Cross-technology Clear Channel Assessment for Low-Power Wide Area Networks

Charalampos Orfanidis*, Laura Marie Feeney†, Martin Jacobsson*, Per Gunningberg†
KTH Royal Institute of Technology, Sweden*
Uppsala University, Sweden†
corf@kth.se, lmfeeney@it.uu.se, martin.jacobsson@sth.kth.se, per.gunningberg@it.uu.se

***Abstract*—Due to their popularity, large coverage areas, and diversity of radio technologies, Low Power Wide Area Networks (LPWAN) will experience high levels of cross-technology interference. In this work, we investigate the impact of LoRa interference on IEEE 802.15.4 (Wi-SUN) networks. Using hardware experiments, we characterize the interaction between these two very different radio technologies. In particular, we observe that the IEEE 802.15.4 Clear Channel Assessment (CCA) mechanism does not reliably detect interfering LoRa transmissions. We therefore propose an enhanced CCA mechanism based on a Multi-Layer Perceptron classifier and show that it significantly reduces the number of unsuccessful transmissions, while remaining compatible with the IEEE 802.15.4 standard.**

***Index Terms*—LPWAN, performance evaluation, LoRa, IEEE 802.15.4, cross-technology interference, clear channel assessment**

I. Introduction

Low-Power Wide Area Networks (LPWAN) will enable new application scenarios for the Internet-of-Things (IoT), due to their long range communication, robustness, and low energy consumption. The popularity of these technologies – combined with their large coverage areas – means that there may be many independently deployed networks operating in the same location. For LPWANs operating in unlicensed spectrum, interference between co-located networks is inevitable and may degrade their performance.

Interactions between networks that use different radio technologies are particularly challenging. The networks may use the shared wireless channel in different ways, making it difficult for radios in one network to detect and avoid transmissions from radios in another network.

Two popular – and very diverse – LPWAN technologies are LoRaWAN [1] and IEEE 802.15.4-SUN (IEEE 802.15.4g) [2], which is the basis of the Wi-SUN standard [3] for Wireless Smart Utility Networks. We identify the case of IEEE 802.15.4 CSMA operating under LoRa interference as particularly interesting, due to expected high levels of LoRa interference and the relatively more vulnerable IEEE 802.15.4 radio.

This paper presents an experimental study of the ability of an IEEE 802.15.4 sender to detect the presence of LoRa transmissions. We observe and explain limitations in the IEEE 802.15.4 Clear Channel Assessment (CCA) mechanism and propose an enhanced CCA mechanism.

The main contributions of this paper are:

- A unique dataset of over 640,000 LoRa–IEEE 802.15.4 interactions, collected in an automated testbed under controlled interference conditions.
- Evidence that IEEE 802.15.4 CCA cannot reliably detect LoRa interference. In our dataset, the CCA decision is only 43% accurate – no better than a coin-flip.
- An enhanced CCA mechanism based on a Multi-Layer Perceptron (MLP) classifier. The enhanced CCA increases the accuracy of the CCA decision by 30 percentage points and reduces the number of unsuccessful IEEE 802.15.4 transmissions by over a factor of three.
- Inter-operability and timing compatibility with the existing IEEE 802.15.4 CCA and CSMA mechanisms. Moreover, the enhanced CCA does not rely on energy-hungry extended sampling of the channel.

II. Background

This section introduces LoRa and IEEE 802.15.4/Wi-SUN.

A. LoRa

LoRa (*Lo*ng *Ra*nge) radios use sub-GHz spectrum and Chirp Spread Spectrum (CSS) [4] to provide robust, low-power communication at low data rates. Depending on transmission parameters, LoRa provides bit rates of 0.3 to 22 kbps and can achieve ranges of 2 to $5\,\mathrm{km}$ in urban areas. LoRa radios are typically deployed as a LoRaWAN, with LoRa devices sending (almost exclusively) uplink traffic to LoRaWAN gateways that provide a coordinated base station infrastructure. The LoRaWAN MAC does not use channel sensing, but instead relies on regulatory limits on the transmit duty-cycle and its own robust modulation to achieve co-existence.

B. IEEE 802.15.4 and Wi-SUN

IEEE 802.15.4 defines an IEEE 802.15.4 PHY layer for Smart Utility Networks (aka IEEE 802.15.4g). The SUN PHY is defined for sub-GHz unlicensed spectrum and uses Frequency Shift Keying (FSK) as its basic modulation. It provides data rates of 50-200 kbps, where support for the former is mandatory and most widely used. Measurement studies [5] report reliable links of up to a few hundred meters under line-of-sight conditions in the $868\,\mathrm{MHz}$ band.

The IEEE 802.15.4-SUN PHY is the basis for the application protocol stack defined in the Wi-SUN (Wireless Smart Utility Network) standard [3]. Wi-SUN uses the IEEE

2155-6814/19/$31.00 ©2019 IEEE
DOI 10.1109/MASS.2019.00032

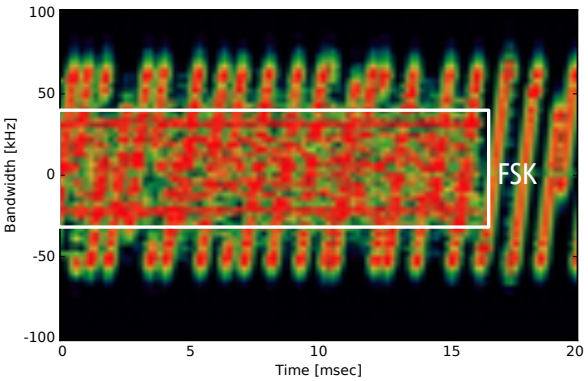

Fig. 1: LoRa (125 kHz BW, SF7) and IEEE 802.15.4 using the same center frequency. The IEEE 8021.5.4 2-FSK trasmission takes ∼16 ms and each LoRa chirp takes 1.024 ms.

802.15.4 upper link layer protocols and the IETF 6LowPAN and RPL protocols to create an IP-based mesh network.

C. CCA

Clear channel assessment (CCA) is a component of numerous protocols: In IEEE 802.15.4, it is essential for the CSMA MAC layer and is also used for the contention access period of the beacon-enabled PAN and for shared links in TSCH networks. CCA is generally based on the amount of RF energy detected on the channel, regardless of its origin. To perform CCA, an IEEE 802.15.4 radio measures the energy present in 98 kHz of filter bandwidth around its center frequency for eight symbol periods (0.160 µs). The channel is CLEAR if the average of these measurements is less than the CCA threshold value, which must be no higher than −90 dBm. (We refer to this as the *default CCA*.)

D. Visualizing LoRa and IEEE 802.15.4 interference

Figures 1 and 2 show spectrum occupied by overlapping IEEE 802.15.4 and LoRa transmissions, captured using a Software Defined Radio. The figures highlight the difference between the two radios.

Each LoRa symbol is a chirp, whose frequency increases at a constant rate as it passes through the LoRa channel bandwidth. The chirps appear as diagonal lines in the spectrogram. The SF determines the rate at which the signal crosses the channel bandwidth (BW). With SF7 (Figure 1), the chirp takes 1.024 ms to cross 125 kHz BW; with SF12 (Figure 2), the chirp takes 32.768 ms to cross the same bandwidth.

The IEEE 802.15.4 signal alternates between two frequencies, deviating 25 kHz from the channel center frequency, that represent the symbols '0' and '1'. In principle, IEEE 802.15.4 transmissions (circled in white) appear as two horizontal lines in the spectrogram (in practice, the trace is noisy). The symbol time is 20 μs (50 kbps bitrate), so the individual symbols are not visible at the depicted time-scale.

These figures provide intuition behind our results: When the IEEE 802.15.4 radio senses the channel, the LoRa chirp might or might not be within its 98 kHz filter bandwidth. The amount of RF energy that the IEEE 802.15.4 radio detects in a sample will depend on the amount of overlap between the IEEE 802.15.4 channel and the bandwidth being used by the LoRa chirp, and on how rapidly the chirp passes through this spectrum. Moreover, the variety of LoRa radio configurations and heterogeneity of the resulting LoRa radio footprint can make specific identification challenging.

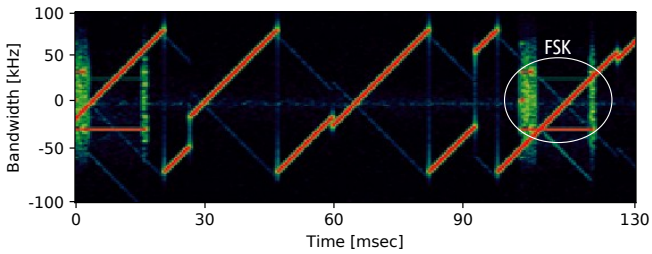

Fig. 2: LoRa (125 kHz BW, SF12) and IEEE 802.15.4 using the same center frequency. Each LoRa chirp takes 32.768 ms.

III. The LPWAN interference environment

In this section, we motivate our focus on LoRa interference and its impact on IEEE 802.15.4. LPWAN is still an emerging technology, with diverse radios, network architectures and business models. Any specific model of the future sub-GHz radio environment is therefore necessarily speculative. Nevertheless, it seems plausible that dense LoRaWAN environments will present substantial interference to co-located IEEE 802.15.4 networks.

A measurement campaign [6] in Aalborg, Denmark reported high levels of both SigFox and LoRa activity in the 868 MHz band. Interference levels of over −105 dBm were observed in up to 33% of measurements in some locations.

The LoRa radio chip manufacturer Semtech describes [7] a large-scale test deployment of a LoRaWAN for a smart metering application. Long-term measurements of data rate distribution, gateway receiver diversity, and packet loss were used to model a high-capacity LoRaWAN network. The report highlights a macro-diversity-based architecture, where frames are transmitted multiple times and are typically received by more than one gateway. In this scenario, maximum capacity is achieved with a channel occupancy of 63%, despite the large number of intra-network collisions this entails.

In this work, we focus on the impact of LoRa on IEEE 802.15.4, rather than the reverse. This is largely because IEEE 802.15.4 is more sensitive to LoRa interference. The LoRa PHY is known to be highly resilient [8], [9] and the IEEE 802.15.4 CSMA MAC layer defers to any ongoing transmissions, while LoRaWAN MAC does not sense the channel before sending.

Moreover, LoRaWAN's infrastructure-oriented architecture results in almost exclusively uplink traffic to LoRaWAN gateways, which typically have high-gain antennas, use mains power and are located on rooftops or even dedicated towers.

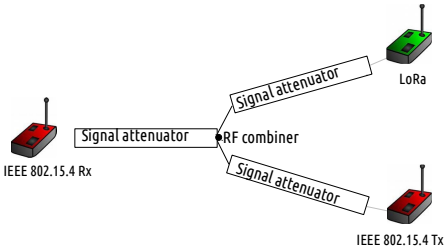

Fig. 3: Two TI CC1310 Launchpad nodes act as sender and receiver and a Netblocks LoRa node acts as interferer. The three nodes are connected via an RF combiner.

By contrast, IEEE 802.15.4-based networks are often used in a mesh configuration, where ordinary devices must be able to both receive and forward traffic and to participate in routing and application protocols.

sectionTestbed

This section describes how the testbed enables large-scale studies of CCA under controlled interference conditions.

A. Hardware

The testbed consists of two IEEE 802.15.4 nodes that act as sender and receiver respectively and one LoRa node that acts as an interferer. The IEEE 802.15.4 nodes are Texas Instruments Launchpad XLs, equipped with a CC1310 radio and running Contiki. The LoRa node is a Netblocks, equipped with a Semtech XRange SX1272 radio and running IBM/Semtech Lmic v1.6 and lorablink. All three nodes are connected to a control PC via USB serial cables.

The three radios are connected to each other using RF cables, fixed attenuators, and an RF combiner (Figure 3). Each radio is shielded from the others by placing it in a heavy stainless steel container, which has small openings for the RF and USB serial cables.

This setup allows the IEEE 802.15.4 sender to send frames to the IEEE 802.15.4 receiver with controlled LoRa interference present at both endpoints. The IEEE 802.15.4 sender first collects 50 Received Signal Strength Indicator (RSSI) values ($\sim$1 ms) and then transmits an 80 B frame, regardless of the channel state. The IEEE 802.15.4 receiver records whether the frame was successfully received.

The LoRa interferer transmits 59 B frames as continuously as possible. The measured gap between frames is 576 µs, giving a channel utilization of $\geq$99.4%. This ensures that there is (almost certainly) a LoRa transmission present on the channel when the IEEE 802.15.4 sender senses the channel and when the transmission arrives at the IEEE 802.15.4 receiver. Because the LoRa frames have a very low data rate and occupy the channel much longer than the IEEE 802.15.4 frames, the usual case is that each IEEE 802.15.4 sensing and transmission sequence overlaps with one LoRa frame.

The experiments and data collection are largely automated: The host PC communicates with the nodes via USB serial to configure their radio parameters and coordinate the timing of the packet transmissions; it also collects and stores RSSI and packet reception information from the nodes.

There are two major advantages of this approach: One, we have control over the signal and interference strengths within the network. Two, we can send a very large number of frames, which is otherwise infeasible given the severely limited transmit duty-cycle (0.1%) in 868 MHz spectrum.

B. Evaluating CCA effectiveness

The testbed allows us to evaluate both existing and experimental CCA algorithms using a consistent data set.

The default CCA measures the channel over eight 20 µs symbol periods (160 µs). In the testbed, this is done by requesting a sequence of RSSI measurements from the radio. Each RSSI measurement takes 21.5 µs, presumably due to a small overhead associated with the API.

For any given CCA algorithm, we can use these RSSI measurements to model the channel sensing and compute whether that algorithm *would have decided* that the channel was BUSY or CLEAR. Since the IEEE 802.15.4 frame is always transmitted, we can also evaluate the correctness of the CCA decision *ex post facto*.

In contention-based networks such as IEEE 802.15.4, a sender performs CCA to detect the presence of transmissions with which it might interfere or be interfered. But whether or not a frame is successfully received is determined by the signal-to-noise+interference ratio (SNIR) *at the receiver* – something that the sender can only partially infer based on its local channel interference conditions. The effectiveness of CCA is therefore inherently limited by the presence of hidden and exposed terminals.

In our experiments, we minimize the role of hidden and exposed terminals by enforcing that the interference conditions are the same at the sender and receiver. In other words, we configure the test network such that the attenuation between the LoRa interferer and the IEEE 802.15.4 sender is the same as the attenuation between the LoRa interferer and the IEEE 802.15.4 receiver. This allows us to isolate the effect of the CCA mechanism.

IV. Dataset

An IEEE 802.15.4 sender's ability to detect an interfering LoRa transmission is determined by three factors: The strength of the interfering LoRa signal at the sender; the overlap between the IEEE 802.15.4 and LoRa channels; and the LoRa transmission parameters. Whether or not an IEEE 802.15.4 frame is received also depends on the relative strength of the IEEE 802.15.4 signal and LoRa interference (SNIR) at the receiver. This section characterizes our dataset in the context of these factors.

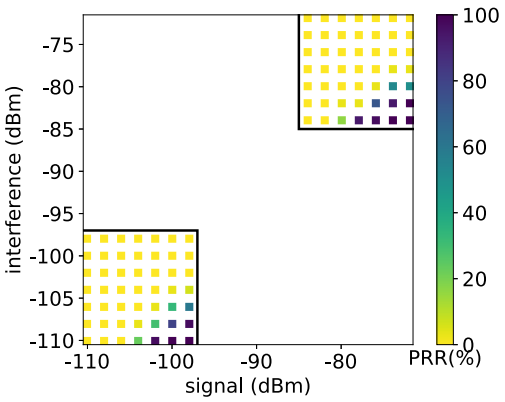

Fig. 4: PRR measurements, with LoRa interference (SF12 and CF offset $-25\,\mathrm{kHz}$) at the IEEE 802.15.4 sender and receiver. Each diagonal contains PRR measurements having the same SNIR. The measurement strategy focuses on regions where the PRR transitions from 0 to 100%. (If the PRR is 100% for a given signal and interference, it will be 100% for all stronger signals and all weaker interferences.)

A. Experiments

Our dataset uses a wide and realistic range of 6468 parameter configurations and measures LoRa-IEEE 802.15.4 interactions for 646,800 frames. The configurations reflect:

1) Signal and interference strength: By varying the radios' output power and the values of the fixed attenuators connecting them, we control the strength of the IEEE 8021.5.4 signal and the LoRa interference arriving at the IEEE 802.15.4 sender and receiver.

Figure 4 illustrates how our measurements focus on regions where the packet reception rate (PRR) transitions from 0% to 100% – challenging conditions where it is uncertain whether the frame will be received. One region has low received signal strengths and low levels of interference (-110 to $-98\,\mathrm{dBm}$ in $2\,\mathrm{dBm}$ increments), where the former value is close to the documented sensitivity of the IEEE 802.15.4 CC1310 radio. The other has high received signal strengths and high levels of interference (-84 to $-72\,\mathrm{dBm}$ in $2\,\mathrm{dBm}$ increments). The dataset includes SNIR values ranging from -12 to $12\,\mathrm{dB}$.

In other words, we exclude scenarios with very low interference values at the sender when the signal strength at the receiver is high (either guaranteed success or a hidden terminal scenario). We similarly exclude scenarios with high interference values at the sender when the signal strength at the receiver is low (either a guaranteed failure or an exposed terminal). This minimizes the impact of hidden and exposed terminals on our estimate of CCA effectiveness. (Recall that we also ensure that the sender and the receiver experience the same level of interference.)

2) Channel offset: The offset between the center frequencies (CF) of an IEEE 802.15.4 channel and a LoRa channel determines how much overlap there is in the spectrum they are using and thus how much of the energy from the LoRa interferer is present in the spectrum that the IEEE 802.15.4 sender and receiver are using.

IEEE 802.15.4 defines 34 fixed channels with $200\,\mathrm{kHz}$ separation in the $868\,\mathrm{MHz}$ band. The most common LoRaWAN bandplan for $868\,\mathrm{MHz}$ uses eight non-overlapping $125\,\mathrm{kHz}$ channels. In a realistic environment, IEEE 802.15.4 devices will therefore experience a variety of CF offsets with respect to LoRa interferers. The precise mix will depend on which channel (or channels in the case of channel hopping) the IEEE 802.15.4 sender is using and how the LoRa channels are defined.

Our dataset therefore uses a uniform mix of channel offsets: To do this efficiently, we fix the IEEE 802.15.4 sender to channel 26 ($CF = 868.325\,\mathrm{MHz}$) and vary the offset of the LoRa CF in $25\,\mathrm{kHz}$ increments from 0 to $\pm 125\,\mathrm{kHz}$, at which point the impact of interference is negligible.

3) LoRa parameters: We use LoRa bandwidth of $125\,\mathrm{kHz}$ (BW125) and spreading factors SF7-SF12. For our purposes, the importance of the SF parameter is that it determines how quickly the LoRa chirp passes through a given portion of spectrum. As with the CF offset, this affects how much energy will be detected by an IEEE 802.15.4 device sensing that spectrum. (The study of 250 and $500\,\mathrm{kHz}$ LoRa configurations and higher IEEE 802.15.4 data rates is future work.)

B. Packet reception rates

Figure 5 further explores the underlying PRR in our dataset. Overall, 33% of the packets are successfully received; this reflects our focus on regions where packet reception is uncertain.

Figures 5a and 5b shows PRR dependence on SNIR, for parameter configurations with lower (Figure 5a) and higher (Figure 5b) signal and interference strengths, corresponding to the two regions shown in Figure 4. The PRR depends on SNIR in the expected way; some artifacts of our dataset design are discussed below.

Figure 5c shows that the PRR strongly depends on the CF offset, which determines how much of the LoRa interference is present in the spectrum used by the IEEE 802.15.4 transmission. At $\pm 125\,\mathrm{kHz}$ CF offset ($\sim$18% of the configurations), the impact of interference on PRR is negligible. This explains why, in Figures 5a and 5b, the PRR remains close to 18%, even at very low SNIR.

At spreading factor SF12, the LoRa symbol time (i.e. one chirp) is longer than the IEEE 802.15.4 frame transmit time. Given fortuitous timing, an IEEE 802.15.4 frame may mostly avoid the LoRa interference (see also Figure 2). This explains why, in Figures 5a and 5b, the IEEE 802.15.4 PRR is higher with LoRa SF12 interference. It also explains why, in Figure 5c, the PRR is non-zero, even at zero CF offset.

C. RSSI measurements

Figure 6 shows a typical trace of 200 RSSI values ($\sim 4\,\mathrm{ms}$) captured by an IEEE 802.15.4 radio in the presence of fairly strong LoRa interference. The narrow interval highlighted in light grey shows one possible sample of eight RSSI values

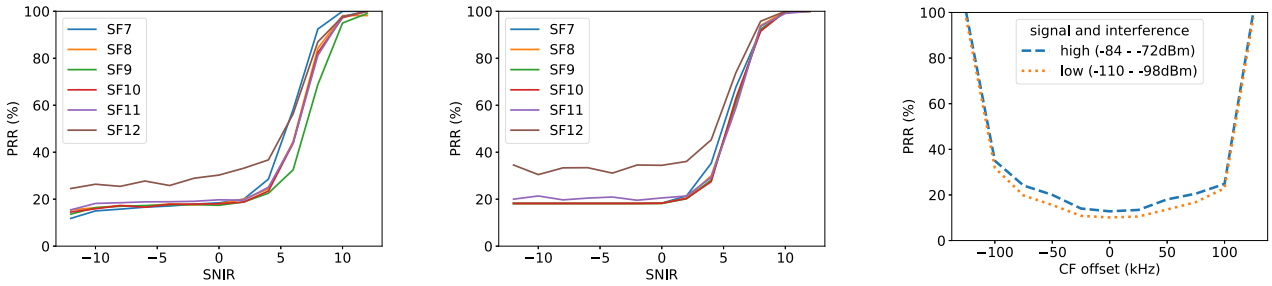

(a) PRR depends on SNIR. Weak signal and interference (–110 - −98 dBm).

(b) PRR depends on SNIR. Stronger signal and interference (–84 - −72 dBm).

(c) PRR depends on the offset between IEEE 802.15.4 and LoRa center frequencies.

Fig. 5: The packet reception rate (PRR) depends heavily on SNIR, but is largely independent of SF, with the exception of SF12 (discussed below). PRR also depends on the CF offset. At ±125 kHz (∼ 18% of the scenarios), impact on PRR is negligible.

(the sample size used in the default CCA). The wider interval highlighted in dark grey shows a sample of fifty RSSI values (the sample size used in our enhanced CCA). Even with this larger sample, it is clear that the average RSSI value depends heavily on when the sample is taken.

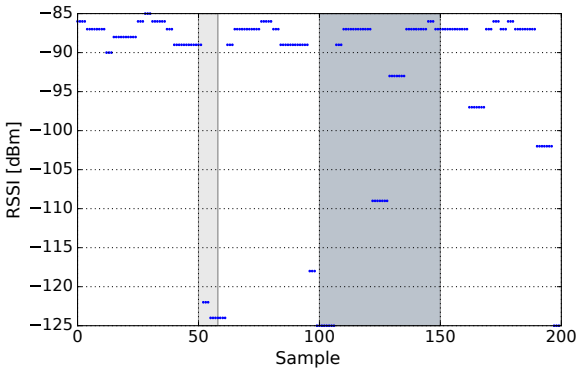

Fig. 6: RSSI values captured at the IEEE 802.15.4 sender. The LoRa interference is −80 dBm at CF offset −25 kHz and SF7.

Figure 7 examines this more systematically, for a range of LoRa spreading factors and CF offsets. A sample represents the input to a CCA decision, since it contains the channel measurements that an IEEE 802.15.4 sender would compare with its CCA threshold value (although here we use a sample size of 50 RSSI values (∼1 ms), rather than the eight used by the default CCA). Each data point in the figure shows the average of a set of 100 uncorrelated samples, along with its standard deviation.

The high variability in the measurements presented in Figure 7 is consistent with the complex interactions suggested by the spectrograms shown in Figures 1 and 2.

When there is maximum overlap between the two radios (offset =0 kHz), the chirp spends most of its time in the IEEE 802.15.4 sender's 98 kHz filter bandwidth, so the RSSI is high and the variation between samples is small.

But when there is an offset between the two radios, the variation between samples is high. In some cases, the LoRa chirp spends most of the sample duration passing through frequencies that overlap with the IEEE 802.15.4 filter bandwidth (resulting in high RSSI). In other cases, the LoRa chirp is mostly using non-overlapping frequencies (resulting in low RSSI). The variation is larger at high LoRa SF values, because the chirp passes through frequencies more slowly, increasing the either-or effect.

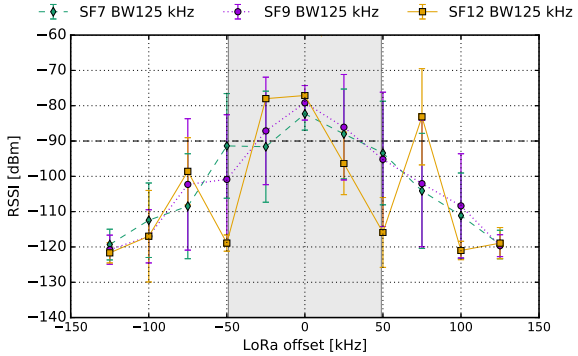

Fig. 7: Each sample represents input to a CCA decision (average and standard deviation of 100 uncorrelated samples).

The standard deviation (shown as error bars) extends well above and well below the 90 dBm CCA threshold (dashed horizontal line). In practice, these reflect cases where the CCA would decide that the channel is BUSY, respective CLEAR. The figure suggests that the decision will be somewhat erratic across a range of CF offset values, especially compared with the PRR behavior in Figure 5. Moreover, there also does not seem to be any obvious choice of a lower CCA threshold that would give better results. This variability suggests that the default IEEE 802.15.4 CCA will not reliably detect the presence of LoRa interference, even when averaging over 50, rather than the standard eight, RSSI values.

V. An Enhanced CCA

We hypothesize that the IEEE 802.15.4 CCA will not consistently detect the presence of LoRa transmissions and that

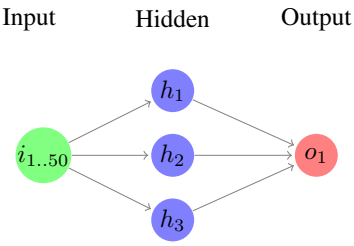

Fig. 8: MLP classifier

simple functions, such as average and threshold, are unlikely to be effective. In this section, we confirm this hypothesis and propose using a Multi-Layer Perceptron classifier (MLP), which is suitable for nonlinear classification, to perform the CCA decision.

A. A Multi-Layer Perceptron Classifier

We propose a MLP classifier that is designed to decide if the channel should be declared BUSY or CLEAR, taking as input a vector of RSSI values. This is achieved by classifying the input into two categories, BUSY or CLEAR. The dataset, which is described in Section IV and corresponds to more than 640000 frames, is divided. 70% of it is used for training the MLP and the remaining 30% is used for testing, with all parameter combinations represented equally in both sets.

Figure 8 presents the proposed MLP classifier. The input $i_{1..50}$ is a vector of 50 one-byte RSSI values. The MLP has one hidden layer with 3 neurons and one output layer, which is responsible for deciding if the channel should be declared BUSY or CLEAR. The number of hidden layers and neurons is chosen using a hyper-parameter tool, GridSearchCV [10], to select optimal values.

The sigmoid function, which performs well for classification purposes [11], is used as the activation function. The MLP classifier is trained by minimizing the Cross-Entropy loss function. Moreover, L2 regularization is used to avoid overfitting by adding a penalty to weights with large magnitude. The GridSearchCV tool [10] is used to select the optimal L2 penalty parameter.

The training phase takes place on a PC and takes a couple of minutes. The weights are numerical parameters which define how much they influence other neurons. After the training phase, the size of the weights is just 1.5 kB. For an input of 50 one-byte RSSI values, 3733 floating point operations are needed to run the MLP classifier. This suggests that it is potentially feasible to implement the MLP itself on relatively constrained hardware, such as the ARM-Cortex-M4.

B. Performance metric

As described in Section III, we evaluate the performance of the default and enhanced CCA by modeling the CCA decision and then checking whether it was correct, based on the eventual fate of the frame. We follow a statistical classification and confusion matrix [12] approach, defining the outcome for each packet according to Table I.

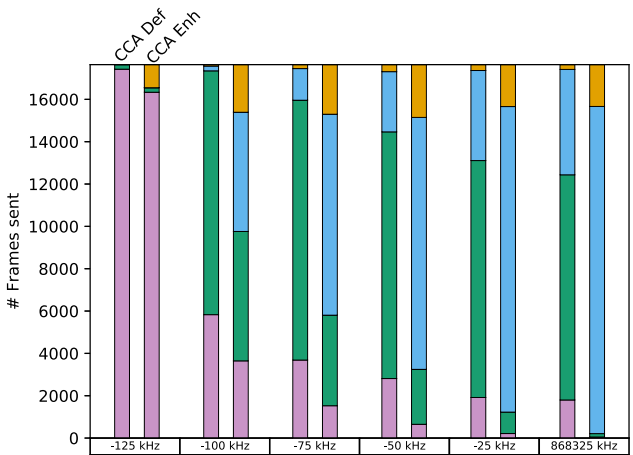

	-125 kHz		-100 kHz		-75 kHz		-50 kHz		-25 kHz		868325 kHz	
CCA	Def	Enh	Def	Enh	Def	Enh	Def	Enh	Def	Enh	Def	Enh
ST	17425	16336	5830	3647	3685	1528	2812	650	1917	214	1799	50
UT	215	214	11517	6112	12274	4277	11652	2595	11192	1012	10637	161
CB	0	1	226	5631	1495	9492	2845	11902	4254	14434	4978	15454
IB	0	1089	67	2250	186	2343	331	2493	277	1980	226	1975

Fig. 9: Two views of the raw data, collated by CF offset. (Positive offsets (25 kHz – 125 kHz) are not shown, as the results are largely symmetric.) The table entries and corresponding bars show the CCA decision and outcome of the transmission, using the definitions in Table I.

	CCA Decision	Ground Truth
Successful Transmission (ST)	CLEAR	Received
Unsuccessful Transmission (UT)	CLEAR	Not received
Correct Backoff (CB)	BUSY	Not received
Incorrect Backoff (IB)	BUSY	Received

TABLE I: Definitions used for evaluation

The CCA decision is **correct** in the **ST** and **CB** cases and **incorrect** in the **UT** and **IB** cases. The simplest metric for the effectiveness of a CCA algorithm is **accuracy**, which is the proportion of decisions that are correct.

C. Performance evaluation

Figure 9 presents a subset of the raw data. In the analysis that follows, we compare the behavior of the two CCA mechanisms. In each figure, we present three views of the full data set, each collated with respect to one of the three factors (from Section IV) that affect the IEEE sender's ability to detect the LoRa interference: the LoRa spreading factor; the CF offset; and the strength of the interfering LoRa signal at the receiver. The latter reflects the regions of low and high signal and interference strengths shown in Figure 4.

Summarizing the results: Figure 10 shows that the enhanced CCA significantly increases the accuracy of the CCA decision. Figure 11 shows that the enhanced CCA is also more conservative, sacrificing some viable transmission opportunities. And Figure 12 shows that this loss is more than compensated for by a reduction in unsuccessful transmissions.

Figure 10 compares the accuracy of the default CCA and enhanced CCA. The enhanced CCA makes a much higher

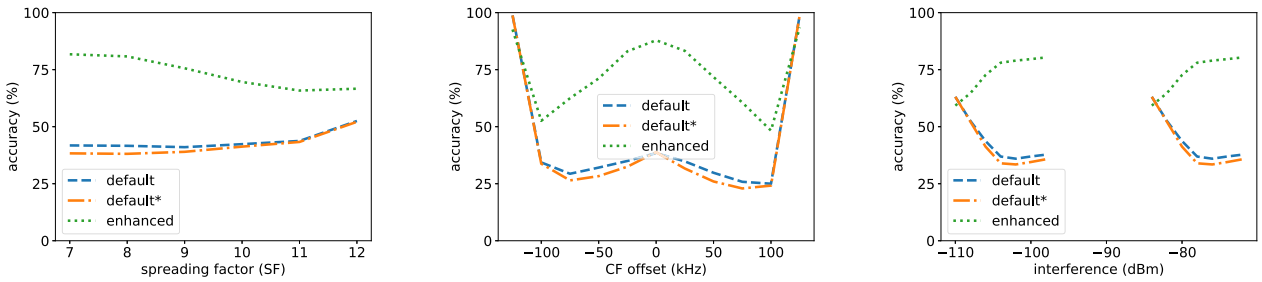

(a) By spreading factor (SF). (b) By channel CF offset. (c) By interfering signal strength at the sender.

Fig. 10: The accuracy of the CCA algorithm is defined as the proportion of decisions that are correct: Either the channel is assessed CLEAR and the transmission succeeds (ST), or the channel is assessed BUSY and the transmission does indeed fail (CB). Each subfigure shows one of the three factors affecting the sender's ability to detect the interfering signal.

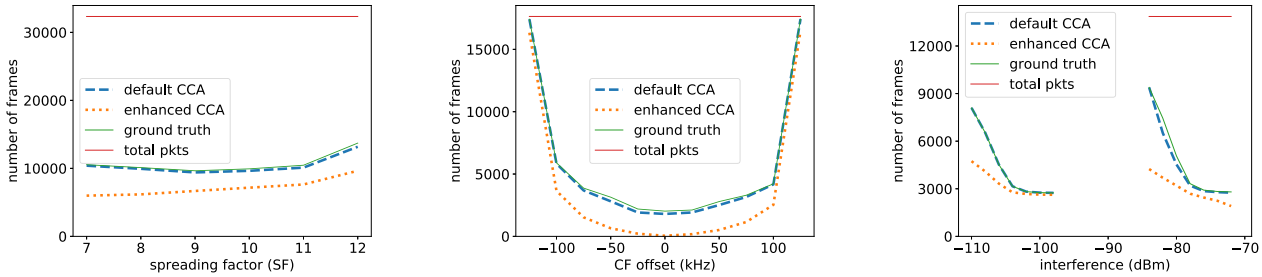

(a) By spreading factor (SF). (b) By channel CF offset. (c) By strength of the interference at the sender..

Fig. 11: The default CCA achieves a successful transmission (ST) in essentially all of the cases where the transmission is feasible (i.e. ground truth is that the frame was successfully received). The total number of packets is also shown.

proportion of correct decisions and the overall accuracy increases from 43% (essentially a coin-flip) to 73%. The figure also includes a straw-man CCA, *default**, which averages the 50 RSSI values provided to the MLP classifier, rather than the eight values used in the default CCA. Default* does not improve the accuracy, confirming the hypothesis suggested by Figure 7. This result further suggests that the MLP-classifier implements some non-trivial function.

However, the huge increase in accuracy contains some real tradeoffs. Figure 11 shows that the default CCA achieves almost all of the transmissions that are feasible (i.e. where ground truth is that the frame was successfully received). The enhanced CCA is more conservative and has fewer successful transmissions (ST) than the default CCA. Across the whole data set, the default CCA successfully transmits 32% of frames, while the enhanced CCA only transmits 22%.

Although the enhanced CCA has one-third fewer successful transmissions than the default CCA, it also has *3.5 times* fewer unsuccessful transmissions (UT). Figure 12 shows the ratio between unsuccessful transmissions (UT) and incorrect backkoffs (IB), for both the default and enhanced CCA.

The two types of errors have different costs: With an unsuccessful transmission (UT), the sender incurs energy and delay costs for both the failed transmission and the subsequent backoff and retransmission. The transmission may also have imposed additional energy costs on the receiver, as well as contributing to overall interference on the channel. With an incorrect backoff (IB), the sender only needs to perform an additional backoff, which has a much smaller delay and energy cost.

Enhanced CCA therefore reduces the number of expensive unsuccessful transmission (UT) errors, at the cost of increasing the number of less costly incorrect backoff (IB) errors. Its more conservative CCA decision may lead to better overall energy and delay performance in complex interference scenarios. Our testbed and data set provide a solid basis for future analysis of these costs, taking into account IEEE 802.15.4 CSMA backoff and retransmission behavior.

Finally, we show that the proposed classifier is not overfitting the dataset. This is illustrated when the offset is -125 kHz. In these cases, interference from LoRa is minimal. For the default CCA case, the CB and IB are 0, (Figure 9) but for the enhanced CCA there are instances where the backoff decisions are triggered. This can be explained from the fact that there are some other cases where LoRa interference is low but enough to cause some frames to drop. These particular RSSI values, captured from IEEE 802.15.4 before sending the frame should have a similar pattern with the cases when the frequency offset is -125 kHz. Nonetheless, these cases are rare.

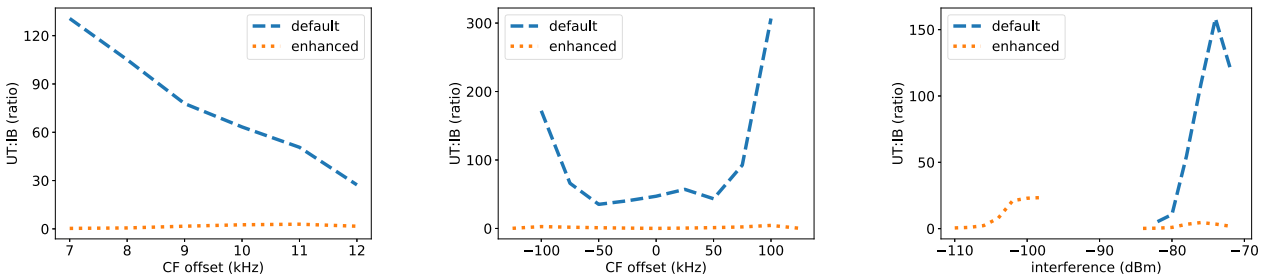

(a) By spreading factor (SF). The default CCA is much more aggressive with lower LoRa spreading factors, while the enhanced CCA has a more consistent and conservative behavior.

(b) By channel CF offset. The default CCA is most aggressive at large offsets, where the LoRa chirp may be quite strong, but is only occasionally present in IEEE 802.15.4 sender's filter bandwidth.

(c) By interfering signal strength at the sender. For low powers, the interference is below the $-90\,\mathrm{dBm}$ CCA threshold, so the default CCA never backs off. All errors are UT errors.

Fig. 12: Types of CCA error. The figure shows the ratio between the two kinds of incorrect decisions: unsuccessful transmissions (UT) and incorrect backoffs (IB).

D. Compatibility with IEEE 802.15.4

The enhanced CCA is inter-operable and timing-compatible with the existing IEEE 802.15.4 CCA and CSMA mechanism.

We have already shown that the enhanced CCA is, on average, more conservative than the default CCA in the presence of LoRa transmissions. In fact, the data set has no instances where the default CCA defers to a LoRa sender and the enhanced CCA does not. Compatibility with the IEEE 802.15.4 standard can therefore be trivially enforced by defining the enhanced CCA to return BUSY if either the default or enhanced CCA returns BUSY.

In addition, the enhanced CCA needs to react correctly to IEEE 802.15.4 transmissions and to operate in an environment that includes devices using the default CCA. This means that it should not be significantly more aggressive or more conservative than the default CCA with respect to other IEEE 802.15.4 transmissions. Table II compares the ability of default and enhanced CCA to detect IEEE 802.15.4 transmissions, using the accuracy criteria defined above. The two behave similarly: The enhanced CCA defers in only 2.6% of instances where the default CCA transmits. This suggests that the former will not be starved of transmission opportunities and that two CCA mechanisms can co-exist.

	Default CCA	Enhanced CCA
ST	4243	4131
UT	383	342
CB	248	289
IB	26	138

TABLE II: The two CCA mechanisms behave similarly detecting IEEE 802.15.4 interference.

Finally, the enhanced CCA should not affect the timing of the IEEE 802.15.4 CSMA algorithm. In the IEEE 802.15.4 standard, the random backoff is always a multiple of the unit backoff period, whose duration is the sum of the CCA sample time and the radio turnaround (rx-to-tx) time. For sub-GHz radios, this value is $1.16\,\mathrm{ms}$. The enhanced CCA duration of $1\,\mathrm{ms}$ is slightly less than the standard unit backoff period. As long as the random backoff interval is at least two unit backoff periods long, part of the backoff can be used for the enhanced CCA without changing its externally visible behavior.

VI. Related Work

Early performance studies of LoRaWAN, including [13] and [14], focus on capacity and scalability limits in the presence of intra-network interference. Low-level studies of LoRa interference include [15], [9], and [16]. Recently, [17] proposed distributed learning to maximize throughput in a LoRaWAN network subject to (potentially adversarial) interference from other LoRa networks.

Cross-technology interference in wireless networks has been widely studied, particularly co-existence among IEEE 802.11, IEEE 802.15.4, and Bluetooth networks operating in 2.4GHz spectrum. Surveys include [18] and [19].

The sub-GHz LPWAN environment is discussed in [20], but has generally received much less attention. A simulation study [21] contrasts the resilience of wideband CSS (LoRa) and ultra-narrow band (SigFox) communication, using an analytic model of the BER obtained by the two modulations. In [22], a controlled interference testbed is used to measure the impact of SigFox, ZWave, and HomeIO transmissions on LoRa packet reception. The authors focus on the timing of the interfering transmissions, demonstrating that the LoRa preamble and header are much more sensitive to interference than the payload.

By contrast, our work addresses the impact of LoRa on IEEE 802.15.4. The resilience of IEEE 802.15.4 to generic interference has been studied in [23]. The testbed study reported in [15] also includes some measurements of interactions between LoRa CSS and GFSK transmissions, where the latter are shown to be much less resilient to interference. Detailed PHY-layer measurements of interactions between LoRa and IEEE 802.15.4 under controlled conditions are presented in [8]. These results also demonstrate that IEEE 802.15.4 is substantially more sensitive to LoRa interference than vice

versa. In addition, the LoRa modulation parameters are shown to affect the IEEE 802.15.4 packet reception.

CCA has long been seen as important for the performance of CSMA-based MAC protocols. Cross-technology interference is a challenge because it can be difficult to detect transmissions from networks that use a lower power or different PHY. Examples of cross-technology CCA for IEEE 802.11 and IEEE 802.15.4 networks operating in 2.4GHz include [24] and [25].

In low-power networks, accurate CCA is also important for power saving methods that are based on devices periodically waking up to sense the channel to detect traffic announcements. In this case, the CCA should avoid expensive spurious wakeups due to misidentifying external transmissions. Examples include [26] and [27], primarily based on adapting the CCA threshold value.

There has been very little work applying machine learning techniques *per se* in CCA. In [28], devices characterize long term observations of the noise floor, the better to detect the presence of a transmission. Classification has also been used to identify specific types of interferers, e.g. [29].

VII. Conclusion and future work

In this work, we identify the impact of LoRa interference on IEEE 802.15.4 as a particularly interesting case, due to the popularity of these technologies, the potentially high levels of LoRa traffic, and the vulnerability of IEEE 802.15.4 transmissions. We present a large dataset of LoRa-IEEE 8021.5.4 interactions, collected in a controlled interference testbed. The data show that the IEEE 802.15.4 CCA mechanism is not able to reliably detect LoRa interference, due to differences in their modulation. We therefore propose using an MLP-classifier as the basis of an enhanced CCA mechanism for IEEE 802.15.4. Our evaluation shows an increase of 30 percentage points in the overall accuracy of the CCA decision. Further analysis reveals that the enhanced CCA is somewhat conservative: It forgoes some viable transmission opportunities, but avoids a much larger number of failed transmissions, which are extremely costly in terms of energy, delay, and interference. We also argue that enhanced CCA is inter-operable and timing compatible with existing IEEE 802.15.4 CCA and CSMA.

ML techniques for resource-constrained devices are an active research topic. Although the MLP used in the enhanced CCA is fairly small, there are still considerable deployment challenges. Our future work therefore focuses on making enhanced CCA more general and practical for use in the wild.

References

[1] LoRa Alliance, https://www.lora-alliance.org, Accessed 05.02.2018.

[2] IEEE, "IEEE standard for low-rate wireless networks," *IEEE Std 802.15.4-2015 (Revision of IEEE Std 802.15.4-2011)*, April 2016.

[3] Wi-SUN Alliance, www.wi-sun.org, Accessed 01.02.2019.

[4] Semtech, http://www.semtech.com/images/datasheet/an1200.22.pdf, Accessed 05.02.2018.

[5] J. Munoz, T. Chang, X. Vilajosana, and T. Watteyne, "Evaluation of IEEE802.15.4g for Environmental Observations," *Sensors*, Oct. 2018.

[6] M. Lauridsen, B. Vejlgaard, I. Z. Kovacs, H. Nguyen, and P. Mogensen, "Interference measurements in the European 868 MHz ISM band with focus on LoRa and SigFox," in *IEEE Wireless Communications and Networking Conference (WCNC)*, 2017.

[7] Semtech Corporation, "Real-world LoRaWAN network capacity for electrical metering applications," Tech. Rep., 2017.

[8] C. Orfanidis, L. M. Feeney, M. Jacobsson, and P. Gunningberg, "Investigating interference between LoRa and IEEE 802.15.4g networks," in *IEEE Int'l Conf on Wireless and Mobile Computing, Networking and Communications (WiMob)*, 2017.

[9] J. Haxhibeqiri, F. Van den Abeele, I. Moerman, and J. Hoebeke, "LoRa scalability: A simulation model based on interference measurements," *Sensors*, vol. 17, no. 6, 2017.

[10] "GridSearchCV," https://scikit-learn.org/, accessed: 2019-05-07.

[11] C. Nwankpa, W. Ijomah, A. Gachagan, and S. Marshall, "Activation functions: Comparison of trends in practice and research for deep learning," *CoRR*, vol. abs/1811.03378, 2018.

[12] S. V. Stehman, "Selecting and interpreting measures of thematic classification accuracy," *Remote Sensing of Environment*, vol. 62, no. 1, 1997.

[13] M. C. Bor, U. Roedig, T. Voigt, and J. M. Alonso, "Do LoRa low-power wide-area networks scale?" in *19th ACM Int'l Conf on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, 2016.

[14] K. Mikhaylov, J. Petaejaejaervi, and T. Haenninen, "Analysis of capacity and scalability of the LoRa low power wide area network technology," in *22th European Wireless Conference*, 2016.

[15] K. Mikhaylov, J. Petajajarvi, and J. Janhunen, "On LoRaWAN scalability: Empirical evaluation of susceptibility to inter-network interference," in *European Conf on Networks and Communications (EuCNC)*, 2017.

[16] A. Rahmadhani and F. Kuipers, "When LoRaWAN frames collide," in *12th Int'l Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (WiNTECH)*, 2018.

[17] A. Azari and C. Cavdar, "Self-organized low-power IoT networks: A distributed learning approach," 2018, arXiv:1807.09333.

[18] N. Baccour, D. Puccinelli, T. Voigt, A. Koubaa, C. Noda, H. Fotouhi, M. Alves, H. Youssef, M. A. Zuniga, C. A. Boano, and K. Römer, "External radio interference," in *Radio Link Quality Estimation in Low-Power Wireless Networks*, 2013.

[19] Y. Han, E. Ekici, H. Kremo, and O. Altintas, "Spectrum sharing methods for the coexistence of multiple RF systems: A survey," *Ad Hoc Networks*, vol. 53, pp. 53–78, 2016.

[20] E. D. Poorter, J. Hoebeke, M. Strobbe, I. Moerman, S. Latré, M. Weyn, B. Lannoo, and J. Famaey, "Sub-GHz LPWAN network coexistence, management and virtualization: An overview and open research challenges," *Wireless Personal Communications*, vol. 95, no. 1, 2017.

[21] B. Reynders, W. Meert, and S. Pollin, "Range and coexistence analysis of long range unlicensed communication," in *23rd Int'l Conf on Telecommunications (ICT)*, 2016.

[22] J. Haxhibeqiri, A. Shahid, M. Saelens, J. Bauwens, B. Jooris, E. De Poorter, and J. Hoebeke, "Sub-gigahertz inter-technology interference. How harmful is it for LoRa?" in *4th IEEE Int'l Smart Cities Conference (ISC2)*, 2018.

[23] C.-S. Sum, F. Kojima, and H. Harada, "Coexistence of homogeneous and heterogeneous systems for IEEE 802.15.4g smart utility networks," in *IEEE Int'l Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2011.

[24] L. Tytgat, O. Yaron, S. Pollin, I. Moerman, and P. Demeester, "Avoiding collisions between IEEE 802.11 and IEEE 802.15.4 through coexistence aware clear channel assessment," *EURASIP Journal on Wireless Communications and Networking*, vol. 2012, no. 1, 2012.

[25] Y. Tang, Z. Wang, T. Du, D. Makrakis, and H. T. Mouftah, "Study of clear channel assessment mechanism for ZigBee packet transmission under Wi-Fi interference," in *IEEE 10th Consumer Communications and Networking Conference (CCNC)*, 2013.

[26] X. Zheng, Z. Cao, J. Wang, Y. He, and Y. Liu, "Interference resilient duty cycling for sensor networks under co-existing environments," *IEEE Transactions on Communications*, no. 7, 2017.

[27] T. Sparber, C. A. Boano, S. Kanhere, and K. Römer, "Mitigating radio interference in large IoT networks through dynamic CCA adjustment," *Open Journal of Internet Of Things (OJIOT)*, vol. 3, no. 1, 2017.

[28] J. Polastre, "Sensor network media access design," University of California, Tech. Rep., 2003.

[29] F. Hermans, O. Rensfelt, T. Voigt, E. Ngai, L.-Å. Nordén, and P. Gunningberg, "Sonic: Classifying interference in 802.15.4 sensor networks," in *12th Int'l Conference on Information Processing in Sensor Networks (IPSN)*, 2013.