable over the interfered channel and thus TIIM recommends channel switching. Identifying such instance is essential for saving energy. We encounter the first uses of FEC from hour 10.5 to hour 11.5. During the same period PM has become more active. It is worth noting that the microwave oven is the dominant interferer during this period. Notice that starting from hour 17 (indicated with a black vertical line), we expose TIIM to condensed corrupted traces. Consequently, TIIM’s countermeasures are dominated by nC-FEC, FEC, and partly PM.

6.3.2 Online Performance

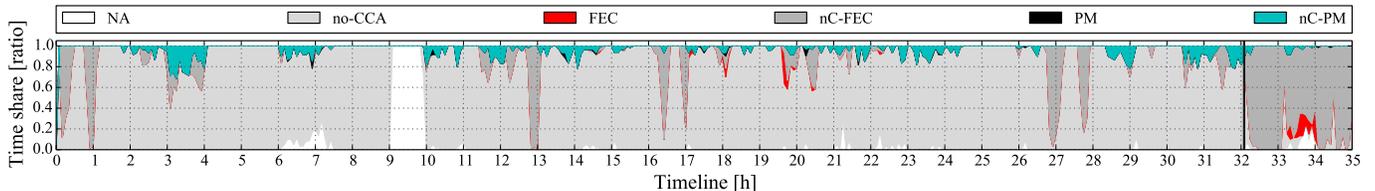
Now, we evaluate the online performance of TIIM. To this end, we run TIIM on extracted features computed over the time window of 5 s. Whenever TIIM recommends a countermeasure, our system applies it.

Figure 13(a) illustrates the base PRR of the traces in black and the additional gain achieved by TIIM in gray. The accuracy of our system in selecting the right countermeasure, at any given time, is plotted as a line. In order to visualize TIIM’s dynamic behaviour, we show the time share of each countermeasure in Figure 13(b), time-synchronized with the PRR in Figure 13(a). Within the first hour, TIIM achieves a gain of 100%, several times enabling PRR to go from 0% to almost 100%. At hour 2, we observe a short drop in the accuracy, but the resulting gain remains high. There are further sharp drops in the accuracy of the system that do not lead to a PRR decrease. These are the cases where the recommended countermeasure still could yield an acceptable gain, but possibly with a higher cost. However, not all inaccuracies remain unpunished by the PRR. For instance, around hours 7, 17, 28, 33, and 24 the PRR drops as a consequence of temporarily decreased accuracy.

TIIM successfully detects the first occurrence of persistent harmful interference (analog phone at location L1) between hours 9 to 10. It detects persistent non-harmful interference (analog phone at location L2) and consequently recommends no-CCA, PM, and nC-PM which yields 100% of



(a) Performance of TIIM while applying recommended countermeasures.



(b) Time share of each countermeasure during the online phase.

Figure 13: Online evaluation of TIIM. Inquiring TIIM for countermeasures while providing the channel conditions as input (resolution 5 min).

gain. However, it fails from hour 10.5 to hour 11 to detect another occurrence of persistent harmful interference, where the system should have recommended channel switching.

Starting from hour 19, TIIM shows its potential in recovering severe performance degradation. TIIM is stressed with a heavy load of concentrated corrupted packets after hour 32. It applies dominantly nC-FEC, FEC, and partly nC-PM to improve a close to 0% PRR to almost 95%.

Reaction time. Traditional classification approaches need a few seconds to detect the type of interference source, e.g., 18.14 s for WiFi [17]. This is relatively a long time for highly dynamic channels. One advantage of TIIM is that it reacts to interference shortly after detecting the degradation caused by interference. Currently, TIIM provides a recommendation 5 s after detecting interference. This allows TIIM to react timely to time-variant interference patterns.

Coexistence with other Radios. We verified empirically that 802.15.4 does not cause harmful interference to high-power wireless devices such as the wireless cameras. Even when disabling CCA, we did not observe any effect on wireless cameras' operation. On the other hand, 802.15.4 can cause harmful interference with coexisting low-power radios, resulting into 802.11 deferrals and packets losses for Bluetooth. TIIM can potentially be trained to make a tradeoff between its performance and the harmful interference it may cause to the low-power networks or be trained to apply no-CCA only for high-power interferers.

7. DISCUSSION

This paper provides a proof of concept on the potential of a CTI-aware and adaptive link-layer solutions. However, more research and experimentation are needed to generalize and realize the full potential of TIIM. Here, we address some practical challenges and research points that assist in evolving TIIM further.

(a) Porting TIIM to other Radios. Although this work focuses on low-power networks, most of the observations can be projected to analogous RF technologies, such as 802.11 radios. The core concept of integrating reasoning to combat interference diversity in the unlicensed bands has not been explored before. We believe that TIIM's core concept can be beneficial to analogous radios, and leave further investigations for future work.

(b) Extending TIIM with new Countermeasures. In this work, we explored the feasibility of addressing the CTI heterogeneity problem by focusing on aspects that are relevant to the set of few mitigation approaches considered in this prototype of TIIM. To extend the system with new mitigation approaches, relevant features need to be redefined and the classifier needs to be retrained. Possible examples of interesting countermeasures that can benefit from a learning module are: Detecting systematically the duration and interspace of interference pluses could be a useful metadata for FEC, to select the right level of redundancy required by FEC, or capturing tendencies in duty cycles can be exploited for a better MAC scheduling.

(c) TIIM's Limitations. TIIM has a narrow view of the RF spectrum that is limited to the 802.15.4 channel width. It focuses on increasing the spectral efficiency over interfered channels, with lack of cognition about the state of the rest of the spectrum. Thus, it lacks a comprehensive view of the RF spectrum to decide whether communication over an interfered channel is preferred over channel switching. Currently, we are tackling how to address this limitation with least possible overhead.

(d) Interference and PHY Layer Information. Over the last few years, researchers advocated a design of wireless systems that allows a better interfacing of physical layer information for higher layers, particularly to cope with interference. We developed this work with legacy systems in mind, thus we were limited to the PHY space provided by these systems. One PHY aspect that can be integrated into TIIM to overcome its limitation of narrow spectrum perception, is the use of cyclostationary analysis for bandwidth estimation of interfering signals as suggested by DOF [29]. DOF estimates the bandwidth of interfering signals solely based on the PHY information retrieved from the channel frequency in use. This allows a better reaction in severe channel conditions, where a less affected channel outer the interferer's bandwidth could be selected.

8. RELATED WORK

We distinguish three major directions adopted to combat interference in the unlicensed bands. The first direction aims at **detecting and avoiding interfered frequencies**

by employing spectrum sensing to identify interference-free channels [18, 23]. Musaloiu et al. [25] propose a distributed algorithm for channel selection and interference estimation using RSSI sampling for 802.15.4 networks. The lack of interference-free channels, and the fast and unpredictable changes in the occupancy state of frequency bands make the sampling overhead of these approaches high, particularly for resource constrained devices.

The second direction aims at **increasing resilience against interference**, by bracing PHY and MAC layers with auxiliary mechanisms. For instance, Liang et al. [15] studied the interplay between 802.11 and 802.15.4 and applied a resilience forward error coding scheme to interference [15]. Analogously, some solutions focused on exploiting the temporal effects of interference induced on the PHY hints, such as variations in soft errors (softPHY) [21] or RSSI variations [9, 20] to recover interfered packets. Others focused on increasing the robustness of existing MAC protocols against interference [13, 14], or considered utilizing multiple radio channels for communication [12, 24] to exploit frequency diversity. Moreover, further PHY solutions have been considered, such as utilizing advancements in MIMO for interference cancellation [7, 28].

The third direction aims at **identifying the type of interference** by employing signal classification techniques [1, 22, 29] or featuring distinct interferer's patterns on corrupted packets [17]. It is, however, not yet clear how the interference classifiers can be utilized in an automated way to mitigate interference, given the diversity of interference technologies. Our work aims at bridging the second and third directions, by featuring classification to recognize interference patterns which can provide useful meta-information about the applicability of a certain mitigation strategy.

9. CONCLUSION

Wireless interference has been a long sought but still crucial problem in wireless communication, notably for systems operating in the unlicensed bands. While most existing solutions focus on careful tuning of signals to realize frequency isolation, less work has focused on thoroughly utilizing the interfered links under heterogeneous interference patterns.

In this paper, we introduced TIIM, an interference mitigation system that proposes countermeasures that work best under the current interference patterns, independent of the particular technology causing it. We leverage previously unconsidered channel attributes and employ a lightweight machine learning classifier to (i) decide whether communication is viable over the interfered link, and (ii) find the best underlying link-layer coexistence scheme. Doing so, TIIM realizes the full potential of interfered wireless links and consequently enhances spectral efficiency. Our evaluation shows that TIIM improves the packet reception ratio under interference by about 30% with only 5.6% additional transmission overhead.

Acknowledgments. We thank Wilhelm Kleiminger, Friedemann Mattern, our shepherd, Chan Mun Choon, and the anonymous reviewers for their insightful comments. This work was partly funded by SSF, the Swedish Foundation for Strategic Research and VINNOVA, Sweden's innovation agency.

10. REFERENCES

- [1] Cisco CleanAir Technology.

- [2] *Texas Instruments. CC2420 datasheet*, 2007.
- [3] Estimating the Utilisation of Key License-Exempt Spectrum Bands, Final Report, Mass Consultants Ltd., Ofcom, 2009.
- [4] Miercom: Cisco CleanAir Competitive Testing, Lab Test Rerpot DR100409D, Miercom, 2010.
- [5] Low-Rate Wireless Personal Area Networks (LR-WPANs), IEEE Std. 802.15.4., 2011.
- [6] *Rulequest Research See5/c5.0.*, 2014.
- [7] Y. Yubo, Y. Panlong, L. Xiangyang, T. Yue, Z. Lan, Y. Lizhao. ZIMO: Building Cross-technology MIMO to Harmonize Zigbee Smog with WiFi Flash Without Intervention. In *ACM MobiCom*, 2013.
- [8] A. Dunkels, B. Gronvall, T. Voigt. Contiki - A Lightweight and Flexible Operating System for Tiny Networked Sensors. In *IEEE LCN*, 2004.
- [9] A. Hithnawi. Exploiting Physical Layer Information to Mitigate Cross-Technology Interference Effects on Low-Power Wireless Networks. In *ACM SenSys*, 2013.
- [10] A. Hithnawi, H. Shafagh, S. Duquenooy. Understanding the Impact of Cross Technology Interference on IEEE 802.15.4. In *ACM WINTeCH*, 2014.
- [11] B. Han, A. Schulman, F. Gringoli, N. Spring, B. Bhattacharjee, L. Nava, L. Ji, S. Lee, R. Miller, Maranello: Practical Partial Packet Recovery for 802.11. In *USENIX NSDI*, 2010.
- [12] B. Nahas, S. Duquenooy, V. Iyer, T. Voigt. Low-Power Listening Goes Multi-channel. In *IEEE DCOSS*, 2014.
- [13] C. A. Boano, T. Voigt, C. Noda, K. Romer, M. A. Zuniga. JamLab: Augmenting Sensornet Testbeds with Realistic and Controlled Interference Generation. In *ACM/IEEE IPSN*, 2011.
- [14] C. A. Boano, T. Voigt, N. Tsiftes, L. Mottola, K. Romer, M. A. Zuniga. Making Sensornet MAC Protocols Robust Against Interference. In *EWSN*, 2010.
- [15] C. Liang, N. Priyantha, J. Liu, A. Terzis. Surviving Wi-Fi Interference in Low-power ZigBee Networks. In *ACM SenSys*, 2010.
- [16] D. Tse, P. Viswanath. *Fundamentals of Wireless Communication*. Cambridge University Press, 2005.
- [17] F. Hermans, O. Rensfelt, T. Voigt, E. Ngai, L. Norden, P. Gunningberg. SoNIC: Classifying Interference in 802.15.4 Sensor Networks. In *ACM/IEEE IPSN*, 2013.
- [18] H. Rahul, N. Kushman, D. Katabi, C. Sodini, F. Edalat. Learning to Share: Narrowband-friendly Wideband Networks. In *ACM SIGCOMM*, 2008.
- [19] Iperf. <http://iperf.sourceforge.net/>.
- [20] J. Hauer, A. Willig, A. Wolisz. Mitigating the Effects of RF Interference through RSSI-Based Error Recovery. In *EWSN*, 2010.
- [21] K. Jamieson, H. Balakrishnan. PPR: Partial Packet Recovery for Wireless Networks. In *SIGCOMM*, 2007.
- [22] K. Lakshminarayanan, S. Sapra, S. Seshan, P. Steenkiste. RFDump: An Architecture for Monitoring the Wireless Ether. In *ACM CoNEXT*, 2009.
- [23] L. Yang, W. Hou, L. Cao, B. Zhao, H. Zheng. Supporting Demanding Wireless Applications with Frequency-agile Radios. In *USENIX NSDI*, 2010.
- [24] R. Gummadi, D. Wetherall, B. Greenstein, S. Seshan. Understanding and Mitigating the Impact of RF Interference on 802.11 Networks. In *ACM SIGCOMM*, 2007.
- [25] R. Musaloiu, A. Terzis. Minimising the Effect of WiFi Interference in 802.15.4 Wireless Sensor Networks. *International Journal of Sensor Networks*, 2008.
- [26] R. O. Duda, P. E. Hart, D. G. Stork. Pattern Classification (2nd Ed.). In *Wiley-Interscience*, 2011.
- [27] S. Gollakota, D. Katabi. Zigzag Decoding: Combating Hidden Terminals in Wireless Networks. In *ACM SIGCOMM*, 2008.
- [28] S. Gollakota, F. Adib, D. Katabi, S. Seshan. Clearing the RF Smog: Making 802.11n Robust to Cross-technology Interference. In *SIGCOMM*, 2011.
- [29] S. Hong, S. Katti. DOF: A Local Wireless Information Plane. In *ACM SIGCOMM*, 2011.
- [30] S. Rayanchu, A. Mishra, D. Agrawal, S. Saha, S. Banerjee. Diagnosing Wireless Packet Losses in 802.11: Separating Collision from Weak Signal. In *IEEE INFOCOM*, 2008.
- [31] S. Rayanchu, A. Patro, S. Banerjee. Airshark: Detecting non-WiFi RF Devices Using Commodity WiFi Hardware. In *ACM IMC*, 2011.
- [32] S. Yun, D. Kim, L. Qiu. Fine-grained Spectrum Adaptation in WiFi Networks. In *MobiCom*, 2013.