

Learning from Errors: Detecting ZigBee Interference in WiFi Networks

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

Abstract—In this work we show how to detect ZigBee interference on commodity WiFi cards by monitoring the reception errors, such as synchronization errors, invalid header formats, too long frames, etc., caused by ZigBee transmissions. Indeed, in presence of non-WiFi modulated signals, the occurrence of these types of errors follows statistics that can be easily recognized. Moreover, the duration of the error bursts depends on the transmission interval of the interference source, while the error spacing depends on the receiver implementation.

On the basis of these considerations, we propose the adoption of hidden Markov chains for characterizing the behavior of WiFi receivers in presence of controlled interference sources (*training phase*) and then run-time recognizing the most likely cause of error patterns. Experimental results prove the effectiveness of our approach for detecting ZigBee interference.

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

I. INTRODUCTION

Recently, the success of ZigBee-based networks has increased the problem of ISM bands overcrowding. Indeed, ZigBee is adopted in many Personal Area Network or Home Area Network 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 granularity [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* cards and can be easily exploited to improve interference detection and troubleshooting algorithms of wireless networks. Specifically, in this paper we model the behavior of the WiFi receiver in presence of non-WiFi interfering sources in order to define a scheme for detecting ZigBee interference.

After a brief review of the some literature solutions (section II), we analyze the theoretical and experimental error rates caused by this interference (sections III-B and III-C). The interference detection model is introduced in section IV, where we also present our implementation choices. 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 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]. Surprisingly, WiFi vulnerabilities arise despite the fact that many mechanisms have been included at the MAC and PHY layer for guaranteeing robustness to interference. 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 and by sequentially moving a WiFi monitoring card to the adjacent channels 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. Complex algorithms are applied to the RSSI 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. ANALYSIS OF RECEIVER ERRORS

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.

Receiver Event	Description
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
Bad PLCP	Parity Check Failure on PLCP Header
Good PLCP	PLCP headers and Parity Check OK
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.

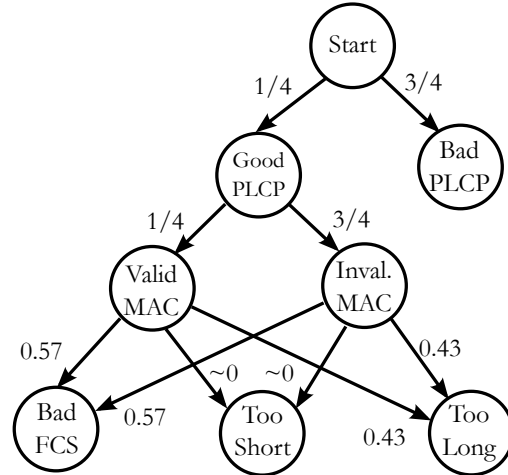


Fig. 1. Error events and relevant probabilities during cross-technology interference.

A. Error Types

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 value in the LENGTH field of the PLCP header is too large or too small compared to the length of a typical WiFi frame or if 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

Name	WiFi ch11		WiFi ch10		WiFi ch8		Microwave		ZigBee HighPW		ZigBee LowPW		Model (%)
	Ev./s	(%)	Ev./s	(%)	Ev./s	(%)	Ev./s	(%)	Ev./s	(%)	Ev./s	(%)	
Bad PLCP	6.5	(0.5)	455.8	(54.8)	1694.2	(75.7)	116.1	(73.6)	266.9	(69.1)	984.4	(72.9)	(75.0)
Good PLCP	1110.0	(99.4)	375.8	(45.2)	542.9	(24.3)	41.7	(26.4)	119.6	(30.9)	366.0	(27.1)	(25.0)
Invalid MAC Header	4.0	(0.4)	286.8	(76.3)	359.1	(66.1)	31.27	(74.9)	84.9	(71.0)	243.1	(66.4)	(75.0)
Good FCS	1067.1	(96.1)	0	(0.0)	0	(0.0)	0	(0.0)	0.0	(0.0)	0.0	(0.0)	(0.0)
Bad FCS	9.0	(0.8)	368.3	(98.0)	285.8	(52.6)	23.1	(55.4)	69.3	(58.2)	147.6	(40.3)	(56.9)
Too Short	0.1	(0.0)	0	(0.0)	1.7	(0.3)	0	(0.0)	0.6	(0.5)	0.2	(0.0)	(0.4)
Too Long	0.2	(0.0)	0.3	(0.1)	251.8	(46.4)	18.5	(44.6)	49.4	(41.3)	218.3	(59.5)	(42.7)

TABLE II

EVENTS CAUSED ON WiFi CHANNEL 11 BY WiFi ON INTERFERING CHANNELS OR DURING ZIGBEE INTERFERENCE (AND NO WiFi TRANSMISSION).

of interference), the probability of having a specific error heavily depends on the format of the expected frame. Figure 1 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 PLCP* error probability increases to $3/4$.

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. Experimental Validation

In order to experimentally validate our theoretical findings, we run some experiments 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 (set on channel 11) 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 dif-

ferent receiver events (summarized in table I) that can be easily mapped in the errors discussed in section III-B. Two types of ZigBee nodes were used in our testbed. Commercial Zolertia Z1 nodes, 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.

We run different experiments by activating a single interference source in each experiment: a WiFi interfering link at channel 11, 10 or 8; a ZigBee interfering link with different transmission powers (0 dBm and -23 dBm); a Microwave oven for generating interference different from ZigBee transmissions. In case of WiFi link on channel 11, all the frames are detected with good PLCP and almost all the frames have also a correct checksum. 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 channel 8, that is 15 MHz apart from the monitoring channel, 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. Now, the error rates follow the statistics of non-WiFi modulated signals and *Too Long* errors appear. Similar statistics are observed for ZigBee and Microwave interference.

IV. INTERFERENCE DETECTION

Although all non-WiFi interfering signals generate receiver errors with similar statistics, a temporal analysis of the receiver events allows to discriminate between different interfering sources. Indeed, multiple events can be generated by the receiver during the same interfering transmission. For example, a checksum failure can follow the detection of a good PLCP, or another (failed or not) synchronization trial can be performed after a bad PLCP event. By organizing the receiver events into bursts generated by the same interfering transmission, it is possible to estimate the timings of the interfering technology.




Fig. 2. Bursts of receiver events corresponding to the reception of ZigBee frames at high power.




Fig. 3. Bursts of receiver events corresponding to the reception of ZigBee frames at low power.

The temporal analysis of the receiver events is also affected by the receiver implementation because the demodulator reset time in case of false or bad preambles depends on the card internal design and results in a different granularity of consecutive events. It follows that the receiver behavior has to be explicitly modeled for recognizing the event patterns due to the effect of different interfering sources.

We propose to use hidden Markov chains for modeling such a receiver behavior and solving the problem of *event pattern* recognition. The observations of the receiver state are given by eight possible receiver events presented in table I. However, the events cannot be read as interrupt signals but need to be indirectly detected by monitoring the card internal registers. We implemented a regular sampling of all the event registers every $250 \mu s$. The sampling interval has been selected as a tradeoff between detection delay and tracking complexity. Because of the periodic sampling, multiple events can occur in the same monitoring interval.

A. Event Patterns

Figures 2 and 3 show two exemplary temporal traces 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

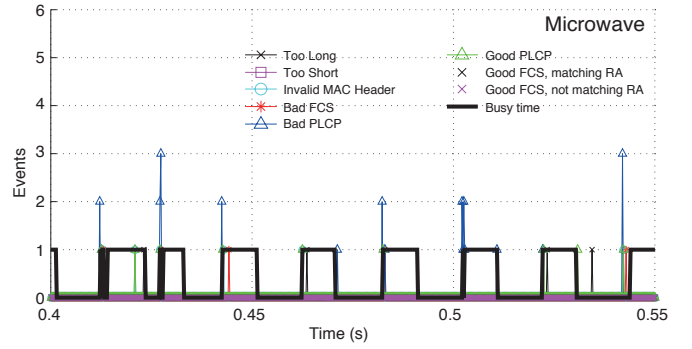


Fig. 4. Bursts of receiver events corresponding to the reception of Microwave interference.

of the ZigBee frame and corresponding acknowledgment (if any), 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 figure 2, 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 (figure 3), 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.

Figure 4 shows a temporal trace of receiver events in case of interference due to a Microwave oven. The oven switches periodically on and off as most Microwave ovens. During the radiation intervals, the WiFi monitoring node senses the channel as busy, as evident from the alternating busy and idle intervals plotted in the figure (whose length is $10 msec$). Event patterns are pretty different from the ones observed in case of ZigBee transmissions: synchronization trials are performed only at the beginning and at the end of the radiation interval (rather than being continuously repeated). This can be due to the power-on and power-down ramp of the Microwave, being the demodulator unable to work when the radiation power is stable.

From both the figures it is evident that receiver events have a different occurrence probability according to the time elapsed from the beginning of the interference or to the type of the previous receiver event (e.g. bad FCS events can occur only after a frame synchronization signaled by a good PLCP event). Since these correlation effects depend on the interference power and interference duration, they can be exploited for classifying the interference sources leading to a given pattern of events.

B. Receiver Model

In order to define the hidden Markov chain modeling the receiver behavior, it is required to specify the receiver observations, the state model and the observation probabilities from each state. While the number of possible events summarized in table I is eight, the overall number of possible observations is higher because multiple events can be triggered during the sampling interval of the card registers. Since the card has a dedicated register for counting the total occurrences of

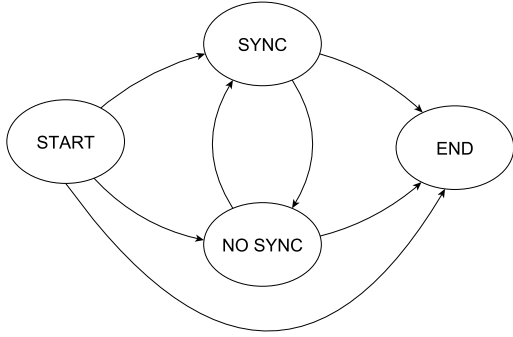


Fig. 5. Generalized state model of the receiver behavior: transition probabilities depend on the interference source.

each event, an observation is given by a vector with eight components, each one specifying the number of occurrences of a given event type during the last observation interval.

The diagram depicted in figure 5 shows the state model triggered at the end of each observation interval (every $250 \mu s$). Although the internal receiver design is not known, the model tries to capture the most evident memory effects, discussed in the previous section, due to the power ramp of the interfering source (START and END states) and to the synchronization of a valid preamble (SYNC and NO SYNC state). Indeed, the occurrence probability of the receiver events, often defined as *emission probabilities*, may vary in each of these states. We assume that the power ramp effects last for one slot only, thus leading to a zero probability to remain in the START and END state. Self transitions to the intermediate states depend on the slot size and on the interference duration.

For tuning the emission and transition probability from each state as a function of a specific source of interference, we implemented a *training phase* of the hidden Markov chain, based on a trace of receiver events acquired in presence of controlled interference. The trace is organized in consecutive event bursts separated by a time interval in which the channel has been sensed as idle. For example, in figure 2 there are four consecutive event patterns, with a last pattern equal to the event sequence {Bad PLCP, Bad PLCP, Bad PLCP, Good PLCP, Too Long}. The state path corresponding to each error pattern can be easily derived by considering that the first and last observations are always performed from the START and END state, while all the others depend on the last preamble synchronization.

We collected three different event traces of $10 s$ under WiFi traffic, ZigBee interference and Microwave interference. By using each trace and corresponding state path, we obtained the maximum likelihood estimates of the emission and transition probabilities devised to characterize the receiver behavior in presence of different signals. Figure 6 visualizes the emission probabilities of the most significant observations for different interference models. It is interesting to observe how the figure quantifies our previous qualitative considerations.

For the WiFi model, most observations result in a synchronized preamble followed by a correct checksum (that can




Fig. 6. Emission Probabilities of most significant observations for different experiments (from top to bottom: WiFi ch. 11, ZigBee and Microwave).

be sampled into the same observation interval or into two consecutive observation intervals due to the short duration of WiFi frames). Packet duration is equal to about $350 \mu s$, because we used frames with 1500 bytes transmitted at 36 Mbps. For the ZigBee model, bad preambles are generated very often: about 60% of error bursts start with such an event, while the other bad preambles are revealed during the intermediate model states. Checksum failures, too long frames or invalid MAC occur at the edge states or when the receiver is synchronized. For the Microwave oven, bad preambles are generated in the START and END states and the no event probability is higher than the previous ones (being the interference interval equal to $10 ms$ and the demodulator active only during the power ramp).

Although the overall occurrence of error rate in presence of non-WiFi signals is known, the approach allows to learn about the implementation-specific reaction times to synchronization errors and sensitivity to narrow-band signals. This allows to define a classification scheme able to work on a generic monitoring node.

C. Classification Scheme

As a result of the training phase, we define three different models for describing the receiver behavior in presence of WiFi, ZigBee and Microwave interference. Being m the number of possible observations, the k -th model is given by the transition probability matrix $P_k^{4 \times 4}$ and emission probability matrix $E_k^{4 \times m}$. For a given event pattern \mathbf{e} , our classification scheme works by selecting the interference model which maximizes the probability of obtaining the sequence \mathbf{e} , i.e. $\text{argmax}_k Pr\{\mathbf{e}|P_k, E_k\}$.

Figure 7 shows the classification results in a temporal trace of 7 consecutive error patterns. The lines plotted on top of the events delimit consecutive errors to be considered as a single error pattern. The pattern delimitation is achieved by monitoring the channel state register and its transition from idle to busy and viceversa. The lines also visualize the

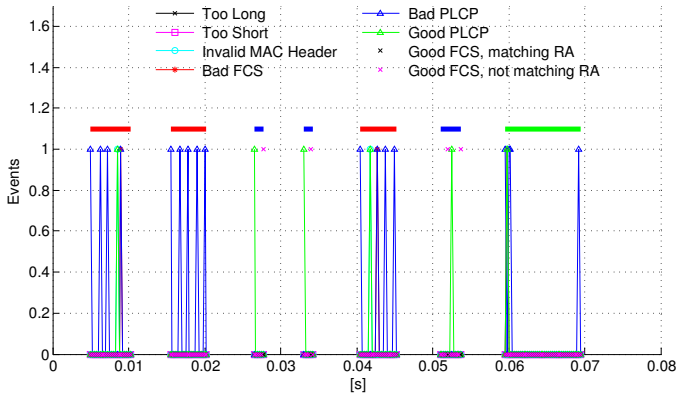


Fig. 7. Temporal trace of consecutive error bursts and classification decisions.

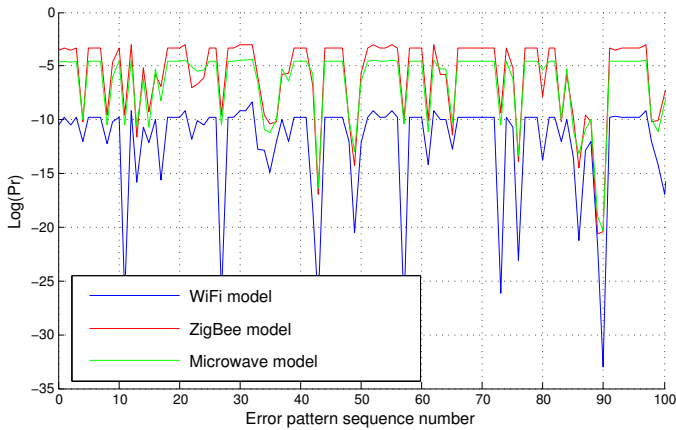


Fig. 8. Comparison between the receiver models working under different interference conditions for a sequence of error bursts due to ZigBee transmissions.

estimated interference source with a different color (blue for WiFi, red for ZigBee and green for the Microwave).

Figure 8 visualizes the effectiveness of the proposed classification scheme by considering a sub-trace of the error patterns corresponding only to ZigBee interference. The figure plots (in logarithmic scale) the occurrence probability of each pattern computed according to the three interfering models. From the figure it is evident that the highest probability corresponds to the ZigBee interference source in almost all the cases. Moreover, the results provided by the Microwave and ZigBee models are much closer than the results provided by the WiFi model.

We tested the classification accuracy by analyzing 10 different temporal traces lasting 100 s where one, two or three interference sources are simultaneously active. We found that the classification accuracy is on average equal to 97% and never lower than 93%. The few decision errors are due to temporal overlapping of multiple interference sources. This type of *combined* interference in principle can be modeled for introducing more advanced interference detection schemes (able for example to recognize WiFi/ZigBee collisions).

V. 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.

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 pattern of the error bursts can be effectively correlated for detecting the presence of non-WiFi signals and identifying the interfering technology. Our solution is based on a simple hidden Markov model characterizing the receiver behavior in presence of interference, whose transition and emission probabilities change as a function of the interference source.

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. Crowddad 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. RF-Dump: An Architecture for Monitoring the Wireless Ether. In Proc. 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.