<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *ACM/IEEE International Conference on Information Processing in Sensor Networks 2016 (IPSN 2016)*.

Citation for the original published paper:

Hithnawi, A., Su, L., Shafagh, H., Duquennoy, S., Gross, J. (2016)

CrossZig: Combating Cros-Technology Interference in Low-Power Wireless Networks.

In: IEEE

<http://dx.doi.org/10.1109/IPSN.2016.7460663>

N.B. When citing this work, cite the original published paper.

© 2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:kth:diva-184883>

CrossZig: Combating Cross-Technology Interference in Low-power Wireless Networks

Anwar Hithnawi, Su Li*, Hossein Shafagh
Department of CS, ETH Zurich, Switzerland

*EPFL Lausanne, Switzerland
{hithnawi, shafagh}@inf.ethz.ch *su.li@epfl.ch

James Gross
KTH Royal Institute of Technology
Stockholm, Sweden
james.gross@ee.kth.se

Simon Duquennoy
INRIA Lille, France
and SICS Swedish ICT
simon.duquennoy@inria.fr

Abstract—Low-power wireless devices suffer notoriously from Cross-Technology Interference (CTI). To enable co-existence, researchers have proposed a variety of interference mitigation strategies. Existing solutions, however, are designed to work with the limitations of currently available radio chips. In this paper, we investigate how to exploit physical layer properties of 802.15.4 signals to better address CTI. We present CrossZig, a cross-layer solution that takes advantage of physical layer information and processing to improve low-power communication under CTI. To this end, CrossZig utilizes physical layer information to detect presence of CTI in a corrupted packet and to apply an adaptive packet recovery which incorporates a novel cross-layer based packet merging and an adaptive FEC coding. We implement a prototype of CrossZig for the low-power IEEE 802.15.4 in a software-defined radio platform. We show the adaptability and the performance gain of CrossZig through experimental evaluation considering both micro-benchmarking and system performance under various interference patterns. Our results demonstrate that CrossZig can achieve a high accuracy in error localization (94.3% accuracy) and interference type identification (less than 5% error rate for SINR ranges below 3 dB). Moreover, our system shows consistent performance improvements under interference from various interfering technologies.

I. INTRODUCTION

In recent years, we have witnessed a proliferation of heterogeneous wireless technologies operating in the unlicensed bands. The escalation of wireless demand has put enormous pressure on available spectrum which exacerbates interference and coexistence problems [1]–[3].

Problem. Embedded computing devices are increasingly integrated in objects and environments surrounding us, paving the way for the Internet of Things’ vision of digitizing the physical world. These devices utilize low-cost sensors for a range of performance-sensitive applications, such as health systems, general monitoring and tracking, home automation, etc. Low-power wireless¹ technologies (e.g., BLE, 802.15.4, and backscatter communication) employed by these applications are expected to endure interference from other radio technologies. The CTI problem is exacerbated for these low-power networks, where energy and complexity constraints prohibit the use of sophisticated interference suppression and cancellation techniques that are finding their ways into unconstrained wireless systems [3], [4].

¹In the context of this paper, we use the term low-power wireless technologies to refer to devices transmitting with less than 0 dBm power.

Due to the inherent application requirements, devices operating in the unlicensed bands transmit at different power levels. Low-power radios typically transmit at less than 1 mW for energy efficiency requirements, others, such as analog phones, can transmit at the maximum allowed power (i.e., 1000 mW). This severe power asymmetry poses significant coexistence problems, where high-power interferers can completely starve low-power technologies. That is because a typical high-power interferer might fail to detect the transmission of a nearby low-power transmitter, thus can interfere with the low-power node’s transmission and monopolize the shared channel.

Wireless systems use a variety of physical layer techniques to combat channel impairments such as attenuation, multipath, or fading. For instance, 802.15.4 employs spread-spectrum modulation and error control coding. However these techniques alone fall short in mitigating the effects of CTI. CTI severely reduces the *Signal-to-Interference-plus-Noise Ratio* (SINR) of the intended transmission, which results in high bit-error rates and limits the effectiveness of these techniques. Most CTI interferers leave the channel recurrently idle for short times [3], [5]. This pattern results in partial overlap with transmitted 802.15.4 packets. The overlap duration is technology/application specific and is reflected on the error characteristics of residual errors in the interfered packets. In this work, we design a recovery mechanism that is aware of the characteristics of the CTI residual errors.

Approach. In this paper, we argue that there exists sufficient unutilized opportunities for low-power wireless networks to coexist in overlapping channels. We thoroughly investigate how to best exploit physical layer (PHY) information to make insightful adaptation decisions. PHY information is available but inaccessible in commercial off-the-shelf low-power radio chips due to conventional network layering abstractions that define the current layer interfaces and the flow of information between these layers. We introduce CrossZig, a cross-layer solution that enables low-power wireless nodes to make informed decisions on their coexistence strategies based on richer physical layer information, thus adapt autonomously to the current interference patterns in the channel. CrossZig achieves this by leveraging two key building blocks:

(a) *CTI Detection*: This component resides in the physical layer. It detects interference in corrupted packets and differentiates its type between *Intra-* and *Cross-Technology Interference*.

(b) *CTI-aware Packet Recovery*: We propose an adaptive recovery scheme, which exploits physical layer hints to adjust the recovery settings. We propose a packet recovery mechanism that exploits time-diversity combining, targeted for low-power single-antenna radios. On top of diversity combining, our scheme applies adaptive *Forward Error Correction* (FEC) coding with redundancy derived from observed error patterns.

We have implemented a prototype of CrossZig for low-power 802.15.4 radios in a *Software-Defined Radio* (SDR) platform. Experimental results show that our approach can substantially improve the performance of 802.15.4 links under various CTI patterns.

Contribution. This paper makes the following contributions:

- We develop a novel lightweight technique that allows low-power nodes to recognize the type of interference in interfered packets. Our design achieves high accuracy in detecting CTI at all SINR ranges where the target signal cannot be decoded correctly and it does not require any prior synchronization between nodes (unlike [6], [7]);
- We propose an adaptive recovery mechanism that exploits both block-based error correction and packet merging through diversity combining at the signal level. We show that both these schemes come at low complexity costs while - if carefully performed - they can effectively alleviate the damage due to *Cross-Technology Interference*;
- We implement and evaluate a prototype of our system in SDR using GNURadio [8] with USRP-N210 [9]. The evaluation results show that our design has consistent improvements in packet reception ratio under interference from various interferer types.

II. BACKGROUND

In order to mitigate the impact of CTI, we need to develop a detailed understanding of how radio technologies interact. In this section, we summarize key features of wireless technologies operating in the ISM bands. Moreover, since CrossZig is in part a physical layer design, we review related aspects of the conventional 802.15.4 physical layer.

A. Cross-Technology Interference

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 problem elevates in the ISM bands where the number and diversity of inhabited wireless networks are continuously increasing. The FCC's light regulations to confine interference have not prevented a range of interference problems across the wireless technologies operating in the ISM bands. Classical coexistence solutions largely focus on avoiding interference by employing carrier sense (a simple access mechanism that supports attentive accommodating of other transmitters) or transmitting over orthogonal channels. These solutions, however, can make low-power technologies more prone to starvation or are no longer feasible given the scarcity of interference free channels, respectively. Moreover, many wireless technologies today have

Fig. 1. IEEE 802.15.4 modulation (O-QPSK).

adopted greedy practices to cope with their high throughput demands and to better exploit the shared spectrum, such as (a) allocating wide-bands, (b) transmitting at high-power, and (c) frequency agility [4], [10]. These solutions exacerbate the CTI problem for low-power networks.

The dominant spectrum access modality that exists in the ISM bands is comprised of non-persistent interferers; where RF technologies exhibit a time-variant ON and OFF pattern of energy emission due to their underlying communication primitive, such as frequency hopping, inter-frame spacing (e.g., SIFS, DIFS), back-off slots, application traffic patterns, and periodic cycles of noise radiation, as for microwave ovens. Therefore, they leave the occupied channel recurrently idle for short times. The ratio of idle periods depends on the interferer technology and the application traffic patterns. Medium access mechanisms utilizing these holes in occupied channels can enhance spatial reuse and network throughput.

Intra- vs. Cross-Technology Interference: We briefly explain why physical layer solutions that are used to tackle intra-technology interference, namely *Interference Cancellation* (IC) [7], [11], [12] are not applicable in case of CTI. The main challenge that hinders the use of IC across different technologies stems from the fact that IC depends on the receiver capability of estimating the channel coefficients between Tx and Rx and understanding the interfered signal structure (i.e., being able of demodulating and decoding). Only then the interference signal can be reconstructed and subtracted. This, consequently, reduces the contribution of interference signal on the SINR to a level where the target signal can be successfully decoded. However, applying IC for the cases where the estimation of the interfering signal is erroneous (e.g., CTI) would worsen the chances of decodability for the target signal.

B. IEEE 802.15.4 Physical Layer

The IEEE 802.15.4 standard [13] allocates 16 channels for devices operating in the 2.4 GHz ISM band, where each channel has a bandwidth of 2 MHz. At the physical layer, data is first grouped into 4-bit symbols and then spread to a specified 32-bit long *Pseudo-random Noise* (PN) sequence ($b_0b_1b_2b_3 \rightarrow c_0c_1c_2 \dots c_{31}$). Each bit (c_i) in a PN sequence is then modulated using Offset Quadrature Phase-Shift Keying (O-QPSK). As shown in Fig. 1, the even chips $c_0c_2c_4 \dots$ are modulated as *In-phase* (I) component of the carrier and the odd indexed chips $c_1c_3c_5 \dots$ are modulated as *Quadrature* (Q) component of the carrier. The time duration of each chip is $1 \mu s$ and there exists a half chip time ($T_s = 0.5 \mu s$) offset between the Q-phase chips and I-phase chips that results in a continuous

Fig. 2. Overview of CrossZig.

phase change and constant envelope. For demodulation, the receiver's radio converts each half-sine pulse signal into a chip. Then these chips are grouped to provide PN sequences. The de-spreading is performed by mapping the PN sequence to the symbol with the highest correlation. Unlike modulation schemes such as QAM or ASK, which operate by varying the amplitude of the carrier wave, 802.15.4 adopts O-QPSK modulation. Hence, the carrier wave amplitude of all chips within one packet is constant and depends on the selected transmission power (i.e., constant envelope). The 802.15.4 chips are shaped by half-sine pulse at the transmitter. While the signal's shape will be distorted by noise in the wireless channel, its basic shape is maintained. The demodulator's output provides an indicator of how close the received signal shape is to the expected shape, we elaborate more on this in §III-B. We leverage these two features (i.e., constant amplitude and signal shape) of the 802.15.4 PHY in the design of CrossZig.

III. ARCHITECTURE AND DESIGN

In this section, we present the detailed design of our approach that extends the 802.15.4 stack to improve medium access efficiency under CTI and recover partially interfered packets. We start by presenting a high-level overview of CrossZig, then present its core components, and finally describe the system integration.

A. Overview

CrossZig is an extension to the standard 802.15.4 that allows low-power wireless nodes to communicate better in interfered environments. Upon the detection of CTI, CrossZig triggers an adaptive recovery scheme. Our extension is accompanied by a CTI-aware medium access mechanism that opportunistically leverages the silence duration in interfered channels. In particular, our extension consists of the following components:

PHY-hints Interface. This interface allows higher layers to access richer physical layer information and consequently use it in a variety of algorithms to boost performance under CTI.

CTI Detection. This component resides in the physical layer. It exploits variations in the PHY hints to detect interference in incoming corrupted packets. More importantly, it differentiates the interference type between *Intra-* and *Cross-Technology Interference*. The error localization and recovery mechanism presented next are enabled by CrossZig only when CTI is detected; otherwise nodes operate in the normal mode, as depicted in Fig. 2.

Symbol Error Localization. This component allows to locate and identify erroneous symbols by solely relying on PHY hints

Fig. 3. Block Diagram of the receiver and corresponding PHY hints.

without additional signaling or redundancy. This is necessary for our introduced recovery mechanisms.

CTI-aware Packet Recovery. Our system recovers from a variety of interference patterns. The receiver estimates errors in interfered packets by relying on physical layer hints. Error information is used to choose a suitable recovery mechanism; currently selecting or combining two recovery mechanisms: (a) Cross-layer based packet merging to recover long error bursts, and (b) Adaptive error-correction coding to deal with transient interference with low BER rates.

B. Physical Layer Hints

When an interfered signal is received, besides standard processing, such as demodulation, chip-to-symbol mapping, and delivering decoded symbols to the data-link layer, the physical layer also accommodates further hints that can be exploited to boost the performance of wireless systems [14]–[17]. In our design, we exploit such hints to detect *Cross-Technology Interference* and to estimate the confidence of received symbols in interfered packets, as depicted in Fig. 3. We consider the following physical layer hints:

Signal Power. When two signals interfere, their energies add up². Therefore interfered segments of the received signal experience larger power than the rest of the signal. The interfered segment of the signal exhibits lower SINR, thus experiences a higher error rate. This insight on additive energy of interfering signals highlights the ambient information the signal carries along and can assist in detecting interference and localizing interfered symbols within interfered packets. Fig. 4(a) plots the signal power of a partially interfered packet. Once exposed to interference, the signal experiences a sudden sharp increase in the signal power.

Hamming Distance. In the 802.15.4 PHY, symbols are spread to a 32-chip codeword before transmission (one of 16 PN codewords). The de-spreading is performed by mapping the received codeword to the symbol with the highest correlation. For an erroneous mapping of a received codeword, many chips have to be flipped. The distance between the input and output codewords of the chip-to-symbol mapper can serve as indicator for the confidence of symbol decoding. Fig. 4(b) plots the Hamming distance within an interfered packet. Large and low Hamming distance values provide a good indicator of corrupted or correct symbols, respectively.

Demodulation Soft Values. *Soft Values* (SV) of demodulated bits are real numbers output by the demodulator. These values are approximations of the transmitted symbols.

²This occur if the signals add up constructively.

Fig. 4. Physical layer hints of a corrupted packet by Interference. The gray area indicates erroneous symbols.

The receiver's demodulator maps the SV to the closest ideal symbol. For instance in case of Binary Phase-Shift Keying, the binary demodulated bits are retrieved after passing the soft values through a binary slicer. The bit is set to 1, if the SV is a positive number, otherwise it is set to 0. However, besides the bit value, SV also contains the confidence information about the demodulation [14].

The confidence information of the SV can be interpreted based on the type of the demodulator. In case the receiver adopts a matched filter-based coherent demodulator, the soft demodulated values indicate the similarity between the received signal and ideal signal shape. Thus the larger SVs, the higher the confidence for the corresponding bit to be correctly demodulated. However, if a non-coherent demodulator is used, the SV carries different information. For example, in our case the receiver uses a quadrature demodulator, which outputs the phase differences between two successive signal samples as SVs and can be computed as:

$$SV(i) = \angle(s(i) \times s^*(i-1)) = \pm \frac{\pi}{4} + \delta, \quad (1)$$

where $\pm \frac{\pi}{4}$ is the ideal value of SV and δ is error caused by interference and noise. Each chip is modulated by a single half-sine pulse in the transmitted signal and is represented by a sequence of four complex samples at the receiver. This implies a total phase change of π for one chip, hence the expected phase change between two signal samples is $\pm \frac{\pi}{4}$. The demodulation confidence does not depend on the absolute value of SV, but the difference between $|SV|$ and $\frac{\pi}{4}$. Chips (i.e., bits of a codeword) with $|SV|$ closer to $\frac{\pi}{4}$ have a higher probability to be correctly demodulated (see Fig. 4(c)).

C. CTI Detection

Performance degradation in wireless systems can be due to *Intra-Technology Interference*, *Cross-Technology Interference*, or *insufficient signal strength*. Determining the cause of performance degradation is essential for the coexistence problem as this defines the corresponding mitigation action to be considered. Discerning the cause of packet corruption while exposing differences between interference and weak signal has been addressed in recent works. This is mainly achieved by relying on soft value jumps or power jumps within received packets. However, differentiating *Intra-* and *Cross-Technology Interference* has not yet been addressed with practical mechanisms. The inability to distinguish the type of interference leads to rather conservative approaches that blindly treat packet losses as collisions (i.e., overlapping

transmissions of the same technology). Thereby exponential backoffs are invoked which can lead to starvation of low-power radios competing with high-power interferers. Moreover, interference cancellation-based solutions (e.g., SIC [11]) would impose undesired overhead and worsen the chances of decodability for the target signal, if applied in the presence of CTI. Generally speaking, bracing wireless nodes with mechanisms that increase their ability to reason about the channel state will allow better adaptation and recovery.

Hence, beyond detecting the presence of interference, we are interested in detecting the presence of CTI. We introduce two complementary CTI detection mechanisms that are utilized in CrossZig: SV-based and correlation-based detections. Introduced recovery mechanisms are only enabled by CrossZig when CTI is detected

1) *SV-based Detection*: We explore the possibility of exploiting variations in demodulated soft values for interference type detection. The core idea is to inspect the modulation and signal shape of the interfered signal, which is reflected by the soft values. While experiencing *Intra-Technology Interference* (interferer is 802.15.4), the demodulator demodulates the stronger signal. Since the interference signal is of the same type (i.e., shape) the variations in the soft values remain small. In contrast, with *Cross-Technology Interference* the signal shape differs from the ideal signal. Thus, the variations in soft values are higher. We take signal samples from the interfered part and compute the variation metric V which we use to determine the received signal type:

$$V = \frac{1}{N} \sum_{i=1}^N \left(|SV(i)| - \frac{\pi}{4} \right)^2 \quad (2)$$

V measures the average distance between the received $|SV|$ and the ideal value $\frac{\pi}{4}$. The smaller V , the higher the chance that the signal is 802.15.4 (i.e., *Intra-Technology Interference*).

This SV-based detection mechanism does not require to compute complex compensation of channel distortions or signal decoding. Moreover, the soft values are readily available which makes this detection mechanism considerably lightweight. Note that any interfering technology using O-QPSK with half-sine pulse shape and similar baseband signal bandwidth other than 802.15.4 is identified as *Intra-Technology Interference* using this mechanism. This technique exploits the capture effect phenomenon, in which a strong interfering signal is successfully demodulated (i.e., of the same technology). It, therefore, works well in the low SINR region, where the target signal is much weaker than the interferer.

Hence, the interference signal dominates the signature shape in the received signal. If this is not the case, we resort to a more costly technique: correlation-based detection. Although we focus our discussion on 802.15.4 PHY, this approach can be adopted to other wireless technologies that provide SVs.

2) *Correlation-based Detection*: The receiver can exploit the fact that 802.15.4 packets start with a predefined preamble and SFD symbols for synchronization, and search for this known signal pattern within the interfered segment by computing the temporal cross-correlation between received signal and the ideal preamble plus the SFD. In case the pattern is present, the receiver can conclude that the interference is of type *Intra-Technology Interference*. Otherwise, it is a *Cross-Technology Interference*. In general, correlation is a typical functionality in standard wireless receivers [18]. To detect the interference type in the received signal, the receiver can perform the cross-correlation between the 802.15.4 preamble (p) and the start of the interfered segment. This approach yields a good performance in theory.

In practice, however, the transmitter and receiver are typically not centered on the same frequency, hence there is a small frequency offset (Δf) between the transmitter and the receiver that causes a linear shift in the phase of the received signal. This frequency offset can distort the correlation and needs to be compensated prior to the correlation process. Standard receivers typically estimate the offset and compensate for it. In the context of CTI, since we are agnostic of the transmission source and do not have access to a decodable pilot or decodable preamble in the interfered signal, it is not possible to compensate the frequency offset of the interference signal, even in case of *Intra-Technology Interference*. This consequently limits the accuracy and usability of this approach.

An alternative approach is applying correlation in the frequency domain. Since frequency offset in the time domain will translate into sampling offset in the frequency domain, it does not affect the value of correlation, but only shift it. The frequency domain correlation with consideration of the frequency offset can be formulated as follows:

$$c(y, \tau, p) = \sum_{n=1}^N P^*(n)Y(n + \tau) \quad (3)$$

$$= \sum_{n=1}^N P^*(n)\mathcal{F}\left((s(k) + i(k) + w(k))e^{j2\pi(\Delta f - \tau)kT}\right) \quad (4)$$

The ideal preamble is independent of transmitted data and the noise, therefore the correlation between the ideal preamble and s and w is about zero.

$$c(y, \tau, p) = \sum_{n=1}^N P^*(n)\mathcal{F}\left(i(kT)e^{j2\pi(\Delta f - \tau)kT}\right) \quad (5)$$

$$= \sum_{n=1}^N P^*(n)I(n + \tau - \Delta f) \quad (6)$$

Where $P(n)$ and $Y(n)$ are the ideal preamble signal and received signal in the frequency domain, respectively. $\mathcal{F}(x)$ means the Fourier transform of x . Moreover, s , i , and w represent signal, interference, and noise in the received signal, respectively. The correlation value is maximized when

Fig. 5. Physical layer hints for correct and corrupted symbols under wireless camera interference. Dark grey dots indicate correct symbols and light grey dots indicate corrupted symbols. The majority of received symbols with Hamming distance under 4 were received correctly and above 10 corrupted. Correct and corrupted symbols overlap in the Hamming distance range in between. Therefore, the Hamming distance alone is not sufficient for the determination of the symbol fate and the power information can help for these symbols.

$\tau = \Delta f$, and the signal is the expected preamble signal ($I(n) = P(n)$). Since Δf is unknown and we cannot compensate it, we compute the correlation for a certain range of τ instead and consider its maximal value as:

$$C(y, p) = \max_{\tau} c(y, \tau, p) \quad (7)$$

The range of τ is not large, given that the frequency offset is typically small.

Complexity. Applying correlation in the frequency domain involves transforming a signal from its time representation to the frequency domain ahead of applying the correlation, which can be an expensive procedure for low-power receivers. CrossZig primarily runs the SV-based mechanism for detection and utilizes the correlation-based technique just for the SINR ranges where the SV-based technique does not yield a good accuracy. The SV-based mechanism in its core examines variations in the SVs which makes it a lightweight mechanism that is practical for low-power radios.

D. Symbol Error Localization (in interfered packets)

As CrossZig involves processing of incomplete packets (partially interfered), it requires the receiver to be able to discern with a high accuracy and without additional feedback from the sender which symbols in a packet are correct and which are not. The physical layer hints described in §III-B expose statistical differences between interfered and non-interfered symbols, which render them suitable candidates to detect erroneous symbols. However, designing practical error detection algorithms based on these PHY hints with acceptable false positive and false negative rates is challenging. As we are interested in per symbol error estimation, we leave out SV (which would introduce more overhead with one SV per chip, i.e., 32 per symbol).

A direct method to estimate the symbol error is setting a threshold on the number of unmatched bits (reflected in Hamming distance) of the decoding results. This indicates the disparity between the chip sequence derived from the received

Fig. 6. Diversity combining of two identical signals under interference.

signal and ideal symbol sequence. However, finding a good threshold is not trivial, as discussed before and illustrated in Fig. 4(b). We propose an error estimation algorithm that jointly uses the number of unmatched bits of decoding results and the received signal power.

Fig. 5 shows the power mean, power variance, and Hamming distance for correct and corrupted symbols for one of our traces (we detail in §V our experiment setup). This plot captures the intuition behind our algorithm; for low and high Hamming distances, we can classify a symbol with a high confidence as successfully decoded or corrupted, respectively. For intermediate values, the Hamming distance alone is not enough (as dark and light gray dots indicating correct and corrupted symbols, respectively, overlap in this Hamming distance range in Fig. 5). The input power though can assist to detect corruption for these cases.

Our symbol error detection algorithm works as follow: A symbol is classified as correct if its Hamming distance is lower than τ_l , whereas those with Hamming distance $\geq \tau_h$ are classified as erroneous symbols. For symbols with a Hamming distance between the decision boundaries τ_l and τ_h , we check the channel SINR. In case SINR is lower than τ_s , we mark the symbol as erroneous. The SINR measures the channel noise and interference, it therefore reflects to what extent the channel preserves the correlation between transmitted and received symbols. We find this joint estimation method to be slightly better and much stabler than just setting a threshold on the number of unmatched bits. The threshold values (τ_l , τ_h , τ_s) are configurable system parameters which we derive empirically.

E. CTI-aware Packet Recovery:

CrossZig mitigates CTI through an adaptive packet recovery scheme. It observes error characteristics and adjusts the recovery mechanism settings accordingly. The recovery scheme integrates two recovery mechanisms, namely cross-layer based packet merging and adaptive RS coding. Cross-layer packet merging tackles long burst errors which are beyond coding recovery capabilities. Adaptive RS coding targets packets that can be recovered with moderate code redundancy, i.e., low bit error ratio. We first explain how these mechanisms work independently and later we describe how they are integrated in CrossZig.

Cross-layer based Packet Merging (CPM). Recovering packets with high error rate with coding is inefficient or even in

Fig. 7. Two consecutive corrupted transmissions of the same packet. L_c highlights the overlapped interfered segment in the two packets which cannot be recovered with basic packet merging. We exploit MRC to recover the segment L_c .

some cases not feasible. Hence, once such pattern is detected, our adaptive mitigation scheme instructs the use of CPM.

Our CPM is realized at two stages; symbol level and signal level. The symbol-level packet merging reconstructs the target packet by combining correct symbols from two packet instances. CrossZig identifies correct symbols using our error localization mechanism. As long as we receive one correct instance of every symbol, this technique allows us to reconstruct the original packet with high confidence. For symbols that are corrupted on all received instances (see Fig. 7), we combine at the signal level by means of *Maximum Ratio Combining* (MRC).

In MRC [19], each signal branch is multiplied by a weight factor that is proportional to the branch SINR. That is, branches with strong SINR have a larger weighted factor and are further amplified. The signal and noise power are computed over interference-free PHY header and SFD symbols. Thus, interference signal power is derived as the difference between the approximated target signal and noise power from interfered signal power. Signals of the first transmission and the corresponding retransmission can be represented as:

$$y_1(t) = s_1(t) + i_1(t) + n_1(t) \quad (8)$$

$$y_2(t) = s_2(t) + i_2(t) + n_2(t) \quad (9)$$

The maximal combined signal can be represented as:

$$y(t) = w_1 y_1(t) + w_2 y_2(t) \quad (10)$$

where the weighted coefficients (w_1 , w_2) are computed by the SINR over the sum of signals of commonly corrupted symbols:

$$w_j = \frac{SINR_j}{\sum SINR_i} \quad (11)$$

Time diversity involves transmitting the same information in two distinct times. In Equations 8 and 9, s_j , i_j , n_j represent signal, interference, and noise components of the received signal at time instant j . The noise in each time instance of the channel is independent of the signal. The signals s_1 and s_2 are essentially identical. In contrast, i_1 and i_2 are not identical and most probably completely uncorrelated.

After combining, target signal s components are amplified. This yields an increase in the power of the target signal, thus increases its decodability chances. Interference i and noise n components can be either canceled, attenuated, or amplified; However on average, the SINR is increased (see Fig. 6).

For the MRC-based combining, co-phasing of all signals is necessary to avoid target signal cancellation. Instead of performing computationally complex frequency and

Fig. 8. Cross-Layer architecture design of CrossZig. Dark gray shaded boxes depict our components added to standard 802.15.4 stack.

initial phase offset compensation for each signal, we estimate the relative phase offset between two signals. Since they are transmitted and received by the same sender and receiver, their frequency offsets are the same. Thus, after compensating the relative phase offset by utilizing the preamble signals in each packet, we can correctly align them.

RS-Adaptive Coding. When the system observes high ratio of corrupted packets, FEC, which adds redundant information to the payload, is used to potentially recover the errors and possibly avoid retransmissions. Although FEC codes are widely used in communications systems, selecting the right coding scheme and setting the right level of redundancy for constrained devices is not trivial. We investigate how to derive an adaptive encoding strategy for low-power devices under various interference patterns, where transmitted redundancy is bounded to the inferred error patterns.

We choose *Reed-Solomon (RS) codes*, which are practical for constrained devices [20], [21]. RS codes are systematic codes, i.e., redundancy data is appended to unaltered source data. This results in no decoding overhead when no error is present. The RS codes are block-based error correcting. The length of the redundant parity (t) defines the maximum number of corrupted blocks a receiver can successfully recover within a partially corrupted packet. RS coding can correct up to $t/2$ and detect up to t block errors. It works well for error patterns that fall under the recovery capacity of the parity check.

The primary goal of our adaptive strategy is to increase packet recovery rates, yet minimize the redundancy overhead on the channel to meet the energy constraints of low-power radios. To realize this, we infer error information from physical layer hints, as discussed in §III-D. This allows us to adaptively derive a redundancy level based on the symbol error rate in the window of received packets. In case the number of corrupted packets is low (i.e., reasonable *Packet Reception Ratio (PRR)*), RS-coding is not triggered which allows CrossZig to avoid introducing redundancy overhead in good links. For the window w_i of observations, we calculate the redundancy level R_i based on the observed degree of corruptions (i.e., R_i is derived from the average number of erroneous symbols per corrupted packet in w_i and is bounded by an upper bound ($R_i \leq \frac{1-PRR_i}{PRR_i} \times \text{packet_length}$)). CrossZig triggers adaptive

Fig. 9. Layout of the online evaluation experiment setup. SDR 802.15.4 Rx-Tx located in an office with a line-of-sight link of 5 m. The interferers are located at locations L1 (1 m), L2 (4 m), and L3 (7 m), where L1 and L2 are in line-of-sight to Rx-Tx and L3 is in non-line-of-sight. Green circles indicate the location of our multiple interferer scenario. Gray squares indicate the location of uncontrolled access points placed in the floor.

coding only if R_i is lower than the calculated upper bound. This allows us to ensure that the introduced redundancy is not significantly higher than potential symbol errors to be recovered. We assume that the corruptions in the upcoming packets follow the trend of our current observations. Hence, the window size should be selected carefully. In our evaluation, we noticed that window sizes of 300 ms to 1 s result in a good and stable performance.

F. System Integration

CrossZig extends the basic 802.15.4 PHY and MAC layers as illustrated in Fig. 8. It provides a single hop reliable delivery mechanism that can counter the CTI effects. The receiver performs packet detection and decoding in a manner similar to the standard 802.15.4. In case a jump in the signal strength is observed during packet reception and the received packet fails the CRC, the receiver initiates the interference detection algorithm discussed in §III-C. If CTI is detected, the transmitter adapts the *Channel Clear Assessment (CCA)* threshold to allow an opportunistic access to the channel. Upon the reception of few partially interfered packets, the system adjusts the initial recovery settings for the next observation window w . Any notable changes on the observed error characteristics trigger the system to adjust the recovery settings.

CrossZig performs the following logic while adjusting the recovery settings. Packet retransmission always carry a fixed, low level of redundancy code. This is used to boost the performance of our CPM: the packets are merged, lowering the BER to a level recoverable with FEC. In case of high packet corruption levels (low PRR, in our settings lower than 75%), RS-adaptive coding is triggered. Here, the redundancy R is derived adaptively based on the observed degree of corruptions. R is derived from the average number of erroneous symbols per corrupted packet in the observation window. Hence, the coding is adapted according to the dynamic interference patterns in the channel. When a packet cannot be recovered with the current redundancy level, CPM is applied.

IV. IMPLEMENTATION

We build a prototype of CrossZig using SDR. For the SDR hardware, we rely on the USRP-N210 [9], equipped with an SBX radio daughterboard [22] as radio front-end.

Fig. 10. Online performance of CrossZig exposed to various types of interferers. The upper plot depicts the goodput ratio for the considered recovery mechanisms. It conveys how much of the transmitted data is useful data, and hence, reflects how well the channel is being utilized. The middle plot depicts the varying redundancy overheads, which affect the radio on time. The bottom plot shows the recovery ratio without considering the cost. CrossZig achieves the highest goodput in most cases, thanks to its adaptivity and a favorable balance between overhead and recovery ratio.

The SBX board incorporates a wide band transceiver that operates from 400 MHz to 4400 MHz, i.e., covers the 2.4 GHz band. For development, we use the GNURadio [8], an open source software toolkit for building software radios.

The transmitter and the non-coherent receiver nodes run 802.15.4 PHY and MAC layers [13]. We modified the receiver PHY to incorporate interference detection logic, error estimation, and channel estimation in our codebase, as described in §III. Moreover, we implement an RS-decoder and the CPM scheme at the receiver side. At the transmitter side, we incorporate the RS-encoder. We implement a virtual feedback channel at host software to carry the receiver feedback to the sender. Note that in our SDR prototype implementation of CrossZig we do not use carrier sense. USRP radios introduce inevitable delays into the processing path of packets, which makes confining with carrier sense strict timing requirements hard to realize [23]. This constraint, however, does not hinder us, as the opportunistic access to the medium in interfered channels is possible without carrier sense. While this is not an optimal solution, it is sufficient to manifest empirically the concepts covered in this paper.

Cross-layer Packet Merging. The standard MRC is carried out on complex signal samples and requires coherent combining at the receiver. In the micro-evaluation of CPM covered in §V-B, we perform the signal alignment offline ahead of the MRC step (trace-based evaluation). This is necessary as our prototype implementation is based on non-coherent receivers. Thus the receiver does not require signals to be synchronized in phase and frequency. To cope with lack of phase offset compensation in our prototype, we carry out MRC on the demodulated SVs instead of the complex signal samples. Given that the quadrature demodulator measures the phase difference of two successive input signal samples, the initial phase offset is no longer an issue. The demodulated soft values of the first transmission and the corresponding retransmission can be represented as follow:

$$y_1[n] = SV_{ideal} + \delta_1[n] \quad (12)$$

$$y_2[n] = SV_{ideal} + \delta_2[n] \quad (13)$$

where $y_1[n]$ and $y_2[n]$ are the soft demodulated values for the n -th symbol in two transmissions, $SV_{ideal} = \pm \frac{\pi}{4}$ is the ideal soft demodulated value for our target signal and $\delta_1[n]$ and $\delta_2[n]$ are the errors caused by interference and noise. Since interference and noise in different transmissions are i.i.d. with zero means, by weighted averaging of the soft value we increase the chances of successful demodulation.

V. EVALUATION

Now we present the experimental evaluation of our prototype implementation on the USRP-N210. In the following, we first define our evaluation objectives and describe the experimental methodology, the considered interferers, the evaluation setup, and the metrics. We continue with a discussion on the system's online performance, followed by a detailed evaluation of the system components, namely CTI detection, error localization algorithm, and the MRC-based packet merging.

Methodology. The ideal experiment setup would evaluate the end-to-end performance of CrossZig using real traffic models with different prominent low-power MAC protocols. However, due to inevitable processing latencies in current software radio platforms, the realization of such an evaluation setup is hard or not feasible with regard to strict time constraint components. Instead, we focus on link performance, by measuring the packet reception rate for various communication links that we subject to external interference sources. Note that the performance degradation under CTI is primarily attributed to starvation or/and discarded corrupted packets. Packet losses, where packets are not successfully detected, account less to the overall performance degradation, and are not directly addressed in this work [2]. Such losses can be resolved by considering better packet detection mechanisms as suggested by [15], [20]. In the second part of this section, we cover the evaluation of individual components of CrossZig. Note that all system components are evaluated empirically. Additionally, in the micro-evaluation, we support part of the empirical results with Matlab simulations, e.g., to show the algorithm's behavior under SINR ranges beyond the empirically-captured ranges.

Fig. 11. The implication of applying fixed levels of redundancy under CTI. The gray-colored area depicts the ranges for recovery ratio under varying interferers per technology.

Interferers. Our set of interference sources includes low/high power, narrow/wide band, channel hopping/fixed frequency, and CSMA/non-CSMA. This represents common underlying properties adopted by most radio technologies. More specifically, as CTI we consider: 802.11 (heavy and light UDP traffic), digital wireless camera, and microwave oven. 802.15.4 is considered as *Intra-Technology Interference*.

Evaluation Setup. The system evaluation is performed in a typical office building. Fig. 9 shows the layout of the experimental setup. Experiments are carried out with controlled single active interferers mentioned above and multiple active interferers. Multiple active interferers are different combinations of single interferers running simultaneously and defined as: *Multiple-1*: microwave oven and wireless camera running simultaneously, *Multiple-2*: microwave oven, wireless camera, and 802.11 with light UDP traffic, and *Multiple-3*: microwave oven, wireless camera, and 802.11 with heavy UDP traffic. The 802.15.4 transmitter-receiver pair was represented by our prototype implementation on USRPs. During the experiments, the 802.15.4 communication link was also exposed to interference from various uncontrolled sources existing in the building. In each experiment, we transmit 6000 packets consecutively with 60 Byte payload, at a 10 ms interval.

Metrics. Within our evaluation we use the following metrics: (a) *Goodput ratio*: defines the ratio of useful received data over total received data. It quantifies the system's efficiency as it reflects both the gain and the transmission overhead together. This metric allows us to observe how well transmitted bytes are utilized. (b) *FEC overhead*: indicates the added transmission overhead which is directly related to energy efficiency, a vital factor in low-power networks. (c) *Packet recovery ratio*: indicates how many of the corrupted packets our recovery mechanisms could recover. The recovery ratio and redundancy overhead show the performance of the considered schemes compared to the baseline where no mitigation scheme took place. Note that in our definition, the basic scheme has 0 recovery ratio and 0 cost. (d) *Precision and Recall*: values are relevant for the performance discussion of symbol error detection, where selection of parameters has an impact on the performance. Precision indicates how many of the identified corrupted symbols are indeed corrupted. Recall indicates how many of the overall corrupted symbols are identified. (e) *Symbol Error Rate (SER)*: is the number of corrupted symbols over the total number of symbols in a received packet.

Fig. 12. Error ratio in discerning the type of interference between CTI and 802.15.4 achieved by the SV-based method (SV) and the correlation-based method (CCF).

A. System Performance

We expose CrossZig first to single active interferers at different distances. The interferers are located first at location L1, then L2, and L3 (see Fig. 9). Second, we consider interference generated from multiple simultaneous sources. Fig. 10 shows the evaluation results achieved by the following recovery schemes: RS-coding with fixed redundancy of 30 Byte, our adaptive coding scheme which selects a redundancy between 0 and 30 Byte based on the average observed SER in the 500 ms window of observations (irrespective of the PRR in the channel), packet merging, and finally CrossZig which combines our cross-layer based packet merging and our adaptive RS-coding scheme.

The error patterns caused by interferers vary as we change the interference types, therefore different experiment settings yield varying performance in terms of goodput ratio, packet recovery ratio, and redundancy overhead.

The RS-fixed scheme achieves the highest packet recovery ratio, but this comes with a fixed 30 Byte redundancy per packet, regardless of channel conditions. This has a negative impact on goodput. This overhead exacerbates for good quality channel conditions which we did not consider in this study. This results in higher in-air time and increased processing for decoding at the receiver side, which are both undesirable for low-power devices. The RS-fixed scheme evaluated in Fig. 10 considers 30 Byte redundancy. Fig. 11 depicts the recovery ratio at different fixed redundancy levels. With a redundancy higher than 20 Byte we do not observe a notable improvement of the recovery ratio. Note that increasing the redundancy has the side-effect of increasing the probability of overlap with interference, hence, reducing the effectiveness. With our RS-adaptive scheme, we observe a similar packet recovery ratio as with the fixed strategy, but at a lower overhead (in average 15 Byte for each packet). This yields a higher goodput.

Packet Merging comes with no FEC overhead because it simply works on the received signal of incoming packets. Its recovery ratio is modest in most cases, except in presence of multiple interferers, because Packet Merging is particularly effective at higher SER levels.

CrossZig improves RS-adaptive which relies only on the observed SER rates for adaptation. In addition CrossZig recovers long error bursts using Packet Merging and is able to keep its cost low under sparse interference. We reach an average packet recovery ratio of 23% overall and up to 50%, for instance for the multiple-3 setup.

Fig. 13. Precision-Recall analysis for symbol error estimation with Power (PW), Hamming Distance (HD), and combination of both (PW_HD). Precision indicates how many of the identified corrupted symbols are indeed corrupted. Recall indicates how many of the overall corrupted symbols are identified.

This is about half of the average packet recovery ratio achieved with the aggressive RS-fixed (40%) over all cases. However, the overall overhead of CrossZig is by a factor of 4.6 lower than the other schemes, and reaches up to a factor of about 20 for the case of WiFi-light. As a result, CrossZig achieves the highest goodput ratios in most scenarios. Note that for fairness we did not compare the performance of CrossZig to the case of no active interferer, where goodput falls drastically for RS-fixed and improves to even higher values for CrossZig.

Conclusion. We show that performing timely adaptation to match induced error patterns from external interference is possible with help of physical layer hints. With this timely adaptation we can achieve better goodput and avoid excessive redundancy which comes at high price for low-power devices.

B. Dive in CrossZig

We carry out an offline micro-benchmark analysis of CrossZig to quantify the performance of its individual components independently. Our traces for this evaluation include the complex signal of 35,875 packets corrupted by interference.

1) *CTI Detection:* We now discuss the performance of our CTI detection scheme introduced in §III-C. We estimate the effectiveness of our scheme in detecting the occurrence of *Intra-* and *Cross-Technology Interference*.

Fig. 12 shows the detection error ratio for both the SV-based and correlation based detection mechanisms. For low SINR ranges under -2 dB, both mechanisms perform well with error rates below 5%. SV-based detection performs well at low SINR because in case of intra-technology interference, the interfering signal can be demodulated by the receiver and this is reflected in lower variations of the soft values. As SV-based detection is the cheaper mechanism, we rely on it for SINR under -2 dB. As the SINR increases (weaker interferer), detecting the source of interference is more challenging. The accuracy of the SV-based scheme degrades sharply, while the correlation-based detection still yields error rates below 10%. Therefore, for SINR greater or equal to -2 dB, we use correlation-based detection. Note that for SINR ranges above 3 dB, the interference signal is very weak and, hence, the target signal is decodable. Consequently, CTI detection is not required for these SINR ranges.

Fig. 14. Simulation performance of Diversity Combining (C) for 802.15.4 signals (S1 and S2) averaging over 1k independent cases subject to Gaussian random interference (GRN), QPSK interference, and 802.15.4 interference.

2) *Symbol Error Localization:* We now discuss the performance of our error localization algorithm introduced in §III-D. Fig. 13 shows the precision and recall of the symbol error detection mechanism using signal power only, decoding Hamming distance only, and using them jointly. This result is aggregated over all the collected traces. The line corresponds to precision and recall for various thresholds (τ_l , τ_h , and τ_s). For our system, we select the thresholds that yield a good balance between precision and recall in the micro analysis, $\tau_l=4$, $\tau_h=10$, and $\tau_s=4$. By combining power and Hamming distance, our symbol error detection approach yields a stable performance with a precision and recall of 82% and 92%, respectively. The average achieved accuracy is $94.3\% \pm 2.4$.

3) *Diversity Combining under CTI:* In this section, we investigate variables that impact MRC performance under CTI which is utilized in our CPM. Moreover, we investigate to what extent MRC can increase the symbol error recovery probability, and put this into the context of recovering packets with bursty errors.

In the context of this work, we exploit time-diversity by combining two interfered copies of the same signal received in different instants of time. We employ the MRC technique for combining the signals. MRC amplifies the SNR of the target signal. The SNR of the combined signal y_c is by factor 2 higher. Therefore, the theoretical SNR gain is 3 dB.

To understand the impact of the interference on the performance of MRC, we first carry out simulations in Matlab. Fig. 14 plots the Bit Error Rate (BER) vs SINR for an interfered 802.15.4 signal before and after MRC. We consider three types of interference here: QPSK signal representing CTI, 802.15.4 representing internal interference, and Gaussian random interference. The time diversity gain from MRC is reflected in the BER drops. As we can see, the MRC gain varies with respect to the type of interference signal. Under interference the gain can exceed the 3 dB expected gain. MRC performs better when the interference signal has an underlying modulation scheme as opposed to noise.

This observation is aligned with our empirical results carried out with the trace-based evaluation. There, the MRC gain for the wireless camera and 802.11 is higher than that for the microwave oven (which is noise radiation). In practice, with MRC we can increase the recovery chance of a symbol by up to 15% which is 2.5x times higher than the random guess. To see how this is reflected in our cross-layer based packet

Fig. 15. Our cross-layer based packet merging mechanism reduces the average SER per packet to 0.11.

merging, we extend our evaluation to packet level. Fig. 15 shows the results of CPM applied to 2 consecutive corrupted transmissions of a packet in our traces. The outcome shows that the mean SER of combined packets is reduced to 0.11 which has good chances to be recovered by low FEC coding on top. Our cross-layer based packet merging can achieve an overall gain of up to 0.34 in SER.

VI. RELATED WORK

Wireless interference is (and has long been) an important topic in wireless communication research. Recent years have seen significant and fundamental contributions to the state-of-the-art interference management, for instance by techniques like interference alignment [24] or joint/coordinated transmission [25], [26]. Nevertheless, these approaches typically require significant computational complexity and/or significant coordination bandwidths, which hinder them applicable for low-power, low-complexity devices of interest in this paper. Hence, in the following we focus on interference mitigation in the unlicensed bands and work related to CrossZig.

Interference Avoidance. Research in this direction aims at *detecting and avoiding interfering signals in space, time, or frequency*. The most common avoidance approach is to employ frequency-based isolation by employing spectrum sensing to identify interference-free channels [27], [28] or adaptive frequency fragmentation [28], [29]. 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. Huang et al. [30] and Boano et al. [31] proposed approaches to avoid interference in time by learning transmission characteristics and the idle cycles of interferers. Radunović et al. [32] proposed an adaptive preamble design to increase the probability of detecting low-power transmissions by high-power competing technologies.

Packet Recovery. Research in this direction aims at *increasing resilience against interference* by bracing PHY and data-link layers with auxiliary mechanisms. For instance, Liang et al. [20] studied the interplay between 802.11 and 802.15.4 and applied a resilience forward error coding scheme against interference. Analogously, some solutions focused on exploiting the temporal effects of interference induced on PHY hints, such as variations in soft errors (softPHY) [14], [15] or RSSI variations [33] to localize interfered segments, hence adapt standard ARQ to retransmit only the interfered segments.

Interference Cancellation. Further physical layer solutions, such as Interference Cancellation, have been considered to combat interference [7], [34], [35]. Here the receiver, with minimal or no coordination from sender, attempts to recover the signal of interest from interference. Halperin et al. [11] utilized Successive Interference Cancellation (SIC) to recover from collisions. The key idea of SIC is that interference signal and target signal are decoded successively. First the receiver decodes the interference signal, i.e., the signal with larger power, afterwards the interference signal is stripped away from the aggregately received signal to get the target signal. Note that these techniques require knowledge about the interfering signal modulation scheme, which makes them not suitable for CTI. Gollakota et al. [3] proposed TIMO, a MIMO design that enables 802.11n to communicate in the presence of CTI. TIMO exploits MIMO capabilities to cancel the interference signal. However, low-power wireless devices are typically single antenna devices, where such approaches are not applicable.

Interference Classification and Signal Detection. Research in this direction aims at *identifying the type of interference technology*. The lack of interference-free channels led researchers to work on novel classification approaches that make networks aware of the type of the existing interference [36]–[39]. It has been shown that when the interference source is known, specialized mitigation approaches can improve the network performance. Researchers explored signal properties by employing signal classification techniques [40] or featuring distinct interferer’s patterns on corrupted packets [37] to build interference classification tools. It is not clear though how these classifiers can be utilized in a systematic way to combat interference. In previous work [41], we address this limitation and propose a system that employs a lightweight machine learning classifier to map the current channel signature to a coexistence strategy. However, this approach requires prior training of the adaptation algorithm which might not always be feasible. Analogously, signal detection techniques [38], [42] for spectrum sensing are important requirements in cognitive radio networks. These techniques enable detection of unused spectrum and sharing of it without causing harm to primary users. This direction has been widely explored in cognitive networks with the focus on detecting known signals in noise. On the contrary, in this work we focus on detecting the type of signal in interfered segments of the packets. Hence, the focus is on signal detection in mixed signals (i.e., interfered signals) where the target signal is mixed with an unknown signal.

Exploiting CTI in Low-power Networks. Recent research efforts focused on exploring opportunities in CTI. For instance [43], [44] harness CTI to beneficially provide security. Others [45], [46] harness channel overlapping between 802.15.4 and 802.11, to allow cross-talk to dispense the role of a dedicated gateway to interconnect these two technologies. CTI is inevitable, hence, utilizing it to provide additional services will enhance spectrum usability. This direction of research is orthogonal to interference mitigation, which is the focus of this work.

Our Approach. Analogously, our work features physical layer hints to infer and recognize interference patterns and harness this to adapt the recovery mechanism. We propose a solution that neither requires interactions with interfered technology nor depends on prior training of the adaptation algorithm, and is agnostic to the interference type. Finally, our system is related to prior work on cross-layer wireless design [3], [15]–[17], [28]. However, our system is optimized to address CTI in low-power and low-complexity radios.

VII. CONCLUSION

Interference is the biggest distress facing wireless networks nowadays, notably in the unlicensed bands. CTI is almost inevitable in these bands and threatens the viability of low-power networks. To address this problem, this paper presents a CTI-aware adaptive recovery mechanism. We investigate how to exploit physical-layer hints to recover from CTI in a low-power environment. Our system combines interference detection, error localization, and an adaptive error recovery mechanism. We do not restrain ourselves with off-the-shelf radios, and resort to SDR for our prototype implementation. Experimental results show that our approach can substantially improve the goodput of 802.15.4 links under various CTI patterns. Moreover, we anticipate that the analysis, insights, and discussions carried out in this paper can inspire further work to address low-power co-existence unconstrained by current chip designs.

Acknowledgments. We thank Friedemann Mattern, our shepherd Tian He, and the anonymous reviewers for their insightful comments. This work was partly supported by a grant from CPER Nord-Pas-de-Calais/FEDER DATA and by VINNOVA, Sweden’s innovation agency.

REFERENCES

- [1] “Estimating the Utilisation of Key License-Exempt Spectrum Bands, Final Report, Mass Consultants Ltd., Ofcom,” 2009.
- [2] A. Hithnawi, H. Shafagh, S. Duquennoy, “Understanding the Impact of Cross Technology Interference on IEEE 802.15.4,” in *WiNTECH*, 2014.
- [3] S. Gollakota, F. Adib, D. Katabi, S. Seshan, “Clearing the RF Smog: Making 802.11n Robust to Cross-technology Interference,” in *ACM SIGCOMM*, 2011.
- [4] S. Hong, J. Mehlman, S. Katti, “Picasso: Flexible RF and Spectrum Slicing,” in *ACM SIGCOMM*, 2012.
- [5] P. Guo, J. Cao, K. Zhang, X. Liu, “Enhancing ZigBee throughput under WiFi interference using real-time adaptive coding,” in *INFOCOM*, 2014.
- [6] S. Souvik, R. Choudhury, S. Nelakuditi, “CSMA/CN: Carrier Sense Multiple Access with Collision Notification,” in *ACM MobiCom*, 2010.
- [7] S. Gollakota, D. Katabi, “Zigzag Decoding: Combating Hidden Terminals in Wireless Networks,” in *ACM SIGCOMM*, 2008.
- [8] GNU Radio Website. [Online]. Available: <http://www.gnuradio.org>
- [9] Universal Software Radio Peripheral, Ettus Inc., www.ettus.com.
- [10] H. Rahul, F. Edalat, D. Katabi, C. Sodini, “Frequency-Aware Rate Adaptation and MAC Protocols,” in *ACM MOBICOM*, 2009.
- [11] D. Halperin, T. Anderson, D. Wetherall, “Taking the Sting out of Carrier Sense: Interference Cancellation for Wireless LANs,” in *MobiCom’08*.
- [12] Kong, Linghe and Liu, Xue, “mZig: Enabling Multi-Packet Reception in ZigBee,” in *ACM MobiCom*, 2015.
- [13] “IEEE 802.15.4 Standard,” IEEE, 2011.
- [14] G. Woo, P. Kheradpour, D. Shen, D. Katabi, “Beyond the Bits: Cooperative Packet Recovery Using Physical Layer Information,” in *ACM MobiCom*, 2007.
- [15] K. Jamieson, H. Balakrishnan, “PPR: Partial Packet Recovery for Wireless Networks,” in *ACM SIGCOMM*, 2007.
- [16] M. Vutukuru, H. Balakrishnan, K. Jamieson, “Cross-Layer Wireless Bit Rate Adaptation,” in *ACM SIGCOMM*, 2009.
- [17] J. Ou, Y. Zheng, M. Li, “MISC: Merging incorrect symbols using constellation diversity for 802.11 retransmission,” in *INFOCOM*, 2014.
- [18] H. Meyr, M. Moeneclaey, S. Fechtel, *Digital Communication Receivers: Synchronization, Channel Estimation, and Signal Processing*, 1997.
- [19] D.G. Brennan, “Linear Diversity Combining Techniques,” *Proceedings of the Institute of Radio Engineers*, vol. 46, no. 1, pp. 1075–1102, 1959.
- [20] C. Liang, N. Priyantha, J. Liu, A. Terzis, “Surviving Wi-Fi Interference in Low-power ZigBee Networks,” in *ACM SenSys*, 2010.
- [21] K. Lin, N. Kushman, D. Katabi, “ZipTx: Harnessing Partial Packets in 802.11 Networks,” in *ACM MobiCom*, 2008.
- [22] SBX, www.ettus.com/product/details/SBX.
- [23] G. Nychis, T. Hottelier, Z. Yang, S. Seshan, P. Steenkiste, “Enabling MAC Protocol Implementations on Software-defined Radios,” in *USENIX NSDI*, 2009.
- [24] V.R. Cadambe, S.A. Jafar, “Interference Alignment and Degrees of Freedom of the K-User Interference Channel,” *IEEE Trans. on Information Theory*, vol. 54, no. 8, pp. 3425–3441, 2008.
- [25] K. Gomadam, V. Cadambe, S. Jafar, “A Distributed Numerical Approach to Interference Alignment and Applications to Wireless Interference Networks,” *IEEE Transactions on Information Theory*, (57:6), 2011.
- [26] D. Gesbert, S. Hanly, H. Huang, S. Shamai, O. Simeone, Y. Wei, “Multi-Cell MIMO Cooperative Networks: A New Look at Interference,” *IEEE Journal on Communications*, vol. 28, no. 9, pp. 1380–1408, 2010.
- [27] H. Rahul, N. Kushman, D. Katabi, C. Sodini, F. Edalat, “Learning to Share: Narrowband-friendly Wideband Networks,” in *SIGCOMM*, 2008.
- [28] L. Yang, W. Hou, L. Cao, B. Zhao, H. Zheng, “Supporting Demanding Wireless Applications with Frequency-agile Radios,” in *NSDI*, 2010.
- [29] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, P. Bahl, “A Case for Adapting Channel Width in Wireless Networks,” in *ACM SIGCOMM*, 2008.
- [30] J. Huang, G. Xing, G. Zhou, R. Zhou, “Beyond Co-existence: Exploiting WiFi White Space for Zigbee Performance Assurance,” in *ICNP*, 2010.
- [31] C. A. Boano, T. Voigt, N. Tsiftes, L. Mottola, K. Roemer, M. Zuniga, “Making Sensor MAC Protocols Robust against Interference,” in *EWSN*, 2010.
- [32] B. Radunović, R. Chandra, D. Gunawardena, “Weeble: Enabling Low-power Nodes to Coexist with High-power Nodes in White Space Networks,” in *ACM CoNEXT*, 2012.
- [33] J. Hauer, A. Willig, A. Wolisz, “Mitigating the Effects of RF Interference through RSSI-Based Error Recovery,” in *EWSN*, 2010.
- [34] D. Divsalar, M. K. Simon, and D. Raphaeli, “Improved parallel interference cancellation for CDMA,” *IEEE Transactions on Communications*, vol. 46, no. 2, pp. 258–268, 1998.
- [35] K. Shankar Kumar and A. Chockalingam, “Parallel interference cancellation in multicarrier DS-CDMA systems,” in *IEEE Communications*, vol. 5, 2004, pp. 2874–2878.
- [36] S. Rayanchu, A. Patro, S. Banerjee, “Airshark: Detecting non-WiFi RF Devices Using Commodity WiFi Hardware,” in *ACM IMC*, 2011.
- [37] 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.
- [38] S. Hong, S. Katti, “DOF: A Local Wireless Information Plane,” in *ACM SIGCOMM*, 2011.
- [39] “Cisco CleanAir,” <http://www.cisco.com/en/US/netsol/ns1070/>.
- [40] K. Lakshminarayanan, S. Sapra, S. Seshan, P. Steenkiste, “RFDump: An Architecture for Monitoring the Wireless Ether,” in *CoNEXT*, 2009.
- [41] A. Hithnawi, H. Shafagh, S. Duquennoy, “TIIM: Technology-Independent Interference Mitigation for Low-power Wireless Networks,” in *ACM IPSN*, 2015.
- [42] W.A. Gardner, “Exploitation of Spectral Redundancy in Cyclostationary Signals,” *IEEE Signal Processing Magazine*, (8:2), pp. 14–36, 1991.
- [43] H. Shafagh, A. Hithnawi, “Poster: Come Closer - Proximity-based Authentication for the Internet of Things,” in *ACM MobiCom*, 2014.
- [44] S. Gollakota, H. Hassanieh, B. Ransford, D. Katabi, K. Fu, “They Can Hear Your Heartbeats: Non-invasive Security for Implantable Medical Devices,” in *ACM SIGCOMM*, 2011.
- [45] Y. Shengrong, L. Qiang, O. Gnawali, “Interconnecting WiFi Devices with IEEE 802.15.4 Devices without Using a Gateway,” in *IEEE DCOSS*, 2015.
- [46] S. M. Kim, T. He, “FreeBee: Cross-technology Communication via Free Side-channel,” in *ACM MobiCom*, 2015.