A Light Weight Approach to Online Detection and Classification of Interference in 802.15.4-Based Sensor Networks

S. Padmaja and J. Uma Maheswari

Abstract--- With a rapidly increasing number of devices sharing access to the 2.4 GHz ISM band, interference becomes a serious problem for 802.15.4-based, low-power sensor networks. Consequently, interference mitigation strategies are becoming commonplace. In this paper, we consider the step that precedes interference mitigation: interference detection. We have performed extensive measurements to characterize how different types of interferers affect individual 802.15.4 packets. From these measurements, we define a set of features which we use to train a neural network to classify the source of interference of a corrupted packet. Our approach is sufficiently light-weight for online use in a resource-constrained sensor network. It does not require additional hardware, nor does it use active spectrum sensing or probing packets. Instead, all information about interferers is gathered from inspecting corrupted packets that are received during the sensor network's regular operation. Even without considering a history of earlier packets, our approach reaches a mean classification accuracy of 79.8%, with per interfereraccuracies of 64.9% for WiFi, 82.6% for Bluetooth, 72.1% for microwave ovens, and 99.6% for packets that are corrupted due to insufficient signal strength.

Keywords--- Interference Mitigation, Network, Interference Detection

I. INTRODUCTION

In this paper, we describe an approach that enables resource-constrained sensor nodes to classify individual corrupted 802.15.4 packets according to the cause of corruption. Using data from our extensive measurements on how different interferers affect802.15.4 communication, we show that each interferer has characteristic patterns that emerge from observing the (1) Link Quality Indicator (LQI) of an interfered 802.15.4packet, (2) the signal strength during packet reception, and (3) information about what parts of the packet are corrupted. We define a set of features on these three observations that extract the essential information. Our features are sufficiently light-weight so that a sensor node can compute them for a given corrupted packet. A neural network maps features to an interference class, i.e., it allows to determine the type of interferer from the data collected about an individual corrupted packet. We implement a fixed-pointneural network on the TelosB platform to demonstrate the feasibility. A key strength of our approach is its resource efficiency: It does not require active spectrum sensing, additional hardware, or probing packets. Instead, it gathers information about interference only during the regular operation of the sensor network.

Assuming that the network uses either forward error correction or retransmissions, our approach does not incur any communication overhead. The main energy cost of our approach comes from turning on the Micro Controller Unit (MCU) during packet reception, whereas usually it would be woken up only when packet reception is completed.

S. Padmaja, Assistant Professor, Vels University, Pallavaram, Chennai-600117. E-mail:padmajasams@gmail.com
J. Uma Maheswari, MCA., M.Phil Scholar, Vels University, Pallavaram, Chennai-600117. E-mail:jk_uma@yahoo.com

II. RELATED WORK

Relevant research in the WiFi domain includes Airshark [16], a system that uses standard802.11 cards to sample the spectrum. The sampled data is analyzed using cyclostationary process methods to detect transmission patterns, which are then used to classify interferers. Another example is RFdump [12], which uses a software-defined radio to detect which devices are accessing the medium. RF dump aims to provide a tcpdump like tool for the wireless communication. Gollakota et al. describe how antenna diversity in 802.11n can be exploited to reconstruct interfered signals [6]. All these approaches have in common that they require advanced signal processing capabilities that are usually not available in sensor networks. Cisco has developed a spectrum analyzer for network analysis that is capable of classifying radio devices [2].

In sensor networks, interference detection can help mitigation, which in turn increases the network lifetime by reducing unsuccessful communication attempts. Chowdhuryet al. describe an approach to interference classification by actively scanning channels for characteristic spectrum usage [1]. In contrast to our work, this approach comes at a higher energy cost, because the radio needs to be turned on even when no sensor network communication is ongoing. Their work is also concerned with interference mitigation. Hauer et al. describe how detection of WiFi interference can be used for interference mitigation [8]. Similar to our approach, they also consider RSSI during packet reception for identifying interference. Their work is concerned with selectively retransmitting parts of a packet that are suspected to be interfered, without having certain knowledge about what caused the corruption. A large body of works considers the effect of interference on sensor networks in terms of high-level metrics such as packet reception rate.

III. TECHNICAL BACKGROUNDS

We briefly summarize the technical aspects of 802.15.4, 802.11b/g, Bluetooth, and microwave ovens that are relevant to our goal of interference classification. 802.15.4 The 802.15.4 standard defines a physical layer and a MAC layer for low power, low-rate wireless networks [10]. We consider 802.15.4 at 2.4 GHz in this paper, because most interference is faced in this popular ISM band. At 2.4 GHz, 16 channels of 2 MHz width are defined with an inter-channel spacing of 3 MHz. A maximum transmission power of 0 dBm is common. The standard implements direct sequence spread spectrum by mapping each four-bit symbol to be transmitted to a pseudo-random32-chip sequence. Offset quadrature phase-shift keying is used for modulation. The data rate is 250 kbps, the symbol period is 16 μs.

The format of an 802.15.4 PHY packet is shown in Fig. 1. Each packet begins with a preamble, which consists of four zero bytes, followed by a one-byte start frame delimiter(SFD) field with a fixed value. The frame length field contains the number of the packet’s payload bytes, which may be up to 127. The length includes the two-byte frame check sequence (FCS) field which trails the packet payload. The FCS field contains a checksum which is calculated over the length field and the payload bytes. A receiver synchronizes to incoming zero-bytes; after receiving four zero-bytes, it scans for an SFD. Only after correctly receiving the SFD, it reads the following payload field and then reads the specified number of payload bytes. If no SFD is received after four zero-bytes, the receiver synchronizes to incoming zero-bytes again. A receiver can detect transmission errors by comparing the FCS against the checksum calculated for the received packet.

IV. INTERFERENCE MEASUREMENTS

We conducted a series of experiments in which an 802.15.4-based sensor network was exposed to radio interference. In each experiment, we activated one interference source and collected data on the corrupted 802.15.4 packets. The purpose of these experiments is twofold. First, we collected corrupted packets to gain a better understanding of the effects that different interferers have on individual 802.15.4 packets. Second, we use subsets of the data for training and evaluating our classification approach. The experiments were carried out in an anechoic chamber, which is shielded from outside radio transmissions. Such a controlled environment gives us high confidence that a corrupted 802.15.4 packet recorded during an experiment was indeed corrupted by the source of interference that we have activated. In a less controlled environment, e.g., our university building, it is virtually impossible to prevent radio devices that are outside of our control from affecting the experiments. The anechoic chamber is also constructed to minimize multipath propagation. This is desirable, because multipath propagation is strongly dependent on the concrete physical layout of an environment, and we do not want to capture environment-specific effects in our measurements.

V. INTERFERENCE DETECTION AND CLASSIFICATION

We describe a classifier that assigns each incorrectly received packet to interference class. Each interference class represents a source of interference: WiFi, Bluetooth, or microwave. We further define an additional class to represent packets that have been received incorrectly in the absence of an interference source due to insufficient signal strength at the receiver. In this section, we first consider what data can be feasibly gathered in a sensor network for such a classification task; we then consider how this data can be condensed into numerical features. We discuss suitable classification algorithms, and finally consider implementation, energy cost and overhead.

VI. CLASSIFICATION ALGORITHM

We use a supervised learning approach to train a classifier to assign each corrupted packet to a class representing either WiFi, Bluetooth or microwave interference, or corruption due to insufficient signal strength. In supervised learning, a classifier is trained on a set of examples for which the correct class (i.e., the interference source in our case) is given. A corrupted packet is represented by the features described in the previous section. The learning phase, which is computationally more costly than classifying individual packets after training is completed, is carried out on a regular PC, whereas the actual classification is performed online in the sensor network. We consider two different classification algorithms: Support Vector Machines (SVMs, [3]) and feed-forward neural networks [18]. An SVM transforms the feature vectors to a high-dimensional space and, during the learning phase, constructs (potentially non-linear) hyper planes to separate the vectors from the learning set according to their classes. In the classification phase, the SVM determines the class of the input vector by considering on which side of the hyper planes the vector lies. Unfortunately, it turned out that online classification in a sensor network with SVMs is not feasible in our case due to the limited amount of RAM available on the sensor nodes. We nevertheless present results for the SVM classification in the evaluation as a reference case, since they are often considered the best “out-of-the box” classification algorithm [15]. Furthermore, SVMs find a global, unique solution, whereas neural networks may get stuck in local minima during the learning phase. Feed-forward neural networks are a class of classification algorithms inspired by biological neural networks. They are represented by directed acyclic graphs, where each node represents a computational unit and weighted edges describe input/output relationships between the nodes. Each node represents an input value (i.e., a component of a feature vector), an output value, or an activation function, which usually is a sigmoid function that takes as input the weighted outputs from incoming edges. During the learning phase, an optimal set of edge weights is found. In the classification phase, an input feature vector is propagated through the network and the output nodes indicate the classification result. After the costly learning phase, a neural network can be represented by a matrix and the classification cost is dominated by matrix multiplications. Thus, with careful implementation, the use of a feed-forward neural network is feasible even on a resource-constrained platform like TeosB.

VII. CONCLUSION AND FUTURE WORK

In this paper, we have investigated the feasibility of a light-weight interference detection and classification approach that only uses data that can be gathered during a sensor network's regular operation. The reason to only use such data is to keep energy consumption as low as possible. We described a set of features that allows us to define the characteristic patterns that we observe for different sources of interference. We also demonstrated that a fixed-point neural network reaches a mean classification accuracy of 79.8% for packets of 64 bytes and more. For this paper, we have gathered training and testing data in an anechoic chamber, which is a highly controlled environment. While we introduce a certain variance into our data by using different hardware models for the interferers, we would expect a less controlled radio environment to contribute further variance due to multipath propagation. Thus, we also plan to perform experiments in the offices of our university to test the robustness of our features. However, we face the practical issue of establishing ground truth in such an environment—we could not with certainty say that a packet is interfered by the interferer that we have activated, because an RF device outside our control (e.g., the university WiFi) may also affect our experiments. Interference detection and classification is an important tool for debugging network problems and mitigation strategies, but it is not an end in itself. Therefore, we plan to integrate our approach into an existing interference mitigation strategy. If our approach helps to make significantly better mitigation decisions in an uncontrolled environment, we may avoid the aforementioned issue of establishing ground truth for measuring the performance of our classification approach. The aim of this work was to assess the feasibility of interference classification with the limited information that can be gathered from corrupted packets. We have focused on classifying the source of interference for individual packets. An interesting track of future work we plan to follow is to incorporate information about previously received, corrupted packets into the classification process. We believe this may yield a significant increase in accuracy. However, incorporating previous information requires careful consideration, especially regarding situations in which multiple interferers are present, or in which the source of interference rapidly changes due to high mobility.

REFERENCE

[1] Chowdhury, K., Akyildiz, I.: Interferer classification, channel selection and transmission adaptation for wireless sensor networks. In: Proc. of ICC '09. pp. 1 –5 (Jun 2009)

[2] Cisco Inc.: Cisco Spectrum Expert Wi-Fi (2012), http://www.cisco.com/en/US/products/ps9393/index.html

[3] Cortes, C., Vapnik, V.: Support-vector networks. Machine Learning 20, 273–297 (1995)

[4] Crossbow Inc.: TelosB datasheet. http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/TelosB_Datasheet.pdf

[5] Gawthrop, P.E., Sanders, F.H., Nebbia, K.B., Sell, J.J.: Radio spectrum measurements ofindividual microwave ovens – volume 1 & 2, NTIA Report TR-94-303-1, NTIA Report TR-94-303-2

[6] Gollakota, S., Adib, F., Katabi, D., Seshan, S.: Clearing the RF smog: making 802.11n robustto cross-technology interference. In: Proc. of SIGCOMM' 11. pp. 170–181 (2011)

[7] Hauer, J.H., Handziski, V., Wolisz, A.: Experimental Study of the Impact of WLAN Interference

[8] Hauer, J.H., Willig, A., Wolisz, A.: Mitigating the Effects of RF Interference through RSSIBased Error Recovery. In: Wireless Sensor Networks, Lecture Notes in Computer Science, vol. 5970, pp. 224–239. Springer (2010)

[9] IEEE Computer Society: Local and metropolitan area networks, specific requirements, part 15.1: Wireless medium access control (MAC) and physical layer (PHY) specifications for wireless personal area networks (WPANs) (2005)

[10] IEEE Computer Society: 802.15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate

Wireless Personal Area Networks (WPANs) (2006)

[11] IEEE Computer Society: Local and metropolitan area networks, specific requirements, part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications (2007)

[12] Lakshminarayanan, K., Sapra, S., Seshan, S., Steenkiste, P.: RFDump: An Architecture for Monitoring the Wireless Ether. In: Procs. of CoNEXT '09 (Dec 2009)

[13] Liang, C.J.M., Priyantha, N.B., Liu, J., Terzis, A.: Surviving Wi-Fi Interference in Low Power ZigBee Networks. In: Procs of SenSys '10. pp. 309–322 (2010)

[14] Musaloiu-E., R., Terzis, A.: Minimising the effect of WiFi interference in 802.15.4 wireless sensor networks. Int. Journal of Sensor Networks 3, 43–54 (2008)

[15] Ng, A.: Support Vector Machines – CS229 Lecture Notes (2011), http://cs229.stanford.edu/notes/cs229-notes3.pdf

[16] Rayanchu, S., Patro, A., Banerjee, S.: Airshark: detecting non-wifirf devices using commodity wifi hardware. In: Procs. of IMC '11. pp. 137–154 (2011)

[17] Rensfelt, O., Hermans, F., Larzon, L.A., Gunning berg, P.: Sensei-UU: a relocatable sensor network tested. In: Procs. of WiNTECH '10. pp. 63–70 (September 2010)

[18] Rojas, R.: Theorie der neuronalenNetze. Springer (1993)