ErrorSense: Characterizing WiFi Error Patterns for Detecting ZigBee Interference

Daniele Croce, Pierluigi Gallo, Domenico Garlisi, Fabrizio Giuliano, Stefano Mangione, Ilenia Tinnirello
Department of Electrical Engineering, Universitá di Palermo, Italy
Email: *name.surname*@unipa.it

***Abstract*—Recent years have witnessed the increasing adoption of heterogeneous wireless networks working in unlicensed ISM bands, thus creating serious problems of spectrum overcrowding. Although ZigBee, Bluetooth and WiFi networks have been natively designed for working in presence of interference, it has been observed that several performance impairments may occur because of heterogeneous sensitivity to detect or react to the presence of other technologies.**

In this paper we focus on the WiFi capability to detect interfering ZigBee links. Despite of the narrowband transmissions performed by ZigBee, in emerging scenarios ZigBee interference can have a significant impact on WiFi performance. Therefore, interference detection is essential for improving coexistence strategies in heterogeneous networks. In our work we show how such a detection can be performed on commodity cards working on time and frequency domain and also analysing data in the *error domain*. Errors are monitored and classified into error patterns observed in the network in terms of occurrence probability and temporal clustering of different error events. Through statistical analysis we are able to detect the presence of ZigBee transmissions measuring the errors raised by the WiFi card.

***Index Terms*—wlan, 802.11, 802.15.4, frame error detection, wireless coexistence.**

I. Introduction

Wireless technologies are changing many aspects of human life. From wireless communications such as satellite or cellular, going to WiFi, Bluetooth and ZigBee, a real "wireless revolution" is transforming and innovating the way technologies are conceived and used. Many of these technologies, especially the ones designed for Local Area Networks (LANs), Home Area Networks (HANs) or Personal Area Networks (PANs), usually make use of free unlicensed ISM radio bands which are becoming increasingly popular and crowded due the widespread dissemination of wireless technologies.

More recently, the success of ZigBee-based networks has increased the problem of cross-technology interference between coexisting wireless applications. Indeed, ZigBee is adopted in many PAN or HAN applications including house and building automation, smart metering systems, surveillance systems, health care monitoring, game remote controllers and so on. With the increased penetration of these new applications, interference will deteriorate radio quality further and, thus, it is important and urgent to provide effective tools which can guarantee a peaceful coexistence of all these applications.

In this paper we specifically deal with ZigBee and WiFi technologies. Despite the fact that many mechanisms have been included in the relevant 802.11 and 802.15.4 standards to cope with interference (e.g. carrier sense, adaptive modulation and coding, signal spreading), both technologies can significantly suffer in presence of the other one [1]. The phenomenon is even more impressive if we consider that the two technologies are pretty heterogeneous in terms of bandwidth (2 MHz for ZigBee and 20 MHz for WiFi) and transmission power (e.g. 0 dBm for ZigBee and 20 dBm for WiFi). Moreover, ZigBee applications are typically low rate, while WiFi networks exhibit abundant channel idle space in time domain [2]. As a matter of fact, the main problems arise because of these heterogeneous features, including frame transmission times and carrier sense capabilities [1].

A critical aspect for improving the spectrum sharing and mitigating the WiFi/ZigBee reciprocal interference, is the correct identification of coexistence problems, which in turn can serve as basis for some inter-technology coordination mechanisms. While state-of-the-art solutions for detecting coexistence problems in WiFi networks have mainly worked on the characterization of RSSI samples observed at different frequencies and with varying temporal gaps, our mechanism is based on the analysis of the *error domain*, i.e. on the classification of error events and on the time intervals between their occurrence. Statistics of these errors are widely available on many WiFi *commodity* NICs and can be easily exploited to improve interference detection and troubleshooting algorithms of wireless networks. In this paper we investigate the feasibility of using these error statistics for building new classifiers (or improving already existing ones), deferring the actual implementation of the classifiers to a future work.

After a brief review of the some literature solutions (section II) and technology features of the WiFi/ZigBee networks (section III), we focus on the possibility to detect ZigBee frames in WiFi networks with the analysis of the error patterns caused by this interference (sections IV and V). Experimental results show that the approach is promising and suitable for further extensions as described in the concluding remarks.

II. Related Work

Several analytical and simulation models, as well as experimental studies, have been proposed for characterizing the

This work has been partially supported by the Italian national research project PON 04 i-NEXT.

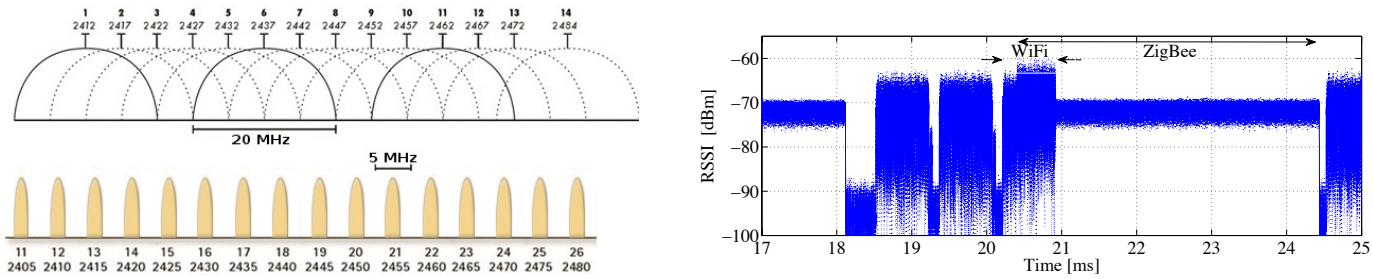

(a) (b)

Fig. 1. Overlapping between WiFi and ZigBee technologies: frequency analysis (available channels) (a) and temporal analysis (RSSI samples) (b).

cross-technology interference in ZigBee and WiFi networks [1], [3]. While early studies mostly focus on the analysis of ZigBee performance degradation in presence of WiFi interference, it has been shown that significant throughput reductions can also be observed in WiFi networks [1], [4]. This phenomenon has been justified by considering two different main reasons: i) an intrinsic reason, due to vendor-dependent implementation choices that in some cases make difficult the detection of non-WiFi modulated signals or introduce latency times in the receiver operations [5]; ii) an extrinsic reason, due to the higher time resolution needed by ZigBee for detecting channel activity and preventing collisions [6], [7].

In such a scenario, it is often required to make orthogonal ZigBee and WiFi transmissions. Early solutions which detect interference and simply choose a better channel to transmit are becoming not viable because of the increasing number of technologies and applications in the market. Other solutions rely on complex and expensive radio transceivers to communicate with multiple protocols and different technologies [8], or increase the robustness of the transmission with use of error correction codes or multiple antennas [9]. Different approaches have considered the possibility to introduce some indirect forms of coordination between the two technologies, based on opportunistic exploitation of WiFi temporal spaces [5], channel reservations [6] by using an additional ZigBee channel for making the channel busy for WiFi stations, or by means of simple forms of adaptive redundancy [7].

Obviously, an important component of any coordination strategy is detecting the coexistence problem, i.e. identifying the presence of two overlapping ZigBee and WiFi networks. The monitoring of heterogeneous RF signals on ISM bands has been specifically addressed in [10], where it is proposed a design of a monitoring module for GNU radio able to quickly identify the transmitting technology and demodulate with the correspondent receiver implementation. Although the approach is very effective, it is based on a dedicated hardware. The possibility to identify WiFi signals by using commodity ZigBee nodes have been explored in [11] and [12]. The approach proposed in [11] is based on the analysis of temporal samples of link quality indicators and RSSI values, as well as on the identification of the portions of ZigBee corrupted packets to be compared with the typical WiFi transmission times. A similar temporal analysis is carried out in [12] with the aim to find periodic interference signatures caused by WiFi beacons and enabling the detection of WiFi networks by using a low-power monitoring interface. Finally, the possibility to detect ZigBee and other interference sources by means of WiFi commodity cards is explored in [13] by using an 802.11n PHY able to read RSSI values at different sub-carriers. Complex algorithms are applied to these samples for characterizing spectral, energy and pulse signals that are mapped into a technology classification scheme. While these previous works rely on the classical analysis of the frequency and time domains, in this paper we study the error domain, i.e. the errors produced by the interfering technologies.

III. Background

In this section we briefly recall some key differences between 802.11g and 802.15.4 MAC/PHY layers that are relevant for understanding our interference detection scheme.

Channels. Both WiFi (802.11g-based) and ZigBee work on the 2.4 GHz ISM band. Each WiFi channel is 20 MHz wide and is spaced of 5 MHz from the adjacent ones. ZigBee channels have only 2 MHz of bandwidth with 3 MHz of inter-channel gap bands (i.e. the center frequencies maintain the spacing of 5 MHz from the adjacent channels, as shown in figure 1-a). In addition, the channels in the two standards match in such a way that one WiFi channel overlaps with exactly four 802.15.4 channels. In practice, since most WiFi networks use channels 1, 6 and 11, few ZigBee channels (15, 20, 25 and 26) are sometimes free from interference. The two technologies also use different transmission powers, since WiFi transmissions are typically performed at 15 or $20dBm$, while ZigBee transmissions can span in the range $[-25, 0]dBm$.

Frames and Rates. Since the two technologies have been defined for different applications (mostly machine-to-machine applications for ZigBee and Internet-based applications for WiFi), the frame size, the coding and the transmission rates considered by the two standards are quite different. ZigBee frames are small, with a maximum payload of only 127 bytes. Bytes are organized into 4-bit symbols that are mapped into 16 pseudo-random sequences of 32-chip transmitted at 2 Mchip/s (i.e. 250 Kbps), which correspond to a frame transmission interval of about $4ms$ for the maximum frame size. Conversely, WiFi frames are much longer, with a maximum frame size of 4096 bytes and multiple OFDM modulations and coding schemes available (from 6 Mbps up to 54 Mbps).

Clear Channel Assessment (CCA). The MAC protocols defined in ZigBee and WiFi are based on CSMA. However, the channel access timings of the two protocols are completely different: the ZigBee backoff slot is set to 320 μs and the WiFi one to 9 μs. This difference is also reflected on a different granularity at which CCA samples are collected. Specifically, during a backoff slot, ZigBee spends 128 μs for detecting the channel activity and 192 μs to switch from reception to transmission mode. If a WiFi transmission is originated during this switching time, it cannot be detected by the ZigBee node (as shown by the USRP trace depicted in figure 1-b).

IV. Detecting Exogenous Interference in WiFi

Our work is motivated by the observation that the receiver errors generated by exogenous RF signals (i.e. non-WiFi modulated signals) exhibit significant differences (in terms of occurrence probability and error intervals) from the ones generated by collisions with other WiFi transmissions. Indeed, in case of coexistence with other technologies, it is possible that the receiver of commodity WiFi cards is triggered by external RF signals. The receiver activation depends on its sensitivity and settings (e.g. the AGC gain) and in some cases is even due to background noise.

A. Classification of Receiver Errors

Regardless of the specific receiver implementation, errors occurring while demodulating a WiFi packet can be categorized into: *i*) an error on the PLCP parity check; *ii*) an error on the FCS checksum of the MAC frame; *iii*) one or more errors in the header fields which make them invalid (either in the PLCP or MAC headers). For example, invalid headers occur if the received frame is too long or too short compared to the value indicated in the LENGTH field of the PLCP header or the protocol version is different from 0 (which is the normal value for current 802.11 standard). These errors have different probabilities to occur depending on the channel conditions and on the power of the received WiFi signal.

B. Error Occurrence Probability

The errors generated by cross-technology interference have much different patterns compared to errors typical of WiFi transmissions. Indeed, in case of wide-band noise and exogenous interference signals, errors may appear randomly at any point during the time the demodulator is active, while for WiFi modulated signals error statistics vary during the frame reception and depend on frame length and rate. For example, PLCP errors have much lower probability to appear compared to bad FCS, because the PLCP transmission is usually more robust and shorter than the rest of the frame. In case the demodulator reveals random bits (i.e. in presence of interference), the probability of having a specific error heavily depends on the format of the expected frame. Figure 2 summarizes the error probability observed when an 802.11g receiver is triggered by non-WiFi modulated signals. Since the PLCP header has one bit only for parity checks, on average one half of the frames should be classified as frames with *Bad PLCP*. However, the receiver can rely also on the RATE field of the header for detecting *Bad PLCP* errors: since the RATE field is 4 bits long while only 8 modulation rates are admitted (out of the 16 possible values), the *Bad PCLP* error probability increases to 3/4.

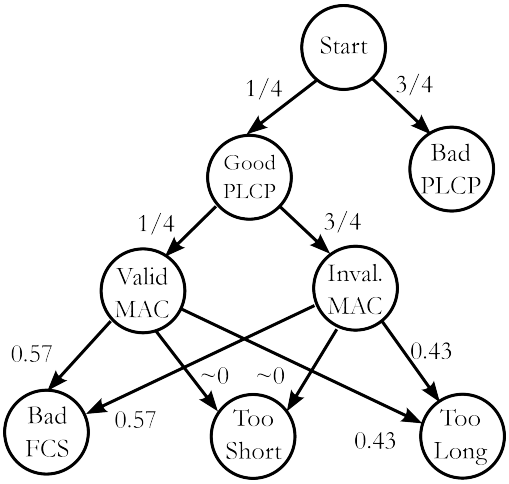

Fig. 2. Error events and relevant probabilities during cross-technology interference.

When a *Bad PLCP* is not detected (25% of the times), the receiver will leave the transceiver on and will continue demodulating until another error is reached, i.e. *Too Long*, *Too Short* or *Bad FCS*. In particular, the LENGTH field in the PLCP header is 12 bits long (values between 0 and 4095) while the length of a WiFi frame is generally between 14 and 2346 Bytes. Therefore, the frame will be considered *Too Long* with probability $1 - 2346/4096 \approx 0.43$ and *Too Short* with probability $14/4096$. The FCS is 32 bits long which means that the probability of having a random sequence with good FCS is only 2^{-32} and, with high probability, a *Bad FCS* error will appear when the frame is not *Too Short* or *Too Long* (~ 0.57).

Finally, an *Invalid MAC Header* error occurs when the 2 bits of the VERSION field in the MAC header are not 0, thus this error occurs 3/4 of the time. However, in this case the transceiver does not suspend the reception but continues until another error is encountered. When the errors detected by a WiFi station closely follow these statistics, it is very likely that interference is generated by non-WiFi modulated signals.

C. Frequency and Time Analysis

Together with the analysis on error statistics, which provides an indication about the existence of RF exogenous signals, in order to classify ZigBee interference it is possible to perform additional tests working on the frequency and time domains.

For example, in [13] it is suggested to sequentially move the WiFi monitoring card to the adjacent channels for scanning the experienced interference with steps of 5 MHz: in case of sudden disappearance of the RF signals when moving from one channel to the next one, it can be assumed that interference was due to a narrow-band ZigBee channel. For example, if the interfering ZigBee node is transmitting on 802.15.4 channel 11, the interference will produce errors on WiFi channel, 1

Receiver Event	Description
Bad PLCP	Parity Check Failure on PLCP Header
Good PLCP	PLCP Header is okay
Too Long	Frame longer than 2346 bytes
Too Short	Frame shorter than 16 bytes
Invalid MAC Header	Protocol Version is not 0
Bad FCS	Checksum Failure on frame payload and MAC Header
Good FCS and RA match	Correct FCS matching the Receiver Address
Good FCS and not RA match	Correct FCS not matching the Receiver Address

TABLE I
RECEIVER EVENTS REPORTED BY BCM4318 CARDS.

but no errors will appear on channels 2, 3, 4 or 5. The results are completely different if the same test is made for other types of interferes such as Bluetooth and microwave ovens. On one side, Bluetooth uses frequency hopping so errors statistically appear on all WiFi channels; on the other side, microwave ovens continuously "sweep" on certain frequencies and have clear ON-OFF patterns.

An alternative solution, also considered in our tests, is performing a time domain analysis of error occurrences. When multiple errors are generated in burst, it is possible to map the error burst duration into an estimation of the interference typical timings. When these timings are compatible with ZigBee access and transmission intervals (e.g. they last for time intervals that may be as long as $4ms$), it is likely that the interference is due to ZigBee frames.

V. EXPERIMENTAL RESULTS

In this section we describe the main results of our experimental campaign devised to prove the feasibility of ZigBee detection with the error statistics of commodity WiFi cards. The experiments have been carried out in our lab at the University of Palermo, in different hours of the day (i.e. under uncontrollable interference from other WiFi networks), by placing a monitoring WiFi card in the same room with two ZigBee nodes and two other WiFi nodes. The transmitting ZigBee and WiFi nodes have been configured for working on different interfering and non-interfering channels, while their reciprocal distance has been set to a few meters.

WiFi monitoring and transmitting nodes employ a Broadcom bcm4318 card, which is able to collect statistics about different receiver events (summarized in table I) that can be easily mapped in the errors discussed in section IV. Although all the required events are tracked by the card, the temporal analysis of these errors is affected by the receiver implementation and in particular by its reaction to the detection of false or bad preambles. Indeed, the time interval in which the demodulator is switched on and off in presence of non-WiFi modulated signals depends on the card internal design. We tried to preliminary characterize these timings under different interference conditions (not only Zigbee frames, but also Bluetooth and microwave ovens) in order to have some preliminary findings about the granularity of consecutive errors. We noticed that, in case of high interfering power (namely, interfering power higher than $-70dBm$), the receiver tries to detect a preamble every 1 ms (channel is continuously sensed busy), while it

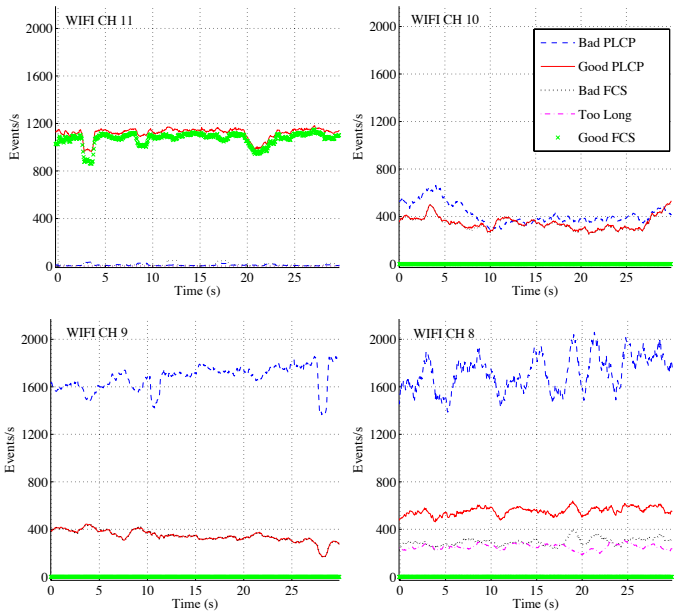

Fig. 3. Receiver events detected by the monitoring WiFi node when another WiFi link is active on the same channel or adjacent channels (10, 9 and 8).

stops and start the demodulator much more frequently in case of interfering signals with lower power (the received signal power is close to the background noise power).

Two types of ZigBee nodes where used in our testbed. Commercial Zolertia Z1 motes, based Texas Instruments CC2420 transceiver, and two self-made nodes based on Microchip MRF24J40 transceiver. Both transceivers are 802.15.4 compatible and, in the experiments, they both generated the same patterns of errors. For ease of presentation, the results shown in the paper are based on the MRF24J40 transceiver only.

A. Error Rates under WiFi interference

Figure 3 summarizes the receiver events detected in our experiments when the monitoring WiFi node is tuned on channel 11 and one WiFi link at 36 Mbps is active (with a saturated traffic source) on varying channels.

In the top-left figure, the link is set-up on the same channel 11. All the frames are detected with good PLCP and almost all the frames have also a correct checksum (i.e. the red good PLCP curve and the green good FCS points almost overlap). When the link is moved on the adjacent channel 10, the monitoring station is able to correctly synchronize about one half of the frames (50% of the PLCP headers pass the parity check and have good rate values) which deterministically result in a failed FCS. Moving the link to the next channel 9 does not affect the frame synchronization probability, but significantly increases the detection of bad PLCP errors which reach over 1700 errors/s. This is due to the fact that when the receiver is not able to correctly synchronize the frame preamble, consecutive trials can be performed during the reception of the same frame and an higher number of error events can be generated for the same frame. In any case, the vast majority of good PLCP frames end up to have bad FCS and the *Too Short* and *Too Long* errors are still close to 0.

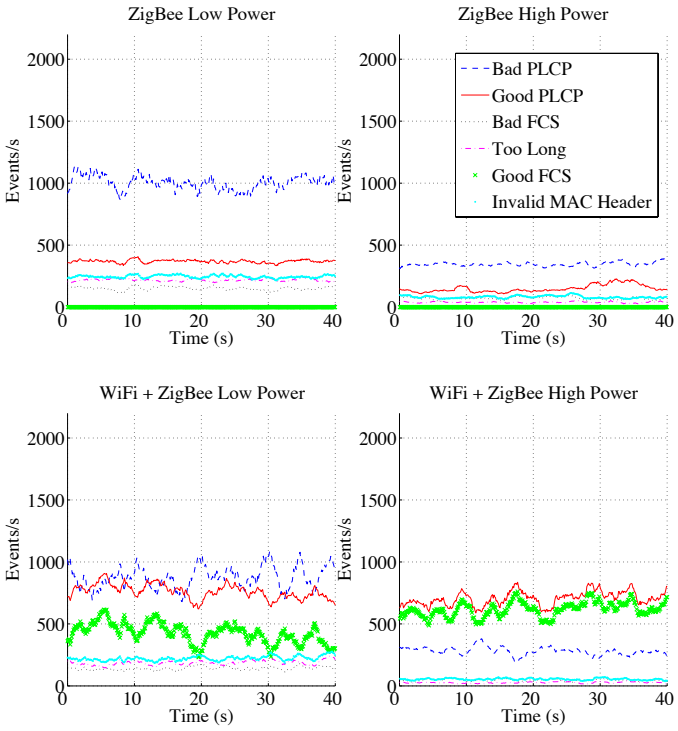

Fig. 4. Receiver events detected by the monitoring WiFi node when a ZigBee link is active on channel 23 (top figures) and an additional WiFi link is active on channel 11 (bottom figures).

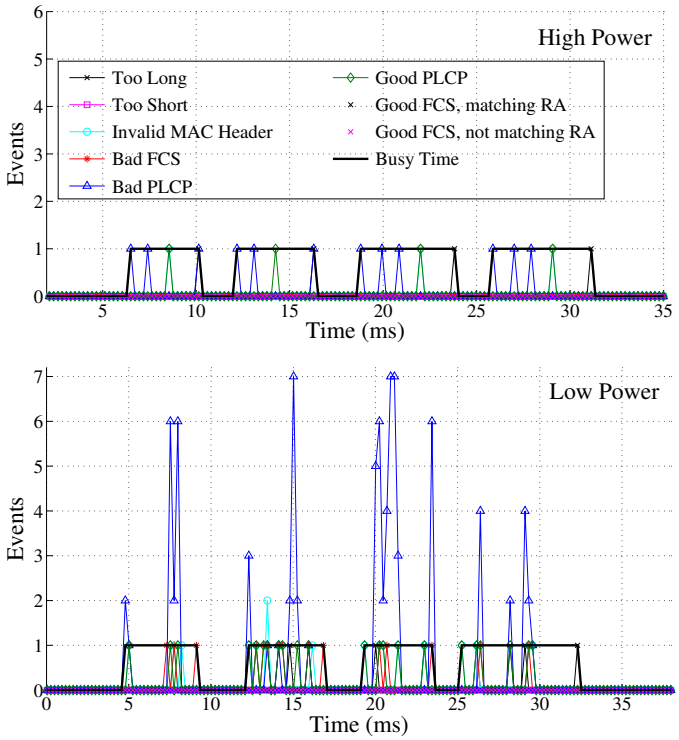

Fig. 5. Clusters of WiFi errors corresponding to the reception of ZigBee frames at high power (top) and low power (bottom).

In the bottom-right figure, the monitoring stations is observing events related to a WiFi link configured on channel 8, i.e. 15 MHz apart from the monitoring channel. From the error statistics, it results that now the WiFi link triggers a number of error events, including the detection of too long frames (46% of the good PLCPs), which follow the statistics described in section IV, as in the case of non-WiFi modulated signals.

B. Error Rates under ZigBee interference

The top plots shown in figure 4 summarize the receiver events detected in our experiments when the monitoring WiFi node is tuned on channel 11 and one ZigBee link with maximum frame size is active on channel 23. The experiments have been repeated by tuning the ZigBee transmission power to $-20dBm$ (low-power case) and $0dBm$ (high power case). An additional WiFi link is considered in the bottom plots. In all the cases, the presence of ZigBee interference can be revealed by the occurrence of too long frames corresponding to the expected ratio of bad PLCP errors.

Although the ZigBee link has been set to a constant traffic in saturation, the total number of events detected in case of low transmission power (top-left figure) is much higher than the corresponding number detected in case of high transmission power (namely, about 1350 events/s in the low power case and 380 events/s in the high power case). This is due to the probability to generate more error signals during the modulation of the same frame, because of more frequent receiver trials to restart the preamble detection. Regardless of the events rate, the ratio between good and bad PCLP in case of ZigBee only transmissions is exactly 1/4 and 3/4 of the total. Similarly, the statics observed on the other type of errors follow the expected characteristics as well.

In case of overlapping WiFi transmissions, the number of good PLCP events is obviously increased of the same amount of injected WiFi frames. A portion of these events, corresponding to the WiFi frames, also result in good FCS events, while the random good PLCP events generated by the ZigBee interference are mapped into too long frames and bad FCS events with probability 0.57 and 0.43. The phenomenon is more evident in the bottom-right figure, where the number of random events is lower and the good PLCP and good FCS curves almost overlap. Similar results were also obtained in hidden-terminal conditions where the number of errors increases due to collisions but the error patterns connected to ZigBee transmissions still remain the same.

C. Error Temporal Analysis

Previous experiments prove that non-WiFi modulated signals can be detected by observing the occurrence rates of different receiver events. In order to classify non-WiFi interfering signals as ZigBee signals, we complement these statistics with the temporal analysis of error bursts. Figure 5 shows an exemplary temporal trace of receiver events in both the cases of high power and low power ZigBee transmissions with maximum payload size. When the interfering signal is high, the receiver employed in the Broadcom card is reset every ms for retrying to synchronize a preamble. At each reset, a good or bad PCLP event occurs with probability 1/4 and 3/4. This implies that during the reception of the ZigBee frame and corresponding acknowledgement (if any),

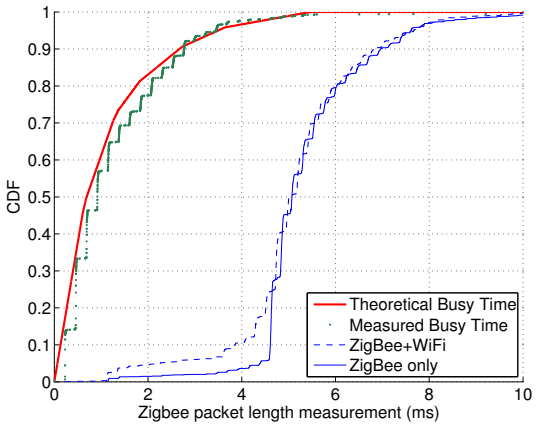

Fig. 6. Cumulative distribution of the ZigBee frame length estimation and comparison with the virtual busy time distribution

the receiver generates a burst of events whose duration is about 4 ms (for unacknowledged frames) or 4.5 ms (for acknowledged frames). For example, in the top part of figure 5, it is possible to easily recognize four consecutive ZigBee frames, with errors spaced about 1 ms from each other. In case of low power transmissions (bottom part of the figure), the demodulator reset is no more regular and more receiver events are generated during each frame transmission. The figure also shows the busy time intervals measured by the monitoring WiFi node. Since the card implements both the actual and virtual carrier sense mechanism, in case of good PLCP events with valid headers, the card will assume that the channel will be busy for a time interval corresponding to: i) a frame length uniformly extracted in the range 14-4096 bytes, and ii) a transmission rate selected with equal probability (namely, 1/8) among the available ones. Specifically, the virtual duration is computed as the number of bytes indicated in the LENGTH field divided by the rate indicated in the RATE field. This explains why, when a good PLCP is raised during the reception of a ZigBee frame, the actual busy time (i.e. the maximum between the frame duration and the virtual busy time generated by the random bits) can exceed $4.5ms$, as in the last frames shown in the figure.

This effect is more evident in figure 6, which shows the cumulative distribution of ZigBee frame length estimates. To isolate the effect of the random virtual duration, the green points in the figure quantify the virtual busy time measured in our experiments from good PCLP events occurred during the ZigBee interference. The curve matches pretty well the theoretical red curve of virtual busy times, corresponding to the random combinations of the RATE and LENGTH fields described above. The blue curves quantify the ZigBee length estimated considering both the actual and virtual carrier sense. Since the probability that the virtual duration is higher than $4.5ms$ is very low, when the good PLCP events occur at the beginning of the ZigBee frame (or when they do not occur at all), the length estimate is equal to $4.5ms$. When the good PLCP event occurs after 2 or 3 ms from the beginning of the frame, the length estimate is equal to the already elapsed time plus the random virtual duration. Similar results have been obtained when ZigBee coexists with a WiFi link or when it transmits at low power.

VI. Conclusions and Future Work

This work has been motivated by the need of introducing novel coordination mechanisms for solving or mitigating the interference suffered by overlapping Zigbee and WiFi networks, in the emerging scenarios of ISM bands overcrowding and increasing ZigBee traffic. In many cases, arbitrators and/or jammers with multiple interfaces have been identified as the only possible approach to support coordination.

To avoid the usage of these arbitrators, we investigated on the possibility to detect ZigBee interference by using commodity WiFi cards. Differently from previous solutions, our approach is based on the analysis of the error signals generated by WiFi receivers when triggered by non-WiFi modulated signals. We prove that the statistics of these signals and the duration of the error bursts can be effectively correlated to the presence of non-WiFi signals and to the typical access and transmission timings of the interfering technology.

Although in this work we just focused on the ZigBee detection problem from WiFi terminals, we are also considering the possibility to conversely detect WiFi transmissions from commodity ZigBee stations. Additionally, we are implementing some forms of inter-technology communication protocols by opportunistically exploiting the generation of error patterns with different durations. Inter-technology communications would allow to easily manage spectrum sharing and channel reservations among overlapping networks.

References

[1] S. Pollin, I. Tan, B. Hodge, C. Chun, and A. Bahai. Harmful Coexistence Between 802.15.4 and 802.11: A Measurement-based Study. In Proc. of CrownCom, 2008.
[2] R. Chandra, R. Mahajan, V. Padmanabhan, and M. Zhang. Crawdad data set microsoft/osdi2006 (v. 2007-05-23), 2007.
[3] Y.S. Soo, S.P. Hong, H.K. Wook. Mutual interference analysis of IEEE 802.15.4 and IEEE 802.11b. In Comp. and Telecomm. Netw., 2007.
[4] R. Gummadi, D. Wetherall, B. Greenstein, S. Seshan. Understanding and Mitigating the Impact of RF Interference on 802.11 Networks. In Proc. of ACM SIGCOMM '07, Pages 385-396.
[5] J. Huang; G. Xing; G. Zhou; R. Zhou. Beyond Co-existence: Exploiting WiFi White Space for ZigBee Performance Assurance. ICNP, 2010.
[6] X. Zhang, K. G. Shin. Enabling Coexistence of Heterogeneous Wireless Systems: Case for ZigBee and WiFi. In Proc. of ACM MobiHoc '11.
[7] C.-J. M. Liang, N. B. Priyantha, J. Liu, and A. Terzis. Surviving Wi-Fi Interference in Low Power ZigBee Networks. In Proc. of SenSys 10, pages 309-322, 2010.
[8] R. Gummadi, H. Balakrishnan, and S. Seshan. Metronome: Coordinating Spectrum Sharing in Heterogeneous Wireless Networks. 1st Int. Workshop on Communication Systems and Networks (COMSNETS), 2009.
[9] S. Gollakota, F. Adib, D. Katabi, and S. Seshan. Clearing the RF smog: making 802.11n robust to cross-technology interference. In Proc. of ACM SIGCOMM 11, pages 170-181, 2011
[10] K. Lakshminarayanan, S. Sapra, S. Seshan, and P. Steenkiste. RFDump: An Architecture for Monitoring the Wireless Ether. In Procs. of CoNEXT 09, Dec. 2009.
[11] F. Hermans, L. Larzon, O. Rensfelt, P. Gunningberg. A Lightweight Approach to Online Detection and Classification of Interference in 802.15.4-based Sensor Networks. In ACM SIGBED Review - CONET 2012, Vol. 9, Issue 3, July 2012, Pages 11-20.
[12] R. Zhou, Y. Xiong, G. Xing. ZiFi: Wireless LAN Discovery via ZigBee Interference Signatures. In Proc. of ACM Mobicom 2010.
[13] S. Rayanchu, A. Patro, and S. Banerjee. Airshark: detecting non-WiFi RF devices using commodity wifi hardware. In Proc. of IMC 2011.