

Chapter 2

External Radio Interference

Abstract An important factor contributing to the degradation and variability of the link quality is radio interference. The increasingly crowded radio spectrum has triggered a vast array of research activities on interference mitigation techniques and on enhancing coexistence among electronic devices sharing the same or overlapping frequencies. This chapter gives an overview of the interference problem in low-power wireless sensor networks and provides a comprehensive survey on related literature, which covers experimentation, measurement, modelling, and mitigation of external radio interference. The aim is not to be exhaustive, but rather to accurately group and summarize existing solutions and their limitations, as well as to analyse the yet open challenges.

2.1 Introduction

The propagation of radio signals is affected by a plethora of variables, such as radio hardware, antenna irregularities, geometry and nature (static or mobile) of the environment, presence of obstacles responsible for shadowing or multipath fading, as well as the environmental conditions (e.g., temperature [1, 2]). These influences can lead to link unreliability and drastically vary the quality of a wireless link over time.

Another important factor that can vary the link quality and cause a degradation of communications is the presence of radio interference. Radio interference is indeed a serious challenge for wireless systems: it is caused by neighbouring devices operating concurrently in the same frequency band, disturbing each other by transmitting unwanted RF signals that play havoc with the desired ones. Interference is a severe problem especially for low-power wireless networks, as the presence of neighbouring devices transmitting at higher power may cause a significant degradation of the overall performance.

The primary outcome of interference is an increase in the packet loss rate, and it is in turn often followed by an increase in the network traffic due to retransmissions,

as well as by a decrease in the performance and efficiency of the overall network. Experiences from several wireless sensor network deployments have shown that an unexpected increase of network traffic compared to the initial calculations may lead to an early battery depletion and/or deployment failure [3, 4].

Interference may also lead to unpredictable medium access contention times and high latencies. This is an important observation for low-power wireless networks used in safety-critical scenarios (e.g., industrial control and automation [5], health care [6], and high-confidence transportation systems [7]), where guaranteeing high packet delivery rates and limited delay bounds is necessary, and where unreliable connections cannot be tolerated.

On a network scale, interference can be both *internal* and *external*, and might affect the totality or only part of the nodes in the network.¹ *Internal interference* is the one generated by other sensor nodes operating in the same network, and it is typically mitigated by either proper placement of nodes and careful topology selection, or an appropriate MAC layer (e.g., making use of time diversity to avoid concurrent activities in the channel). *External interference* is caused by other wireless appliances operating in the same frequency range of the network of interest using other radio technologies. In the context of wireless networks composed of low-power sensor nodes, several devices operating at higher powers can be source of external interference (e.g., Wi-Fi access points and microwave ovens).

Because of its strong impact on the quality of wireless links, it is important to understand how interference affects the communications among wireless sensor nodes, and how to develop techniques that can properly mitigate its impact. Making wireless communications robust and reliable in the presence of interference is not an easy task. While internal interference can be minimized by means of a proper configuration of the network and protocol selection, the mitigation of external interference is often more complex for several reasons. Firstly, it is hardly possible to know in advance all potential sources of interference in a given environment and to predict their behaviour. Secondly, interference is often intermittent and highly dynamic, therefore it is difficult to create solutions that guarantee a reliable communication over time. Furthermore, fading due to multi-path propagation or shadowing from obstacles in the surroundings can affect the quality of wireless communications unpredictably, and things get even more erratic in the presence of interference because of the superposition of the unwanted radio signals [9].

In the next sections of this chapter we focus on wireless sensor networks, and we report on experimentation, measurement, modelling, and avoidance of external interference. We start reporting on the increasing congestion in the unlicensed frequency bands used by wireless sensor networks, as it makes their communications vulnerable to the interference generated by other wireless appliances.

¹ Some works also define *protocol interference* as the one occurring when multiple local protocols send conflicting commands to the radio transceiver [8].

2.2 Crowded Spectrum

Wireless communication technology has become increasingly popular in the last decades, because of the greater flexibility and remarkable reduction of costs for installation and maintenance compared to traditional wired solutions. This triggered a massive proliferation of wireless devices in our everyday life: the world of telecommunications and networking has experienced a large-scale revolution, and wireless systems have become ubiquitous, especially in residential and office buildings.

This proliferation has caused the radio spectrum to become a very expensive resource, therefore many standardized technologies operate in increasingly crowded and lightly regulated Industrial, Scientific and Medical (ISM) radio bands. The latter are freely-available portions of the radio spectrum internationally reserved for industrial, scientific and medical purposes other than communications.

When several technologies operate in the same ISM radio band, many devices concurrently share the same frequencies, and coexistence may become problematic. This applies especially to low-power wireless networks, as the presence of neighbouring devices transmitting at higher power may cause a significant degradation of the link quality, as well as a decrease in the packet delivery rates. Also wireless sensor networks compliant to the IEEE 802.15.4 standard are vulnerable to the external interference generated by neighbouring devices coexisting on the same frequencies, because the standard specifies operations in unlicensed and crowded ISM bands [10].

The IEEE 802.15.4 standard. The IEEE 802.15.4 protocol specifies the two lowest layers of the protocols stack for low-rate wireless personal area networks, namely the medium access control and physical layers. The physical layer (PHY) manages the physical RF transceiver, and provides the data transmission service: its main responsibilities are the data transmission and reception according to specific modulation techniques, as well as the channel frequency selection and management of energy and signal functions (e.g., LQI and energy detection). The first standard released in 2003 [11] specified only two PHY layers based on direct sequence spread spectrum (DSSS): the 868/915 PHY, employing binary phase-shift keying (BPSK) modulation, and the 2450 PHY, employing offset quadrature phase-shift keying (O-QPSK) modulation. Wireless sensor networks compliant to this standard were operating on one of three possible unlicensed ISM frequency bands:

- 868.0–868.6 MHz, available in Europe, one communication channel with center frequency $F_c = 868.3$ MHz;
- 902–928 MHz, available in North America, up to ten communication channels with center frequency $F_c = 906 + 2 \cdot (k - 1)$ MHz, for $k = 1, 2, \dots, 10$;
- 2400–2483.5 MHz, available worldwide, up to sixteen communication channels with center frequency $F_c = 2405 + 5 \cdot (k - 11)$ MHz, for $k = 11, 12, \dots, 26$.

The standardization process has been very active in the last decade, and important updates and amendments have been made to the first version of the IEEE 802.15.4 standard. A revision in 2006 [12] defined two optional 868/915 PHY layers employing a different modulation scheme, namely an 868/915 direct sequence spread spectrum PHY employing offset quadrature phase-shift keying (O-QPSK) modulation,

and an 868/915 parallel sequence spread spectrum (PSSS) PHY employing a combination of binary keying and amplitude shift keying (ASK) modulation.

Because of the increasing congestion of the original three ISM band used, several amendments were defined to support new bands. The last version of the standard released in 2011 [13] encompasses the amendments defined by task groups IEEE 802.15.4.a/c/d and adds new physical layers, some of which make use of UltraWideBand (UWB) and Chirp Spread Spectrum (CSS) modulation techniques, as shown in Table 2.1.

UWB is a radio technology that employs high-bandwidth communications and uses a large portion of the radio spectrum, offering significant advantages with respect to robustness, energy consumption, and location accuracy compared to narrow-band DSSS. Because of the large bandwidth communications, UWB achieves a higher robustness against interference and fading, which makes this technology promising to achieve robust communications [14].

Despite the introduction of new ISM bands, the vast majority of wireless sensor networks deployed nowadays still makes use of the three original unlicensed ISM bands specified in the first edition of the IEEE 802.15.4 standard [11]. The main reason for this trend is essentially the availability of several off-the-shelf inexpensive hardware platforms developed in the beginning of the 21st century, such as the Moveiv TelosB and Sentilla Tmote Sky motes. Based on the available literature, we now briefly summarize the most prominent sources of interference in the three ISM bands specified in [11].

Interference in the 868/915 MHz ISM bands. Although the 868 and 915 MHz frequency bands are known to be relatively interference-free [15], several radio technologies have proven to cause significant problems to deployed wireless sensor networks. Barrenetxea et al. [16] have highlighted how cellular phones can be an interference source for sensor nodes operating in the 868 MHz band. The proximity of the European GSM band (that uses 890–915 MHz to send information from the mobile station to the base station and 935–960 MHz for the other direction), causes an impact on mote transmissions during the first seconds of an incoming call. Kusy et al. [17] have described the impact of in-band and out-of-band interference in the 900 MHz frequencies: in-band interference is caused by telemetry networks and cordless telephones, whereas mobile phones and pagers often cause out-of-band interference. Furthermore, several wireless devices marketed in Europe, including wireless domestic weather stations, car alarms, garage openers, and residential electronic alarms, use the 868 MHz frequency and are therefore potential sources of interference for wireless sensor networks operating in the 868/915 MHz ISM bands.

Interference in the 2.4 GHz ISM band. Most of the wireless sensor networks deployed nowadays use the 2.4 GHz frequencies, because they are available worldwide, and because several popular off-the-shelf sensor nodes embed radio transceivers operating in the 2450 PHY layer defined by the IEEE 802.15.4 standard (e.g., Moteiv TelosB, Crossbow MicaZ, and Sentilla Tmote Sky motes). However, to date, the 2.4 GHz is by far the most congested ISM band, because of the pervasiveness of devices operating in those frequencies, and their high transmission power. Sensor nodes must indeed compete with the communications of Wi-Fi (IEEE 802.11) and

Table 2.1 Frequency bands, modulation techniques, and data rates specified in the IEEE 802.15.4 standard [13]

PHY Layer (MHz)	Frequency Band (MHz)	Modulation	Bit Rate (kb/s)	Symbol Rate (ksymbol/s)	Symbols	802.15.4 Standard
780	779–787	O-QPSK	250	62.5	16-ary orthog.	2011 [13]
780	779–787	MPSK	250	62.5	16-ary orthog.	2011 [13]
868/915	868–868.6	BPSK	20	20	Binary	2003 [11]
	902–928		40	40		
868/915*	868–868.6	ASK	250	12.5	20-bit PSSS	2006 [12]
	902–928		250	50	5-bit PSSS	
868/915*	868–868.6	O-QPSK	100	25	16-ary orthog.	2006 [12]
	902–928		250	62.5	16-ary orthog.	
950	950–956	GFSK	100	100	Binary	2011 [13]
950	950–956	BPSK	20	20	Binary	2011 [13]
2450 DSSS	2400–2483.5	O-QPSK	250	62.5	16-ary orthog.	2003 [11]
2450 CSS*	2400–2483.5	DQCSK	250	167	64-ary orthog.	2011 [13]
			1000	167	8-ary orthog.	
UWB*	250–750	BPM-BPSK	Variable parameters			2011 [13]
	3244–4742					
	5944–10234					

*(optional)

Bluetooth (IEEE 802.15.1) devices, as well as with the noise generated by microwave ovens and other domestic appliances such as cordless phones, baby monitors, game controllers, presenters, and video-capture devices [10, 18–20].

As the 2.4 GHz band is, by far, the most crowded ISM band nowadays, we describe in the remainder of this section the characteristics of the most important sources of interference in this band, namely Wi-Fi devices, Bluetooth devices, and microwave ovens. We then conclude with a short discussion on the impact of adjacent channels interference from co-located IEEE 802.15.4 networks.

2.2.1 Coexistence Between IEEE 802.15.1/15.4 Devices

The IEEE 802.15.1 (Bluetooth) standard specifies 79 channels, spaced 1 MHz, in the range 2402–2480 MHz, with center frequency $F_c = 2402 + k$, with $0 \leq k \leq 78$. Bluetooth uses the Frequency Hopping Spread Spectrum (FHSS) technology to combat interference and fading: it hops 1600 times per second, and therefore it remains at most $625 \mu\text{s}$ in the same channel. Given that only 79 channels are available, on average, one channel is used approximately 20 times each second: this makes interference generated by Bluetooth devices uniformly distributed across the whole 2.4 GHz band.

As of version 1.2, several Bluetooth stack implementations apply an Adaptive Frequency Hopping (AFH) mechanism to combat interference and fading, in which the hopping sequence is modified to avoid interfered channels. However, because the low-power sporadic communications of sensor nodes do not constitute a real threat compared to the communications between more powerful transmitters (e.g., Wi-Fi), Bluetooth devices will still make use of channels in which wireless sensor networks are operating.

The interference produced by IEEE 802.15.1 devices is, however, not so problematic for wireless sensor networks, because of the randomness and adaptiveness of Bluetooth's random frequency hopping [21]. Several experimental works have been carried out to study the impact of IEEE 802.15.1 communications on the reliability of sensor network transmissions. The packet loss rate of a wireless sensor network operating in the presence of Bluetooth interference is often between 3% (as reported by Bertocco et al. [22] and Penna et al. [23]) and 5% (as reported by Huo et al. [24]) up to a maximum of 9–10% (as shown in the experimental results of Boano et al. [18] and Sikora and Groza [20]), hence not too critical.

2.2.2 Coexistence Between IEEE 802.11/15.4 Devices

IEEE 802.11, better known as Wi-Fi, is a set of standards for wireless local area network (WLAN) communications. The standard divides the ISM bands into channels: the 2.4 GHz band (2400–2483.5 MHz), for example, is divided into up to 14

channels,² each of which has a bandwidth of 22 MHz. Channels are therefore partially overlapping, and there is only room for three orthogonal channels. The standard evolved significantly in the last decade (the first version was released in 1997), with data rates increasing from the original 2 Mbit/s to the 11 Mbit/s of 802.11b (1999), 54 Mbit/s of 802.11g (2003), up to the 150 Mbit/s of 802.11n (2009); and it is still undergoing changes, with the new high-throughput 802.11ac protocol currently under development.

The coexistence between IEEE 802.11b/g/n and IEEE 802.15.4 devices represents a challenge for several reasons. Firstly, Wi-Fi devices are nowadays ubiquitous, especially in residential and office buildings where many Access Points (AP) are installed. Secondly, IEEE 802.11 devices operate at significantly higher power (≈ 24 dBm) than traditional low-power sensor nodes. Thirdly, Wi-Fi channels have a bandwidth of 22 MHz and can therefore interfere with multiple IEEE 802.15.4 channels at the same time. Fourthly, IEEE 802.11 supports high-throughput transmissions that generate interference patterns that are difficult to predict, as they depend on factors such as the number of active users, their activities, the protocols they use (UDP or TCP), or the traffic conditions in the backbone.

Several works in the literature investigate the impact of IEEE 802.11 communications on the reliability of IEEE 802.15.4 transmissions, and show that wireless sensor networks suffer from high packet loss rates in the presence of Wi-Fi interference [19, 20, 22, 23, 25, 26]. Under certain conditions, also IEEE 802.11 devices can suffer from the interference of nearby wireless sensor networks. The communications between sensor nodes can indeed trigger a nearby Wi-Fi transmitter to back off: if this happens, only the header of IEEE 802.15.4 packets is typically corrupted [8, 26].

The actual packet loss rate that IEEE 802.15.4 networks experience in the presence of IEEE 802.11 transmissions depends on the Wi-Fi activity, as well as the location of the nodes. Boano et al. [18] have varied the Wi-Fi traffic pattern, showing that activities such as continuous radio streaming are not too critical for sensor network communications, as it results in approximately 15 % packet loss rate. On the contrary, activities such as video streaming (≈ 30 % packet loss rate) and file transfers (≈ 90 % packet loss rate) can destroy the majority of wireless sensor networks transmissions and cause long delays, drastically decreasing the performance of the network.

2.2.3 Coexistence Between Microwave Ovens and IEEE 802.15.4 Devices

Microwave ovens are a common kitchen appliance used to cook or warm food by exposing food to non-ionizing microwave radiations to make water and other polarized molecules oscillate, usually at a frequency of 2.45 GHz. Therefore, microwave ovens are a potential source of interference for wireless sensor networks operating in

² The availability of channels is regulated by each country, e.g., channel 14 is currently only available in Japan.

the 2.4 GHz spectrum. There are two main categories of microwave ovens: the ones designed for domestic use and the ones designed for commercial purposes, with the former being much more common [27].

The characteristics of the interference patterns emitted by domestic microwave ovens depend on the model; nevertheless all the ovens present the same basic properties. Firstly, with respect to the frequency spectrum, microwave ovens can potentially interfere on a large portion of IEEE 802.15.4 channels. The frequency of emitted microwaves depends on the sizes of the cavities of the magnetron, and varies also with changes in load impedance, supply current, and temperature of the tube. Also other factors, including the oven content, the amount of water in the food, and the position within the oven can affect the temperature of the magnetron [28]. Therefore, it is not possible to state with certainty which channel(s) will be mostly affected by the interference generated by microwave ovens. Secondly, with respect to the temporal behaviour, the noise generated by microwave ovens is rigorously periodical, as ovens continuously turn on and off according to the frequency of the AC supply line. The duration of a single period, called power cycle, hence mostly depends on the power grid frequency, but can also slightly vary depending on the oven model. Works in the literature report a power cycle of roughly 20 ms (at 50 Hz) or 16 ms (at 60 Hz) with an active period of at most 50% of the power cycle [27, 29]. During the active period, the communications of low-power sensor nodes in proximity of a microwave oven are likely destroyed (because microwave ovens operate at up to ≈ 60 dBm), but during the inactive period several consecutive packet can be scheduled, as shown in [30]. The interference generated by microwave ovens can therefore be easily modelled as a deterministic sequence of interference pulses.

2.2.4 Coexistence Between IEEE 802.15.4 Devices Operating in Adjacent Channels

Several studies have highlighted that IEEE 802.15.4 channels in the 2.4 GHz ISM band are not orthogonal to each other, and hence wireless sensor networks operating on adjacent channels may interfere with each other [31–36].

Wu et al. [34] have shown that the number of orthogonal IEEE 802.15.4 channels in the 2.4 GHz ISM band is only eight, despite the actual number of channels with 5 MHz spacing available is sixteen. The authors carry out experiments using MicaZ nodes (that employ the CC2420 radio chip) and show that transmissions in adjacent channels decrease the packet reception rate, whereas transmissions generated at least two channels away from the one of interest do not harm the reception (that remains basically unaffected). The interference generated in the adjacent channel can decrease the packet reception rate as much as 50% when using a low transmission power, leading to a potential disruption of connectivity that cannot be neglected. These results were confirmed by the experiments by Ahmed et al. [31]: the packet reception

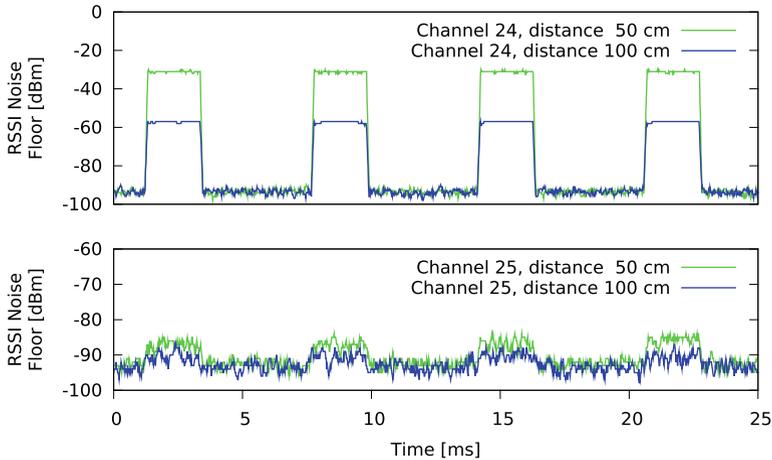


Fig. 2.1 IEEE 802.15.4 transmissions create noise on adjacent channels

rate decreases significantly in the presence of activities in adjacent IEEE 802.15.4 channels.

Figure 2.1 shows the impact of IEEE 802.15.4 packet transmissions on adjacent channels as measured by a Sentilla Tmote Sky node as follows. Two nodes \mathcal{A} and \mathcal{B} are placed in an interference-free environment at a given distance from each other (50 and 100 cm in two successive experiments). Node \mathcal{A} listens first on channel 24 (center frequency at 2470 MHz) and later on channel 25 (center frequency 2475 MHz), recording the signal strength at the antenna pins. At the same time, node \mathcal{B} continuously transmits packets at a rate of $\frac{1}{128}$ packets/s on channel 24. Figure 2.1 shows that the packet transmissions of node \mathcal{B} on channel 24 generate also noise on channel 25 (bottom figure) that are high enough to affect communication in low-quality links. Hence, if node \mathcal{A} would use channel 25 for its communications, it may experience packet loss due to the adjacent channel interference on channel 24.

2.3 Interference Measurement and Modeling

In the presence of external interference, the properties of a wireless channel can change unpredictably over time in both indoor and outdoor environments. Interference can be sporadic, causing only a temporary impact on communications, or persistent, causing a channel to experience heavy interference and become unavailable for long periods of time. Wireless sensor nodes may therefore need to adapt dynamically to changeable interference patterns and adjust their behaviour at runtime in order to maximize the reliability of their communications.

To adapt dynamically to various interference patterns, wireless sensor nodes firstly need to acquire a detailed understanding about the surrounding interference by means of accurate measurements. The latter must be carried out in a simple and energy efficient fashion, in order to meet the constrained capabilities of wireless sensor nodes. Secondly, efficient metrics to estimate the presence of interference need to be obtained from the measurements, in order to adapt dynamically to changing interference patterns, for example by carrying out a dynamic channel selection in multichannel protocols, or by ranking the available channels. Thirdly, there is a need to derive lightweight interference models that can be parametrized at runtime, in order to carry out a dynamic protocol selection or a dynamic adjustment of protocol parameters as soon as certain properties in the environment have changed.

In this section, we describe the most popular ways to accurately measure interference using off-the-shelf wireless sensor nodes (Sect. 2.3.1), and show how these measurements can be used to identify the active sources of interference in the surroundings (Sect. 2.3.2). We then describe simple interference models that can be implemented on resource-constrained wireless sensor nodes, and explain how they can be parametrized at runtime to achieve interference mitigation (Sect. 2.3.3).

2.3.1 Measuring Interference Using Sensor Nodes

The most common way to measure interference using wireless sensor nodes is the so called RF noise (floor) measurement, i.e., an estimate of the received signal power within the bandwidth of an IEEE 802.15.4 channel in absence of sensornet transmissions [37]. RF noise measurements are typically retrieved using the energy detection (ED) feature available in IEEE 802.15.4-compliant radios as part of a channel selection algorithm [11], and typically return an RSSI value that can be converted in dBm.

The key difference between RF noise measurements and traditional RSSI and LQI values is that the former can be obtained anytime, whereas the latter are generated only upon packet reception, and hence cannot describe the interference in a fine-grained way. Indeed, RSSI and LQI, as well as packet reception rate, can be used to derive the presence of interference and react accordingly (e.g., by switching channel when high packet losses arise [38]), but do not unequivocally identify and quantify the presence of interference in the environment. For example, the LQI describes the chip error rate for the first eight symbols following the SFD field, which obviously has a correlation with the amount of interference. Nevertheless, low LQI values may also result from unreliable links in absence of external interference. The same applies to the packet reception rate, as a low reception ratio may be caused by other factors than external interference, including routing issues, and software bugs.

Figure 2.2 shows the RSSI values returned by a Maxfor MTM-CM5000MSP sensor mote measuring RF noise in the presence of sensornet transmissions and external interference in the 2.4 GHz ISM band [39]. The mote employs the CC2420 radio and samples the RSSI at approximately 50 kHz, enough to capture the interference of

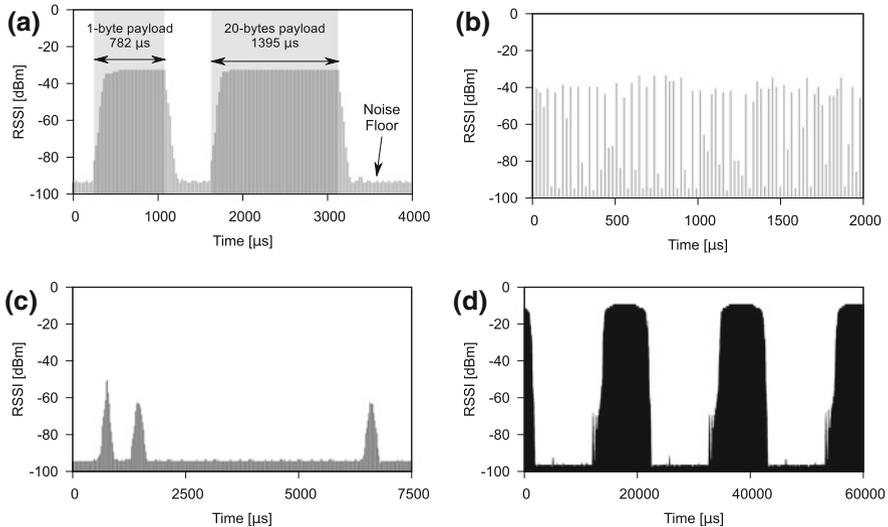


Fig. 2.2 RF noise values measured at a speed of 50 kHz using Maxfor MTM-CM5000MSP sensor motes operating in the 2.4 GHz ISM band [39]

common sources of interference such as microwave ovens, as well as several Bluetooth and Wi-Fi devices. In case interference is absent, the RSSI values are typically close to the sensitivity threshold of the radio and identify the noise floor (see Fig 2.2a). The latter has typically values between -100 and -95 dBm when using sensor nodes equipped with the CC2420 radio transceiver. As shown in Fig. 2.2, the interference generated from different devices produces different types of RSSI traces: one can easily identify the periodicity of microwave ovens, as well as the short and bursty transmissions of Wi-Fi devices.

Based on RF noise measurements, several works have proposed empirical metrics to derive the presence of interference. Musaloïou et al. [40] have measured the RF noise in each IEEE 802.15.4 channel at a rate of 20 samples/s in the presence of Wi-Fi activity and proposed three metrics that can be implemented on resource-constrained motes using off-the-shelf radios. The first metric is based on the cardinality of RSSI values, and counts the amount of unique RSSI values collected. Channels with high cardinality are likely to be rich of interference, as non-interfered channels have a small RSSI variation [18, 41]. The second metric is based on the maximum and mean RSSI value: channels with high levels of interference will record high maximum RSSI values as well as high average RSSI readings. The third metric counts the number of RSSI measurements above a given threshold: channels with the most measurements above such thresholds are considered to be highly interfered. This technique essentially consists in counting the amount of failed clear channel assessments carried out using a given RSSI threshold. The authors propose a threshold of -90 dBm, but this value may not be optimal because off-the-shelf sensors are typically uncalibrated [42].

Also Hauer et al. [25] suggest to predict the degradation of the link quality caused by interference using RF noise measurements. According to their experimental results gathered in an urban residential area, a substantial increase in RSSI noise floor values anticipates heavy packet loss, in particular when using very low transmission power levels.

Although very popular, RF noise measurements carried out using energy detection have several drawbacks:

1. Energy detection comes at a very high energy cost: the radio needs to be turned on in listening mode during the whole duration of the measurements.
2. A suitable RSSI sampling rate needs to be selected, a trade-off between energy-efficiency and accuracy. On the one hand, a low sampling rate may not capture the interference produced by devices transmitting at high frequency, e.g., Wi-Fi devices, and therefore assessing a medium as idle even though transmissions actually occurred between consecutive RSSI samples. On the other hand, a high sampling frequency corresponds to a significant energy expenditure that battery-powered sensor nodes may not be able to afford. Several works have tried to reach the limit of sensor nodes, and obtained RSSI sampling rates up to 62.5 kHz [41, 18]. This enables a quite accurate understanding of interference: Hauer et al. [41] exploit the RSSI profiles to estimate the bit error positions inside corrupted frames; Boano et al. [18] exploit the high-frequency samples to get a precise trace of interference for a later regeneration. However, even a resolution of 62.5 kHz is not sufficient to capture all possible interference sources: one can detect for example all IEEE 802.11b frames, but not all IEEE 802.11g/n frames.
3. Energy detection is a technique known to be brittle. RSSI readings have a low accuracy (± 6 dBm in the CC2420 radio, for example), and the readings may not be accurate in the presence of narrow unmodulated carriers because of saturation of the Intermediate Frequency (IF) amplifier chain, as highlighted in [18].
4. The sensor nodes often need to dedicate all their resources to the energy detection, and hence cannot receive and send packets or carry out other tasks, especially when sampling the RSSI with a very high frequency.

Because of the inefficiency of energy detection, a few works in the context of multichannel protocols do not make use of RF noise measurements to measure or detect the presence of interference. These works rely on either the packet reception rate [38, 43] or on the amount of failed transmissions and CCA failures [44, 45] to escape from congested channels. However, as mentioned above, these techniques have two main drawbacks: (i) they take effect only after the interference affected the communications within the sensor nodes, (ii) they do not unequivocally identify the presence of external interference.

2.3.2 Interferer Identification

An accurate measurement of interference can be used to identify which devices are generating the unwanted signals, a precious information that can be exploited to adapt packet transmissions and increase the robustness of communications. Interferers are typically classified based on RF noise measurements, but there are also a number of approaches based on the hardware estimators available in common IEEE 802.15.4-compliant radio transceivers such as RSSI and LQI (see Sect. 3.4).

Chowdhury and Akyildiz [30] describe a classification mechanism in which a sensor node actively scans channels in the 2.4 GHz ISM band to detect the interference generated by microwave ovens and IEEE 802.11b devices. In this approach, RF noise readings are used to obtain a spectral signature in all the available IEEE 802.15.4 channels. This signature is then used by the sensor nodes to derive the type of interferer by matching the observed spectral pattern with a stored reference shape. The knowledge about the type of interfering source is then exploited to construct an adaptive scheme for channel selection and scheduling of packets. Thanks to the correct classification, the authors show that the protocol can significantly reduce the packet loss rate in the presence of interference. However, as with other solutions exploiting RF noise measurements, this techniques may require a high amount of energy to scan all the available IEEE 802.15.4 channels, and cause the nodes in the network to be unreachable while performing energy detection.

Zacharias et al. [46] propose a classification algorithm that uses RSSI noise floor readings to monitor a single IEEE 802.15.4 channel in the 2.4 GHz ISM and understand which device is the source of interference. Compared to the approach by Chowdhury and Akyildiz, they also recognize Bluetooth interferers in addition to Wi-Fi devices and microwave ovens, but they do not exploit the knowledge to adapt packet transmissions. In particular, the authors exploit one second of RF noise readings (which results in approximately 11300 RSSI values) and classify the interferer depending on a given RSSI threshold, the channel usage, and the duration of interference over time. For example, their algorithm classifies as Wi-Fi interference all traces in which the usage of the channel is between 1 and 30% and the maximum time of a clear channel is less than 100ms (with an RSSI threshold of -85 dBm). Although the proposed duration of the scan (1 s) seems to be relatively energy efficient, it may not be enough to obtain a clear picture of the ongoing interference. Also, this approach cannot detect multiple sources at the same time and it is strongly dependent on empirical values that can vary depending on the calibration of different sensor nodes [42].

Differently from the previous two approaches, Hermans et al. [47] propose a method that is not based on energy detection. The authors combine the properties of LQI and RSSI during packet transmissions to investigate the feasibility of a lightweight interference detection and classification approach that only uses data that can be gathered during a sensor network's regular operation. Similarly to the work of Zacharias et al. [46], the authors try to differentiate between microwave ovens, Bluetooth and Wi-Fi devices, but without additional overhead due to scanning of

the medium in absence of packet transmissions. In particular, they use LQI (that represents the chip error rate) to identify the cases in which a packet is received with a high LQI, but the packet fails the CRC check, which implies that channel conditions were good when reception started, but then deteriorated. The authors claim that such a sudden change in channel conditions can be observed when an interferer starts emission during packet reception. They further exploit the RSSI during packet reception to get a series of RSSI values for each received packet, containing about one sample per payload byte. Using this, the authors analyse the corruption of some 802.15.4 packets in the presence of different interferers and derive which parts of a packet have been corrupted, and use a supervised learning approach to train a classifier to assign each corrupted packet to a class representing the interfering device.

A work different in spirit from the above ones is that of Boers et al. [48, 49], which does not aim to identify the type of device, but rather the characteristic of interference. The authors distinguish interference with infrequent spikes, periodic spikes, periodic and frequent spikes in the RF noise measurements, as well as interference exhibiting a value constantly higher than the RSSI sensitivity threshold, or bimodal interference (such as the one produced by microwave ovens). In case a channel shows a mixture of these characteristics, the channel is classified according to the dominant pattern. The authors use decision trees as classifier, since, after training, the classification of new cases is simple, and can therefore meet the requirements of constrained wireless sensor nodes.

2.3.3 Modeling Interference

The creation of precise and lightweight models of common interference sources is not an easy task, because of the severe hardware limitations of wireless sensor nodes, and their reduced energy budget.

Very popular is the modelling of the channel occupancy as a two-state semi-Markov model, in which, at a given time instant, a channel is defined as busy if any interfering devices is transmitting packets and defined as idle otherwise [50]. The advantage of this simple model is that it can be easily parametrized at runtime using RF noise measurements: several works have defined a channel as busy when the RSSI values returned by RF noise measurements are above a certain threshold, and idle when the RSSI values are below such threshold [18, 51].

The parametrization of the lightweight two-state semi-Markov interference model at runtime has been used in the context of interference estimation and runtime adaptation of protocol parameters. Noda et al. [51] have presented a channel quality metric based on the availability of the channel over time, which quantifies spectrum usage. Distinctive feature of this metric is the ability to distinguish between a channel in which interference is bursty with large inactive periods and a channel that has very high frequency periodic interference with the same energy profile. As the first case is more favourable for having successful packet transmissions, the proposed metric ranks in a more favourable way channels with larger inactive periods or vacancies.

Using the two-state semi-Markov model of channel occupancy, Boano et al. [39] make use of RF noise measurements to measure the duration of the idle and busy instants, and compute the probability density function of idle and busy periods. Based on the duration of the longest busy period, the authors derive protocols parameters such that certain QoS requirements are met even in the presence of external interference.

The two-states model has also been extensively used with the purpose of generating interference. Boano et al. [18, 52] have implemented a bursty interferer model that sends continuous blocks of interference with uniformly distributed duration and spacing, which can be easily implemented on off-the-shelf sensors nodes to emulate, for example, the bursty transmission caused by Wi-Fi or Bluetooth devices. Interference follows continuous off/on periods and the transitions between the two states are specified by a Bernoulli random variable. A uniformly distributed random variable is further used to control the burstiness and duration of the interference. A second model, called semi-periodic interferer model, produces continuous blocks of interference as well, but the duration of the periods and their spacing have smaller variance, in order to emulate, for example, the type of interference generated by a sensor performing periodic data collection.

In their follow-up work [18], the authors have followed a different approach: since the patterns generated by external interference differ substantially depending on the interfering source (as highlighted in Sect. 2.2), external interference was instead modelled per device. Because of its characteristics, microwave oven interference is simple to model, as it follows a deterministic on/off sequence. Hence, the interference is a function of three parameters τ , δ , ρ , where τ is the period of the signal, δ is the duty cycle (fraction of time the oven is “on”) and ρ is the output power of the microwave, which determines the strength of the interference signal. The authors have used this model to generate a series of on/off signals resembling microwave oven interference using sensor nodes and the CC2420 radio transceiver [18]. A similar model was used by Chowdhury and Akyildiz [30] for interference-aware scheduling of packets. To model Wi-Fi traffic in a simple and lightweight fashion, Boano et al. [18] resorted to an analytical model for saturated traffic sources, and derived models from empirical data for non-saturated traffic. In particular, for the latter, they denoted a random variable X as the *clear* time between two consecutive *busy* times, and obtained the probability mass function $p(x) = P_r\{X = x\}$ from the empirical sampling of the channel, where x is the time in number of slots (each slot is 1 ms). For saturated traffic, the authors have exploited the analytical model for the Distributed Coordination Function (DCF) mode of 802.11 proposed by Bianchi [53], and its simplification proposed by Garetto and Chiasserini [54].

2.4 Mitigating Interference

Several techniques have been proposed to tolerate external interference in wireless sensor networks [55]. In this section, we review related literature and propose a taxonomy that classifies existing interference mitigation techniques and aims to help

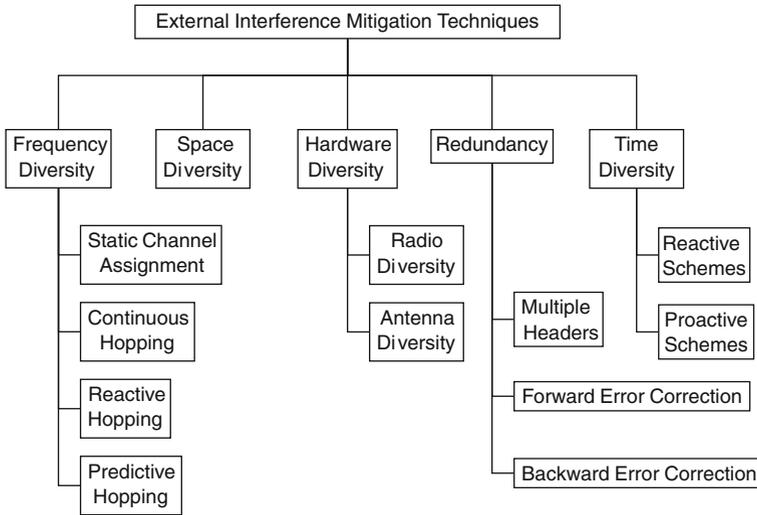


Fig. 2.3 Taxonomy for external interference mitigation techniques

protocol designers in identifying relevant mechanisms that make wireless sensor networks communications robust to external interference.

Most techniques exploit the availability of several radio channels in a given ISM band, for example by continuously hopping among channels over time. We group these applications in the frequency diversity class. Another popular interference mitigation technique consists in avoiding interference by routing packets through different links: we describe this solution in the space diversity class. Protocols deferring transmissions or scheduling packets in such a way to avoid interference are described in the time diversity class, whereas solutions exploiting multiple radios and antennas are grouped in the hardware diversity class. Finally, we define a class named redundancy, in which we group solutions that add redundancy to the transmitted information, such as the use of multiple headers and forward error correction techniques, as well as protocols making use of retransmissions. Figure 2.3 summarizes the proposed taxonomy of external interference mitigation techniques.

2.4.1 Frequency Diversity Solutions

One of the most common ways to mitigate external interference is to exploit the availability of several radio channels in a given ISM band.

Static channel assignment. The easiest way to pursue interference avoidance by exploiting multiple radio frequencies consists in statically assigning channels depending on the expected interference sources. This represents a very primordial

way, that we name *static channel assignment*, in which the designer of the network essentially assigns each node to one or more fixed IEEE 802.15.4 channels that are supposedly interference-free for extended periods of time or throughout the lifetime of the network. In the easiest scenario, all nodes communicate on the same channel: this is the case for several real-world wireless sensor networks assigned to channel 26 in order to escape Wi-Fi interference in the 2.4GHz ISM band [6, 8, 56]. This technique assumes that, in addition to Wi-Fi, no other interference source will ever interfere in that specific channel with the communications of the wireless sensor nodes, and hence it is highly unreliable. In more complex scenarios, involving dense wireless sensor networks covering large areas, different interference sources may be present throughout the network, and the quality of channels may differ from node to node. Several works have created multichannel protocols and statically assigned portions of the network to specific channels. However, this was mostly done with the intention of maximizing the bandwidth available for communications by increasing the number of channels for unicast transmissions [34, 57–61]. Hence, these multichannel solutions are not meant to provide any guarantee against external interference, and do not enhance coexistence among devices operating in the same frequency range.

Continuous hopping. Another way to pursue interference avoidance resembles the operations of the IEEE 802.15.1 standard (Bluetooth) and consists in continuously hopping among channels according to the same pseudo-random sequence. This technique, that we name *continuous hopping*, can be blind (i.e., nodes hop among all available channels) or adaptive, i.e., nodes carry out some form of blacklisting of undesirable channels where high traffic loads or excessive interference is present [62]. Examples of protocols adopting adaptive continuous hopping are the Time Synchronized Mesh Protocol (TSMP) [63], the Wireless-HART standard [64], the protocol developed by Du et al. [65] and Yoon et al. [66], and EM-MAC [45]. The latter is a typical example of adaptive continuous hopping: it uses a penalty system with channel blacklisting based on the results of the clear channel assessment (CCA) operation. A node switches among channels based on its pseudo-random channel schedule, except that, if the next pseudo-randomly chosen channel is on the node's channel blacklist, the node stays on its current channel.

Continuous hopping exploits the potential of frequency diversity and hence can reduce the impact of narrow-band interference and persistent multipath fading. Furthermore, channel hopping also ensures fairness among the chosen channels. However, channel hopping requires a tight time synchronization across the network in order to work properly. Also, the seed and the list of blacklisted channels needs to be shared in a reliable fashion, a critical operation in the presence of interference. It is important to note that in case of blind continuous channel hopping, the interference avoidance is only passive, i.e., by continuously hopping, a pair of nodes will sooner or later pick a good communication channel: this might not necessarily happen in a short time interval. Another drawback of continuous channel hopping is the additional energy required to continuously switch channels [67] on the long run, as well as the protocol overhead. Compared to protocols switching on demand, this represent a non-negligible burden. Also, adaptive continuous hopping protocols

require to continuously update and spread the list of blacklisted channels, which may cause a significant energy consumption.

Reactive hopping. In order to avoid the burden of continuously switching the channel, several protocols switch or blacklist channels only once performance has degraded, e.g., in only a specific part of the network [68]. Indeed, these approaches continuously monitor the quality of the current channel and check whether it is satisfactory: if too much interference is detected, a frequency switch is carried out. We call these protocols *reactive hopping* protocols, as, to mitigate interference by frequency diversity, they first need to experience interference and a performance degradation. In this category fall protocols such as CoReDac [69], Chrisso [44], and ARCH [38]. The latter uses the expected number of transmissions (ETX) to monitor the quality of the link, and as soon as the ETX values collected in a given observation window exceed a certain threshold, a new channel is selected. The authors show that 15 minutes of observation are enough to predict channel reliability, and an interesting method is suggested to select the next channel. After blacklisting the current channel, ARCH hops to a new channel that is further away from the previous one. This has two benefits: on the one hand it avoids wideband interferers, on the other hand it avoids deep fades, as highlighted by Watteyne et al. [70].

The advantages of reactive protocols are, as mentioned above, the significant energy saving compared to hopping continuously and maintaining time synchronization among nodes. However, such protocols may not suit safety-critical systems, as they need to experience packet loss before performing a channel switch. Moreover, the switch is performed without knowledge about the stability of the other channels.

Proactive hopping. In order to overcome this problem, some other works try to avoid packet loss by predicting when the channel conditions will degrade and by hopping before this happens. We name protocols falling in this category *predictive hopping* protocols. This type of protocol typically uses channel quality estimation metrics to detect an early degradation of the channel [25, 40, 51]. However, because of the overhead required to continuously estimate the channel quality, no well-established proactive protocol has been developed yet. The work by Kerkez et al. [71] uses a sort of proactive strategy, as it keeps track of the quality of all channels by periodically forcing a channel switch. Another example is MuZi [72], in which as soon as the performance of a channel degrades, all the channels are scanned and a new reliable channel is selected. A fundamental role in the development of practical proactive protocols is played by proper interference metrics that can reliably assess the degradation of channels, as discussed in Sect. 2.3.1, as well as by an efficient link quality ranking algorithm [73].

2.4.2 Space Diversity Solutions

A solution widely investigated in the context of large and dense wireless sensor networks consists in avoiding interference by routing packets through different links [74]. Adaptive routing has been studied by several works that do

not explicitly target the presence of external interference, but rather aim for an effective link estimation in order to achieve reliable communication.

Alizai et al. [75] proposed to apply a bursty routing extension to detect short-term reliable links. Their approach allows a routing protocol to forward packets over long-range bursty links in order to minimize the number of transmissions in the network. Liu and Cerpa [76] have developed a receiver-driven estimator based on a machine learning approach to predict the short temporal quality of a link. Their estimation is based on trained models that predict the link quality using both packet reception rate and other physical layer parameters, such as RSSI, SNR and LQI.

Gonga et al. [77] have carried out a comparison between multichannel communication and adaptive routing, in order to determine which one guarantees more reliable communications in the presence of external interference and high link dynamics. On the one hand, the authors have shown the good performance of adaptive routing in dense wireless sensor networks. The key reason behind this is the selection of good long-term stable links, which avoids low-quality links that may be more vulnerable to external interference. When external interference is present, one could indeed try to route a packet towards a closer node, so that the probability that a stronger signal corrupts the packet is smaller. On the other hand, in sparse networks, where the choice of forwarding links is rather limited, adaptive routing loses its flexibility, and multichannel solutions yield better performance in terms of both average end-to-end delay and reliability.

2.4.3 Hardware Diversity Solutions

In case an ISM band is highly congested and it is hardly possible to find a reliable channel for communications, or in case it is not possible to route a packet through a different link, a new strategy needs to be selected in order to mitigate external interference. Several works have made use of wireless sensor network platforms equipped with dual radios to communicate in multiple ISM bands. Other works have proposed the use of directional antennas or spatially separated antennas to achieve more robust communications even in the presence of interference. We group these solutions in the hardware diversity class, and describe them in more detail in the remainder of this section.

Radio diversity. Kusy et al. [17] have shown that radio transceivers operating at dual widely spaced radio frequencies and through spatially separated antennas offer robust communication, high link diversity, and better interference mitigation. Using dual radios, the authors have shown, through experiments, a significant improvement in the end-to-end delivery rates and network stability, at the price of a slight increase in energy cost compared to a single radio approach. This solution is rather effective, but it is however only feasible for wireless sensor network platforms equipped with dual radios, such as the BTnode, the Mulle node, and the Opal node.

Antenna diversity. Rehmani et al. [78] have envisioned the possibility to design and implement a software-defined intelligent antenna switching capability for

wireless sensor nodes. More precisely, the authors have attached an inverted-F antenna to a TelosB mote in addition to the built-in antenna in order to achieve antenna diversity. Based on the wireless link condition, and in particular on physical layer measurements, the sensor node should then dynamically switch to the most appropriate antenna for communication.

Another options are dynamically steerable directional antennas, as shown experimentally by Giorgetti et al. [79]. Such antennas are able to dynamically control the gain as a function of direction, and, because of these properties, they can be very useful to increase the communication range and reducing the contention on the wireless medium. The authors have proposed a four-beam patch antenna and showed interference suppression from IEEE 802.11g systems. They have further discussed the use of the antenna as a form of angular diversity useful to cope with the variability of the radio signal. Other examples of directional antennas applied on wireless sensor nodes [80–82], in which a sensor node can concentrate the transmitted power towards the intended receiver dynamically.

2.4.4 Solutions Based on Redundancy

Several solutions exploit redundancy to mitigate the impact of external interference, such as the use of multiple headers, retransmissions, as well as forward error correction techniques. In the remainder of this section, we discuss those in detail.

Multiple headers. Liang et al. have experimentally shown that IEEE 802.11 transmitters can back off due to elevated channel energy when nearby IEEE 802.15.4 nodes start sending packets [8, 26]. When this happens, the IEEE 802.15.4 packet header is often corrupted, but the rest of the packet is still intact. Based on this observation, the authors have proposed the use of multi-headers to protect the IEEE 802.15.4 packets from the corruption generated by Wi-Fi interference. The authors have suggested that two additional headers represent a good trade-off between overhead and performance. It is important to notice that multi-headers are only effective in the so called symmetric region, i.e., when an IEEE 802.15.4 transmission is able to affect the behaviour of a IEEE 802.11 transmitter, because the bit errors occur mainly in the beginning of the packet. In contrast, in the asymmetric region, i.e., when the IEEE 802.15.4 signal is too weak to affect IEEE 802.11 behaviour, the bit errors are distributed across the packet in a uniform way.

Forward error correction techniques. In order to mitigate interference in the asymmetric region, Liang et al. [8, 26] have also proposed the use of forward error correction (FEC) techniques to recover from corrupted packets. FEC techniques use extra redundancy added to the original information frame to enable the correction of errors directly at the receiving node without the need for retransmission. When using forward error correction, the original message is encoded into a larger message by using an error correction code, which implies a longer time in which the radio is switched on, and a longer computation time for encoding and decoding the packet. The receiver then decodes the original message by applying the reverse

transformation of the error correction code. The redundancy in the encoded message allows the receiver to recover the original message in the presence of a limited number of bit errors. The authors demonstrated that Reed-Solomon (RS) correcting codes perform well while recovering packets corrupted by the activities of IEEE 802.11 [26].

However, FEC techniques pose a trade-off between data recovery capacity and its inherent payload and computation overhead. Forward error correction indeed creates overhead both on the receiver and the transmitter, and therefore requires a significant amount of energy as well as powerful nodes. Liang et al. [26] have shown that the time required to encode an original 65-byte message into an RS-encoded message with 30-byte parity is approximately 36 ms, whereas the decoding of the message depends on the presence of errors and can vary between 100 and 200 ms.

Backward error correction techniques. An often used alternative to forward error correction is the use of acknowledgement (ACK) or negative-acknowledgement (NACK) packets to trigger a retransmission of the corrupted frames. This solution may not necessarily lead to a good result in the presence of external interference, as retransmitted packets are prone to corruption as well as the original packet. Furthermore, when sending ACK/NACK packets, one may increase the channel congestion and the energy consumption of the nodes.

In order to minimize the energy consumption required for retransmissions, Hauer et al. [41] have developed an Automatic Repeat reQuest (ARQ) scheme that minimizes the amount of data that the sender needs to retransmit. In their scheme, a receiving node records the RSSI of the received packet during reception at high frequency, and tries to estimate the position of the error within the packet. The RSSI-based recovery mechanism is effective also in the presence of external interference, because collisions of frames with the transmissions generated by other devices such as IEEE 802.11 or IEEE 802.15.4 can be detected through an increase in the RSSI profile, which would otherwise be very stable (typically ± 1 dBm), as one can see in Fig. 2.4. The authors then propose an RSSI-based recovery mechanism, in which the receiving node triggers only the retransmission of the damaged portion of the packet, which implies a significant amount of energy in case of packets with relatively long data payloads.

Exploiting the stability of RSSI over time in absence of interference, Boano et al. [39] have proposed a novel protocol that uses jamming signals instead of message transmissions as the last step of a packet handshake between two nodes. This permits the two nodes to reliably agree on whether a packet was correctly received even in the presence of external interference. Agreement is indeed an issue in the presence of external interference, as ACK packets may be lost as well as original packets, and there is no way to guarantee the actual reception of a given packet. If two sensor nodes need to agree, for example, on a new time slot or frequency channel, message loss caused by external interference may break agreement in two different ways: none of the nodes use the new information (time slot, channel) and stick with the previous assignment, or—even worse—some nodes use the new information and some do not, leading to reduced performance and failures [39]. To get around this problem, the authors have proposed a jamming-based agreement protocol that con-

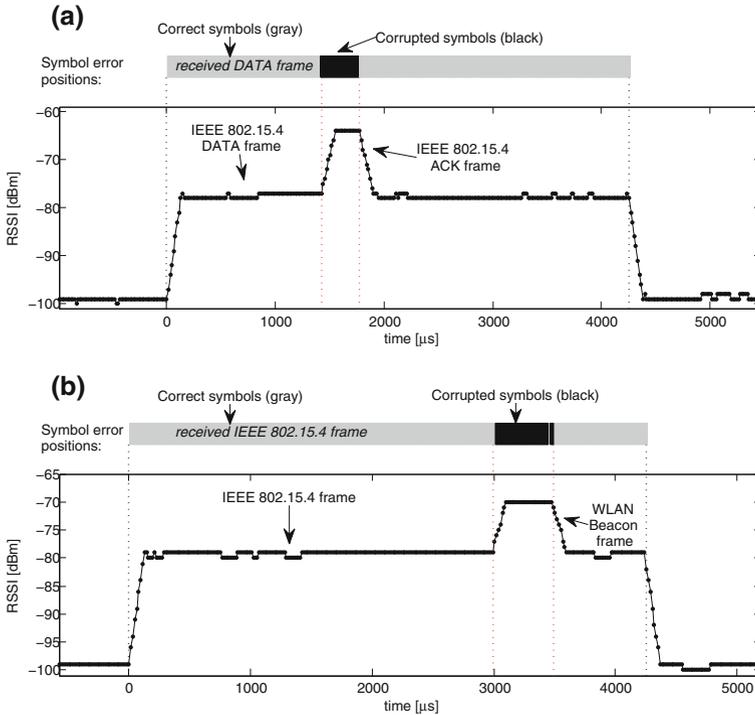


Fig. 2.4 Error positions and RSSI profiles of an IEEE 802.15.4 frame (133-byte PHY Protocol Data Unit) colliding with an IEEE 802.15.4 packet (*top*) and IEEE 802.11 beacon frame (*bottom*). The RSSI profile is measured on a Tmote Sky mote placed at less than 5 m from the interferer [41]

sists in a three-way handshake in which the last acknowledgement packet is sent in the form of a jamming signal. The transmission of a jamming signal has the property of being easily recognizable even in the presence of interference. As shown in Fig. 2.5, a jamming sequence results in a stable RSSI value above the sensitivity threshold of the radio, whereas in the presence of additional external interference, the RSSI register will return the maximum of the jamming signal and the interference signal due to the co-channel rejection properties of the radio. However, typical interference sources—in contrast to a jamming signal—do not produce continuous interference for long periods of time, rather they alternate between idle and busy. Therefore, by sending a jamming signal lasting long enough (i.e., longer than the longest busy period of the interference signal), one can unequivocally detect the absence of the jamming signal, by checking if any of the RSSI samples equals the sensitivity threshold of the radio. The authors have further shown that by carefully selecting a proper jamming time window, one can guarantee the identification of the ACK despite the presence of external interference, which makes this approach suitable for safety-critical systems.

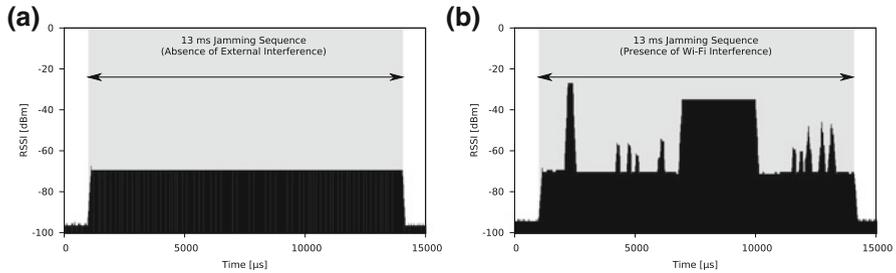


Fig. 2.5 RSSI values recorded during the transmission of a jamming sequence without external interference (a), and with external Wi-Fi interference (b) [39]

2.4.5 Time Diversity Solutions

Another class of external interference mitigation techniques is time diversity, which consists in either deferring transmissions, or scheduling them in such a way to avoid interference.

Reactive schemes. A basic way to mitigate interference is to defer transmission until interference clears, using for example the Carrier Sense Multiple Access with Collision Avoidance (CSMA-CA) technique. Those solutions, that we name *reactive*, react to the given interference pattern by adjusting protocol or parameter settings accordingly. In this context, the role of congestion back-off and Clear Channel Assessment (CCA) has been studied extensively in the presence of external interference. Boano et al. [52] have investigated the selection of a suitable congestion back-off scheme when using CCA and detecting a busy channel. They also investigate protocols or parameter settings that enable potentially more handshakes in case some fail due to interference, and extend X-MAC, showing an improved robustness to interference.

An important role is also played by the CCA threshold. Yuan et al. [83] have proposed a decentralized approach in which sensor nodes adaptively and distributively adjust their CCA thresholds in the presence of external heavy interference. The authors show that this approach substantially reduce the amount of discarded packets due to channel access failures, and hence increases the performance of sensornet protocols under interference.

Bertocco et al. [84] study the performance of different CCA modes in the presence of in-channel wide-band additive white gaussian noise (AWGN). Similarly, Petrova et al. [19] investigate the three CCA modes defined by the IEEE 802.15.4 PHY standard (energy above threshold, carrier sense only, and carrier sense with energy above threshold). They have observed that dynamic CCA thresholds can improve the performance of sensornet communications both in the overlapping and non-overlapping channels with IEEE 802.11n.

Proactive schemes. Another class of protocols is the one in which the sensor node tries not to defer transmissions, but rather to schedule them in a way to avoid

the interference of other devices. An example from this class, that we name *proactive schemes*, is the scheduling of packets proposed by Chowdury and Akyildiz [30]. In their protocol, the authors analyze the cases in which microwave ovens and Wi-Fi device are operating, and propose a scheme in which the sensor nodes transmit whenever the channel is predicted to be free based on the Wi-Fi traffic or microwave oven duty cycle. In the case of microwave oven interference, the sensor nodes can align their own sleep cycles with the duty cycle of the microwave oven, and synchronize their transmissions with the beginning of the off-time. In the case of Wi-Fi transmissions, the sensor nodes exploit the detection of Short Inter-Frame Space (SIFS) and Distributed Inter-Frame Space (DIFS). As also highlighted by Liang et al. [26], the peaked power pulses emitted by the sensor nodes interrupt the DIFS carrier sense and force a back-off among the contending Wi-Fi devices, leaving the channel free for the sensor nodes to complete their transmissions.

2.5 Experimenting with Interference

As interference can severely affect the reliability of wireless communications, there is a strong need for understanding the performance of existing sensornet protocols under interference, as well as designing and validating novel protocols that can deliver high and stable performance despite changing interference patterns. Furthermore, testing of the correct functionality of the system prior deployment is of fundamental importance in wireless sensor networks: several reports from real-world deployments have highlighted how the lack of testing can lead to a partial or complete system failure [3, 4].

An accurate validation and testing requires a proper infrastructure, in which realistic interference patterns can be created in an easily controllable, inexpensive, accurate, and repeatable way. This is in practice very difficult to obtain, and especially when it comes to radio interference and wireless communications, experimentation can be frustrating and may require a large amount of time and resources.

On the one hand, wireless propagation is extremely complex and dependent on a plethora of variables, such as radio hardware, antenna irregularities, geometry and nature (static or mobile) of the environment, presence of obstacles responsible for shadowing or multipath fading, as well as the environmental conditions (e.g., temperature [2]). These influences can only be modelled to a limited extent in simulators, and cannot be easily controlled when experimenting using real-hardware.

On the other hand, radio interference can be produced by several devices, and the generated patterns can be highly diverse, causing a given protocol to perform differently in different scenarios. Hence, one would need to verify several different possibilities and control several variables, which is in practice very difficult to obtain. Indeed, all possible scenarios cannot be exhaustively verified using real-hardware or simulators, as their number would be prohibitively large [85].

For these reasons, experimentation with radio interference, whether carried out in simulation, in a laboratory testbed, or in a real-world deployment, can be frustrating

and time-consuming, and needs to address several key aspects, such as accuracy and repeatability, that we analyse in the remainder of this section.

2.5.1 Requirements

We now summarize a list of key properties that experiments (carried out both in simulation and using real-hardware) involving radio interference should satisfy.

Realism and accuracy. When testing the reliability and robustness of a protocol or an application against external interference in a systematic fashion, one needs to set up realistic and credible experiments. The interference patterns used in the experiments must be accurate and be a representative set of how interference appears in reality. Having a device that is permanently interfering for long periods of time would not represent a realistic scenario, as it hardly occurs in practice (interference is instead typically bursty).

Device diversity. Experiencing packet loss due the presence of external interference is rather common nowadays, as it is typically enough to operate a sensor network in the presence of active Wi-Fi transmissions. In order to achieve a complete investigation, however, it is important to use different interference sources. A given protocol may be resilient to the periodic interference generated by microwave ovens, but may suffer the randomness of Bluetooth or the bursty nature of Wi-Fi transmissions, hence testing the protocol using only one interfering device is often not optimal. One should also make use of several heterogeneous devices at the same time, especially when testing protocols that exploit frequency diversity. For example, a multichannel protocol that blacklists congested channels would provide very good results in the presence of only one interfering device operating permanently on a given channel. However, the use of several wide-band devices operating concurrently using different frequencies may cause the protocol to perform poorly.

Spatial diversity. Because of the complexity of wireless propagation, and the intrinsic properties of low-power hardware such as antenna irregularities (see Chap. 1), it is also important to vary the location of both interfering sources and network nodes. Very often one aims to generate the “worse case” and hence places the interfering device very close to the sensor nodes, so to block their communications. This is often referred to as binary interference [18]: if a device is active, it automatically interferes the operations of a given sensor node. However, this may not necessarily lead to the worse case setting, as the most challenging scenarios may often be the ones in which interference affects the communications between a pair of nodes only intermittently, or when it only affects parts of the nodes in the network.

Temporal diversity. Different devices can interfere in different ways, and it is therefore important to try different patterns for each interference source. For example, in the case of a Wi-Fi device, if one varies the user activity, the transport protocol, or the packet size and other low-level parameters, one may obtain significantly different interference patterns that may lead to a different performance. Similarly, it is important to vary the usage pattern of a given device, as they often differ on the long-term.

As an example, the amount of traffic generated by a Wi-Fi station typically varies significantly from day to night and from weekday to weekends.

Scalability and controllability. The infrastructure for interference generation should support as many network nodes and interfering devices as possible, without requiring an excessive amount of time and resources, or additional costs. It is also very important to be able to easily control the interfering devices, as well as to ensure that the interference patterns that are generated correspond to the expected ones. This may not be an easy task in practice, as several devices often need to be set manually. For example, the activation of several microwave ovens in a large-scale network would be hardly feasible if the devices cannot be programmed remotely.

Repeatability. Being able to repeat an experiment under the same interference patterns is one of the most challenging aspects of experimenting with interference, especially when comparing the performance of different protocols on real hardware. On the one hand, wireless propagation is extremely complex and may be affected by several external factors that are usually unknown to the experimenter. On the other hand, most experiments are carried out in “RF jungles” such as office environments and residential buildings [86], where one does not usually have control on the background noise level or the presence of people, nor one can make sure that the experimental setup did not change between consecutive experiments, as discussed in Sect. 2.5.3.

2.5.2 *Experimenting Using Simulators*

Even though simulation environments offer high levels of scalability, controllability, and repeatability, they still lack realistic and accurate behaviour, especially when it comes to simulating wireless communications and interference. Although several solutions have been proposed in the literature to make simulation of wireless communications and interference more accurate [87–89], there are still several fundamental and implementation issues that affect the accuracy of simulators.

A fundamental issue is represented by mathematical models, such as the unit disk model, and the log-normal shadowing model, which do not capture the actual behaviour of wireless sensor nodes accurately [90]. This problem neutralizes in a sense the simplicity with which the experimenter can relocate network nodes and interfering devices in a simulation environment, as the inaccuracy of the propagation model may not be true to reality. Furthermore, the parametrization of the simulation models is still a major challenge to make simulation experiments realistic, as their choice is often rather arbitrary. Trace-based models such as the one offered by TOSSIM have several limitations, including the fact that the noise is established separately and randomly for each of the nodes, failing in representing link burstiness [88].

There are also a number of implementation issues. On the one hand, the repeatability across different simulators may be limited due to the usage of random number generators, as highlighted by Garg et al. [90]. On the other hand, simulators often come with a limited amount of models, and implementing a new solution (e.g.,

extending an existing simulator with impulsive interference in wireless sensor networks [91]) may not be trivial. Especially if one needs to simulate the behaviour of different devices generating interference at the same time, designing and implementing new models may also require a substantial amount of time.

As the creation of models that accurately reflect reality is a rather difficult task, some works propose to augment existing simulation tools with the playback of realistic interference traces [92]. An example of a simulator providing such capabilities is COOJA [93]: traces recorded using off-the-shelf sensor nodes can be incorporated directly into the simulation environment, improving the level of realism.

Despite the above limitations of existing simulators, a substantial number of experiments and evaluations regarding the impact of external radio interference have been carried out using simulation. Nethi et al. [94] have implemented a multichannel protocol to avoid interference and simulated its performance under interference using the ns-2 simulator, but did not provide details about the interference generation. Yuan et al. [83] have proposed a decentralized approach that adjusts CCA thresholds of IEEE 802.15.4 nodes in the presence of heavy interference, and validated it using OPNET using two scenarios involving IEEE 802.11b nodes. Several parameters of the simulation (transmission power, receiver sensitivity, data rate, CCA threshold, payload size) are given, along with a description of the position of the interfering nodes. OPNET has also been used by Shin et al. [95] to investigate the coexistence of ZigBee and Bluetooth devices exploiting the Suiteetooth package, and several parameters of the simulation (line-of-sight distance, path loss exponent, payload size, transmission power) are reported. Voigt et al. [69] evaluate a multichannel protocol in COOJA by adding disturber nodes generating interference that prevents the use of selected channels. Similarly, Iyer et al. [96] use jamming nodes on specific portions of the network in TOSSIM.

Due to the aforementioned limitations of existing simulators, some works have extended existing simulators with new interference patterns or implemented a new custom simulator. Bertocco et al. have extended the OMNeT++ simulator and supported the analysis of interference between IEEE 802.11b and IEEE 802.15.4-based networks [97]. Dominicis et al. have used this simulator to investigate coexistence issues when using the WirelessHART protocol [98]. Chowdhury et al. have written a custom C++ simulator to simulate the transmissions of IEEE 802.11b within a wireless sensor network, and varied the number of Wi-Fi nodes between 10 and 20, each of which generated packets at different rates [30]. Boers et al. have extended a SIDE-based emulator by adding the ability to define scripted external impulsive interference by creating a user-specified configuration, a new node type and running threads producing the specified interference [91].

2.5.3 *Experimenting Using Real Hardware*

Experiments on real-hardware, whether carried out in a real-world deployment or in a laboratory testbed, offer a higher level of accuracy and realism compared to simulation, but they do not scale well, and they often come at a very high cost.

Existing sensornet testbeds lack capabilities for interference generation, and upgrading them with additional heterogeneous devices in order to introduce interference sources is a costly, inflexible, and labor-intensive operation. One cannot easily place bulky equipment such as Wi-Fi access points and microwave ovens into a laboratory testbed, or, even worse, bring several devices to the deployment location, and ensure a stable power supply. Also, changes in the experimental setup, including the relocation or addition of interference sources, may need to be done manually. Changes in the behaviour of the interfering sources require manual or remote activation and programming of each device, which creates a significant overhead (especially in the case of unconventional equipment such as microwave ovens).

To avoid this problem, a number of sensornet testbeds are equipped with heterogeneous devices in order to enable the generation of different types of interference. One example is EasiTest [99], in which high-speed multi-radio nodes and low-speed single-radio nodes have been used to study the co-existence problem between IEEE 802.11 and IEEE 802.15.4 devices. However, the number and the size of such heterogeneous testbeds is still rather limited.

In the last years, many researchers have also proposed the use of Software Defined Radio (SDR) devices, such as USRP and WARP, to generate dynamic interference patterns. This choice is mainly driven by the easy reconfiguration and adaptivity of software defined radios. Following this idea, Sanchez et al. [100] have envisioned a novel testbed federation incorporating SDR devices, which would facilitate recording and playback of interference patterns. However, the cost of SDR hardware is still very high, and hence this approach does not scale to large testbeds.

A solution that was proposed in order to avoid the need of heterogeneous devices and speed up the setup time of experiments in an inexpensive way consists in using sensor motes as interfering devices [18, 69, 101, 102]. Although sensor motes can only interfere on a single channel and with limited transmission power, and hence the accuracy of the generated patterns is rather low, the clear advantage of this approach are the limited setup time, and the use of sensor nodes (no additional hardware required).

A big concern when experimenting using real-hardware experiments is repeatability. Burchfield et al. [86] have described the environments in which most wireless experiments are carried out as real “RF jungles”, as way too many assumptions are made on the environment, and there is hence a high risk of misinterpreting the data obtained from such experiments. For example, it is not necessarily true that experiments conducted at night are interference-free. In the same way, one can hardly make sure that, when generating specific interference patterns, no other interference source is present in the environment, as wireless propagation is affected by several external factors that are usually unknown to the experimenter and that can severely

bias experimental results. Indeed, very often, sensornet testbeds are located in office and residential buildings rich of activities of other wireless devices [103], and only very few studies have been carried out in special environments, e.g., anechoic chambers [47, 104] or shielded Faraday cages [30].

Another problem comes from the fact that real-world interference cannot be easily repeated. Gnawali et al. have highlighted that evaluating protocols by running them one at a time on real-hardware is not optimal since no experiment is absolutely repeatable [105], and this applies especially to experiments involving wireless systems, as wireless propagation depends on a myriad of factors. For example, experiments exploiting “ambient interference” surrounding the wireless sensor network testbed may not be a suitable option to compare several protocols or applications under realistic interference, as the interference patterns are not fully controllable and cannot be recreated precisely. This may not enable a fair comparison among different protocols.

Based on the available literature, we now classify the existing works studying external interference experimentally on real-hardware in four different categories, in which the experiments are carried out by:

- exploiting noisy environments;
- generating specific interference patterns exploiting existing equipment;
- generating specific interference patterns using a custom setup;
- using sensor nodes to generate interference.

Exploiting noisy environments. Several experiments exploit the “ambient interference” surrounding the wireless sensor network to evaluate their protocols or applications under realistic interference patterns.

Gonga et al. [77] run experiments in which the IEEE 802.11 access points co-located with the sensor network in an office environment act as sources for narrow band interference. The authors also report the presence of microwave ovens and people moving throughout the day, which makes the experiment more realistic due to the shadowing and multipath fading being introduced. Open office environments have also been used in several other experiments [65, 106–110].

Also university campuses have hosted several experiments, such as the ones by Zhou et al. [111] and Hauer et al. [41]. In the latter, the authors run their experiments on publicly accessible indoor sensor network testbeds deployed in university buildings (TWIST and MoteLab) surrounded by Wi-Fi access points. Noda et al. [51] have analyzed the bursty distribution of interference in a university library, and reported heavy traffic from co-located Wi-Fi networks.

Finally, Sha et al. [112] have analysed the spectrum usage in residential environments and collected several observations used to derive a new frequency-adaptive MAC protocol [38].

Exploiting noisy environments has two key advantages: the realism of real-world interference patterns, and the rather low-cost and effort required to set up the experiments, as there is no need to explicitly generate interference. This comes at the price of a complete uncontrollability and non-repeatability of the experiments: one

cannot have full knowledge of the devices that are actually interfering, nor can one differentiate their impact.

Generating specific interference patterns exploiting existing equipment. In several works, the authors generate specific interference patterns exploiting existing infrastructures.

Dutta et al. [113] evaluate their protocol on a university testbed with 94 TelosB nodes, and actively verify the performance while “a file transfer is in progress using a nearby 802.11 access point”. No information is reported on the protocol used and on the size of IEEE 802.11 packets. Similarly, Moeller et al. [114] perform experiments on a 40-motes indoor wireless sensor network testbed, and operate two 802.11 radios on a given channel, transmitting UDP packets of size 890 bytes. Also Lao et al. [115] and Kang et al. [116] exploit existing Wi-Fi access points in the proximity of an indoor testbed and download, from a laptop, a large amount of files from a server using the FTP protocol. Liang et al. [26] evaluated their protocol on a 57-node testbed and tested its performance using the interference originating from a co-located IEEE 802.11g testbed on the same building floor. Full details on the 802.11g testbed and the devices downloading are given. Musaloïou and Terzis [40] run six consecutive UDP transfers at different rates and report the impact on the Zigbee network, along with a detailed map of the testbed and the location of the access points. Rohde et al. [117] performed a measurement to quantify interference in a 2-story residential building. The authors used an iPhone in active connection with a Wi-Fi access point. Zhao et al. [99] directly program the existing infrastructure of Easitest and generate 1480-byte-long IEEE 802.11 UDP packets to study the coexistence with co-located wireless sensor networks.

Exploiting existing infrastructure to specifically generate certain interference patterns is a popular approach, and offers several advantages and disadvantages. On the one hand, the time required to set up the experiments is minimal, since one can exploit the existing infrastructure. On the other hand, the experimenter is limited by the amount, location, and type of the devices, and often cannot make sure that no other interference source was active in addition to the one being generated. Interestingly, most of the works following this approach consist in “continuous file transfer using Wi-Fi”, but without any further documentation or justification.

Generating specific interference patterns using a custom setup. The most popular way to experiment with interference seems, by far, to be the creation of a custom experimental setup. The latter is typically small-scale and involves a small amount of sensor nodes and the interfering device of interest (typically one or two units).

Examples of this approach are the works by Won et al. [68], that created an ad-hoc network of two 802.11b devices and run a file transfer. The same approach was used by Hauer et al. [25], Sikora and Groza [20], Ahmed et al. [31, 118, 119], Tang et al. [45] (that uses the iperf network testing tool instead of a file transfer), Ansari et al. [120], Petrova et al. [19], Hossian et al. [121], Huang et al. [122], Xu et al. [72], Shuaib et al. [123], and Jeong et al. [124].

Several studies experiment with multiple interfering devices, e.g., IEEE 802.11 and Bluetooth interference. This is the case for the studies of Penna et al. [23],

Bertocco et al. [22], and Arkoulis et al. [125]. Similar experiments were also carried out by Hou et al. [24] and Boano et al. [18], but in this case, also microwave ovens were used to generate interference. A distinctive feature of the experiments of Boano et al. [18] with Wi-Fi devices is the differentiation of the interference depending on the user activity. In their experiments, the authors present different results depending on the activity of the user on the machine (file transfer, video streaming, radio streaming).

Common features of experiments carried out using a custom setup are the rather detailed description of the experimental facility that, however, typically uses only a few (often just one) interfering device in a single configuration of speed, power and protocol. Especially in studies involving more advanced equipment, such as signal generators, and software-defined radios [126, 127], the setup is limited to a couple of nodes, but offers high degrees of controllability.

Using sensor nodes to generate interference. Several works [69, 101, 102] use sensor nodes as jammers and continuously transmit random packets, introducing noise into the frequencies of interest and evaluating the impact of such interference on several protocols. Although the transmission of packets is not fully controllable (e.g., inter-packet times [128]) and does not resemble the interference produced by devices using other technologies, the clear advantage of this approach are the limited setup time, and the use of sensor nodes (no additional hardware required). This kind of solution can be very useful if the experimenter is primarily interested in binary interference generation, i.e., in a setup in which the interfering nodes either blocks the communication of the nodes in their surroundings by emitting a strong-enough interference signal, or by not interfering at all.

Boano et al. [18] have enhanced this methodology and proposed to augment existing sensor network testbeds with JamLab, a low-cost infrastructure for the creation or playback of realistic and repeatable interference patterns. With JamLab, either a fraction of the existing nodes in a testbed are used to record and playback interfer-

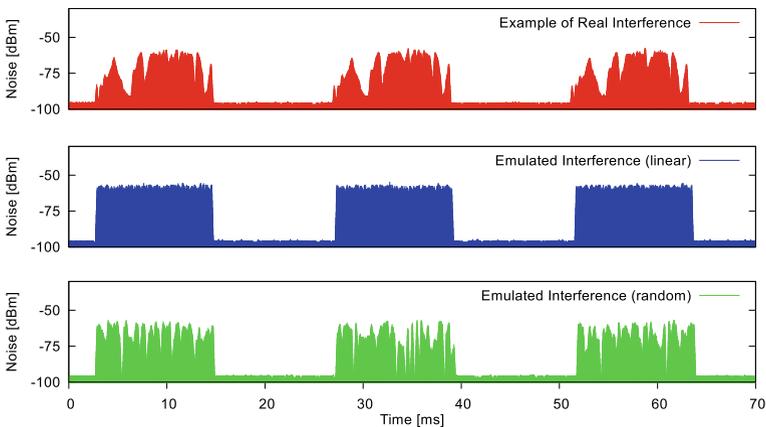


Fig. 2.6 Emulation of microwave oven interference (*top*) with fixed (*middle*) and random power (*bottom*) using JamLab [18]

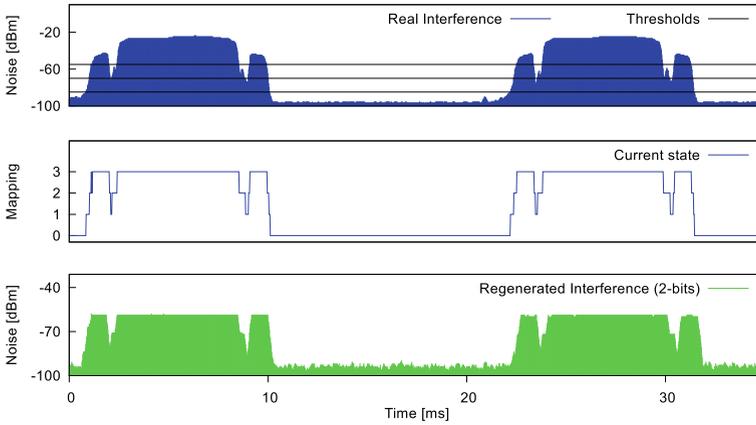


Fig. 2.7 Regenerated interference of a microwave oven using JamLab [18]

ence patterns, or a few additional notes are placed in the testbed area; hence the installation overhead is minimal. The nodes selected to generate interference can have two modes of operation: emulation, where a simplified model is used to generate interference patterns that resemble those generated by a specific appliance (such as a WiFi device or a microwave oven); and regeneration, where each interfering node autonomously samples the interference in the environment, compresses and stores it locally, and regenerates the recorded patterns later. The latter mode offers the possibility to record realistic interference patterns at the deployment site, and bring them back to the laboratory testbed for a more detailed study of their impact. Figures 2.6 and 2.7 show an example of emulated and regenerated interference using JamLab [18]. On the one hand, the interference is an accurate representation of a given model and fully repeatable (given that no other uncontrollable source of interference is present in the environment), as it repeats over time continuously. On the other hand, the hardware used to generate the interference are off-the-shelf sensor nodes, and therefore one can only emulate the behaviour of other devices with all the limitations of the sensor nodes, e.g., limited transmission power, operations on a single IEEE 802.15.4 channel. Several works have exploited JamLab to generate interference in existing sensor network testbeds, including [44, 52, 129–131].

2.5.4 Observations

Based on a careful analysis of existing published work in the area, we now illustrate our observations regarding experimentation with interference.

1. **Most experiments are carried out using real-hardware.** Because of the inaccuracy of existing simulators and because of the complexity of deriving reliable and precise models of different interference sources, the majority of experiments are carried out using real devices.

2. **Table experiments are very popular.** Most experiments involving real hardware are carried out using customized setups involving only few sensor nodes and one interfering device “on top of a table”. This is probably due to the necessity of selecting and controlling the interference patterns, and due to the fact that interference would be difficult to control on a large scale. On the other hand, a small-scale setup cannot capture situations in which different portions of a large network experience different interference patterns. Also in most “table experiments” the transmissions/noise generated by the interfering devices are typically strong enough to destroy the sensor network communication, given the short distances. It is hence not possible to capture cases in which the interference destroys only some of the packets sent by sensor nodes.
3. **Lack of frequency diversity.** It is perhaps not surprising that most experiments are carried out using sensor nodes operating in the 2.4GHz ISM band, given the high number of devices and technologies sharing these frequencies. However, very few studies are carried out in other frequency bands [15], and few studies are known about potential interference patterns caused by medical and RFID devices, as well as cellular phones in other frequencies.
4. **Lack of device diversity.** Most of the experiments target a specific interference source, the latter being often Wi-Fi, with a few experiments carried out under Bluetooth and microwave oven interference. It follows that several interference sources have still not been considered, and that external interference in wireless sensor networks is often conceptually referred to as Wi-Fi interference, or that Wi-Fi interference implicitly captures the worst-case of external interference.
5. **Lack of temporal diversity.** Apart from being carried out using Wi-Fi devices, the majority of experiments involving external interference often makes use of “continuous file transfers”. This choice is probably driven by the intention of generating a worse-case scenario with high levels of interference, which is typically the case of heavy file transfers. However, this is not necessarily a good choice. A multichannel protocol in the presence of a heavily congested channel would likely blacklist the channel and not experience any problem, compared to a multichannel protocol in the presence of intermittent interference among different channels. Also, the choice of file transfer often comes without a justification, and no further configurations of the interfering device are tested.
6. **Insufficient descriptions.** A fundamental task of scientists is the meticulous description of the setup and the surrounding environment in scientific papers. Although the typical size of a manuscript limits the space available for describing the setup of an experiment, some information is sometimes indispensable for a deep understanding of the results. For example, a fundamental piece of information is often the actual congestion of the medium during the experiments, which helps in understanding the packet loss rates. Only a few works, such as [65, 112, 114] provide the spectral traces of the available channels, and it is often difficult to understand the reasons for a certain performance. Without resorting to additional devices to capture the spectrum, one may just run one of the existing interference estimation metrics, such as [51], and report its results together with the performance of protocols, as done in [39]. Alternatively, one may use the Expected

Network Delivery [132], a metric that quantifies the delivery performance that a collection protocol can be expected to achieve given the network topology.

7. **Monitoring background activity.** A recurring assumption is the absence of unwanted signals during an experiment. Even the smallest and most controlled setup may however experience unwanted interfering sources, often unknown to the experimenter. It is a common practice in several works to use spectrum analyzers to verify the actual absence of unwanted interfering sources. Very popular is also the use of portable cheaper spectrum analyzers such as the Wi-Spy [25, 26, 32, 41, 65, 86, 107, 112].
8. **Comparability among experiments is difficult.** Despite several experiments are carried out using Wi-Fi and file transfers, it is difficult to compare them. Firstly, different experiments often have a different setup. Secondly, devices can be configured in several ways (for example, Wi-Fi access points) and the configuration information is rarely reported. Furthermore, the interference generated by file transfers using Wi-Fi devices often depends on the congestion of the backbone or on the actual distance from the access point. Therefore, results should be taken with a grain of salt, and their comparison is rather difficult.

Although experimentation with interference evolved significantly in the last years, these observations show that further research is necessary. Experiments involving sensor nodes and interfering devices “on top of a table” still represent the most common experimental setup, and there are very few studies that analyse the performance of a given protocol under interference generated using different devices and varying interference patterns.

Also, the congestion of the 2.4 GHz band has now been extensively studied, and there is a need to understand the performance of communication protocols in specific settings. For example, in hospital and clinical structures, it is important to make sure that sensor networks do not interfere with medical equipment and vice-versa, as alarm procedures must be triggered immediately when monitoring life-threatening deteriorations in vital signs of hospitalized patients.

Furthermore, a large number of protocols have been developed and proposed to mitigate the impact of interference, but there is no comprehensive study comparing their performance under different interference patterns. Experiments similar to [52] would significantly help in understanding which approach (i.e., frequency hopping, forward or backward error correction, routing through different links) offers the best performance in the presence of a given interference pattern.

2.6 Conclusion

Radio interference is receiving increasing attention in wireless sensor network research. As more and more wireless devices are being deployed, there is a strong need to increase the robustness of communications carried out in unlicensed frequencies, as they are vulnerable to the interference generated by other wireless appliances.

In order to cope with the massive proliferation of wireless devices in our everyday life, new standards are being introduced to increase the number of unlicensed frequencies that can be used for communication. Recent amendments, for example, have standardized the use of the UWB technology for wireless sensor networks, as it offers excellent interference immunity and low complexity at rather low costs. New frequency bands will be released in the near future to mitigate coexistence problems in existing bands (e.g., in the IEEE 802.11ac standard, Wi-Fi devices will not operate in the 2.4 GHz band anymore, and therefore a big portion of potential interfering devices will move to other frequencies). However, the increasing rate in which wireless devices are being deployed hints that there will still be several devices sharing the same frequencies, and that hence will need to coexist. This requires the design of protocols that guarantee reliable and robust communications among wireless sensor networks, especially for real-time and safety-critical applications that can hardly tolerate high packet loss rates and long latencies.

In the last decade, the research community has actively tackled the problem, and provided several solutions aimed to mitigate external interference in wireless sensor networks. The use of multichannel protocols, radio diversity, redundancy in all its forms (e.g., forward error correction techniques, multiple packet headers) was shown to increase significantly the robustness of communications in the presence of external interference.

However, the dynamism of interference requires solutions that can adapt at runtime to the changing interference patterns. In their renowned hitch-hiker's guide to successful wireless sensor network deployments, Barrenetxea et al. [16] recommend to make sure that the radio frequency used by the sensor nodes is not already in use. This task is hard to fulfil, as interference changes dynamically over time, and therefore there is a need for lightweight and efficient solutions that adapt to interference at runtime. The research community still needs to come up with self-learning solutions that can efficiently adapt the behaviour of sensor nodes at runtime, such as a dynamic protocol selection, or the dynamic adjustment of model parameters. The main obstacle towards this goal is currently the inefficiency of energy detection, as the radio transceiver needs to be turned on in listening mode for extended periods of time.

Recent advances in cognitive radio technologies have highlighted the possibility to apply dynamic spectrum access techniques in order to get access to less congested spectrum [133]. Cognitive radio capable sensor nodes can adapt to varying channel conditions and adapt their parameters at runtime, in order to increase the transmission efficiency and save energy. The research community has hence started to explore this promising paradigm, in order to adopt cognitive radio capabilities in wireless sensor networks in the near future.

References

1. Boano CA, Brown J, He Z, Roedig U, Voigt T (2009) Low-power radio communication in industrial outdoor deployments: the impact of weather conditions and ATEX-compliance. In: Proceedings of the 1st international conference on sensor networks applications, experimentation and logistics (SensAppeal), pp 159–176
2. Boano CA, Brown J, Tsiiftes N, Roedig U, Voigt T (2010) The impact of temperature on outdoor industrial sensor network applications. *IEEE Trans Industr Inform* 6(3):451–459
3. Beutel J, Römer K, Ringwald M, Woehrle M (2009) Deployment techniques for sensor networks. In: Ferrari G (ed) *Sensor networks*. Springer, Berlin, pp 219–248
4. Langendoen K, Baggio A, Visser O (2006) Murphy loves potatoes: experiences from a pilot sensor network deployment in precision agriculture. In: Proceedings of the 14th international workshop on parallel and distributed real-time systems (WPDRTS), pp 174–181
5. GINSENG: Performance control in wireless sensor networks. <http://www.ict-ginseng.eu/>
6. Chipara O, Lu C, Bailey TC, Roman GC (2010) Reliable clinical monitoring using wireless sensor networks: experiences in a step-down hospital unit. In: Proceedings of the 8th ACM international conference on embedded networked sensor systems (SenSys), pp 155–168
7. Jeong J (2009) Wireless sensor networking for intelligent transportation systems. Ph.D. thesis, University of Minnesota, MN, USA
8. Liang CJM (2011) Interference characterization and mitigation in large-scale wireless sensor networks. Ph.D. thesis, John Hopkins University, Baltimore
9. Son D, Krishnamachar B, Heidemann J (2006) Experimental study of concurrent transmission in wireless sensor networks. In: Proceedings of the 4th international conference on embedded networked sensor systems (SenSys '06). ACM, New York, pp 237–250
10. Zhou G, Stankovic JA, Son SH (2006) Crowded spectrum in wireless sensor networks. In: Proceedings of the 3rd workshop on embedded networked sensors (EmNets)
11. IEEE 802.15.4 Working Group (2003) IEEE Standard for Local and Metropolitan Area Networks—Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs), IEEE std 802.15.4-2003 edn
12. IEEE 802.15.4 Working Group (2006) IEEE Standard for Local and Metropolitan Area Networks—Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs), IEEE std 802.15.4-2006 edn
13. IEEE 802.15.4 Working Group (2011) IEEE Standard for Local and Metropolitan Area Networks—Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs), IEEE std 802.15.4-2011 edn
14. Zhang J, Orlik PV, Sahinoglu Z, Molisch AF, Kinney P (2009) UWB systems for wireless sensor networks. *Proc IEEE* 97(2):313–331
15. Woehrle M, Bor M, Langendoen KG (2012) 868 MHz: a noiseless environment, but no free lunch for protocol design. In: Proceedings of the 9th international conference on networked sensing systems (INSS)
16. Barrenetxea G, Ingelrest F, Schaefer G, Vetterli M (2008) The Hitchhiker's guide to successful wireless sensor network deployments. In: Proceedings of the 6th ACM international conference on embedded networked sensor systems (SenSys), pp 43–56
17. Kusy B, Richter C, Hu W, Afanasyev M, Jurdak R, Brüning M, Abbott D, Huynh C, Ostry D (2011) Radio diversity for reliable communication in WSNs. In: Proceedings of the 10th IEEE international conference on information processing in sensor networks (IPSN), pp 270–281
18. Boano CA, Voigt T, Noda C, Römer K, Zúñiga MA (2011) Jamlab: augmenting sensor network testbeds with realistic and controlled interference generation. In: Proceedings of the 10th IEEE international conference on information processing in sensor networks (IPSN), pp 175–186
19. Petrova M, Wu L, Mähönen P, Riihijärvi J (2007) Interference measurements on performance degradation between colocated IEEE 802.11g/n and IEEE 802.15.4 networks. In: Proceedings of the international conference on networking (ICN), pp 93–98

20. Sikora A, Groza VF (2005) Coexistence of IEEE 802.15.4 with other systems in the 2.4 GHz-ISM-band. In: Proceedings of the IEEE conference on instrumentation and measurement technology (IMTC), pp 1786–1791
21. Farahani S (2008) ZigBee wireless networks and transceivers. Elsevier Inc., Amsterdam
22. Bertocco M, Gamba G, Sona A (2008) Is CSMA/CA really efficient against interference in a wireless control system? an experimental answer. In: Proceedings of the 13th IEEE international conference on emerging technologies and factory automation (ETFA), pp 885–892
23. Penna F, Pastrone C, Spirito M, Garelli R (2009) Measurement-based analysis of spectrum sensing in adaptive WSNs under Wi-Fi and bluetooth interference. In: Proceedings of the 69th IEEE vehicular technology conference (VTC), pp 1–5
24. Huo H, Xu Y, Bilen CC, Zhang H (2009) Coexistence issues of 2.4 GHz sensor networks with other RF devices at home. In: Proceedings of the 3rd international conference on sensor technologies and applications (SENSORCOMM), pp 200–205
25. Hauer JH, Handziski V, Wolisz A (2009) Experimental study of the impact of WLAN interference on IEEE 802.15.4 body area networks. In: Proceedings of the 6th European conference on wireless sensor networks (EWSN), pp 17–32
26. Liang CJM, Priyantha NB, Liu J, Terzis A (2010) Surviving Wi-Fi interference in low power zigbee networks. In: Proceedings of the 8th ACM conference on embedded networked sensor systems (SenSys '10). ACM, New York, pp 309–322
27. Kamerman A, Erkocevic N (1997) Microwave oven interference on wireless LANs operating in the 2.4 GHz ISM band. In: Proceedings of the 8th IEEE international symposium on personal, indoor and mobile radio communications (PIRMC), vol. 3, pp. 1221–1227
28. Vollmer M (2004) Physics of the microwave oven. *Phys Education* 39:74–81
29. Taher TM, Misurac MJ, LoCicero JL, Ucci DR (2008) Microwave oven signal modeling. In: Proceedings of the IEEE wireless communications and networking conference (WCNC), pp 1235–1238
30. Chowdhury KR, Akyildiz IF (2009) Interferer classification, channel selection and transmission adaptation for wireless sensor networks. In: Proceedings of the IEEE international conference on communications (ICC), pp 1–5
31. Ahmed N, Kanhere S, Jha S (2009) Multi-channel interference in wireless sensor networks. In: Proceedings of the 8th IEEE international conference on information processing in sensor networks (IPSN), poster session, pp 367–368
32. Bello LL, Toscano E (2009) Coexistence issues of multiple co-located IEEE 802.15.4/zigbee networks running on adjacent radio channels in industrial environments. *IEEE Trans Ind Inform* 5:157–167
33. Incel ÖD, Dulman S, Jansen P, Mullender S (2006) Multi-channel interference measurements for wireless sensor networks. In: Proceedings of the 31st IEEE international conference on communications (LCN), pp 694–701
34. Wu Y, Stankovic JA, He T, Lin S (2008) Realistic and efficient multi-channel communications in wireless sensor networks. In: Proceedings of the 27th IEEE international conference on computer communications (INFOCOM), pp 1193–1201
35. Xing G, Sha M, Huang J, Zhou G, Wang X, Liu S (2009) Multi-channel interference measurement and modeling in low-power wireless networks. In: Proceedings of the 30th IEEE international real-time systems symposium (RTSS), pp. 248–257
36. Xu X, Lao J, Zhang Q (2010) Design of non-orthogonal multi-channel sensor networks. In: Proceedings of the 30th IEEE international conference on distributed computing systems (ICDCS), pp 358–367
37. Srinivasan K, Dutta P, Tavakoli A, Levis P (2010) An empirical study of low-power wireless. *ACM Trans Sens Netw* 6:1–49
38. Sha M, Hackmann G, Lu C (2011) ARCH: practical channel hopping for reliable home-area sensor networks. In: Proceedings of the 17th IEEE international real-time and embedded technology and applications symposium (RTAS), pp 305–315

39. Boano CA, Zúñiga MA, Römer K, Voigt T (2012) JAG: reliable and predictable wireless agreement under external radio interference. In: Proceedings of the 33rd IEEE international real-time systems symposium (RTSS), pp 315–326
40. Musaloiu-ER, Terzis A (2007) Minimising the effect of WiFi interference in 802.15.4 wireless sensor networks. *Int J Sens Netw (IJSNet)* 3(1):43–54
41. Hauer JH, Willig A, Wolisz A (2010) Mitigating the effects of RF interference through RSSI-based error recovery. In: Proceedings of the 7th European conference on wireless sensor networks (EWSN), pp 224–239
42. Chen Y, Terzis A (2010) On the mechanisms and effects of calibrating RSSI measurements for 802.15.4 radios. In: Proceedings of the 7th European conference on wireless sensor networks (EWSN). LNCS 5970, pp 272–288
43. Doddavenkatappa M, Chan MC, Leong B (2011) Improving link quality by exploiting channel diversity in wireless sensor networks. In: Proceedings of the IEEE 32nd real-time systems symposium (RTSS), pp 159–169
44. Iyer V, Woehrle M, Langendoen K (2011) Chryso: a multi-channel approach to mitigate external interference. In: Proceedings of the 8th IEEE communications society conference on sensor, mesh, and ad hoc communications and networks (SECON)
45. Tang L, Sun Y, Gurewitz O, Johnson DB (2011) EM-MAC: a dynamic multichannel energy-efficient MAC protocol for wireless sensor networks. In: Proceedings of the 12th ACM international symposium on mobile ad hoc networking and computing (MobiHoc), pp 23:1–23:11
46. Zacharias S, Newe T, O’Keeffe S, Lewis E (2012) Identifying sources of interference in RSSI traces of a single IEEE 802.15.4 channel. In: Proceedings of the 8th international conference on wireless and mobile communications (ICWMC)
47. Hermans F, Rensfelt O, Voigt T, Ngai E, Larzon LA, Gunningberg P (2013) SoNIC: classifying interference in 802.15.4 sensor networks. In: Proceedings of the 12th ACM/IEEE international conference on information processing in sensor networks (IPSN)
48. Boers NM, Nikolaidis I, Gburzynski P (2010) Patterns in the RSSI traces from an indoor urban environment. In: Proceedings of the 15th international workshop on computer aided modeling, analysis and design of communication links and networks (CAMAD), pp 61–65
49. Boers NM, Nikolaidis I, Gburzynski P (2012) Sampling and classifying interference patterns in a wireless sensor network. *ACM Trans Sens Netw (TOSN)* 9(1):1–19
50. Stabellini L, Zander J (2010) Energy-efficient detection of intermittent interference in wireless sensor networks. *Int J Sens Netw (IJSNET)* 8(1):27–40
51. Noda C, Prabh S, Alves M, Boano CA, Voigt T (2011) Quantifying the channel quality for interference-aware wireless sensor networks. *ACM SIGBED Rev* 8(4):43–48
52. Boano CA, Voigt T, Tsiftes N, Mottola L, Römer K, Zúñiga MA (2010) Making sensor MAC protocols robust against interference. In: Proceedings of the 7th European conference on wireless sensor networks (EWSN). LNCS 5970, pp 272–288
53. Bianchi G (2000) Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE J Sel Areas Commun* 18(3):535–547
54. Garetto M, Chiasserini CF (2005) Performance analysis of 802.11 WLANs under sporadic traffic. In: Proceedings of the 4th IFIP-TC6 networking conference (NETWORKING), Waterloo, Canada
55. Yang D, Xu Y, Gidlund M (2011) Wireless coexistence between IEEE 802.11- and IEEE 802.15.4-based networks: a survey. *Int J Distrib Sens Netw (IJDSN)* 2011:17pp
56. Angelopoulos CM, Nikolettseas S, Theofanopoulos GC (2011) A smart system for garden watering using wireless sensor networks. In: Proceedings of the 9th ACM international symposium on mobility management and wireless access (MobiWac), pp 167–170
57. Incel ÖD, Jansen P, Mullender S (2011) MC-LMAC: a multi-channel MAC protocol for wireless sensor networks. *J Ad Hoc Netw* 9:73–94
58. Kim Y, Shin H, Cha H (2008) Y-MAC: an energy-efficient multi-channel MAC protocol for dense wireless sensor networks. In: Proceedings of the 7th IEEE international conference on information processing in sensor networks (IPSN), pp 53–63

59. Salajegheh M, Soroush H, Kalis A (2007) HyMAC: hybrid TDMA/FDMA medium access control protocol for wireless sensor networks. In: Proceedings of the 18th IEEE international symposium on personal, indoor and mobile radio communications (PIMRC)
60. So HSW, Walrand J, Mo J (2007) McMAC: a parallel rendezvous multi-channel MAC protocol. In: Proceedings of the IEEE wireless communications and networking conference (WCNC), pp 334–339
61. Wu Y, Keally M, Zhou G, Mao W (2009) Traffic-aware channel assignment in wireless sensor networks. In: Proceedings of the 4th international conference on wireless algorithms, systems, and applications (WASA), pp 479–488
62. Watteyne T, Mehta A, Pister K (2009) Reliability through frequency diversity: Why channel hopping makes sense. In: Proceedings of the 6th international symposium on performance evaluation of wireless ad hoc, sensor, and ubiquitous networks (PE-WASUN)
63. Pister K, Doherty L (2008) TSMP: time synchronized mesh protocol. In: Proceedings of the IASTED international symposium on distributed sensor networks (DSN), pp 391–398
64. Song J, Han S, Mok A, Chen D, Lucas M, Nixon M, Pratt W (2008) WirelessHART: applying wireless technology in real-time industrial process control. In: Proceedings of the 14th IEEE international real-time and embedded technology and applications symposium (RTAS), pp 377–386
65. Du P, Roussos G (2011) Adaptive channel hopping for wireless sensor networks. In: Proceedings of the IEEE international conference on selected topics in mobile and wireless networking (iCOST), pp 19–23
66. Yoon SU, Murawski R, Ekici E, Park S, Mir ZH (2010) Adaptive channel hopping for interference robust wireless sensor networks. In: Proceedings of the IEEE international conference on communications (ICC), pp 432–439
67. Chen H, Cui L, Lu S (2009) An experimental study of the multiple channels and channel switching in wireless sensor networks. In: Proceedings of the 4th international symposium on innovations and real-time applications of distributed sensor networks (IRADSN), pp 54–61
68. Won C, Youn JH, Ali H, Sharif H, Deogun J (2005) Adaptive radio channel allocation for supporting coexistence of 802.15.4 and 802.11b. In: Proceedings of the 62nd IEEE vehicular technology conference (VTC), pp 2522–2526
69. Voigt T, Österlind F, Dunkels A (2008) Improving sensor network robustness with multi-channel convergecast. In: Proceedings of the 2nd ERCIM workshop on e-Mobility
70. Watteyne T, Lanzisera S, Mehta A, Pister KS (2010) Mitigating multipath fading through channel hopping in wireless sensor networks. In: Proceedings of the IEEE international conference on communications (ICC), pp 1–5
71. Kerkez B, Watteyne T, Magliocco M, Glaser S, Pister K (2009) Feasibility analysis of controller design for adaptive channel hopping. In: Proceedings of the 4th international conference on performance evaluation methodologies and tools (VALUETOOLS)
72. Xu R, Shi G, Luo J, Zhao Z, Shu Y (2011) MuZi: multi-channel zigbee networks for avoiding WiFi interference. In: Proceedings of the 4th international conference on cyber, physical and social computing (CPSOCOM), pp 323–329
73. Zúñiga MA, Irzynska I, Hauer JH, Voigt T, Boano CA, Römer K (2011) Link quality ranking: getting the best out of unreliable links. In: Proceedings of the 7th IEEE international conference on distributed computing in sensor systems (DCOSS), pp 1–8
74. Radi M, Dezfouli B, Bakar KA, Lee M (2012) Multipath routing in wireless sensor networks: survey and research challenges. *Sensors* 12(1):650–685
75. Alizai MH, Landsiedel O, Bitsch Link JA, Götz S, Wehrle K (2009) Bursty traffic over bursty links. In: Proceedings of the 7th ACM conference on embedded networked sensor systems (SenSys), pp 71–84
76. Liu T, Cerpa A (2011) Foresee (4c): wireless link prediction using link features. In: Proceedings of the 10th IEEE international conference on information processing in sensor networks (IPSN), pp 294–305
77. Gongga A, Landsiedel O, Soldati P, Johansson M (2012) Revisiting multi-channel communication to mitigate interference and link dynamics in wireless sensor networks. In: Proceedings of the 8th IEEE international conference on distributed computing in sensor systems (DCOSS)

78. Rehmani MH, Alves T, Lohier S, Rachedi A, Poussot B (2012) Towards intelligent antenna selection in IEEE 802.15.4 wireless sensor networks. In: Proceedings of the 13th ACM international symposium on mobile ad hoc networking and computing (MobiHoc), pp 245–246
79. Giorgetti G, Cidronali A, Gupta SK, Manes G (2007) Exploiting low-cost directional antennas in 2.4 GHz IEEE 802.15.4 wireless sensor networks. In: Proceedings of the 10th European conference on wireless technologies (ECWT), pp 217–220
80. Nilsson M (2009) Directional antennas for wireless sensor networks. In: Proceedings of the 9th Scandinavian workshop on wireless adhoc network (Adhoc)
81. Öström E, Mottola L, Nilsson M, Voigt T (2010) Smart antennas made practical: the SPIDA way. In: Proceedings of the 9th ACM/IEEE international conference on information processing in sensor networks (IPSN), demo session, pp 438–439
82. Voigt T, Hewage KC, Mottola L (2013) Understanding link dynamics in wireless sensor networks with dynamically steerable directional antennas. In: Proceedings of the 10th European conference on wireless sensor networks (EWSN)
83. Yuan W, Linnartz JPM, Niemegeers IG (2010) Adaptive CCA for IEEE 802.15.4 wireless sensor networks to mitigate interference. In: Proceedings of the IEEE wireless communication and networking conference (WCNC), pp 1–5
84. Bertocco M, Gamba G, Sona A (2007) Experimental optimization of CCA thresholds in wireless sensor networks in the presence of interference. In: Proceedings of the IEEE Europe the workshop on electromagnetic compatibility (EMC)
85. Woehrle M (2010) Testing of wireless sensor networks. Ph.D. thesis, Eidgenössische Technische Hochschule (ETH), Zürich, Switzerland
86. Burchfield R, Nourbakhsh E, Dix J, Sahu K, Venkatesan S, Prakash R (2009) RF in the jungle: effect of environment assumptions on wireless experiment repeatability. In: Proceedings of the IEEE international conference on communications (ICC), pp 4993–4998
87. Kamthe A, nán MACP, Cerpa AE (2009) M&M: multi-level Markov model for wireless link simulations. In: Proceedings of the 7th ACM conference on embedded networked sensor systems (SenSys), pp 57–70
88. Lee H, Cerpa A, Levis P (2007) Improving wireless simulation through noise modeling. In: Proceedings of the 6th international conference on information processing in sensor networks (IPSN '07). ACM, New York, pp 21–30
89. Zhou G, He T, Krishnamurthy S, Stankovic J (2006) Models and solutions for radio irregularity in wireless sensor networks. *ACM Trans Sens Netw (TOSN)* 2:221–262
90. Garg K, Förster A, Puccinelli D, Giordano S (2011) Towards realistic and credible wireless sensor network evaluation. In: Proceedings of the 3rd international conference on ad hoc networks (ADHOCNETS), pp 49–64
91. Boers NM, Nikolaidis I, Gburzynski P (2012) Impulsive interference avoidance in dense wireless sensor networks. In: Proceedings of the 11th international conference on ad-hoc networks and wireless (AdHocNow)
92. Boano CA, Römer K, Österlind F, Voigt T (2011) Realistic simulation of radio interference in COOJA. In: Adjunct proceedings of the 8th European conference on wireless sensor networks (EWSN), demo session, pp 36–37
93. Österlind F (2011) Improving low-power wireless protocols with timing-accurate simulation. Ph.D. thesis, Uppsala University, Uppsala, Sweden
94. Nethi S, Nieminen J, Jäntti R (2011) Exploitation of multi-channel communications in industrial wireless sensor applications: avoiding interference and enabling coexistence. In: Proceedings of the IEEE wireless communications and networking conference (WCNC), pp 345–350
95. Shin SY, Kang JS, Park HS (2009) Packet error rate analysis of zigbee under interferences of multiple bluetooth piconets. In: Proceedings of the 69th IEEE vehicular technology conference (VTC)
96. Iyer V, Woehrle M, Langendoen K (2010) Chamaeleon: exploiting multiple channels to mitigate interference. In: Proceedings of the 7th international conference on networked sensing systems (INSS), pp 65–68

97. Bertocco M, Gamba G, Sona A, Tramarin F (2008) Investigating wireless networks coexistence issues through an interference aware simulator. In: Proceedings of the 13th IEEE international conference on emerging technologies and factory automation (ETFA), pp 1153–1156
98. Dominicis CMD, Ferrari P, Flammini A, Sisinni E, Bertocco M, Giorgi G, Narduzzi C, Tramarin F (2009) Investigating WirelessHART coexistence issues through a specifically designed simulator. In: Proceedings of the IEEE international instrumentation and measurement technology conference (I2MTC), pp 1085–1090
99. Zhao Z, Yang GH, Liu Q, Li V, Cui L (2010) Easitest: a multi-radio testbed for heterogeneous wireless sensor networks. In: Proceedings of the IET international conference on wireless sensor network (IET-WSN), pp 104–108
100. Sanchez A, Moerman I, Bouckaert S, Willkomm D, Hauer JH, Michailow N, Fettweis G, Dasilva L, Tallon J, Pollin S (2011) Testbed federation: an approach for experimentation-driven research in cognitive radios and cognitive networking. In: Proceedings of the of the 20th future network and mobile summit
101. Xu W, Trappe W, Zhang Y (2008) Defending wireless sensor networks from radio interference through channel adaptation. *ACM Trans Sens Netw (TOSN)* 4:1–34
102. Zhou G, Lu J, Wan CY, Yarvis MD, Stankovic JA (2008) BodyQoS: adaptive and radio-agnostic QoS for body sensor networks. In: Proceedings of the 27th IEEE international conference on computer communications (INFOCOM), pp 565–573
103. Sakamuri D (2008) NetEye: a wireless sensor network testbed. Master's thesis, Wayne State University, Detroit, Michigan
104. Khaleel H, Pastrone C, Penna F, Spirito M, Garellò R (2009) Impact of Wi-Fi traffic on the IEEE 802.15.4 channels occupation in indoor environments. In: Proceedings of the 11th international conference on electromagnetics in advanced applications (ICEAA)
105. Gnawali O, Guibas L, Levis P (2010) Case for evaluating sensor network protocols concurrently. In: Proceedings of the 5th ACM international workshop on wireless network testbeds, experimental evaluation and characterization (WiNTECH), pp 47–54
106. Ortiz J, Culler D (2008) Exploring diversity: evaluating the cost of frequency diversity in communication and routing. In: Proceedings of the 6th ACM international conference on embedded networked sensor systems (SenSys), poster session, pp 411–412
107. Ortiz J, Culler D (2010) Multichannel reliability assessment in real world WSNs. In: Proceedings of the 9th IEEE international conference on information processing in sensor networks (IPSN), pp 162–173
108. Stabellini L (2010) Energy-aware channel selection for cognitive wireless sensor networks. In: Proceedings of the 7th international symposium on wireless communication systems (ISWCS), pp 892–896
109. Gnawali O, Fonseca R, Jamieson K, Moss D, Levis P (2009) Collection tree protocol. In: Proceedings of the 7th ACM conference on embedded networked sensor systems (SenSys '09). ACM, New York, pp 90–100
110. Stabellini L, Parhizkar MM (2010) Experimental comparison of frequency hopping techniques for 802.15.4-based sensor networks. In: Proceedings of the 4th international conference on mobile ubiquitous computing, systems, services and technologies (UBICOMM), pp 110–116
111. Zhou G, Lu L, Krishnamurthy S, Keally M, Ren Z (2009) SAS: self-adaptive spectrum management for wireless sensor networks. In: Proceedings of the 18th international conference on computer communications and networks (ICCCN), pp 1–6
112. Sha M, Hackmann G, Lu C (2011) Multi-channel reliability and spectrum usage in real homes: empirical studies for home-area sensor networks. In: Proceedings of the 19th IEEE international workshop on quality of service (IWQoS), pp 39:1–39:9
113. Dutta P, Dawson-Haggerty S, Chen Y, Liang CJM, Terzis A (2010) Design and evaluation of a versatile and efficient receiver-initiated link layer for low-power wireless. In: Proceedings of the 8th ACM international conference on embedded networked sensor systems (SenSys), pp 1–14

114. Moeller S, Sridharan A, Krishnamachari B, Gnawali O (2010) Routing without routes: the backpressure collection protocol. In: Proceedings of the 9th international conference on information processing in sensor networks (IPSN), pp 279–290
115. Lau SY, Lin TH, Huang TY, Ng IH, Huang P (2009) A measurement study of zigbee-based indoor localization systems under RF interference. In: Proceedings of the 4th ACM international workshop on experimental evaluation and characterization (WINTTECH), pp 35–42
116. Kang MS, Chong JW, Hyun H, Kim SM, Jung BH, Sung DK (2007) Adaptive interference-aware multi-channel clustering algorithm in a zigbee network in the presence of WLAN interference. In: Proceedings of the 2nd international symposium on wireless pervasive computing (ISWPC)
117. Rohde J, Toftegaard TS (2011) Mitigating the impact of high interference levels on energy consumption in wireless sensor networks. In: Proceedings of the 2nd international conference on wireless communication, vehicular technology, information theory and aerospace and electronic systems technology (Wireless VITAE), pp 1–5
118. Ahmed N, Kanhere SS, Jha S (2010) Mitigating the effect of interference in wireless sensor networks. In: Proceedings of the 35th IEEE international conference on local computer networks (LCN), pp 160–167
119. Ahmed N, Kanhere SS, Jha S (2010) Experimental evaluation of multi-hop routing protocols for wireless sensor networks. In: Proceedings of the 9th ACM/IEEE international conference on information processing in sensor networks (IPSN), pp 416–417
120. Ansari J, Ang T, Mähönen P (2011) WiSpot: fast and reliable detection of Wi-Fi networks using IEEE 802.15.4 radios. In: Proceedings of the 9th ACM international workshop on mobility management and wireless access (MobiWac), pp 35–44
121. Hossain MA, Mahmood A, Jäntti R (2009) Channel ranking algorithms for cognitive coexistence of IEEE 802.15.4. In: Proceedings of the 20th IEEE international symposium on personal, indoor and mobile radio communications (PIMRC), pp 112–116
122. Huang GJ, Xing G, Zhou G, Zhou R (2010) Beyond co-existence: exploiting WiFi white space for zigbee performance assurance. In: Proceedings of the 18th IEEE international conference on network protocols (ICNP), pp 305–314
123. Shuaib K, Boulmal M, Sallabi F, Lakas A (2006) Co-existence of zigbee and WLAN: a performance study. In: Proceedings of the IEEE international wireless telecommunications symposium (WTS)
124. Jeong Y, Kim J, Han SJ (2011) Interference mitigation in wireless sensor networks using dual heterogeneous radios. *Wirel Netw* 17(7):1699–1713
125. Arkoulis S, Spanos DE, Barbounakis S, Zafeiropoulos A, Mitrou N (2010) Cognitive radio-aided wireless sensor networks for emergency response. *Meas Sci Technol* 21(12):124002
126. Ansari J, Ang T, Mähönen P (2010) Spectrum agile medium access control protocol for wireless sensor networks. In: Proceedings of the 7th IEEE international conference on sensor mesh and ad hoc communications and networks (SECON), pp 1–9
127. Qin Y, He Z, Voigt T (2011) Towards accurate and agile link quality estimation in wireless sensor networks. In: Proceedings of the 10th IFIP annual mediterranean ad-hoc networking workshop (Med-Hoc-Net), pp 179–185
128. Boano CA, He Z, Li Y, Voigt T, Zúñiga M, Willig A (2009) Controllable radio interference for experimental and testing purposes in wireless sensor networks. In: Proceedings of the 4th international workshop on practical issues in building sensor network applications (SenseApp), pp 865–872
129. Duquenois S, Österlind F, Dunkels A (2011) Lossy links, low power, high throughput. In: Proceedings of the 9th ACM conference on embedded networked sensor systems (SenSys), pp 12–25
130. Fotouhi H, Zúñiga MA, Alves M, Koubâa A, Marrón PJ (2012) Smart-HOP: a reliable hand-off mechanism for mobile wireless sensor networks. In: Proceedings of the 9th European conference on wireless sensor networks (EWSN), pp 131–146
131. Österlind F, Mottola L, Voigt T, Tsiftes N, Dunkels A (2012) Strawman: resolving collisions in bursty low-power wireless networks. In: Proceedings of the 11th international conference on information processing in sensor networks (IPSN), pp 161–172

132. Puccinelli D, Gnawali O, Yoon S, Santini S, Colesanti U, Giordano S, Guibas L (2011) The impact of network topology on collection performance. In: Proceedings of the 8th European conference on wireless sensor networks, EWSN' 11. Springer, Berlin, pp 17–32
133. Akan OB, Karli OB, Ergul O (2009) Cognitive radio sensor networks. *IEEE Network: The Magazine of Global Internetworking* 23(4):34–40



<http://www.springer.com/978-3-319-00773-1>

Radio Link Quality Estimation in Low-Power Wireless Networks

Baccour, N.; Koubâa, A.; Noda, C.; Fotouhi, H.; Alves, M.; Youssef, H.; Zúñiga, M.A.; Boano, C.A.; Römer, K.; Puccinelli, D.; Voigt, T.; Mottola, L.

2013, XIII, 147 p. 42 illus., 29 illus. in color., Softcover
ISBN: 978-3-319-00773-1